



National Cybersecurity Protection System Cloud Interface Reference Architecture

Volume Two: Reporting Pattern Catalog

May 14, 2021

Version 1.1

Cybersecurity and Infrastructure Security Agency
Cybersecurity Division
Capability Delivery Subdivision
NCPS Program Management Office

Revision/Change Record

Version	Date	Revision/Change Description	Section/Page Affected
Version 1.0	10/16/2020	Initial Release Version	All
Version 1.1	5/14/2021	Response to Comments and Feedback	Added Sections 2.4, 2.5, 2.6, 4.4; expanded Section 3; moved Section on Reporting Pattern-Level Characteristics to Volume One; moved all appendices to Volume One; added Acronyms table; minor graphic and text revisions throughout.

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 Cybersecurity and Infrastructure Security Agency (CISA) analysts can continue to provide situational awareness and support to the agencies. To support this goal, CISA is developing a cloud-based architecture to collect and analyze agency cloud security data. This reference architecture 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 reporting patterns to explain how that information can be sent to CISA.

The *NCPS Cloud Interface Reference Architecture* is being released as two individual volumes. The first volume provides an overview of changes to NCPS to accommodate the collection of relevant data from agencies' cloud environments and provides general reporting patterns for sending cloud telemetry to CISA. This second volume builds upon the concepts presented in *NCPS Cloud Interface Reference Architecture: Volume One* and provides an index of common cloud telemetry reporting patterns for how agencies can send cloud-specific data to the NCPS cloud-based architecture. It also discusses the various characteristics in these reporting patterns that agencies should consider as they select and implement patterns. Individual cloud service providers (CSPs) can leverage the reporting patterns in this volume to offer guidance on the solutions they provide that enable agencies to send cloud telemetry to CISA in fulfillment of NCPS requirements.

A cloud-based NCPS architecture is currently in development at CISA. This *NCPS Cloud Interface Reference Architecture* is being released to Federal Civilian Agencies in advance of a production system to accomplish the following:

- 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.....	8
1.1	Document Organization.....	8
1.2	Purpose.....	8
1.3	Document Guide.....	9
2	GENERIC REPORTING PATTERNS	10
2.1	GN-NNNN-SS: Agency CSP Cloud-Native Source Data Push to CLAW	14
2.2	SN-NNNN-LS: CLAW Pull from Agency CSP Cloud-Native Source.....	17
2.3	GV-NNAN-SS: Agency Aggregated Data Push to CLAW.....	20
2.4	SP-NNAN-LS: CLAW Pull of Agency Aggregated Service Data.....	23
2.5	SA-SDNN-SS: Agency Filtered Data Push to CLAW	26
2.6	NN-SDNI-LS: CLAW Pull of Agency Filtered Data.....	29
2.7	SF-NDNJ-SS: Agency CSP SECaaS Data Push to CLAW.....	33
2.8	SA-SANC-SS: CSP SECaaS Data, Agency Processing, and Push to CLAW.....	36
3	COMBINATION REPORTING PATTERNS.....	40
3.1	Differentiated Processing of Multi-Account Data (GV-NNAN-SS + SN-NNNN-LS).....	42
3.2	Per-Region Processing of Multi-Region Data	44
3.3	Push from Integrated Sharing Solution.....	46
3.4	Push to Local Regional CLAW in Multiple CSPs.....	48
4	CONCLUSION	49

LIST OF FIGURES

Figure 1: Reporting Pattern Structure	9
Figure 2: Reporting Pattern Identifier Format	10
Figure 3: Stage A Reporting Pattern Flow Chart.....	12
Figure 4: Stage B Reporting Pattern Flow Chart.....	13
Figure 5: Stage C Reporting Pattern Flow Chart.....	13
Figure 6: Roles and Telemetry Flow – GN-NNNN-SS.....	14
Figure 7: Visual Pattern Summary – GN-NNNN-SS.....	15
Figure 8: Roles and Telemetry Flow – SN-NNNN-LS.....	17
Figure 9: Visual Pattern Summary – SN-NNNN-LS.....	18
Figure 10: Roles and Telemetry Flow – GV-NNAN-SS.....	20
Figure 11: Visual Pattern Summary – GV-NNAN-SS.....	21
Figure 12: Roles and Telemetry Flow – SP-NNAN-LS.....	23
Figure 13: Visual Pattern Summary – SP-NNAN-LS.....	24
Figure 14: Roles and Telemetry Flow – SA-SDNN-SS.....	26
Figure 15: Visual Pattern Summary – SA-SDNN-SS.....	27
Figure 16: Roles and Telemetry Flow – NN-SDNI-LS.....	29
Figure 17: Visual Pattern Summary – NN-SDNI-LS.....	30
Figure 18: Roles and Telemetry Flow – SF-NDNJ-SS.....	33
Figure 19: Visual Pattern Summary – SF-NDNJ-SS.....	34
Figure 20: Roles and Telemetry Flow – SA-SANC-SS.....	36
Figure 21: Visual Pattern Summary – SA-SANC-SS.....	37
Figure 22: Visual Pattern Summary – Differentiated Processing of Multi-Account Data	42
Figure 23: Visual Pattern Summary – Per-Region Processing of Multi-Region Data	44
Figure 24: Visual Pattern Summary – Push from Integrated Sharing Solution.....	46
Figure 25: Visual Pattern Summary – Push to Local Regional CLAW in Multiple CSPs.....	48

LIST OF TABLES

Table 1: Reporting Pattern Identification	10
Table 2: Index of Reporting Patterns.....	11
Table 3: Pattern Summary Table – GN-NNNN-SS.....	15
Table 4: Pattern Summary Table – SN-NNNN-LS.....	18
Table 5: Pattern Summary Table – GV-NNAN-SS.....	21
Table 6: Pattern Summary Table – SP-NNAN-LS.....	24
Table 7: Pattern Summary Table – SA-SDNN-SS.....	27
Table 8: Pattern Summary Table – NN-SDNI-LS.....	30
Table 9: Pattern Summary Table – SF-NDNJ-SS.....	34
Table 10: Pattern Summary Table – SA-SANC-SS.....	37
Table 11: Pattern Summary Table – Differentiated Processing of Multi-Account Data.....	42
Table 12: Pattern Summary Table – Per-Region Processing of Multi-Region Data	44
Table 13: Pattern Summary Table – Push from Integrated Sharing Solution.....	46

ACRONYMS AND DEFINITIONS

Abbreviation	Definition
AI	Artificial Intelligence
API	Application Programming Interface
AWS	Amazon Web Services
C2	Command & Control
CASB	Cloud Access Security Broker
CEF	Common Event Format
CIO	Chief Information Officer
CISA	Cybersecurity and Infrastructure Security Agency
CISO	Chief Information Security Officer
CLAW	Cloud Log Aggregation Warehouse
CSP	Cloud Service Provider
DDOS	Distributed Denial of Service
DGA	Domain Generation Algorithms
DHS	Department of Homeland Security
DNS	Domain Name System
FedCIRC	Federal Computer Incident Response Capability
FIPS	Federal Information Processing Standards
FISMA	Federal Information Security Management Act
FOUO	For Official Use Only
GCP	Google Cloud Platform
GMT	Greenwich Mean Time
HIDS	Host-Based Intrusion Detection System
HR	Human Resources
HTTP	Hypertext Transfer Protocol
IaaS	Infrastructure as a Service
IAM	Identity and Access Management
ICRF	International Celestial Reference Frame
IDS	Intrusion Detection System
IERS	International Earth Rotation and Reference Systems
IOC	Indicators of Compromise
IP	Internet Protocol
IPFIX	Internet Protocol Flow Information Export
IPS	Intrusion Prevention System
IT	Information Technology
JSON	JavaScript Object Notation
LEEF	Log Event Extended Format
ML	Machine Learning
MOU	Memorandum of Understanding
MTIPS	Managed Trusted Internet Protocol Services
NAT	Network Address Translation
NCIRA	NCPS Cloud Interface Reference Architecture

Abbreviation	Definition
NCPS	National Cybersecurity Protection System
NIST	National Institute of Standards and Technology
NOC	Network Operations Center
OMB	Office of Management and Budget
PaaS	Platform as a Service
PIA	Privacy Impact Assessments
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
SaaS	Software as a Service
SECaaS	Security as a Service
SIEM	Security Information and Event Management
SOAR	Security Orchestration, Automation and Response
SOC	Security Operations Center
TIC	Trusted Internet Connections
TLS	Transport Layer Security
URL	Uniform Resource Locator
UT	Universal Time
UTC	Coordinated Universal Time
VM	Virtual Machine
VPC	AWS Virtual Private Cloud
VPN	Virtual Private Network

1 INTRODUCTION

Federal civilian departments and agencies¹ must participate in the National Cybersecurity Protection System (NCPS).² CISA analysts use this data for 24/7 situational awareness, analysis, and incident response. Traditionally, NCPS sensors located at Trusted Internet Connections (TIC) and Managed Trusted Internet Protocol Service (MTIPS) gateways capture security information as traffic passes between the agency and the Internet. As agencies move their information technology (IT) infrastructure to the cloud, some 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 CISA analysts can continue to provide situational awareness and support to the agencies. To support this goal, CISA is deploying a cloud-based architecture, the Cloud Log Aggregation Warehouse (CLAW), to collect and analyze agency cloud security data. CISA has released the *NCPS Cloud Interface Reference Architecture (NCIRA)* as a two-volume document set to explain how agencies can provide cloud-generated security information to the CLAW. Volume One introduces fundamental concepts about cloud data aggregation and reporting patterns (including attributes and options for how agencies can send cloud telemetry to NCPS). Volume Two (i.e., this document) provides a catalog of common reporting patterns based on the reporting pattern framework developed in Volume One. It also discusses the various characteristics in these reporting patterns that agencies should consider as select and implement patterns.

NCIRA Volume Two (this document) is a continuation of NCIRA Volume One and builds on the concepts presented in that document. In order to understand and implement the reporting patterns presented in this document, agencies must be familiar with the concepts introduced in NCIRA Volume One.

1.1 Document Organization

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

- Section 1 provides a document overview and a guide on how to use this volume in conjunction with Volume One.
- Section 2 contains a catalog of simple reporting patterns that can be mapped to common agency cloud use cases and includes reader aids (an index of the patterns and flow charts).
- Section 3 is a catalog of more complex reporting patterns that combine one or more of the individual patterns developed in Section 2.
- Section 4 discusses conclusions and future work.

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

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. See <https://cyber.dhs.gov/agencies/> for additional information.

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

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 CISA. This reference architecture is divided into two volumes.

1. Volume One of the *NCPS Cloud Interface Reference Architecture* provides general guidance for agencies on participating in NCPS in the cloud. The information provided includes an introduction to general reporting patterns. The discussion in Volume One is vendor-agnostic and not specific to any particular CSP.
2. Volume Two of the *NCPS Cloud Interface Reference Architecture* contains a catalog of reporting patterns for how agencies can participate in NCPS in the cloud under different cloud service models. The catalog includes individual reporting patterns (typical of an agency using a single CSP) as well as complex reporting patterns (illustrating how an agency can use several cloud service models and providers and send cloud security data to NCPS in the cloud). Volume Two also discusses the various characteristics in these reporting patterns that agencies should consider as they select and implement patterns.

1.3 Document Guide

Section 2 of this document discusses the NCPS cloud telemetry characteristics that are common across the reporting patterns in Sections 3 and 4. Sections 3 and 4 of this document are intended to serve as a catalog of common reporting patterns; it is not necessary for the document to be read in its entirety. Agencies should identify which reporting patterns apply to their cloud use cases and use these patterns to implement NCPS in the Cloud. As shown in Figure 1, each reporting pattern in this document is presented in the following format.

1. **Identifier and Title:** The naming scheme and title description provides a high-level summary of the reporting pattern and the attribute options leveraged.
2. **Overview:** An overview that provides the reader with a brief summary of the reporting pattern, including information used to understand its context and application.
3. **Roles and Telemetry Flow Figure:** This figure depicts which entity is responsible for each of the three telemetry reporting stages and what functions each entity performs.
4. **Stage Summary:** The stage summary provides an explanation about how each of the three telemetry reporting stages are performed.
5. **Visual Pattern Summary Figure:** This figure provides a visual summary of options selected for each of the attributes in each stage of the pattern.
6. **Pattern Summary Table:** This table articulates the option selected for each of the attributes in each stage of the pattern.
7. **Pattern Characteristics:** Additional details describe the pattern-level characteristics for full agency cloud telemetry sharing.

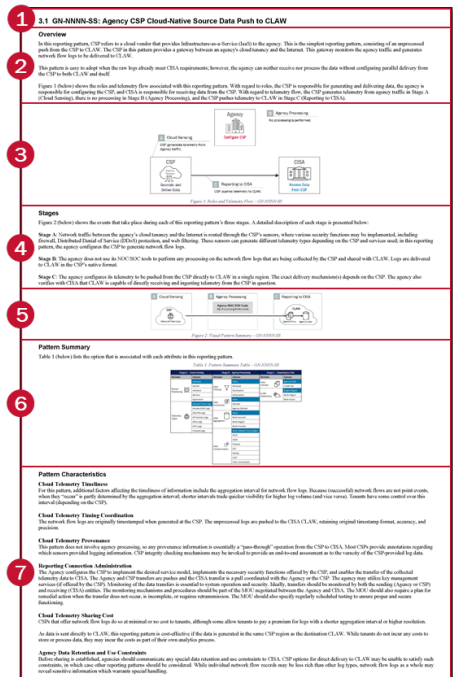


Figure 1: Reporting Pattern Structure

2 GENERIC REPORTING PATTERNS

The selection of a combination of Stage A, B, and C options constitutes a reporting pattern. Stage A addresses Cloud Sensing, Stage B addresses Agency Processing, and Stage C covers Reporting to CISA. Because there are three stages in any reporting pattern, and each stage has multiple attributes and options (as listed in Table 1), there are many possible reporting patterns. Therefore, it is desirable to have a scheme for easily identifying generic reporting patterns. To address this need, each generic reporting pattern will be identified by an eight-character identifier in the format shown in Figure 2.



Figure 2: Reporting Pattern Identifier Format

The acceptable values for each character position, and its corresponding option, are listed in Table 1.

Table 1: Reporting Pattern Identification

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Option and Letter Code	Attribute	Option and Letter Code	Attribute	Option and Letter Code
Sensor Positioning (S ₁)	Gateway (G)	Data Filtering (P ₁)	None (N)	Data Transfer (R ₁)	Agency Push (S)
	Subnet (N)		Removal (R)		CLAW Pull (L)
Telemetry Types (S ₂)	Interface (I)	Data Enrichment (P ₂)	Sanitization (S)	CLAW Distribution (R ₂)	Single Region (S)
	Service (S)		Obfuscation (O)		Multi-Region (R)
	Application (A)	None (N)	Multi-Cloud (C)		
	Network Flow Logs (N)	Derived (D)			
	Access/Auth Logs (A)	Agency-Defined (A)			
Telemetry Types (S ₂)	IDS/IPS Logs (I)	Data Aggregation (P ₃)	None (N)		
	API Activity Logs (P)		Multi-Account (A)		
	DNS Logs (D)		Multi-Region (R)		
	VPN Logs (V)	Multi-Provider (P)			
	Firewall Logs (F)	Data Transformation (P ₄)	None (N)		
			IPFIX (I)		
			JSON (J)		
Parquet (P)					
CEF (F)					
Syslog (S)					
LEEF (L)					
Firewall Logs (F)	CISA Coordinated (C)				

For example, the identifier “GN-NNNN-SS” indicates that there is a gateway sensor sending network flow logs (Stage A), with no additional processing (Stage B), and an agency push to its regional CLAW delivery point (Stage C).

Each generic reporting pattern also includes a short name (e.g., “Agency CSP Cloud-Native Source Data Push to CLAW”) to better accommodate conversation. These short names will be provided as part of the reporting pattern title.

Based on the variety of options available, there are many reporting pattern permutations, and it is not practical to discuss every possible permutation within this document. Instead, this document will focus on a small set of common reporting patterns. Patterns not shown here may still be viable alternatives and should be discussed with CISA on a case-by-case basis for adoption and possible inclusion in future versions of this volume.

While these patterns focus on the process of delivering telemetry to CLAW, agencies may use the same telemetry in their own analytics process. When considering how well each pattern would satisfy an agency’s need to share telemetry with CISA, agencies should also consider how well the pattern overlaps with their existing analytics process, as leveraging this overlap may result in significant cost savings.

Reporting Pattern Index

As previously stated, NCIRA Volume Two provides a catalog of common reporting patterns based on the reporting pattern framework developed in Volume One. As a result, there is no expectation that agencies review each individual reporting pattern at length, as an agency may only need to reference one of the reporting patterns covered in this catalog. Table 2 (below) provides a high-level index of all reporting patterns covered in NCIRA Volume Two. If an agency’s desired reporting pattern does not appear in this table, it does not mean that the agency cannot report its data to NCPS with a pattern that is not in this document. Instead, it means that the agency must develop a new reporting pattern, using one or more similar patterns, and discuss the resulting reporting pattern with CISA.

Table 2: Index of Reporting Patterns

Pattern Identifier	Cloud Sensing		Agency Processing				Reporting to CISA	
	Sensor Positioning	Telemetry Types	Data Filtering	Data Enrichment	Data Aggregation	Data Transformation	Data Transfer	CLAW Distribution
2.1 GN-NNNN-SS	Gateway	Network Flow Logs	None	None	None	None	Agency Push	Single Region
2.2 SN-NNNN-LS	Service	Network Flow Logs	None	None	None	None	CLAW Pull	Single Region
2.3 GV-NNAN-SS	Gateway	VPN Logs	None	None	Multi-Account	None	Agency Push	Single Region
2.4 SP-NNAN-LS	Service	API Activity Logs	None	None	Multi-Account	None	CLAW Pull	Single Region
2.5 SA-SDNN-SS	Service	Access/Auth Logs	Sanitization	Derived	None	None	Agency Push	Single Region
2.6 NN-SDNI-LS	Subnet	Network Flow Logs	Sanitization	Derived	None	IPFIX	CLAW Pull	Single Region
2.7 SF-NDNJ-SS	Service	Firewall Logs	None	Derived	None	JSON	Agency Push	Single Region
2.8 SA-SANC-SS	Service	Access/Auth Logs	Sanitization	Agency-Defined	None	CISA Coordinated	Agency Push	Single Region

Reporting Pattern Reader’s Guide

The reporting pattern flow charts in this section can help agencies determine which of the presented reporting pattern(s) apply to their cloud deployment(s) (if any). Each flow chart is specific to one of the stages in a reporting pattern. The first flow chart shows options for Stage A, the second shows options for Stage B, and the final shows options for Stage C. In each flow chart, the attributes are used to develop decision points and the options are shown as choices to move to the next selection point. Once a chart has been traversed, the resulting destination will either reference one or more specific reporting patterns in this catalog or will note that the agency will have to develop a unique reporting pattern (“Tailored Reporting Pattern”). Because this document only discusses eight reporting patterns and there are many other possible combinations of attributes and options, Tailored Reporting Pattern should always be considered as a potential option at the end of each flow chart.

The Stage A reporting pattern flow chart is shown in Figure 3.

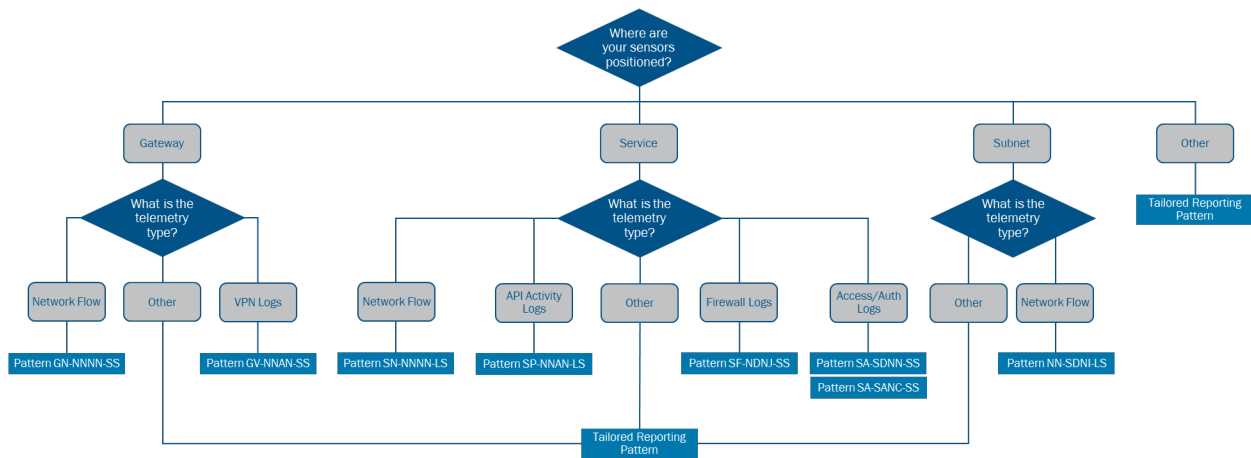


Figure 3: Stage A Reporting Pattern Flow Chart

The Stage B reporting pattern flow chart is shown in Figure 4.

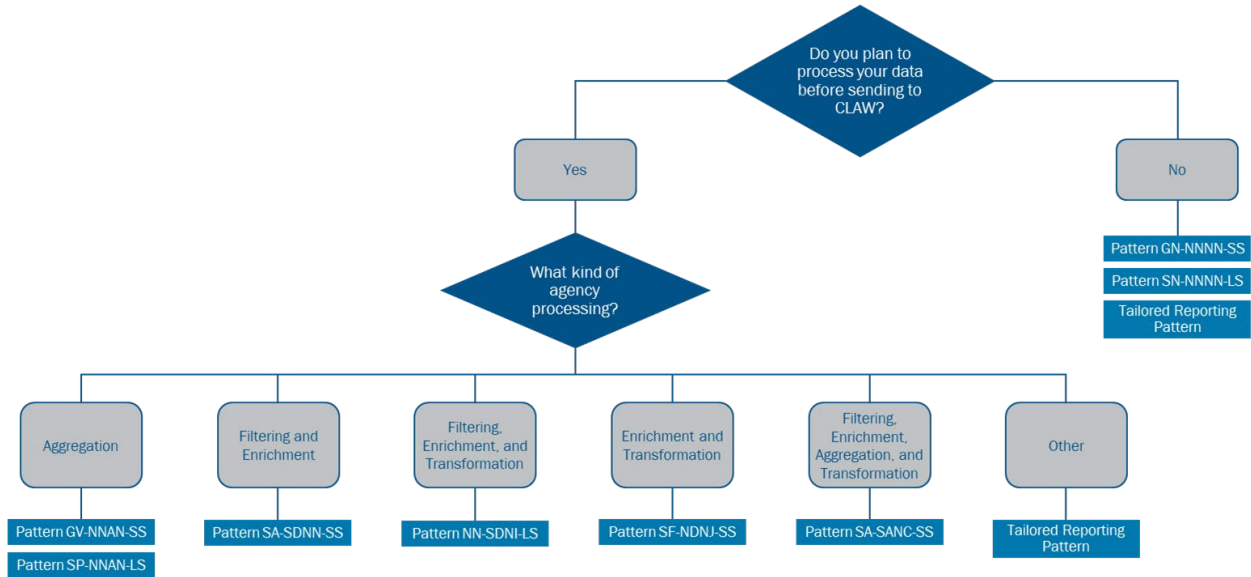


Figure 4: Stage B Reporting Pattern Flow Chart

The Stage C reporting pattern flow chart is shown in Figure 5.

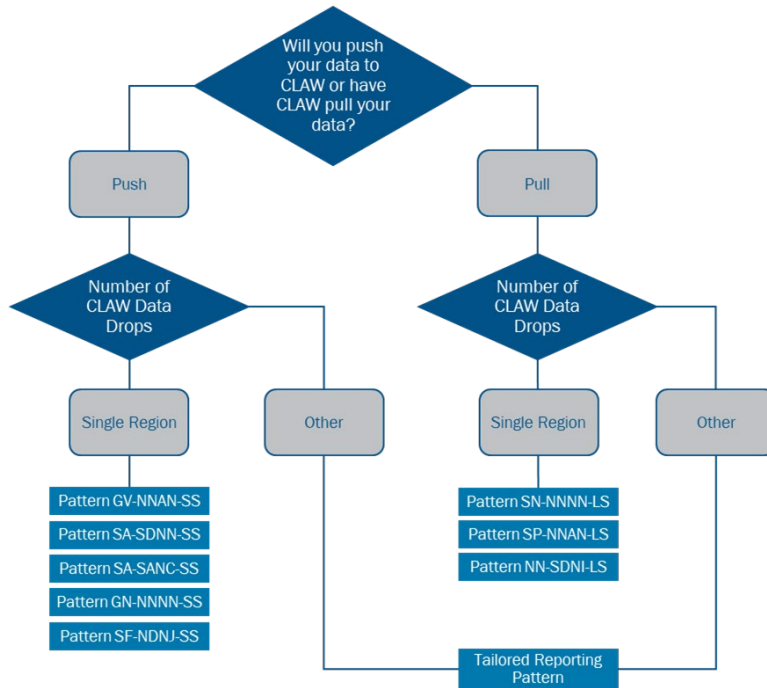


Figure 5: Stage C Reporting Pattern Flow Chart

2.1 GN-NNNN-SS: Agency CSP Cloud-Native Source Data Push to CLAW

Overview

In this reporting pattern, CSP refers to a cloud vendor that provides Infrastructure as a Service (IaaS) to the agency. This is the simplest reporting pattern, consisting of an unprocessed push from the CSP to CLAW. The CSP in this pattern provides a gateway between an agency's cloud tenancy and the Internet. This gateway monitors the agency traffic and generates network flow logs to be delivered to CLAW.³

This pattern is easy to adopt when the raw logs already meet CISA requirements; however, the agency can neither receive nor process the data without configuring parallel delivery from the CSP to both CLAW and itself.

Figure 6 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP, and CISA is responsible for receiving data from the CSP. With regard to telemetry flow, the CSP generates telemetry from agency traffic in Stage A (Cloud Sensing), there is no processing in Stage B (Agency Processing), and the CSP pushes telemetry to CLAW in Stage C (Reporting to CISA).

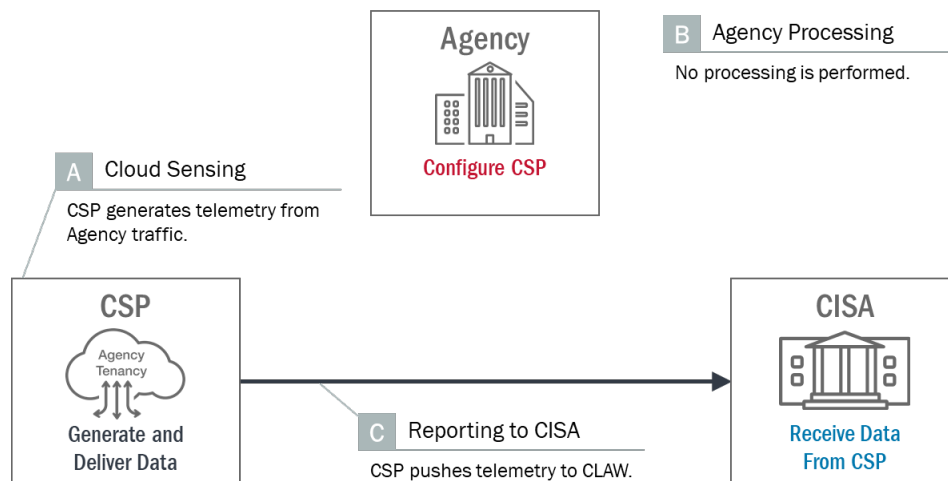


Figure 6: Roles and Telemetry Flow – GN-NNNN-SS

Stages

Figure 7 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: Network traffic between the agency's cloud tenancy and the Internet is routed through the CSP's sensors, where various security functions may be implemented, including firewall, Distributed Denial of Service (DDoS) protection, and web filtering. These sensors can generate different telemetry types depending on the CSP and services used; in this reporting pattern, the agency configures the CSP to generate network flow logs.

³ The CSP may also provide gateways between the agency's cloud tenancy and external networks that are not the Internet, such as the agency's on-premise network. These gateways provide similar monitoring capabilities.

Stage B: The agency does not use its Network Operations Center/Security Operations Center (NOC/SOC) tools to perform any processing on the network flow logs that are being collected by the CSP and shared with CLAW. Logs are delivered to CLAW in the CSP’s native format.

Stage C: The agency configures its telemetry to be pushed from the CSP directly to the regional CLAW. The exact delivery mechanism(s) depends on the CSP. The agency also verifies with CISA that CLAW is capable of directly receiving and ingesting telemetry from the CSP in question.⁴












Figure 7: Visual Pattern Summary – GN-NNNN-SS

Pattern Summary

Table 3 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 3: Pattern Summary Table – GN-NNNN-SS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning 	Gateway	Data Filtering 	None	Data Transfer 	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization	CLAW Distribution 	Single Region
	Service		Obfuscation		Multi-Region
	Application		None		Multi-Cloud
Telemetry Types 	Network Flow Logs	Data Enrichment 	Derived		
	Access/Auth Logs		Agency-Defined		
	IDS/IPS Logs	Data Aggregation 	None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
	Firewall Logs		None (Native Forms Align)		
		Data Transformation 	IPFIX		
			JSON		
			Parquet		
		CEF			
		Syslog			
		LEEF			
		CISA Coordinated			

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, additional factors affecting the timeliness of information include the aggregation interval for network flow logs. Because (successful) network flows are not point events, when they “occur” is partly determined by the aggregation interval; shorter intervals trade quicker visibility for higher log volume (and vice versa). Tenants have some control over this interval (depending on the CSP).

⁴ If this is not the case (e.g. for CSPs that are not commonly used by agencies), the agency may request CISA to add support for this CSP. Alternatively, the agency themselves can transform the data into IPFIX or another format that it negotiates with CISA; refer to Patterns 6-8 for examples.

Cloud Telemetry Timing Coordination

The network flow logs are originally timestamped when generated at the CSP. The unprocessed logs are pushed to the CISA CLAW, retaining original timestamp format, accuracy, and precision.

Cloud Telemetry Provenance

This pattern does not involve agency processing, so any provenance information is essentially a “pass-through” operation from the CSP to CISA. Most CSPs provide annotations regarding which sensors provided logging information. CSP integrity checking mechanisms may be invoked to provide an end-to-end assessment as to the veracity of the CSP-provided log data.

Reporting Connection Administration

The agency configures the CSP to push the telemetry using connection security parameters and credentials coordinated with CISA. The agency may utilize key management services (if offered by the CSP). Monitoring of the data transfer health can only be performed by CISA and within CSP-native functions. Ideally, transfers should be monitored by both the sending (CSP) and receiving (CISA) entities. Any retransmission of telemetry, for whatever reason, would take place within CSP-native functionality. Agencies should ensure sufficient telemetry cache durations and retransmission capabilities can satisfy CISA preferences.

Cloud Telemetry Sharing Cost

CSPs that offer network flow logs do so at minimal or no cost to tenants, although some allow tenants to pay a premium for logs with a shorter aggregation interval or higher resolution.

As data is sent directly to CLAW, this reporting pattern is cost-effective if the data is generated in the same CSP region as the destination. While tenants do not incur any costs to store or process data, they may incur the costs as part of their own analytics process.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. CSP options for direct delivery to CLAW may be unable to satisfy such constraints, in which case other reporting patterns should be considered. While individual network flow records may be less rich than other log types, network flow logs as a whole may reveal sensitive information which warrants special handling.

2.2 SN-NNNN-LS: CLAW Pull from Agency CSP Cloud-Native Source

Overview

In this reporting pattern, CSP refers to an agency's Platform as a Service (PaaS) cloud vendor. CLAW sends requests to pull network flow log data from the CSP, which in turn responds with the desired telemetry. The CSP in this pattern may provide various services, such as load balancing, network/application firewalls, Domain Name System (DNS), identity/authentication, key management, web hosting, etc. These services each generate telemetry, which is made available through an API (either the service's own, or, if the telemetry is exported to a CSP storage service, then that service's API).

This pattern is applicable when the raw network flow logs already meet CISA requirements. In addition, the agency must ensure that the mechanism used by CLAW to pull agency data from the PaaS vendor cannot be abused by other third parties.

Figure 8 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and storing data, the agency is responsible for configuring the CSP, and CISA is responsible for retrieving data from the CSP. With regard to telemetry flow, the CSP generates telemetry from agency traffic in Stage A (Cloud Sensing), there is no processing in Stage B (Agency Processing), and CLAW pulls telemetry from the CSP in Stage C (Reporting to CISA).

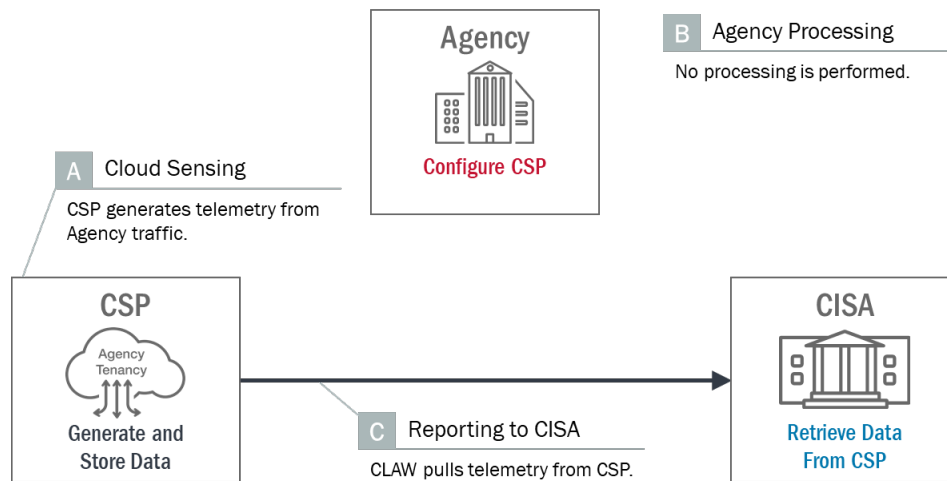


Figure 8: Roles and Telemetry Flow – SN-NNNN-LS

Stages

Figure 9 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: Network traffic between the agency's cloud tenancy and the Internet is routed through the CSP services; the agency configures one or more of these services to generate network flow logs.

Stage B: The agency does not use its NOC/SOC tools to perform any processing on the network flow logs that are being collected by the CSP and shared with CLAW. Logs are delivered to CLAW in the CSP's native format.

Stage C: The agency configures its telemetry to be supplied from the CSP service directly to the regional CLAW. This involves configuring permissions on the CSP such that CLAW⁵ has the proper pull credentials to make the necessary requests and pull the network flow logs from the CSP. The exact delivery mechanism(s) depends on the CSP.




Figure 9: Visual Pattern Summary – SN-NNNN-LS

Pattern Summary

Table 4 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 4: Pattern Summary Table – SN-NNNN-LS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization		Single Region
	Service		Obfuscation	CLAW Distribution	Multi-Region
	Application		None		Multi-Cloud
Telemetry Types	Network Flow Logs	Data Enrichment	Derived		
	Access/Auth Logs	Data Aggregation	Agency-Defined		
	IDS/IPS Logs		None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs	Multi-Provider			
	Firewall Logs	Data Transformation	None (Native Forms Align)		
			IPFIX		
	JSON				
	Parquet				
		CEF			
		Syslog			
		LEEF			
		CISA Coordinated			

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the aggregation interval for network flow logs and the polling frequency of CLAW. Because (successful) network flows are not point events, when they “occur” is partly determined by the aggregation interval; shorter intervals trade quicker visibility for higher log volume (and vice versa). Tenants have some control over this interval (depending on the CSP). The frequency with which CLAW checks for and pulls new data adds delay and is limited by mechanisms such as API request throttling.

Cloud Telemetry Timing Coordination

The network flow logs are originally timestamped when generated at the CSP. The unprocessed logs are pulled by CLAW, retaining original timestamp format, accuracy, and precision.

⁵ An authenticated identity principal corresponding to the CLAW instance in the selected region.

Cloud Telemetry Provenance

This pattern does not involve agency processing, so any provenance information is essentially a “pass-through” operation from the CSP to CISA. Most CSPs provide annotations regarding which sensors provided logging information. CSP integrity checking mechanisms may be invoked to provide an end-to-end assessment of the veracity of the CSP-provided log data.

Reporting Connection Administration

The agency prepares for the CISA pull transfer by provisioning credentials and establishing reachability to the telemetry source from CLAW. The agency may utilize key management services (if offered by the CSP). The frequency of CLAW telemetry pull transactions, buffer sizes for individual telemetry items, notifications for successful receipt, and other parameters should be negotiated prior to telemetry sharing initiation. Monitoring of the data transfers should be monitored by both the sourcing (CSP) and receiving (CISA) entities. The monitoring mechanisms and procedures are limited to CSP-native and CLAW functionality, as the agency does not perform any supplemental processing.

Cloud Telemetry Sharing Cost

CSPs which offer network flow logs do so at minimal or no cost to tenants, although some allow tenants to pay a premium for logs with a shorter aggregation interval or higher resolution.

As data is sent directly to CLAW, this reporting pattern is most cost-effective if the data is generated in the same CSP region as the destination. While tenants do not incur any costs to store or process data, they may incur the costs as part of their own analytics process.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. Options for pulling data directly from the CSP may be unable to satisfy such constraints, in which case other reporting patterns should be considered. While individual network flow records may be less rich than other log types, network flow logs as a whole may reveal sensitive information which warrants special handling.

2.3 GV-NNAN-SS: Agency Aggregated Data Push to CLAW

Overview

In this reporting pattern, CSP refers to an agency's IaaS cloud vendor. The agency has multiple accounts with the CSP and aggregates data from each before sending to CLAW. The CSP in this pattern provides VPN gateways agency's cloud tenancies with their on-premises network and/or remote endpoints. These gateways generate VPN logs to be delivered to CLAW. The log format is the same for each agency tenancy.

This reporting pattern follows naturally from an agency's own log aggregation and analytics and it simplifies connection administration by consolidating similar telemetry streams. However, if operations in each account are concentrated in separate regions, then egress data transfer costs increase.

Figure 10 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP and combining data, and CISA is responsible for receiving data from the agency. With regard to telemetry flow, the CSP generates telemetry from agency activity on multiple accounts in Stage A (Cloud Sensing), the agency aggregates telemetry received from the CSP in Stage B (Agency Processing), and the agency pushes telemetry to CLAW in Stage C (Reporting to CISA).

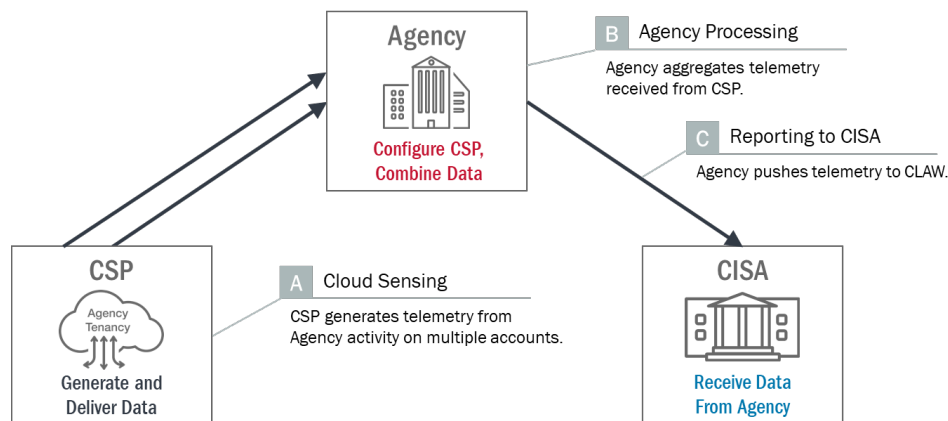


Figure 10: Roles and Telemetry Flow – GV-NNAN-SS

Stages

Figure 11 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: VPN traffic between the agency's cloud tenancies and its on-premises network and/or remote endpoints is routed through the CSP's VPN gateways, where various security functions may be implemented, including authentication and security posture checks. These gateways can generate VPN logs which contain the history of remote accesses; in this reporting pattern, the agency configures the CSP to generate such logs for each account.

Stage B: The agency uses its NOC/SOC tools to aggregate the VPN logs from multiple CSP accounts into a single stream. The format of the logs (i.e., the CSP's native format) is preserved and the agency does not perform any filtering or enrichment.

Stage C: The agency pushes the aggregated telemetry to its regional CLAW. The exact delivery mechanism(s) depends on the CSP.












Figure 11: Visual Pattern Summary – GV-NNAN-SS

Pattern Summary

Table 5 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 5: Pattern Summary Table – GV-NNAN-SS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning 	Gateway	Data Filtering 	None	Data Transfer 	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization	CLAW Distribution 	Single Region
	Service		Obfuscation		Multi-Region
Telemetry Types 	Application	Data Enrichment 	None		Multi-Cloud
	Network Flow Logs		Derived		
	Access/Auth Logs		Agency-Defined		
	IDS/IPS Logs	Data Aggregation 	None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
Firewall Logs	Data Transformation 	None (Native Forms Align)			
		IPFIX			
		JSON			
		Parquet			
		CEF			
		Syslog			
		LEEF			
		CISA Coordinated			

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the agency’s own policy for delivery to CLAW. Agencies can delay delivering individual records/objects (e.g., as part of a batching policy) and may do so if they do not exceed the maximum processing delay parameters.

Agency processing itself should not significantly affect timeliness; aggregating a common log type from multiple sources is expected to be a low complexity operation, facilitating rapid execution.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the VPN logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are aggregated. However, the original log entries’ timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves agency processing on multiple log streams across multiple tenancies. As a common log type and format is assumed across each data source, the agency is able to aggregate the sources either by interleaving or combining them in some other fashion (e.g., data from one tenancy might precede that from another). In this case, provenance claims are likely to be made by the agency. In particular, although multiple streams may arrive at the agency labeled and integrity-protected, the process of interleaving would create a new stream that itself requires provenance metadata. In short, the agency would be responsible for asserting that it provided the aggregation of the multiple streams, and constituent streams may retain sufficient provenance information to be checked end-to-end by CISA.

Reporting Connection Administration

The agency aggregator, as the sender of the telemetry to CISA, utilizes the CISA-provided credentials and security parameters to establish the data sharing connection. Transfer system health should be monitored by both the sending agency and CISA entities. The monitoring mechanisms and procedures may be able to leverage agency aggregation system native functionality, reducing cost and complexity. In addition, a plan for remedial action when the transfer does not occur, is incomplete, or requires retransmission may also leverage agency aggregation system native functionality.

Cloud Telemetry Sharing Cost

Note that multi-account aggregation is not necessarily exclusive with multi-region aggregation; if operations (and thus telemetry) for the two accounts are in separate CSP regions, egress data transfer costs will apply when the data is aggregated.

Implementations of this reporting pattern may involve persistent compute resources to perform the agency push to CLAW, in which case agencies incur the cost to maintain these resources.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. Any special constraints on the aggregate data stream are driven by those of the constituent telemetry sources, so agencies should be careful when combining data of different sensitivities. Furthermore, aggregation may produce sensitive information not deducible from either telemetry source alone.

2.4 SP-NNAN-LS: CLAW Pull of Agency Aggregated Service Data

Overview

In this reporting pattern, CSP refers to an agency's IaaS cloud vendor. The CSP in this pattern provides various services through one or more APIs and logs all requests made to these interfaces. There are separate telemetry streams of API activity logs for each account. The agency gathers data from the multiple accounts and aggregates it. CLAW sends requests to pull data from the agency, which in turn responds with the desired telemetry.

This reporting pattern follows naturally from an agency's own log aggregation and analytics, and it simplifies connection administration by consolidating similar telemetry streams. However, if operations in each account are concentrated in separate regions, then egress data transfer costs increase. In addition, the agency must ensure that the mechanism used by CLAW to pull agency data from the IaaS vendor cannot be abused by other third parties.

Figure 12 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP and combining data, and CISA is responsible for retrieving data from the agency. With regard to telemetry flow, the CSP generates telemetry from agency activity on multiple accounts in Stage A (Cloud Sensing), the agency aggregates telemetry received from the CSP in Stage B (Agency Processing), and CLAW pulls telemetry from the agency in Stage C (Reporting to CISA).

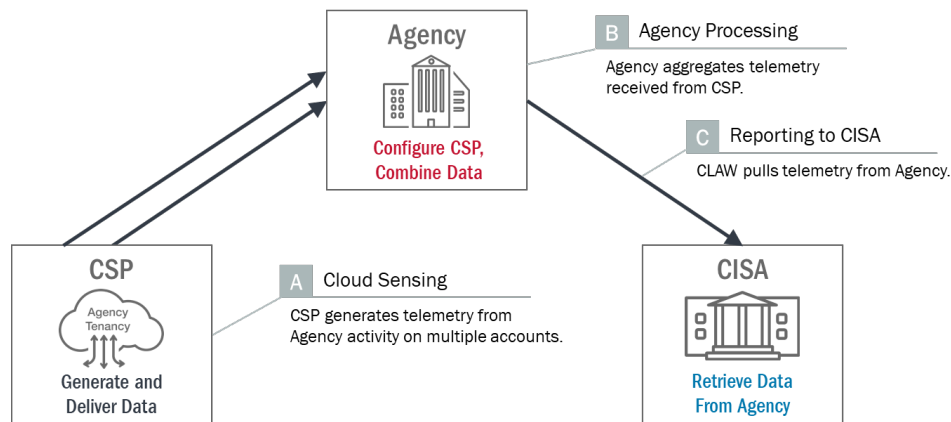


Figure 12: Roles and Telemetry Flow – SP-NNAN-LS

Stages

Figure 13 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: API calls to CSP services made by the agency, as well as API calls made by other entities on agency resources, are logged by the CSP. The agency configures the appropriate CSP auditing service so it receives a copy of these logs.

Stage B: The agency uses its NOC/SOC tools to aggregate the API activity logs from multiple CSP accounts. Once merged into a single stream, the aggregated logs may then be stored for later retrieval by CLAW. The aggregation may take place on agency premise equipment or may occur on agency-

configured CSP infrastructure. The format of the logs (i.e., the CSP’s native format) is preserved and the agency does not perform any filtering or enrichment.

Stage C: The agency supplies the aggregated telemetry to be pulled by the CLAW in the same region. This involves configuring pull credentials such that CLAW⁶ is authorized to make the necessary requests. The exact delivery mechanism(s) depends on the CSP.




Figure 13: Visual Pattern Summary – SP-NNAN-LS

Pattern Summary

Table 6 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 6: Pattern Summary Table – SP-NNAN-LS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization		Single Region
	Service		Obfuscation		Multi-Region
Telemetry Types	Application	Data Enrichment	None	CLAW Distribution	Multi-Cloud
	Network Flow Logs		Derived		
	Access/Auth Logs	Data Aggregation	Agency-Defined		
	IDS/IPS Logs		None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
	Firewall Logs		None (Native Forms Align)		
			Data Transformation	IPFIX	
				JSON	
	Parquet				
	CEF				
			Syslog		
			LEEF		
			CISA Coordinated		

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the polling frequency of CLAW. The frequency with which CLAW checks for and pulls new data adds delay and is limited by mechanisms such as API request throttling.

Agency processing itself should not significantly affect timeliness; aggregating a common log type from multiple sources is expected to be a low complexity operation, facilitating rapid execution.

⁶ An authenticated identity principal corresponding to the CLAW instance in the selected region.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the API activity logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are aggregated. However, the original log entries' timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves agency processing on multiple log streams across multiple tenancies. As a common log type and format is assumed across each data source, the agency is able to aggregate the sources either by interleaving or combining them in some other fashion (e.g., data from one tenancy might precede that from another). In this case, provenance claims are likely to be made by the agency. In particular, although multiple streams may arrive at the agency labeled and integrity-protected, the process of interleaving would create a new stream that itself requires provenance metadata. In short, the agency would be responsible for asserting that it provided the aggregation of the multiple streams, and constituent streams may retain sufficient provenance information to be checked end-to-end by CISA because no agency filtration is performed in this pattern. In addition, as the agency is not necessarily guaranteed to receive incoming telemetry requests at a predetermined rate, the agency may need to decide which data to retain or discard. Should it be necessary for the agency to discard data, this fact should be noted and integrity protected as part of the provenance claims.

Reporting Connection Administration

The agency prepares for the CISA pull transfer by provisioning credentials and establishing reachability to the aggregation source from CLAW. The agency may utilize key management services (if offered by the CSP). The frequency of CLAW telemetry pull transactions, buffer sizes for individual telemetry items, notifications for successful receipt, and other parameters should be negotiated prior to telemetry sharing initiation. Monitoring of the data transfers should be monitored by both the sourcing (agency aggregation service) and receiving (CISA) entities. The monitoring mechanisms and procedures may leverage agency aggregation system native functionality.

Cloud Telemetry Sharing Cost

Note that multi-account aggregation is not necessarily exclusive with multi-region aggregation; if operations (and thus telemetry) for the two accounts are in separate CSP regions, egress data transfer costs will apply when the data is aggregated.

Agencies will incur the cost to keep processed data in storage until pulled by CLAW, though the frequency of pulls should allow data to be quickly transitioned to cheaper tiers of storage.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. Any special constraints on the aggregate data stream are driven by those of the constituent telemetry sources, so agencies should be careful when combining data of different sensitivities. Furthermore, aggregation may produce sensitive information not deducible from either telemetry source alone.

2.5 SA-SDNN-SS: Agency Filtered Data Push to CLAW

Overview

In this reporting pattern, CSP refers to an agency's Software as a Service (SaaS) cloud vendor. The CSP provides various application services, such as customer relations, email, or support service delivery tracking, where users login to perform certain actions. These services each generate access and authentication logs, either standalone or as part of the application's general logging output. Especially in the latter case, the telemetry may contain sensitive information; the agency gathers the telemetry and then filters and enriches it before sending to CLAW.

Access and authentication logs can be rich and useful in improving CISA's situational awareness and this reporting pattern allows agencies to share such logs while protecting agency-sensitive information; however, depending on the log format, it may not be trivial to discover and sanitize all instances of sensitive data.

Figure 14 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP and filtering data, and CISA is responsible for receiving data from the agency. With regard to telemetry flow, the CSP generates telemetry from agency applications in Stage A (Cloud Sensing), the agency filters telemetry received from the CSP in Stage B (Agency Processing), and the agency pushes telemetry to CLAW in Stage C (Reporting to CISA).

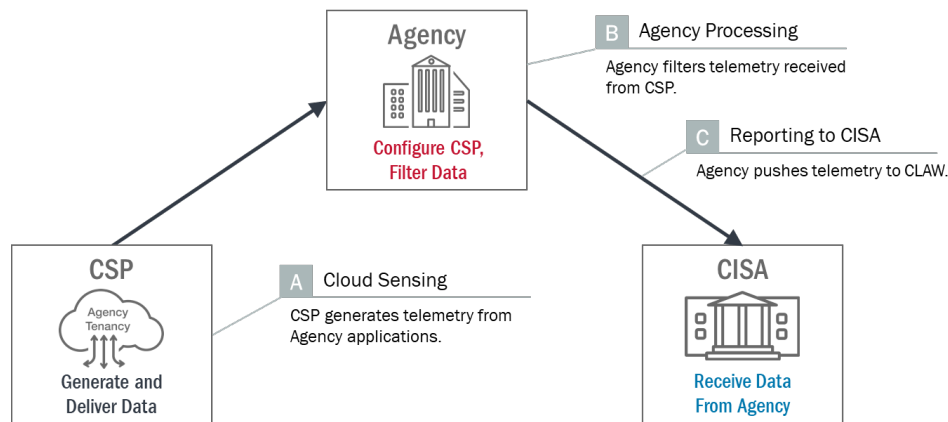


Figure 14: Roles and Telemetry Flow – SA-SDNN-SS

Stages

Figure 15 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: Network traffic between the agency's cloud tenancy and the Internet is handled by the CSP services. The agency configures one or more of these services to generate access and authentication logs.

Stage B: The agency uses its NOC/SOC tools to perform data sanitization and enrichment functions to process the raw data. The raw data may be filtered to remove agency "private/internal" sources, personally identifiable information (PII), or other sensitive information in conformance with agency

sanitization and sharing requirements. The agency then enriches some fields with derived information (e.g., destination country). The agency does not perform any aggregation or transformation.

Stage C: The agency pushes the processed telemetry to the regional CLAW. The exact delivery mechanism(s) depends on the CSP.




Figure 15: Visual Pattern Summary – SA-SDNN-SS

Pattern Summary

Table 7 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 7: Pattern Summary Table – SA-SDNN-SS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization		Single Region
Telemetry Types	Service	Data Enrichment	Obfuscation	CLAW Distribution	Multi-Region
	Application		None		Multi-Cloud
	Network Flow Logs	Data Aggregation	Derived		
	Access/Auth Logs		Agency-Defined		
	IDS/IPS Logs		None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
Firewall Logs	None (Native Forms Align)				
		Data Transformation	IPFIX		
			JSON		
			Parquet		
			CEF		
			Syslog		
			LEEF		
			CISA Coordinated		

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the application and the agency’s own policy for delivery to CLAW. Agencies should consult application-specific documentation and determine which fields might have sensitive information that is not trivial to detect and remove (which may introduce processing delay). Agencies can delay delivering individual records/objects (e.g., as part of a batching policy) and may do so if they do not exceed the maximum delay parameters.

Agency processing may significantly affect timeliness. Some log fields may be sanitized by withholding them while others may require deep scanning (e.g., PII embedded in a Uniform Resource Locator (URL) field). Agencies should also characterize the performance of different methods of cross-referencing the relevant data for enrichment.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the access and authentication logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are processed. However, the original log entries' timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves agency processing on log content, including data removal and addition. In this case, the agency is an author of log information, as it is providing enrichment and editing. Provenance claims in this context are three-fold: the origin of the information from the SaaS service, the origin of the information used in performing the enrichment, and resulting stream provided to CISA by the agency. Agency processing should be arranged to convey both the nature of the modifications (e.g., enrichment) performed, the type of information removed, and the processing mechanisms (e.g., software artifacts) used in performing the processing.

Reporting Connection Administration

The agency processing service, as the sender of the telemetry to CISA, utilizes the CISA-provided credentials and security parameters to establish the data sharing connection. Transfer system health should be monitored by both the sending agency and CISA entities. The monitoring mechanisms and procedures may be able to leverage agency processing system native functionality, reducing cost and complexity. In addition, a plan for remedial action when the transfer does not occur, is incomplete, or requires retransmission may also leverage agency processing system native functionality.

Cloud Telemetry Sharing Cost

The agency processing in this reporting pattern, especially the sanitization of sensitive information like PII, may not be trivial, and agencies may consider PaaS or SaaS capabilities when implementing a solution. Data used for enrichment during processing may be open-source or provided by vendors which charge for the service.

Implementations of this reporting pattern may involve persistent compute resources to perform the agency push to CLAW, in which case agencies incur the cost to maintain these resources.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. This reporting pattern allows agencies to filter (sanitize) raw data that would otherwise require special handling into an output stream that can be shared with less or no such constraints.

2.6 NN-SDNI-LS: CLAW Pull of Agency Filtered Data

Overview

In this reporting pattern, CSP refers to an agency’s IaaS cloud vendor. The agency configures sensors for specific subnets within its cloud tenancy. These sensors monitor the agency traffic to/from those subnets and generate network flow logs, which are processed extensively by the agency prior to being retrieved by CLAW. The agency processing is done at a single location, so retrieval by CLAW in the local region is utilized.

This reporting pattern also allows an agency to configure sensing for all its subnets – which it may already do for its own analytics – and only share logs for higher-risk segments; however, the agency is responsible for more extensive processing. In addition, the agency must ensure that the mechanism used by CLAW to pull agency data from the IaaS vendor cannot be abused by other third parties.

Figure 16 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP and filtering data, and CISA is responsible for retrieving data from the agency. With regard to telemetry flow, the CSP generates telemetry from agency traffic in Stage A (Cloud Sensing), the agency filters telemetry received from the CSP in Stage B (Agency Processing), and CLAW pulls telemetry from the agency in Stage C (Reporting to CISA).

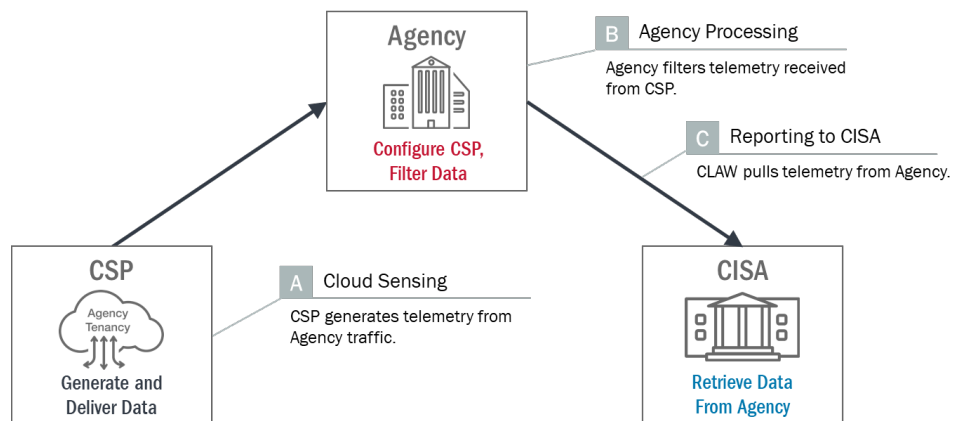


Figure 16: Roles and Telemetry Flow – NN-SDNI-LS

Stages

Figure 17 (below) shows the events that take place during each of this reporting pattern’s three stages. A detailed description of each stage is presented below:

Stage A: Network traffic to and from the agency’s chosen⁷ subnets pass through the CSP’s sensors, where security functions (e.g., firewall) are implemented. In addition to executing their security functions, these subnet-level sensors also generate network flow logs.

⁷ One possible selection of subnets consists of just those that are publicly accessible from the Internet; this allows the agency to filter out much of the data corresponding to “private/internal” sources even before the Agency Processing stage. CISA is primarily interested in this data (as opposed to private traffic between internal components).

Stage B: The agency uses its NOC/SOC tools to perform data sanitization, enrichment, and transformation functions to process the raw data. The raw data is filtered to remove agency “private/internal” sources, PII, and other sensitive information in conformance with agency sanitization requirements. The agency may perform filtering before or after other processing. The data is also transformed to the Internet Protocol Flow Information Export (IPFIX) format (although CLAW is likely capable of ingesting the data in the CSP’s native format, the agency may prefer IPFIX for its own analytics). The agency enriches some fields with derived information (e.g., destination country) in the IPFIX format.⁸ The agency does not perform any aggregation.

Stage C: The agency supplies the filtered telemetry to be pulled by the regional CLAW. This involves configuring pull credentials such that CLAW⁹ is authorized to make the necessary requests. The exact delivery mechanism(s) depends on the CSP.

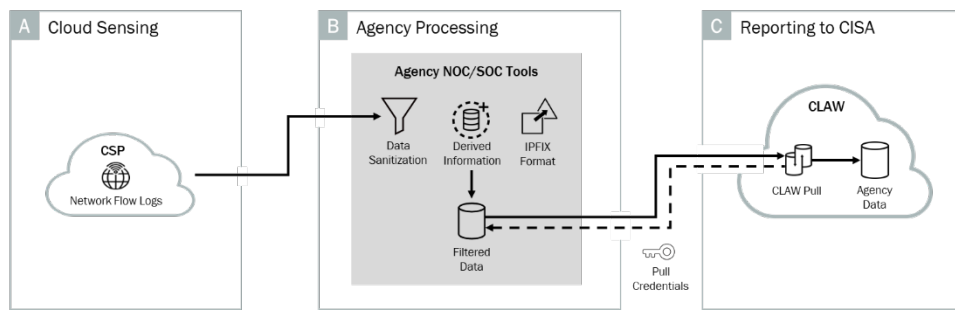


Figure 17: Visual Pattern Summary – NN-SDNI-LS

Pattern Summary

Table 8 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 8: Pattern Summary Table – NN-SDNI-LS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA		
Attribute	Options	Attribute	Options	Attribute	Options	
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push	
	Subnet		Removal		CLAW Pull	
	Interface		Sanitization		Single Region	
	Service		Obfuscation		Multi-Region	
	Application		None		Multi-Cloud	
Telemetry types	Network Flow Logs	Data Enrichment	Derived	CLAW Distribution		
	Access/Auth Logs	Data Aggregation	Agency-Defined			
	IDS/IPS Logs		None			
	API Activity Logs		Multi-Account			
	DNS Logs		Multi-Region			
	VPN Logs		Multi-Provider			
	Firewall Logs		Data Transformation			None (Native Forms Align)
						IPFIX
						JSON
						Parquet
CEF						
			Syslog			
			LEEF			
			CISA Coordinated			

⁸ Potentially in the form of enterprise-specific information elements.

⁹ An authenticated identity principal corresponding to the CLAW instance in the selected region.

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the aggregation interval for network flow logs and the polling frequency of CLAW. Because (successful) network flows are not point events, when they “occur” is partly determined by the aggregation interval; shorter intervals trade quicker visibility for higher log volume (and vice versa). Tenants have some control over this interval (depending on the CSP). The frequency with which CLAW checks for and pulls new data adds delay and is limited by mechanisms such as API request throttling.

Agency processing may significantly affect timeliness. As there is more extensive processing than in other patterns, agencies should test and document the end-to-end processing time for logs, ideally under realistic workloads.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the service application logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are processed. However, the original log entries’ timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves potentially significant processing on log content by the agency, including data removal, data transformation, and enrichment. Agency processing should be arranged to convey both the nature of the modifications (e.g., enrichment) performed, the type of information removed, and the processing mechanisms (e.g., software artifacts) used in performing the processing. In this case, the agency is an author of log information or metadata. Provenance claims in this context are three-fold: the origin of the information from the IaaS service, the origin of the information used in performing the enrichment, and the resulting stream provided to CISA by the agency. The report stream has undergone a transformation so the nature and entity author(s) of the transformations should be captured in the provenance claims.

Reporting Connection Administration

The agency prepares for the CISA pull transfer by provisioning credentials and establishing reachability to the processing source from CLAW. The agency may utilize key management services (if offered by the CSP). The frequency of CLAW telemetry pull transactions, buffer sizes for individual telemetry items, notifications for successful receipt, and other parameters should be negotiated prior to telemetry sharing initiation. Monitoring of the data transfers should be monitored by both the sourcing (agency processing service) and receiving (CISA) entities. The monitoring mechanisms and procedures may leverage agency processing system native functionality.

Cloud Telemetry Sharing Cost

CSPs that offer network flow logs do so at minimal or no cost to tenants, although some allow tenants to pay a premium for logs with a shorter aggregation interval or higher resolution.

The agency processing in this reporting pattern, especially the sanitization of sensitive information like PII, may not be trivial, and agencies may consider PaaS or SaaS capabilities when implementing a

solution. Data used for enrichment during processing may be open-source or provided by vendors which charge for the service.

Agencies will incur the cost to keep processed data in storage until pulled by CLAW, though the frequency of pulls should allow data to be quickly transitioned to cheaper tiers of storage.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. This reporting pattern allows agencies to filter (sanitize) raw data that would otherwise require special handling into an output stream that can be shared with less or no such constraints. While individual network flow records may be less rich than other log types, network flow logs as a whole may reveal sensitive information which warrants special handling.

2.7 SF-NDNJ-SS: Agency CSP SECaaS Data Push to CLAW

Overview

In this reporting pattern, the CSP provides Security as a Service (SECaaS) to the agency. In the SECaaS model, the sensors that generate telemetry are managed by the CSP and configured by the agency. This reporting pattern outlines a basic case where telemetry generated by the CSP is delivered directly to CLAW.

In addition to sensing, this reporting pattern allows an agency to use the telemetry processing and delivery capabilities of a SECaaS vendor; however, not all vendors may provide the features necessary to support all components of this reporting pattern (e.g., transformation to a CISA-acceptable format, delivery to CLAW).

Figure 18 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP, and CISA is responsible for receiving data from the CSP. With regard to telemetry flow, the CSP generates telemetry from agency traffic in Stage A (Cloud Sensing), the CSP processes telemetry based on agency settings in Stage B (Agency Processing), and the CSP pushes telemetry to CLAW in Stage C (Reporting to CISA).

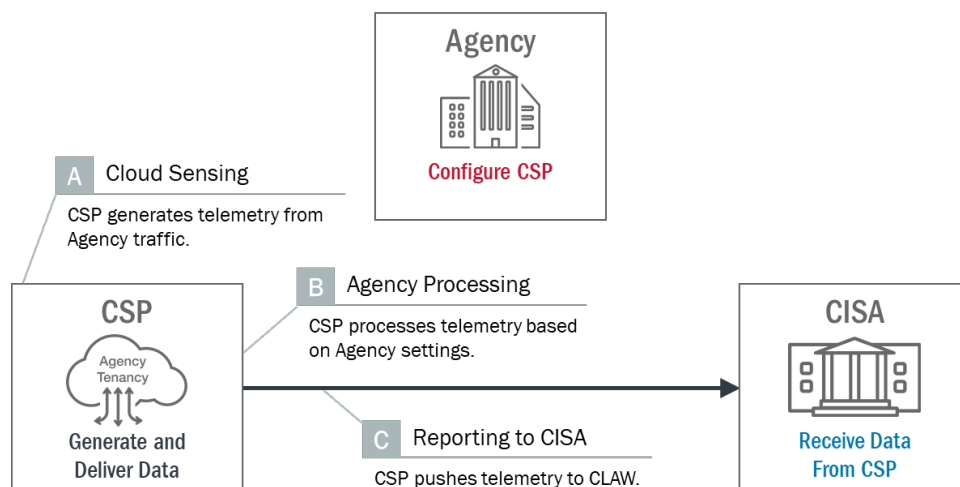


Figure 18: Roles and Telemetry Flow – SF-NDNJ-SS

Stages

Figure 19 (below) shows the events that take place during each of this reporting pattern’s three stages. A detailed description of each stage is presented below:

Stage A: Network traffic between the agency and the Internet is routed through the CSP’s services, where various security functions are implemented, which may include firewall, DDoS protection, or web filtering. These services can generate different telemetry types depending on the CSP and services used. In this reporting pattern, the agency configures the CSP to generate firewall logs.

Stage B: The agency configures the CSP service using NOC/SOC tools. The agency does not configure any filtering¹⁰ but does configure enrichment and transformation. The agency configures the CSP option to include some enrichment fields with derived information (e.g., destination country). Finally, data is transformed from its native format into JSON.¹¹ The agency does not perform any aggregation.

Stage C: The agency configures its telemetry to be pushed from the NOC/SOC tools to the regional CLAW. The exact delivery mechanism(s) depends on the CSP; while coordinating on the telemetry format, the agency and CISA also work together to ensure that CLAW is capable of directly receiving telemetry from the third-party.




Figure 19: Visual Pattern Summary – SF-NDNJ-SS

Pattern Summary

Table 9 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 9: Pattern Summary Table – SF-NDNJ-SS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization	CLAW Distribution	Single Region
	Service		Obfuscation		Multi-Region
Telemetry Types	Application	Data Enrichment	None	Multi-Cloud	
	Network Flow Logs		Derived		
	Access/Auth Logs	Agency-Defined			
	IDS/IPS Logs	Data Aggregation	None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
	Firewall Logs	Data Transformation	None (Native Forms Align)		
			IPFIX		
			JSON		
	Parquet				
	CEF				
	Syslog				
	LEEF				
	CISA Coordinated				

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the agency’s own policy for delivery to CLAW. Agencies can delay delivering individual records/objects (e.g., as part of a batching policy) and may do so if they do not exceed the maximum delay.

¹⁰ Similar to pattern NN-SDNI-LS, the agency can configure the routing so that only traffic between public-facing components and the Internet is routed through the CSP’s sensors, removing one of the common drivers of filtering.

¹¹ Before telemetry is sent, the agency shares schema information (field names, field types, and other constraints) with CISA.

Agency processing may significantly affect timeliness. Agencies should characterize the performance of different methods of cross-referencing the relevant data for enrichment.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the service application logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are processed. However, the original log entries' timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves the agency applying processing to transform data from a SECaaS vendor format to a CISA-acceptable format, including potential data enrichment. In this case, the agency is the primary author of log information. Provenance claims in this context are three-fold: the origin of the information from the SECaaS service, the origin of the information used in performing the enrichment, and the information regarding the resulting stream provided to CISA and authored by the agency. The stream is being freshly authored based on information provided by enrichment and the SECaaS provider and is not limited to simple transformations. In this pattern, agency processing should be arranged to convey both the nature of the original sources, the processing mechanisms (e.g., software artifacts) used in performing the processing, and an indicator of the agreement between the agency and CISA governing the streams provided.

Reporting Connection Administration

The agency processing service, as the sender of the telemetry to CISA, utilizes the CISA-provided credentials and security parameters to establish the data sharing connection. Transfer system health should be monitored by both the sending agency and CISA entities. The monitoring mechanisms and procedures may be able to leverage agency processing system native functionality, reducing cost and complexity. In addition, a plan for remedial action when the transfer does not occur, is incomplete, or requires retransmission may also leverage agency processing system native functionality.

Cloud Telemetry Sharing Cost

The agency processing in this reporting pattern may not be trivial, and agencies may consider PaaS or SaaS capabilities when implementing a solution. Data used for enrichment during processing may be open-source or provided by vendors which charge for the service.

Implementations of this reporting pattern may involve persistent compute resources to perform the agency push to CLAW, in which case agencies incur the cost to maintain these resources.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. CSP options for direct delivery to CLAW may be unable to satisfy such constraints, in which case other reporting patterns should be considered.

2.8 SA-SANC-SS: CSP SECaaS Data, Agency Processing, and Push to CLAW

Overview

In this reporting pattern, the CSP provides SECaaS to the agency. Telemetry generated by the CSP is sent to the agency, which processes the data prior to sending it to CLAW.

In addition to sensing, this reporting pattern allows an agency to use the telemetry processing capabilities of the SECaaS vendor, augmented with its own processing; however, the pattern may be difficult to implement for an agency that does not already have its own mature analytics process. The processed data format does not align with established CISA supported structures and requires CISA-coordination prior to delivery.

Figure 20 (below) shows the roles and telemetry flow associated with this reporting pattern. With regard to roles, the CSP is responsible for generating and delivering data, the agency is responsible for configuring the CSP and processing data, and CISA is responsible for receiving data from the agency. With regard to telemetry flow, the CSP generates telemetry from agency traffic in Stage A (Cloud Sensing), the agency processes telemetry received from the CSP in Stage B (Agency Processing), and the agency pushes telemetry to CLAW in Stage C (Reporting to CISA).

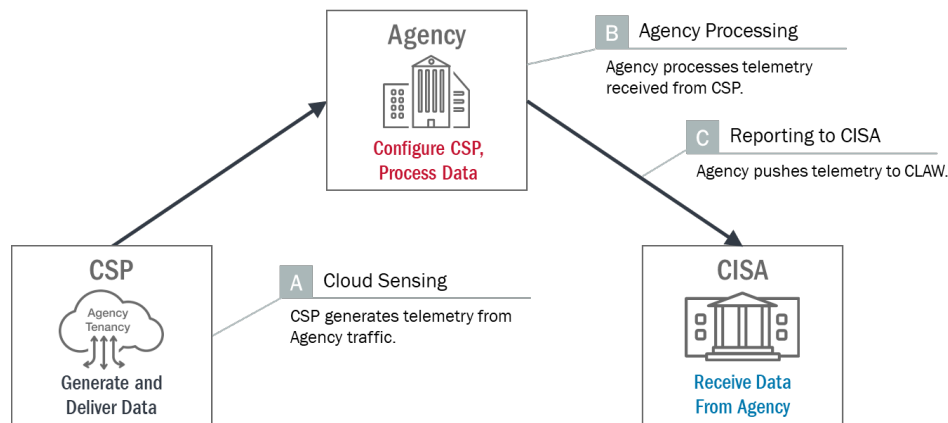


Figure 20: Roles and Telemetry Flow – SA-SANC-SS

Stages

Figure 21 (below) shows the events that take place during each of this reporting pattern's three stages. A detailed description of each stage is presented below:

Stage A: The agency configures the CSP services to generate access and authentication logs.

Stage B: The agency performs data filtering, enrichment, and transformation, first by using functions provided by the CSP services (to perform data sanitization and enrichment) and then by using its own NOC/SOC tools (to perform further data sanitization and enrichment, as well as transformation). Factors for selection of where processing occurs include performance, cost, and privacy. The agency processing may include capabilities implemented through self-hosted services or from an agency's cloud telemetry processing service. The agency uses the CSP's service capabilities to pre-process the telemetry to include certain enrichment fields with agency-defined information and exclude certain fields with sensitive information that should not be shared (i.e., data sanitization). Optionally, the agency may also

configure the CSP services to output the telemetry in an intermediate format convenient for its own processing. After the processing at the CSP service, the telemetry is delivered to the agency for additional processing. For example, the agency further sanitizes web transaction telemetry by scanning for sensitive data embedded within the URL field, and further enriches firewall transaction logs with agency-defined data (such as labels identifying resources within its cloud tenancy). As final processing, the agency transforms the data into a format agreed upon with CISA.¹² No data aggregation is performed.

Stage C: The agency pushes the processed logs to the regional CLAW.

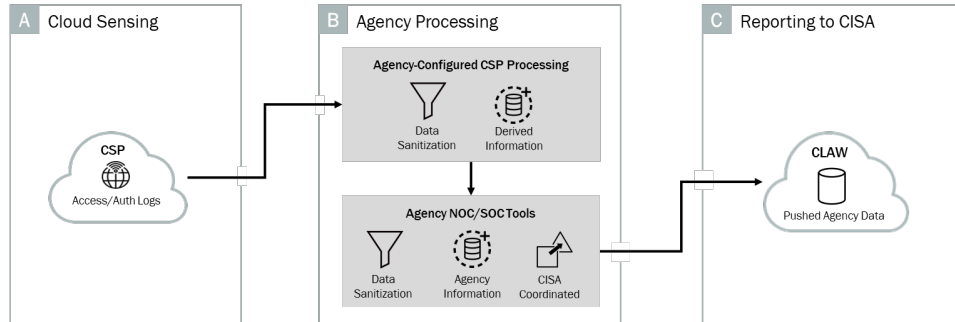


Figure 21: Visual Pattern Summary – SA-SANC-SS

Pattern Summary

Table 10 (below) lists the option that is associated with each attribute in this reporting pattern.

Table 10: Pattern Summary Table – SA-SANC-SS

Stage A – Cloud Sensing		Stage B – Agency Processing		Stage C – Reporting to CISA	
Attribute	Options	Attribute	Options	Attribute	Options
Sensor Positioning	Gateway	Data Filtering	None	Data Transfer	Agency Push
	Subnet		Removal		CLAW Pull
	Interface		Sanitization		Single Region
Telemetry Types	Service	Data Enrichment	Obfuscation	CLAW Distribution	Multi-Region
	Application		None		Multi-Cloud
	Network Flow Logs	Data Aggregation	Derived		
	Access/Auth Logs		Agency-Defined		
	IDS/IPS Logs		None		
	API Activity Logs		Multi-Account		
	DNS Logs		Multi-Region		
	VPN Logs		Multi-Provider		
Firewall Logs	Data Transformation	None (Native Forms Align)			
		IPFIX			
		JSON			
		Parquet			
		CEF			
		Syslog			
		LEEF			
		CISA Coordinated			

¹² CISA may request modifications to the CSP’s default format in order to include required information or to improve ingestion processing. Once CISA decides on a format for a given vendor and service, subsequent agencies that use the same CSP/service may use it as a standard.

Pattern Characteristics

Cloud Telemetry Timeliness

For this pattern, factors affecting the timeliness of information include the agency's own policy for delivery to CLAW. Agencies can delay delivering individual records/objects (e.g., as part of a batching policy) and may do so if they do not exceed the maximum delay parameters.

Agency processing may significantly affect timeliness. As there is more extensive processing than in other patterns, agencies should test and document the end-to-end processing time for logs, ideally under realistic workloads. However, it is expected that the CSP can perform any pre-processing configured by the agency in real-time.

Cloud Telemetry Timing Coordination

In this case, the processing during aggregation will have an opportunity to introduce its own timestamps as provenance claims. However, the network flow logs, and access and authentication logs are still timestamped when the log entry is generated at the CSP. Additional timestamps may be added at the time when the log entries are processed. However, the original log entries' timestamps should be preserved.

Cloud Telemetry Provenance

This pattern involves the agency applying processing to transform data from a SECaaS vendor format to a CISA-acceptable format, along with arbitrary data transformations, filtration, and enrichment decided by the agency. In this case, the agency is the primary author of log information. Provenance claims in this context are multiple (depending on the complexity of the agency processing performed) but include: the origin of the information from the SECaaS and other services, the origin of the information used in performing the enrichment, and information regarding the resulting stream provided to CISA and authored by the agency. The stream is being freshly authored based on information provided by enrichment and the SECaaS provider. In this pattern, agency processing should be arranged to convey the provenance of all original sources, all processing mechanisms (e.g., software artifacts and services) used in performing the processing, and an indicator of the agreement between the agency and CISA demonstrating how the stream provided to CISA is sufficient for NCPS operations.

Reporting Connection Administration

The agency processing service, as the sender of the telemetry to CISA, utilizes the CISA-provided credentials and security parameters to establish the data sharing connection. Transfer system health should be monitored by both the sending agency and CISA entities. The monitoring mechanisms and procedures may be able to leverage agency processing system native functionality, reducing cost and complexity. In addition, a plan for remedial action when the transfer does not occur, is incomplete, or requires retransmission may also leverage agency processing system native functionality.

Cloud Telemetry Sharing Cost

Configuring the CSP that generates the data to perform some pre-processing can result in significant cost reduction, as the agency processing in this reporting pattern (especially the sanitization of sensitive

information like PII) is not trivial. Agencies may not have to pay any additional cost (over what they already pay for the SECaaS provided by the CSP) to have this pre-processing done.

Implementations of this reporting pattern may involve persistent compute resources to perform the agency push to CLAW, in which case agencies incur the cost to maintain these resources.

Agency Data Retention and Use Constraints

Before sharing is established, agencies should communicate any special data retention and use constraints to CISA. This reporting pattern allows CSPs and agencies to filter (sanitize) raw data that would otherwise require special handling into an output stream that can be shared with less or no such constraints. On the other hand, data enrichment with agency-defined information may introduce sensitive information into the output stream which warrants special handling.

3 COMBINATION REPORTING PATTERNS

A combination reporting pattern is when two or more existing reporting patterns are selected to be applied in concert. Combination patterns tend to arise when there are multiple sources of raw telemetry and one reporting pattern is not appropriate for all of them. As with Section 2, this document will only focus on a small set of possible combinations. Combinations not shown here may still be viable alternatives and should be discussed with CISA on a case-by-case basis.

A short description is provided for each combination reporting pattern, along with pros, cons, and alternatives to guide characteristics. For brevity, familiarity with Section 2 is assumed and discussion about the attributes and options of each constituent pattern is omitted.

Combination Reporting Pattern Characteristics

Cloud Telemetry Timeliness

The combination reporting patterns mix various details of the previously discussed timeliness characteristic. Agencies should not expect or try to achieve “uniform” timeliness from all sources. Instead, they should make sure that the delay from event occurrence to delivery to CLAW is within CISA preferences in all cases. This may require extensive testing.

Cloud Telemetry Timing Coordination

In the case of combination reporting patterns, the processing stage will have an opportunity to introduce its own timestamps into the overall chain that originates from the source and terminates at the CLAW. However, the cloud telemetry logs are still timestamped when they are generated at the CSP. Additional timestamps may be added at the time when the log entries are processed. However, the original log entries’ timestamps should be preserved.

Cloud Telemetry Provenance

The combination reporting patterns mix various details of the other generic reporting patterns and provenance concerns vary depending on the specific details. The combination reporting pattern scenarios may be more complex, as data formats and provenance from different types of systems (e.g., SaaS, IaaS) and locations or administrative controls may be interleaved, each of which may have different levels of abstraction/granularity and reporting capabilities (e.g., time, identity). In cases where multiple different log types can be aggregated and processed, a common field is typically used to correlate information. A timestamp or transaction identifier is commonly used;¹³ note that time should be of sufficient precision and accuracy to make such log aggregation possible.

Reporting Connection Administration

The need for effective connection administration increases as an agency shares more output streams with CLAW. Automated or semi-automated processes for key management, health monitoring (including completeness and timeliness of data transfer), issue remediation (including re-delivering data in case of a visibility gap or terminating transfer in case of a suspected compromise), and other aspects of connection administration will significantly reduce the burden on agencies, especially as more streams are shared with CLAW.

¹³ See, for example, minimum requirements for 1msec granularity in the financial industry (Consolidated Audit Trail National Market System; available at <https://www.sec.gov/rules/proposed/2010/34-62174.pdf>).

Cloud Telemetry Sharing Cost

When applied in the right context, the combination reporting patterns can be more cost-efficient than their individual counterparts. Reasons include applying less/no agency processing to some input streams, keeping data streams within their respective CSP regions, and utilizing an agency's existing analytics process. Some combination reporting patterns require more cloud resources (compute, storage, etc.) and/or more staffing hours to develop and maintain, offsetting some of the savings.

Agency Data Retention and Use Constraints

A distinguishing feature of combination reporting patterns is that, in many of the patterns, the agency shares multiple output streams to CLAW. When telemetry sources (or agency processing) produce data at varying levels of sensitivity, agencies can group data of like sensitivity into their own stream; by separating out lower sensitivity data, special handling constraints are applied only to the higher sensitivity data to require it. Many options exist to determine the content of each stream and agencies should weigh the costs and benefit of this approach.

3.1 Differentiated Processing of Multi-Account Data (GV-NNAN-SS + SN-NNNN-LS)

Description

This combination pattern can be used by agencies that have multiple accounts, where telemetry from different accounts may have different attributes. Telemetry from Accounts 1/2 are handled as in pattern GV-NNAN-SS, with some additional processing, and data from multiple accounts is aggregated prior to delivery to CLAW. Telemetry from Account 3 is handled independently, pulled directly from the CSP by CLAW (just as in pattern SN-NNNN-LS). This approach can support various use cases, such as bypassing agency processing for streams that do not require it (e.g., no sanitization required for Account 3) or for sending multiple streams to CLAW based on sub-organizations within the agency (e.g., one group owns Accounts 1/2, and another owns Account 3).

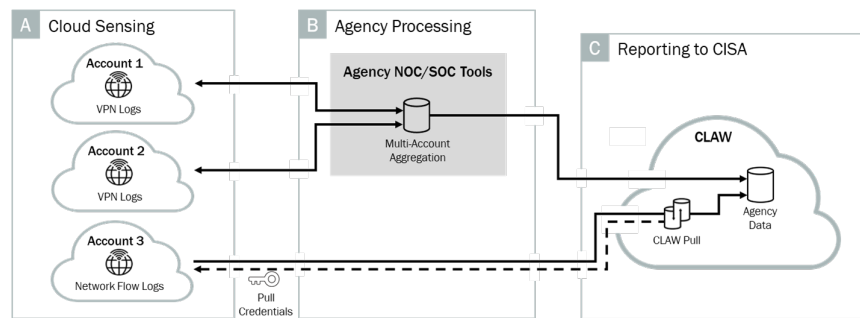


Figure 22: Visual Pattern Summary – Differentiated Processing of Multi-Account Data

Pattern Summary

Table 11: Pattern Summary Table – Differentiated Processing of Multi-Account Data

Stage	Attribute	Account 1 Option	Account 2 Option	Account 3 Option
Stage A: Sensing	Sensor Positioning	Gateway	Gateway	Subnet
	Telemetry Types	VPN Logs	VPN Logs	Network Flow Logs
Stage B: Agency Processing	Data Filtering	None	None	None
	Data Enrichment	None	None	None
	Data Aggregation	Multi-Account	Multi-Account	None
Stage C: Reporting to CISA	Data Transformation	None	None	None
	Data Transfer	Agency Push	Agency Push	CLAW Pull
	CLAW Distribution	Single region	Single region	Single region

Pros

- Different input streams are handled naturally according to their needs.
- A "sub-agency" can be assigned to each output stream sent to CLAW, allowing CISA to conduct both whole-agency and more granular analysis.
- Issues pushing Account 1/2 data to CLAW do not necessarily affect CLAW's ability to pull Account 3 data.

Cons

- Account-level granularity may not be enough when differentiating streams.

- Multiple groups may be responsible for sending data to CLAW.
- Without additional configuration, the agency NOC/SOC does not have visibility into Account 3.

Alternatives

In the simplest alternative, Account 3 telemetry is aggregated along with Account 1/2 data, reducing this combination pattern into a variant of pattern GV-NNAN-SS. This approach largely inverts the pros/cons listed above.

In another alternative, Account 3 telemetry undergoes a separate and minimal processing pipeline, resulting in a push to CLAW independent of the Account 1/2 telemetry. This approach alleviates some of the cons listed above but results in additional complexity in the Agency Processing stage.

3.2 Per-Region Processing of Multi-Region Data

Description

This combination pattern can be used by agencies using a single CSP in multiple regions. Like pattern SA-SDNN-SS, the agency has logs that are sanitized prior to delivery to CLAW. This time, the logs are network flow logs and originate from two different regions, which the agency handles entirely in-region; the agency provisions identical processing pipelines in both regions and send the output of each to the local regional CLAW. In other words, two instances of pattern SA-SDNN-SS are combined to handle two regions. This approach can be generalized to any number of regions and can be applied in any instance where similar telemetry is generated in multiple regions.

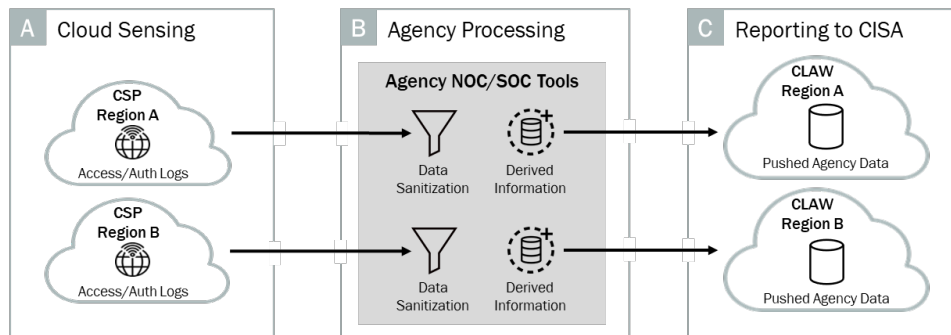


Figure 23: Visual Pattern Summary – Per-Region Processing of Multi-Region Data

Pattern Summary

Table 12: Pattern Summary Table – Per-Region Processing of Multi-Region Data

Stage	Attribute	Region A Option	Region B Option
Stage A: Sensing	Sensor Positioning	Service	Service
	Telemetry Types	Access/Auth Logs	Access/Auth Logs
	Data Filtering	Sanitization	Sanitization
Stage B: Agency Processing	Data Enrichment	Derived	Derived
	Data Aggregation	None	None
	Data Transformation	None	None
Stage C: Reporting to CISA	Data Transfer	Agency Push	Agency Push
	CLAW Distribution	Multi-Region	Multi-Region

Pros

- Data is kept within one region, minimizing data transfer costs.
- Infrastructure-as-code services can be used so the agency only implements a pipeline template once.
- Issues in one pipeline do not necessarily affect others.

Cons

- Cost of operating multiple pipelines may exceed the cost of a single pipeline capable of handling all the data.
- In the absence of infrastructure-as-code services, changes need to be applied independently to each pipeline.

Alternatives

Agencies may instead conduct multi-region aggregation to produce a single stream of data, processed by a single pipeline and delivered to a single regional CLAW. This approach largely inverts the pros/cons listed above.

3.3 Push from Integrated Sharing Solution

Description

In the integrated sharing solution, an agency is already performing robust cloud telemetry processing and is extending the output of their tools to now include reporting to CISA via CLAW. This pattern takes an “all of the above” approach to the breadth of input and processing. Input sources may include telemetry from the local CSP, other CSPs, on-premise analytics, mobile device management systems, and CSP or third-party threat intelligence. The cloud sensing may include multiple CSP sensor positions with multiple telemetry types. The resulting information is aggregated together with other (possibly non-security) information for subsequent filtration, enrichment, transformation, and export as selected by the agency. CISA is one consumer; others may include the agency’s own risk management, security, and operational personnel.

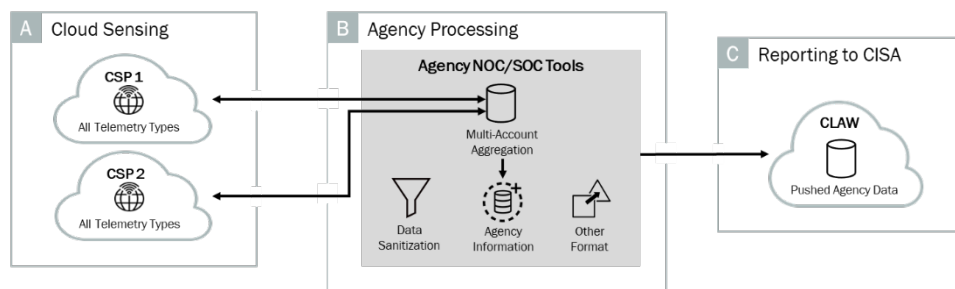


Figure 24: Visual Pattern Summary – Push from Integrated Sharing Solution

Pattern Summary

Table 13: Pattern Summary Table – Push from Integrated Sharing Solution

Stage	Attribute	Tenancy Options
Stage A: Sensing	Sensor Positioning	Gateway, Subnet, Interface, Service, Application
	Telemetry Types	All Types
	Data Filtering	Removal, Sanitization, Obfuscation
Stage B: Agency Processing	Data Enrichment	Derived, Agency-Defined
	Data Aggregation	Multi-Account, Multi-Region, Multi-Provider
	Data Transformation	CISA Coordinated
Stage C: Reporting to CISA	Data Transfer	Agency Push
	CLAW Distribution	Single-Region

Pros

- Visibility is broad due to multiple input streams.
- Existing agency capabilities and integration are leveraged.
- CLAW attribution and coordination is simplified, as all telemetry for the agency is originating from a single source system.

Cons

- Provenance claims must rely on complex mechanisms to ensure unique identifiers for all physical and logical resources in both cloud environments, including computing resources, person and non-person accounts, and IP addressable infrastructure components.
- Diversity in underlying data types and attributes between cloud environments can increase difficulty in transferring a normalized, repeatable telemetry set.

- Deployment may include unique requirements and supported capabilities for ingestion of information from multiple CSPs.
- Telemetry transfer health monitoring for all parties may also require sophisticated retransmission accommodations when gaps in data are observed from only one source.

Alternatives

Agencies may instead determine that CLAW telemetry sharing requirements align with an existing output consumer (permitting reuse).

3.4 Push to Local Regional CLAW in Multiple CSPs

Description

This combination reporting pattern is for agencies with operations in multiple IaaS CSPs or in multiple regions within a CSP that wish to minimize costs associated with data egress. As with the combination reporting pattern “Per-Region Processing of Multi-Region Data” (see Section 3.2), the agency processes logs in the same region and CSP where they originated. While there is additional complexity for the agency implementing processing pipelines in multiple CSPs, this avoids the egress costs associated with performing multi-CSP aggregation. After processing the logs, the agency pushes the data to the local regional CLAW.

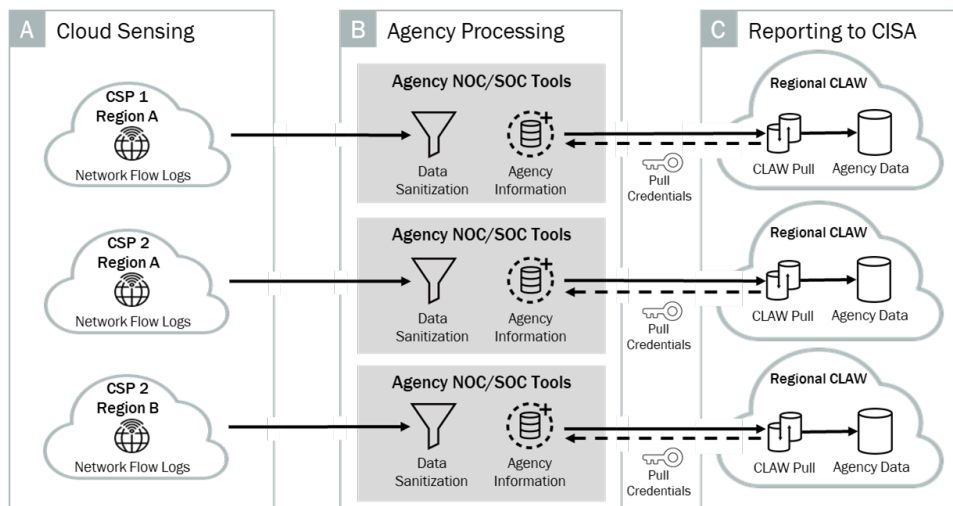


Figure 25: Visual Pattern Summary – Push to Local Regional CLAW in Multiple CSPs

Pros

- Data is kept within one region, minimizing data transfer costs.
- Infrastructure-as-code services can be used so the agency only implements a pipeline template once.
- Problems in one pipeline do not necessarily impact other pipelines.
- Reporting pattern is generalizable to any number of CSPs and regions.

Cons

- Cost of operating multiple pipelines may exceed the cost of a single pipeline capable of handling all the data.
- In the absence of infrastructure-as-code services, changes need to be applied independently to each pipeline.
- The agency processing pipeline implementation and associated infrastructure-as-code templates may need to be customized for each CSP.

Alternatives

Agencies may instead conduct multi-region and/or multi-CSP aggregation to produce a single stream of data, processed by a single pipeline and delivered to a single regional CLAW. This approach largely inverts the pros/cons listed above.

4 CONCLUSION

As agencies move more of their applications and services to cloud, the NCPS Program is evolving to ensure that security information for cloud-based traffic can be captured and analyzed and that CISA analysts can continue to provide situational awareness and support to the agencies. The *NCPS Cloud Interface Reference Architecture: Volume One* document introduces a framework for developing reporting patterns for how cloud logs will be collected and transferred to CLAW. This companion document (*NCPS Cloud Interface Reference Architecture: Volume Two*) provides a catalog of generic reporting patterns that match common agency cloud use cases and shows how more complex reporting patterns can be developed to describe use cases with a combination of attributes and options. Together, these two documents provide guidance for how an agency can adapt its cloud environments to allow for security data to be sent to CLAW.

Individual CSPs can use these documents to provide vendor solutions that match reporting patterns. Vendors are encouraged to develop overlays that identify how their agency customers can comply with NCPS visibility requirements while using the CSP's products and services. While CISA will not provide formal authorization or approval of a vendor overlay solution, CISA may provide input to the vendor on a case-by-case basis to convey desired approaches and intent.