



Brazil Telecommunications Service Operations Management

Last updated: 10/08/2026

Some examples and graphics depicted herein are provided for illustration only. No real association or connection to ServiceNow products or services is intended or should be inferred.

ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Please read the ServiceNow Website Terms of Use at www.servicenow.com/terms-of-use.html

Company Headquarters
2225 Lawson Lane
Santa Clara, CA 95054
United States
(408) 501-8550

Table of Contents

Telecommunications Service Operations Management.....	4
Explore.....	5
Telecom data model.....	7
Telecom Visibility.....	16
Telecom Assurance.....	56
Configure.....	69
Configure Telecom Visibility.....	70
Configure Telecom Assurance.....	135
Use.....	156
Use Telecom Discovery patterns.....	156
Run SGC imports.....	160
Validate JSON payloads.....	164
Run Telecom Discrepancy audit.....	166
Attribute value discrepancy report.....	169
Reference.....	170
Arista VeloCloud installed integrations.....	170
Fortinet installed integrations.....	170
Fortinet extension points reference.....	171
Cisco Meraki installed integrations.....	174
Telecom API notifications roles.....	175
Nokia Altiplano system components.....	176
Discrepancy and Reconciliation components.....	177
Fortinet Service Graph Connector API endpoints.....	180
Cisco Meraki Service Graph Connector API endpoints.....	181
Cisco Meraki extension points.....	187
Arista VeloCloud Service Graph Connector API endpoints.....	191
MPN Formulas table.....	193
MPN latency KPIs.....	194
MPN health status rule reference.....	195
Equinix connector metrics.....	197
Retrieve data from Nokia Altiplano via REST API.....	198
Page limits for Meraki pull connector endpoints.....	199
Pull connector granularity constraints.....	200
Metric aggregation configuration reference.....	200
Metric type IDs.....	202

Telecommunications Service Operations Management

Learn how Telecommunications Service Operations Management (TSOM) empowers communication service providers (CSPs) to proactively monitor, analyze, and resolve network and service issues before they impact customers. Built on the ServiceNow AI Platform, TSOM delivers a unified operations view across distributed, multi-domain telecom environments, helping teams improve service availability, operational efficiency, and customer satisfaction.

Get started

<p>Explore</p>  <p>Learn about how telecom service providers use Telecommunications Service Operations Management.</p>	<p>Configure</p>  <p>Plan and configure your Telecommunications Service Operations Management.</p>
<p>Use</p>  <p>Use Telecommunications Service Operations Management to track comprehensive telecom service operations.</p>	<p>Reference</p>  <p>Get Telecommunications Service Operations Management reference information.</p>

Additional resources

- Learn more about what's new and changed, see the <https://www.servicenow.com/docs/access?context=telecommunications-service-operations-management-rn&version=brazil&pubname=brazil-release-notes&ft:locale=en-US>.
- Review the multiple TM Forum (TMF) Open APIs supported for the product, see [TMF APIs for TMT](#).
- Log in to your ServiceNow® account and find additional information about implementing and deploying Telecommunications Service Operations Management features at the [Best Practices Library](#).
- Access real time courses, self-paced training, and career resources at [ServiceNow University](#).

- Find useful resources related to your role and explore general guidelines at [ServiceNow Impact](#).
- Connect with other Telecommunications Service Operations Management users at [Now Community](#).


Exploring Telecommunications Service Operations Management

Learn how Telecommunications Service Operations Management (TSOM) empowers communication service providers (CSPs) proactively monitor, analyze, and resolve network and service issues before they impact customers. Built on the ServiceNow AI Platform, TSOM delivers a unified operations view across distributed, multi-domain telecom environments, helping teams improve service availability, operational efficiency, and customer satisfaction.

Telecommunications Service Operations Management is a solution for telecom providers that offers complete visibility into network and service health. TSOM collects, correlates, and prioritizes events from across network domains—such as access, transport, and core—using standard APIs and the power of the ServiceNow AI Platform. By providing real-time, actionable insights, TSOM enables both frontline and back-office teams to address service-impacting issues and maintain consistent performance.

How TSOM works

TSOM architecture and Telecom Applications



Telecommunications Service Operations Management simplifies telecom operations by connecting with existing monitoring and telemetry platforms, identifying actionable patterns, and automating resolution workflows. It leverages:

- External event management via Telecommunications API notifications for standardized alarm ingestion.
- Event Management for event correlation, suppression, and prioritization.
- Metric Intelligence to detect threshold breaches, performance degradation, and anomalous behavior in real time.
- Health Log Analytics to analyze log data, triage related issues, and identify root causes before users are impacted.

- Service Impact Analysis to assess and trace service disruptions based on impacted infrastructure and business services.
- Service Graph Connectors and Discovery to build a dynamic, telecom-aware CMDB.

Key features

TSOM key capabilities

Capability	Description
Real-time event monitoring	Ingest alarms and events from multi-domain network monitoring systems using External event management via Telecommunications API notifications.
Event correlation & analysis	Use Event Management and Metric Intelligence to correlate related events, reduce noise, and detect anomalies. For more information, see Event Notification Management Open API .
Telecom Visibility	Gain end-to-end visualization of network and service health, including impact traceability.
Service Impact Analysis	Understand how network or infrastructure issues affect services and prioritize remediation based on business impact.
Metric Intelligence	Monitor performance trends, detect threshold breaches and anomalies in metrics to proactively identify issues. For more information, see Metric Intelligence .
Health Log Analytics	The ServiceNow Health Log Analytics application helps prevent IT issues before your users are affected. It helps you identify the root cause of an issue by enabling you to triage related logs and analyze the raw data. For more information, see Health Log Analytics .
Automated remediation	Use guided workflows and playbooks to drive fast, consistent, and auditable issue resolution.
Telecom-aware CMDB	Link infrastructure, services, and physical/logical configuration items (CIs) using a telecom-specific model for accurate root cause analysis.
Alert Management	Manage alerts efficiently with correlation, grouping, and automated response actions. For more information, see Manage and monitor alerts and Fault Management: Events and alerts .

Key benefits

- MTTR and MTTD reduction: Identify and respond to service-impacting issues before escalation.
- Improve operational efficiency: Noise reduction and streamline operations with guided workflows.
- Gain End-to-End Visibility: Understand how infrastructure issues impact customer-facing services.
- Integrate seamlessly: Ingest data from legacy and modern NMS/EMS tools with open standards.
- Verify compliance: Align with TM Forum standards for telecom service assurance.

Key personas

- Network Operations Center (NOC) agents: Monitor network events and execute guided resolutions.
- Service Assurance managers: Analyze service trends and track resolution metrics.
- System Integrators / admins: Configure data ingestion, correlation rules, and workflows.
- Back-office analysts: Investigate root causes and track CI relationships in the CMDB.

Integration with the ServiceNow AI Platform

TSOM is embedded into the ServiceNow AI Platform and integrates seamlessly with core capabilities such as:

- Discovery
- Service Mapping
- CMDB
- Flow Designer and Playbooks

This integration confirms consistent workflows, accurate service models, and unified operations management across telecom domains.

Related topics

[External event management via Telecommunications API notifications](#)

[Telecom Visibility](#)

Telecom data model

The Telecom data model defines a structured framework for representing telecommunications networks within the CMDB. It extends the core configuration item (CI) model to describe telecom-specific infrastructure, connections, and relationships in a consistent and scalable way.

At its core, the Telecom data model organizes network elements such as equipment, interfaces, ports, sites, and connections, and defines how these elements relate to one another. The model captures both physical and logical aspects of the network and uses containment and consumption relationships to represent how network components are assembled and interconnected. This structure enables accurate modeling of complex telecom environments while remaining aligned with the standard CMDB architecture.

The Telecom data model is extensible and technology-agnostic. It provides a common foundation that supports multiple telecommunications domains by reusing existing CMDB classes where possible and introducing new classes only when telecom-specific representation is required.

Key components

- Network equipment and hardware: Physical and logical equipment, along with the components they contain, such as interface cards and ports.
- Interfaces and ports: Network interfaces and ports that enable connectivity between devices and services.

- Sites and locations: Logical and physical locations used to group and organize network infrastructure.
- Links, circuits, and connections: Physical and logical connections that describe how network elements communicate and exchange data.

Network equipment classes

The Telecom data model includes network equipment classes that represent specialized telecommunications hardware as children of the Telco Equipment (cmdb_ci_telco_equipment) class. The following table describes these classes.

Network equipment classes

Class	Description
Network Packet Broker (cmdb_ci_network_packet_broker)	A specialized device that sits between network TAPs or SPAN ports and security and monitoring tools. A network packet broker aggregates, filters, and distributes traffic. Each connected tool receives only the data it needs, which optimizes monitoring performance across high-speed, complex, and hybrid cloud networks. Example devices include the Iris Packet Broker IPB220 and IPB420, and APCON IntellaFlex XR monitoring switches. This class is a child of the Telco Equipment (cmdb_ci_telco_equipment) class.

Benefits

By providing a consistent way to model telecom infrastructure and relationships, the Telecom data model enables accurate network inventory, impact analysis, and operational insight. The shared structure supports integration across systems, improves data consistency, and simplifies analysis of network topology and dependencies. This common modeling approach helps teams manage complex networks more effectively and supports use cases such as service visibility, operational monitoring, and lifecycle management.

Domain-specific data models

Domain-specific data models build on the Telecom data model by refining and specializing its classes for particular network technologies. These models reuse the core structure of the Telecom data model while introducing domain-specific attributes and relationships as needed. The SD-WAN data model is one example of a domain-specific data model. It focuses on the subset of Telecom data model concepts that are relevant to SD-WAN environments. It applies them to represent SD-WAN topology, services, and operations while remaining consistent with the broader Telecom data model.

Related topics

[SD-WAN data model](#)

[Data model](#) 

SD-WAN data model

The ServiceNow AI Platform[®] uses a custom data model that defines how SD-WAN connectors discover and retrieve device information.

The high-level Telecom Data Model, part of the Common Service Data Model (CSDM), does not define how specific technologies such as SD-WAN should be modeled. To address this gap,

the SD-WAN data model provides a dedicated set of classes and relationships that define the structure of an SD-WAN network within the Configuration Management Database (CMDB).

The model documents comprehensive details about network assets, including configuration data, available ports, and bandwidth allocations across sites and services. This enables centralized management of SD-WAN repositories, streamlined provisioning and monitoring, and efficient resource allocation for building and maintaining network infrastructure.

The Discovery Builder framework supports the SD-WAN data model. You can use the framework to create and customize discovered objects that map directly to SD-WAN-specific classes and relationships, such as controllers, edge devices, tunnels, overlays, and underlays. Discovered data is mapped to the appropriate SD-WAN classes rather than to generic network CI types, supporting downstream assurance workflows such as fault management and performance monitoring."

ServiceNow AI Platform[®] SD-WAN data model architecture

The following table describes the classes used in the SD-WAN data model and the type of data each class stores.

SD-WAN data model classes

Class	Description
Group	Discovered organizations or administrative domains (ADOMs). Required
Company	Discovered organizations or administrative domains
Network site	Network site, including geographical information
Network service instance	Logical grouping of devices and their associated configurations
Network equipment models	Product model of the equipment
IP router or any other supported equipment classes	Hardware or network gear classification, used for storing information about network equipment
Key values	Additional attributes or fields related to a CI that aren't covered by predefined CI columns
Network interface	Network interface on a piece of equipment
Network adapter	Network adapter on a piece of equipment
IP address	Management IP address associated with a device
Per-device license	License assigned to a specific device. Optional
Location	Location of a network site

MPN data model

The Mobile Private Network (MPN) data model extends the Telecom data model to represent the physical hardware and virtual network functions that make up a mobile private network. It also represents the relationships that connect them.

A mobile private network contains a mix of physical and virtual objects. Physical objects include servers, firewalls, routers, switches, antennas, and radio units. Virtual objects include the network functions that run on those servers, such as the User Plane Function (UPF) and the Unified Data Management (UDM) function. These virtual objects also include their associated licenses. The MPN data model provides dedicated classes and relationships in the Configuration Management Database (CMDB) for these objects. This enables the physical and virtual components of an MPN to be modeled accurately and consistently. Accurate inventory and relationship data enables faster troubleshooting, supports impact analysis when faults occur, and provides the foundation for automated event and performance monitoring across MPN network elements.

The model spans two domains of the mobile network: the Radio Access Network (RAN) and the Mobile Core. The RAN connects user equipment to the network over the air interface. The Mobile Core provides packet routing, session management, authentication, and policy control. The model also defines upper-layer logical objects that group RAN and Core components into end-to-end network service instances and assign them to network sites.

Physical hardware classes

Physical hardware classes represent the tangible equipment deployed in the RAN and the Mobile Core. The following table describes the physical classes used in the MPN data model.

MPN physical hardware classes

Domain	Class	Description
RAN	Remote Radio Unit	Radio frequency unit, typically mounted close to the antenna, that transmits and receives signals over the air interface.
RAN	Antenna	Antenna assembly used to radiate and receive signals to and from user equipment.
RAN	Baseband Unit	Equipment that processes baseband signals and connects radio units to the transport network.
RAN	Baseband Control and Interface Cards	Cards inside a Baseband Unit that provide control plane processing and interfaces to radio and transport networks.
RAN	Server	Compute platform that hosts virtualized 5G Distributed Units (DUs) and Centralized Units (CUs).
RAN	BBU chassis (cldb_ci_chassis)	Physical enclosure that houses one or more Baseband Units and their interface cards.
RAN	Antenna Tower (cldb_ci_antenna_structure)	Physical tower structure that supports one or more antennas and their associated radio units.
RAN	Indoor Radio Unit (IRU) (cldb_ci_indoor_radio_unit)	Radio frequency unit deployed indoors to provide cellular coverage inside buildings.
RAN	Radio Dot/ Distributed IRU (cldb_ci_distributed_indoor_radio_unit)	Distributed indoor radio unit that extends indoor coverage from a central Indoor Radio Unit.
RAN and Core	MPN box/ cabinet (cldb_ci_container_chassis)	Outdoor or indoor cabinet that houses the BBU chassis and other RAN equipment at the cell site.
RAN	Cable (cldb_ci_cable)	Physical cable connection between MPN hardware components.
Core	Firewall	Security appliance that enforces traffic policies at the boundaries of the Mobile Core network.
Core	IP Router	Layer 3 device that forwards IP traffic between network segments in the Mobile Core.
Core	IP Switch	Layer 2 device that forwards traffic between hosts and routers in the Mobile Core.
Core	Time sync source	Timing source that provides accurate clock synchronization to the mobile network.

MPN physical hardware classes (continued)

Domain	Class	Description
		Common sources include GPS receivers and other PTP grandmaster clocks.
Core	Server	Compute platform that hosts 4G and 5G virtualized network functions (VNFs).
RAN and Core	Network Adapter	Physical network interface hardware (cmdb_ci_network_adapter) installed in a Server, Router, Switch, or Firewall.
RAN and Core	Network Interface	Logical interface (cmdb_ci_network_interface) exposed by a Network Adapter and used to terminate IP connectivity.
RAN and Core	IP Address	IP address (cmdb_ci_ip_address) assigned to a Network Adapter or Network Interface and used for routing traffic across the MPN.

Virtual network function classes

Virtual network function classes represent the software components that run on the MPN servers. RAN virtual functions handle radio-side processing, while Core virtual functions provide subscriber, session, and policy services. The following table describes the virtual classes used in the MPN data model.

MPN virtual network function classes

Domain	Class	Description
RAN	Mobile Cell	Logical representation of a radio cell served by a Baseband Unit and its associated radio units.
RAN	vDU (virtual Distributed Unit)	5G virtual service that performs lower-layer baseband processing and connects to one or more radio units.
RAN	vCU-CP (virtual Centralized Unit, Control Plane)	5G virtual service (cmdb_ci_5g_cu_control_plane_network_function) that performs higher-layer control plane processing and signaling between the RAN and the Mobile Core.
RAN	vCU-UP (virtual Centralized Unit, User Plane)	5G virtual service (cmdb_ci_5g_cu_user_plane_network_function) that performs higher-layer user plane processing and forwards subscriber traffic to the Mobile Core.
Core	4G and 5G VNFs	Virtual services that represent Mobile Core network functions. The 5G Core includes AMF (Access and Mobility Management), AUSF (Authentication Server), SMF (Session Management), and UPF (User Plane). It also includes UDM (Unified Data Management), PCF (Policy Control), NRF (Network

MPN virtual network function classes (continued)

Domain	Class	Description
		Repository), NEF (Network Exposure), and NSSF (Network Slice Selection). The 4G EPC (Evolved Packet Core) includes MME (Mobility Management Entity), HSS (Home Subscriber Server), and PCRF (Policy and Charging Rules). It also includes AAA (Authentication, Authorization, and Accounting), SGW (Serving Gateway), and PGW (Packet Data Network Gateway).
Core	VNF license	License assigned to a specific virtual network function. Tracks entitlement and consumption per VNF.
RAN	Virtual Machine	Compute instance (cmdb_ci_vm_object) that runs on a RAN Server and hosts a vDU or vCU virtual service.
Core	Virtual Machine	Compute instance (cmdb_ci_vm_object) that runs on a Core Server and hosts a 4G or 5G VNF.
Core	IMS network functions	Shared 3GPP network functions (cmdb_ci_network_function_instance) that provide IP Multimedia Subsystem (IMS) services for voice and multimedia sessions across 4G and 5G.
Core	MPN O&M functions	Operations and management applications (cmdb_ci_app) such as ProbeApp and Prometheus that monitor and manage the MPN.

MPN-specific classes

The MPN data model introduces custom classes that represent assets and resources unique to mobile private networks, including radio spectrum, network slices, data network names, and SIM provisioning. The following table describes the MPN-specific classes used in the MPN data model.

MPN-specific classes

Class	Description
Spectrum	Frequency band (u_mpn_spectrum) allocated to the mobile private network and consumed by Mobile Cells for radio transmission.
SIM range	Block of subscriber identities (u_mpn_sim_range) provisioned to the MPN for allocation to SIM cards.

CI attributes

For each CI class in the MPN data model, the pull connector captures and stores relevant data attributes in the CMDB. This provides detailed information for management and troubleshooting. The following table lists the key attributes captured per CI type.

Key attributes captured per MPN CI type

CI type	Attribute	Description
---------	-----------	-------------

Upper-layer logical classes

Upper-layer logical classes group the physical and virtual components of an MPN into end-to-end service constructs and assign them to a geographic site. The following table describes the upper-layer logical classes used in the MPN data model.

MPN upper-layer classes

Class	Description
Network service instance	End-to-end logical service that represents an entire MPN deployment, including its RAN and Mobile Core components.
Mobile Core network service instance	Logical service that groups the Mobile Core VNFs, supporting hardware, and licenses that deliver core packet, session, and policy functions.
RAN network service instance	Logical service that groups the RAN virtual functions, baseband and radio hardware, and Mobile Cells that deliver radio access.
Network Site	Physical or logical location to which the MPN network service instance and its components are assigned.
Slice	Logical network slice (u_mpn_slice) that partitions the MPN to deliver tailored connectivity for specific use cases or tenants.
DNN	Data Network Name (u_mpn_slice_dnn) that identifies the external data network reachable through a slice.
SIM card	Individual subscriber identity module (cmdb_ci_sim_card) assigned from a SIM range and consumed by user equipment connecting to a Mobile Cell.

Relationships between physical and virtual objects

The MPN data model uses CMDB relationships to express how physical equipment hosts virtual functions and how virtual functions consume each other. These relationships also show how individual components roll up into network service instances and sites. They make dependencies and connectivity visible in the CMDB and support downstream assurance workflows such as fault management, performance monitoring, and impact analysis.

MPN CMDB relationships

Parent CI	Relationship type	Child CI	CMDB type
RAN Server	Contains	Virtual Machine (vDU / vCU)	

MPN CMDB relationships (continued)

Parent CI	Relationship type	Child CI	CMDB type
Core Server	Contains	Virtual Machine (4G / 5G VNF)	
Baseband Unit	Contains	Baseband Control and Interface Cards, Remote Radio Units, Antennas	
Baseband Unit and radio units	Provides	Mobile Cell	
VNF license	Is consumed by	VNF	
MPN network service instance	Contains	MPN RAN network service instance, MPN Mobile Core network service instance	
MPN Network Site	Contains	MPN network service instance and components	

Related topics

[Configure elastic connectors for MPN alarm collection](#)

[Configure elastic event pull connectors for MPN](#)

[MPN Formula Engine processing](#)

[Metric-to-CI binding](#)

Telecom Visibility

Telecom Visibility extends the discovery and reconciliation capabilities of ITOM Visibility to meet the specific needs of telecom service providers. It enables accurate, telecom-aware network inventory by discovering real-time network data, reconciling mismatches, and maintaining a consistent telecom model in your CMDB and TNI (Telecommunications Network Inventory).

Discover physical and logical network components across telecom domains and reconcile discrepancies between discovery and inventory data. Maintain a telecom-specific CMDB structure to support advanced automation and assurance use cases.

Telecom Visibility overview

Telecom Visibility is a unified solution for discovering and managing telecom network resources. It leverages proven discovery technologies such as Horizontal Discovery, Telecommunications Discovery Patterns, and Service Graph Connectors (SGC). These technologies bring network data from various systems (for example, CLI/SNMP-based standalone devices or API-driven EMS/NMS/controllers) into your CMDB.

Telecom Visibility helps communication service providers (CSPs) and telecom operators:

- Discover and maintain telecom-specific network inventory
- Reconcile mismatches between planned/design inventory and actual network state
- Support Autonomous Network Operations (ANO), service fulfillment, and assurance workflows

Key Capabilities

Key capabilities of TSOM

Capability	Description
Telecom Discovery with Patterns	Discover standalone network functions (xNFs) such as routers and switches using SNMP and CLI protocols via predefined Telecom Discovery Patterns.
Service Graph Connector Integration	Bring inventory and configuration data from external EMS/NMS/controllers into CMDB using northbound APIs.
Discrepancy Identification & Reconciliation	Detect and resolve mismatches between inventoried and discovered data (for example, missing CIs, model mismatches, stale entries).
Telecom-Centric Data Modeling	Maintain structured and hierarchical telecom data including interface cards, slots, subslots, VLANs, and ports.
CMDB Integration	Use CMDB 360 and Discovery Admin Workspace to monitor discovery sources, attribute changes, and reconciliation outcomes.

Applications and Plugins

- 1. Telecommunications Horizontal Discovery Patterns:** Discovers telecom devices (xNFs) such as routers and switches using SNMP/CLI. Enables accurate inventory updates by scanning live network infrastructure. Supports vendor-specific patterns including:
 - Cisco (ASR1K, 7613, Nexus 9000, Nexus 3548)
 - Juniper (MX80, MX104, MX240, MX460)
- 2. Service Graph Connectors (SGC s) for TSOM:** Prebuilt integrations with EMS/NMS/controllers to import inventory data via APIs. Supports:
 - Scheduled or on-demand discovery
 - Multi-instance configurations
 - Integration with systems like Nokia Altiplano (GPON networks)
- 3. Telecom Visibility Plugin:** Provides shared logic and foundational components for both discovery and discrepancy reconciliation. Ensures:
 - Consistent behavior across TSOM workflows
 - Reusable components for reconciliation and remediation
- 4. CMDB CI Class Models for Telecom (Required version: 1.82 or later):** Extends the CMDB to include telecom-specific classes and relationships, including:
 - Interface cards
 - Slots and subslots
 - Network interfaces

Telecom Visibility architecture

Telecom Visibility combines direct and indirect discovery methods:

- Direct Discovery using SNMP/CLI (for example, standalone routers/switches)
- Indirect Discovery using APIs from EMS/NMS/controllers (for example, Nokia Altiplano)

These feed into the CMDB and TNI, which then support decision-making, automation, and network operations. The following infographic demonstrates the architecture of Telecom Visibility.

Use Cases

Telecom Visibility use cases

Use Case	Outcome
Autonomous Network Operations	Reliable network inventory supports AI/ML-driven automation.
Service Fulfillment	Accurate CMDB/TNI ensures order accuracy and rapid provisioning.
Service Assurance	Reconciled inventory reduces troubleshooting time and increases uptime.
Audit & Compliance	Discovery Admin Workspace and CMDB 360 support audit readiness and traceability.

Key personas

Telecom Visibility is used by:

- Network and infrastructure teams: Discover and manage real-time network resource data across vendors and technologies.
- Inventory and fulfillment teams: Maintain an accurate and complete CMDB/TNI to support service fulfillment, order management, and provisioning.
- Platform administrators: Configure plugins, manage discovery sources, validate reconciliation rules, and monitor discovery performance.

Monitoring tools

Discovery Admin Workspace - Central console to:

- Monitor discovery tasks
- Validate reconciliation processes
- Tune schedules and view diagnostics

CMDB 360 - Visualize:

- Attribute-level update history
- Source attribution
- Reconciliation rules in action

Key benefits

Key benefits of Telecom Visibility

Feature	Benefit
CLI/SNMP Discovery	Expands network coverage without relying on external APIs.
Nokia Altiplano Integration	Supports GPON network discovery across multiple controller instances.
Discrepancy Reconciliation	Keeps inventory clean, actionable, and aligned with real network state.

Related topics

[Configure Telecom Visibility](#)

[Discovery Admin Workspace](#) 

Telecom Visibility vs. ITOM Visibility

As telecom networks evolve and become more hybrid and complex, visibility into infrastructure is more critical than ever. To meet the unique needs of different environments, ServiceNow AI Platform offers two purpose-built visibility solutions: ITOM Visibility and Telecom Visibility.

While they're built on the same foundation, each is tailored to serve a different world—ITOM for traditional IT infrastructure and TSOM for telecom networks.

Both ITOM Visibility and Telecom Visibility are built on the same underlying capabilities of the ServiceNow Discovery engine, Identification and Reconciliation Engine (IRE), and CMDB. They both provide:

- Agent-based and agentless discovery.
- Horizontal and pattern-based discovery.
- Reconciliation of discovered data into CMDB.
- Dependency mapping and CI relationship creation.
- Integration with the Discovery Admin Workspace and CMDB 360.

Despite this shared architecture, the scope, use cases, and data models they serve are different.

Key differences between TSOM and ITOM Visibility

Telecom Visibility vs. ITOM Visibility

Feature / Focus Area	ITOM Visibility	Telecom Visibility
Target Environment	Traditional IT infrastructure (servers, applications, databases, cloud resources)	Telecom infrastructure (xNFs, network elements, EMS/NMS controllers)
Discovery Method	Horizontal Discovery with IT patterns	Horizontal Discovery with Telecommunications Discovery Patterns (SNMP/CLI) and Service Graph Connectors for telecom

Telecom Visibility vs. ITOM Visibility (continued)

Feature / Focus Area	ITOM Visibility	Telecom Visibility
CMDB Model	ITOM CMDB classes (for example, Windows Server, Application, Network Adapter)	Telecom-aware CMDB classes and Telecom Network Inventory (TNI) (for example, Interface Cards, Slots, LAGs, Subslots, VLANs)
Plugins required	com.snc.discovery and com.snc.itom.visibility	sn_tsom_core, sn_tsom_patterns, and telecom-specific SGC plugins (for example, sn_sgc_altiplano_connector)
Use Case Focus	Application dependency mapping, service modeling, cloud infrastructure discovery	Telecom network inventory discovery, reconciliation, autonomous network operations
Discrepancy Handling	General IRE reconciliation rules	Telecom-specific discrepancy identification and reconciliation (for example, hierarchy mismatches, attribute-level conflicts)
Vendor Data Ingestion	Primarily via discovery patterns	Strong emphasis on northbound API integrations using SGCs (EMS/NMS/controllers)
Network Types Supported	Enterprise networks, datacenters, cloud	Multi-vendor, multi-domain telecom networks (RAN, Core, Transport, Access)

IT Discovery vs Telecom Discovery

The following table helps you understand the differences between ITOM Visibility and TSOM Discovery.

ITOM Discovery	Telecom Discovery
Flat, simple or no hierarchy	Hierarchical, alignment to telecom models
Basic attributes	Advanced attributes
The network is the trusted source of truth	Inventory/CMDB design is the trusted source of truth
Simple CI identification and reconciliation	Complex telecom CI identification and reconciliation
Populates the CMDB with the CIs found in the network (everything found is written into the CMDB)	Validates that the network implementation is in sync with CMDB/TNI records as designed/planned
Never miss a new CI for monitoring	Start monitoring new CIs after they have been validated

Key benefits

Use ITOM Visibility when:

- Discover IT infrastructure components (for example, Windows servers, cloud VMs, databases, load balancers).
- Primary goals include service mapping, operational resilience, or cloud optimization.
- Focus on ITSM, ITOM, or DevOps use cases.

Use Telecom Visibility when:

- Discovering telecom network infrastructure, including devices not managed by traditional IT systems.
- Dealing with telecom-specific network hierarchies like cards, ports, subslots, and LAGs.
- Relying on EMS/NMS/controllers as authoritative data sources.
- Need discrepancy detection and reconciliation tailored for telecom inventory models.
- Aligning with TM Forum standards, supporting autonomous network operations, or enabling closed-loop assurance.

Examples

Use Case	ITOM Visibility	Telecom Visibility
Discover a fleet of virtual machines in AWS	Yes	No
Ingest router and switch data from an EMS using APIs	No	Yes
Identify application-to-application dependencies	Yes	No
Detect and reconcile mismatches in telecom card hierarchies	No	Yes

While ITOM Visibility and Telecom Visibility both serve to populate and maintain an accurate CMDB, they are optimized for different domains. ITOM Visibility is geared toward enterprise IT environments, while Telecom Visibility is tailored for the specialized needs of telecom infrastructure discovery, discrepancy management, and inventory reconciliation.

Choosing the right visibility solution—or using both in tandem—confirms that you maintain trusted, domain-specific operational visibility across both IT and telecom landscapes.

Telecom Discovery

ServiceNow AI Platform Telecom Discovery provides visibility into your telecom network infrastructure by extending ITOM Visibility to support telecom-specific use cases.

Built for communication service providers (CSPs), this solution enables the discovery and mapping of network elements across multi-vendor environments using standardized protocols and integration with network management systems.

Telecom Discovery plugins work with Service Graph Connectors and Discovery Patterns to automatically populate and maintain accurate records of your telecom resources in the CMDB. This provides a unified view of both IT and network infrastructure.

With Telecom Discovery, you can:

- Discover physical and logical telecom network resources across domains and vendors.
- Integrate with Element Management Systems (EMS), Network Management Systems (NMS), and software-defined networking (SDN) controllers.
- Automatically populate and update CMDB and TNI records based on real-time network data.
- Discover standalone xNFs using SNMP and Command-Line Interface (CLI).
- Enrich CMDB data using Service Graph Connectors and specialized discovery patterns.
- Identify discrepancies between discovered network data and inventory records.
- Support automation use cases through consistent and accurate infrastructure visibility.

Integration with ITOM Visibility

Telecom Discovery complements existing ITOM Visibility features. You can:

- Use Horizontal Discovery and ITOM capabilities alongside TSOM plugins.
- Maintain consistent discovery practices for IT and telecom resources.
- Use the same CMDB data model to manage cross-domain service visibility.

This integration supports unified asset management, faster issue resolution, and streamlined operations across both IT and network domains.

Customization with low-code and no-code tools

Telecom Discovery provides design tools to extend discovery logic without writing code. You can:

- Build or modify custom Service Graph Connectors.
- Extend Telecommunications Discovery Patterns to match vendor-specific requirements.
- Accelerate onboarding of new device types and network domains.

This approach helps CSPs adapt and reduce onboarding time when expanding their discovery coverage.

Key components

- Telecommunication Discovery Patterns (sn_tsom_patterns): Provides patterns for SNMP-based discovery of standalone routers, switches, and xNFs. Includes Cisco and Juniper-specific discovery logic.
- Service Graph Connector for Nokia Altiplano (sn_sgc_altiplano_connector): Enables data collection from the Nokia Altiplano Access SDN Controller via REST APIs.
- Telecom Core (sn_tsom_core): Delivers foundational capabilities such as discrepancy identification, remediation logic, and shared telecom discovery features.

Related topics

[Configure Telecom Visibility](#)

[Use Telecom Discovery patterns](#)

Direct Discovery using Discovery Patterns

The Telecommunications Discovery Patterns plugin (also known as TSOM Patterns) extends ServiceNow AI Platform Telecom Discovery. It supports direct discovery of standalone network elements—such as routers and switches—without relying on traditional network management

systems. These patterns enable Communication Service Providers (CSPs) to identify and map multivendor xNFs using protocols like SNMP and CLI.

Telecommunications Discovery Patterns provide a powerful, pattern-based approach to discover and manage telecom network resources directly from network elements. These patterns are especially valuable for discovering standalone xNFs (such as routers and switches) that aren't managed through traditional EMS/NMS systems.

This capability enhances visibility across telecom infrastructures and facilitates both vendor-neutral and vendor-specific device data is captured and reflected in the CMDB and Telecom Network Inventory (TNI), following the TNI data model.

Note: Telecommunications Discovery Patterns are part of the TSOM Visibility subscription and are available as a customer-visible plugin from the ServiceNow Store.

Key capabilities

- **Direct Network Element Discovery**
 - Use SNMP and CLI to communicate directly with physical network elements.
 - Discover physical network inventory such as interface cards, slots, ports, and devices without relying on EMS/NMS mediation.
- **Automated CI Mapping and CMDB Integration**
 - Discovered data is automatically mapped to telecom-aligned configuration item (CI) classes.
 - Integrates directly with the Identification and Reconciliation Engine (IRE) to promote accurate, non-duplicate CI records in the CMDB and Telecom Network Inventory (TNI).
- **CMDB Accuracy and Compliance**
 - Patterns trigger CMDB Compliance Certification Audits to detect mismatches or outdated records.
 - Supports discrepancy identification and remediation, helping maintain data integrity.
- **Low-Code Pattern Customization:** Extend or customize discovery patterns using low-code pattern designers to meet your specific network architecture and business needs.

How it works

Telecommunications Discovery Patterns are executed through Horizontal Discovery using the Nebula Discovery Language (NDL). A pattern is a sequence of steps that:

1. Establishes a connection to a target device.
2. Executes commands using SNMP, CLI, or both.
3. Extracts CI attributes and relationships.
4. Sends results to IRE for reconciliation and insertion into the CMDB and TNI.

The discovery logic follows the TNI data model, which restructures complex CI relationships (e.g., card-on-card scenarios) to align with telecom modeling best practices. For example, If a child card is discovered inside a parent card, Telecom Discovery synthesizes a subslot to insert the child, avoiding an invalid card-on-card configuration.

For every discovered network interface, both logical (such as VLANs) and physical (such as Gigabit Ethernet ports), Telecom Discovery sets the Equipment reference field to the parent equipment. This applies even when the underlying port resides on a card or on a card within a card. The interface's relationship to that card is captured separately in the CI relationship

hierarchy. The Equipment reference field consistently resolves to the parent equipment. This alignment follows the TNI data model.

Discovery payload cleanup before reconciliation

Before sending discovered data to the Identification and Reconciliation Engine (IRE), Telecom Discovery Patterns clean up two common data issues:

- **Duplicate interface cards.** If two interface cards report the same serial number, Telecom Discovery keeps only the card with more discovered sub-components (such as ports or subslots). It removes the other card, along with anything discovered underneath it. If both cards have the same number of sub-components, the pattern keeps whichever card it discovered first.
- **Missing containment relationships.** If a discovered network interface, slot, or subslot isn't already linked to its parent router or switch, Telecom Discovery automatically adds that containment relationship. This happens before the data reaches IRE.

Both of these actions are recorded in the discovery log, so you can confirm whether a device's structure was adjusted during discovery.

Default life cycle values for discovered CIs

When Telecom Discovery Patterns create or update a CI, they set default life cycle values. The **Life cycle stage** is set to **Operational** and the **Life cycle stage status** is set to **In Use**. Your instance customizations can return different values. Discovery doesn't overwrite either field if it's already set.

If you use life cycle fields to distinguish newly discovered, not-yet-provisioned equipment from equipment already in production, confirm this default matches your operational process before relying on it.

Architecture using Horizontal Discovery and Telecommunications Discovery Patterns

The following infographic is an example of the implementation for standalone SNMP or/and CLI

xNFs.

Horizontal Discovery Application

The Horizontal Discovery application in ServiceNow is a versatile and highly scalable discovery engine. It operates effectively across network, IT, and cloud environments, collecting data across multiple layers to provide a holistic view of the infrastructure.

For more information, see [Horizontal discovery process flow with patterns](#) .

Supported discovery patterns

A pattern is a sequence of commands to detect attributes of a configuration item (CI) and its outbound connections. Telecom Discovery provides a set of preconfigured Patterns that cover a wide range of network elements. ServiceNow offers several predefined TSOM patterns, including:

- Telecom Router Pattern – Generic SNMP-based router discovery.
- Telecom Cisco 7613 Router Pattern – For Cisco 7613 routers using SNMP.
- Telecom Juniper MX SSH Router Pattern – SNMP + CLI discovery of Juniper MX routers.
- Telecom Cisco Switch Pattern – SNMP-based Cisco switch discovery.
- Telecom Switch Pattern – SNMP-based generic switch discovery.

TNI entity creation

If your instance has the **Telecom Network Inventory (TNI)** plugin installed:

- Every discovered CI automatically results in a TNI entity record.
- The IRE payload includes mapping to both `cmdb_ci` and `tni_entity` tables.

This promotes seamless alignment between operational and inventory systems, which is essential for order fulfillment, assurance, and network planning.

MID Server

MID Server is a Java application that runs as a Windows service or UNIX daemon on a server within your local network. The ServiceNow[®] MID Server facilitates communication and data transfer between a ServiceNow instance and external applications, data sources, and services.

For more information, see [MID Server](#) .

Identification & Reconciliation Engine (IRE)

IRE offers a centralized framework for identifying and reconciling data from multiple sources. It verifies the integrity of the CMDB and some non-CMDB tables when various data sources are used to create or update CI records.

Related topics

[Install Horizontal Discovery and set up Discovery Patterns](#)

Telecom Router Pattern

The ServiceNow[®] Telecom Discovery application uses the Telecom Router discovery pattern to find SNMP-based routers in the network. Discovering some of these resources requires updating the Telecommunications Discovery Patterns (TSOM Patterns) from the ServiceNow[®] Store.

Telecom Discovery uses the Telecommunications Discovery Patterns to run Horizontal Discovery. This Telecommunications Discovery Pattern uses a set of SNMP requests to find, classify, and discover network elements.

Telecom Router pattern is part of the Telecommunications Discovery Patterns application (`sn_tsom_patterns`), which is part of TSOM Visibility.

Request apps on the Store

Visit the [ServiceNow Store](#) to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Prerequisites

- Subscription to TSOM.
- Confirm that your network router devices have SNMP access.
- On the ServiceNow instance, configure SNMP credentials. For more information, see [SNMP support for Discovery](#).
- For setting up TSOM Patterns, see [Install Horizontal Discovery and set up Discovery Patterns](#).

Impacted CMDB CIs and CI Relationships (Physical Layer)

CIs	CI Relationships
IP Router CI	<p>IP Router Device is represented by the IP Router CI.</p> <p>Table name: cmdb_ci_ip_router</p> <p>IP Router CI contains Slots or Network Interfaces.</p>
Slot CI	<p>Slots are represented by the Slot CI.</p> <p>Table name: cmdb_ci_container_slot</p> <p>Slot is contained by the IP Router.</p> <p>Slot CI contains the Interface Card.</p>
Subslot CI	<p>Subslots are represented by the Subslot CI.</p> <p>Table name: cmdb_ci_container_subslot</p> <p>Subslot is contained by Interface Card CI.</p> <p>Subslot CI contains the Interface Card CI.</p>
Interface Card CI	<p>Different types of cards are represented by the Interface Card CI.</p> <p>Fan and Power Supply Units are also represented by the Card CI.</p> <p>Table name: cmdb_ci_interface_card</p> <p>Interface Cards are contained by Slots or Subslots.</p>

CIs	CI Relationships
	Interface Cards can contain Network Interface or Subslots.
Network Interface CI	<p>Any type of Network Interface is represented by the Network Interface CI.</p> <p>Table name: cmdb_ci_ni_interface</p> <p>Network Interface is contained by Interface Card, IP Router.</p>

Telecom Cisco 7613 Router Pattern

The ServiceNow® Telecom Discovery application uses the Telecom Cisco 7613 Router discovery pattern to find SNMP-based Cisco 7613 in the network. Discovering some of these resources requires updating the Telecommunications Discovery Patterns (TSOM Patterns) from the ServiceNow® Store.

Telecom Discovery uses the Telecommunications Discovery Patterns to run Horizontal Discovery. This Telecommunications Discovery Pattern uses a set of SNMP requests to find, classify, and discover network elements.

Telecom Cisco 7613 Router Pattern is part of the Telecommunications Discovery Patterns application (sn_tsom_patterns), which is part of TSOM Visibility.

Request apps on the Store

Visit the [ServiceNow Store](#) to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Prerequisites

- Subscription to TSOM.
- Verify that your network router devices have SNMP access.
- On the ServiceNow instance, configure SNMP credentials. For more information, see [SNMP support for Discovery](#).
- For setting up TSOM Patterns, see [Install Horizontal Discovery and set up Discovery Patterns](#).

Impacted CMDB CIs and CI Relationships (Physical Layer)

CIs	CI Relationships
IP Router CI	<p>IP Router Device is represented by the IP Router CI.</p> <p>Table name: cmdb_ci_ip_router</p> <p>IP Router CI contains Slots or Network Interfaces.</p>

CIs	CI Relationships
Slot CI	<p>Slots are represented by the Slot CI.</p> <p>Table name: cmdb_ci_container_slot</p> <p>Slot is contained by the IP Router.</p> <p>Slot CI contains the Interface Card.</p>
Subslot CI	<p>Subslots are represented by the Subslot CI.</p> <p>Table name: cmdb_ci_container_subslot</p> <p>Subslot is contained by Interface Card CI.</p> <p>Subslot CI contains the Interface Card CI.</p>
Interface Card CI	<p>Different types of cards are represented by the Interface Card CI.</p> <p>Fan and Power Supply Units are also represented by the Card CI.</p> <p>Table name: cmdb_ci_interface_card</p> <p>Interface Cards are contained by Slots or Subslots.</p> <p>Interface Cards can contain Network Interface or Subslots.</p>
Network Interface CI	<p>Any type of Network Interface is represented by the Network Interface CI.</p> <p>Table name: cmdb_ci_ni_interface</p> <p>Network Interface is contained by Interface Card, IP Router.</p>

Telecom Juniper MX SSH Router Pattern

The ServiceNow® Telecom Discovery application uses the Telecom Juniper MX SSHRouter discovery pattern to find SNMP and CLI -based Juniper MX Series routers in the network. Discovering some of these resources requires updating the Telecommunications Discovery Patterns (TSOM Patterns) from the ServiceNow® store.

Telecom Discovery uses the Telecommunications Discovery Patterns to run Horizontal Discovery. This Telecommunications Discovery Pattern uses a set of SNMP requests to find and classify CLI over SSH to discover network elements.

Telecom Juniper MX SSH Router pattern is part of the Telecommunications Discovery Patterns application (sn_tsom_patterns), which is part of TSOM Visibility.

Request apps on the Store

Visit the [ServiceNow Store](#) to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Prerequisites

- Subscription to TSOM.
- Verify that your network router devices have SNMP access.
- On the ServiceNow instance, configure SNMP credentials. For more information, see [SNMP support for Discovery](#).
- For setting up TSOM Patterns, see [Install Horizontal Discovery and set up Discovery Patterns](#).

Impacted CMDB CIs and CI Relationships (Physical Layer)

CIs	CI Relationships
IP Router CI	<p>IP Router Device is represented by the IP Router CI.</p> <p>Table name: cmdb_ci_ip_router</p> <p>IP Router CI contains Slots or Network Interfaces.</p>
Slot CI	<p>Slots are represented by the Slot CI.</p> <p>Table name: cmdb_ci_container_slot</p> <p>Slot is contained by the IP Router.</p> <p>Slot CI contains the Interface Card.</p>
Subslot CI	<p>Subslots are represented by the Subslot CI.</p> <p>Table name: cmdb_ci_container_subslot</p> <p>Subslot is contained by Interface Card CI.</p> <p>Subslot CI contains the Interface Card CI.</p>
Interface Card CI	<p>Different types of cards are represented by the Interface Card CI.</p> <p>Fan and Power Supply Units are also represented by the Card CI.</p> <p>Table name: cmdb_ci_interface_card</p> <p>Interface Cards are contained by Slots or Subslots.</p>

CIs	CI Relationships
	Interface Cards can contain Network Interface or Subslots.
Network Interface CI	<p>Any type of Network Interface is represented by the Network Interface CI.</p> <p>Table name: cmdb_ci_ni_interface</p> <p>Network Interface is contained by Interface Card, IP Router.</p>

Telecom Cisco Switch Pattern

The ServiceNow® Telecom Discovery application uses the Telecom Cisco Switch discovery pattern to find SNMP-based Cisco switches in the network. Discovering some of these resources requires updating the Telecommunications Discovery Patterns (TSOM Patterns) from the ServiceNow® Store.

Telecom Discovery uses the Telecommunications Discovery Patterns to run Horizontal Discovery. This Telecommunications Discovery Pattern uses a set of SNMP requests to find, classify, and discover network elements.

Telecom Cisco Switch pattern is part of the Telecommunications Discovery Patterns application (sn_tsom_patterns), which is part of TSOM Visibility.

Request apps on the Store

Visit the [ServiceNow Store](#) to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Prerequisites

- Subscription to TSOM.
- Confirm that your network router devices have SNMP access.
- On the ServiceNow® instance, configure SNMP credentials. For more information, see [SNMP support for Discovery](#).
- For setting up TSOM Patterns, see [Install Horizontal Discovery and set up Discovery Patterns](#).

Impacted CMDB CIs and CI Relationships (Physical Layer)

CIs	CI Relationships
IP Switch CI	<p>IP Switch Device is represented by the IP Switch CI.</p> <p>Table name: cmdb_ci_ip_switch</p> <p>IP Switch CI contain Slots or Network Interfaces.</p>

CI	CI Relationships
Slot CI	<p>Slots are represented by the Slot CI.</p> <p>Table name: cmdb_ci_container_slot</p> <p>Slot is contained by the IP Switch.</p> <p>Slot CI contains the Interface Card.</p>
Subslot CI	<p>Subslots are represented by the Subslot CI.</p> <p>Table name: cmdb_ci_container_subslot</p> <p>Subslot is contained by Interface Card CI.</p> <p>Subslot CI contains the Interface Card CI.</p>
Interface Card CI	<p>Different types of cards are represented by the Interface Card CI.</p> <p>Fan and Power Supply Units are also represented by the Card CI.</p> <p>Table name: cmdb_ci_interface_card</p> <p>Interface Cards are contained by Slots or Subslots.</p> <p>Interface Cards can contain Network Interface or Subslots.</p>
Network Interface CI	<p>Any type of Network Interface is represented by the Network Interface CI.</p> <p>Table name: cmdb_ci_ni_interface</p> <p>Network Interface is contained by Interface Card, IP Switch.</p>

Telecom Switch Pattern

The ServiceNow® Telecom Discovery application uses the Telecom Switch discovery pattern to find SNMP-based Telecom switches in the network. Discovering some of these resources requires updating the Telecommunications Discovery Patterns (TSOM Patterns) from the ServiceNow® Store.

Telecom Discovery uses the Telecommunications Discovery Patterns to run Horizontal Discovery. This Telecommunications Discovery Pattern uses a set of SNMP requests to find, classify, and discover network elements.

Telecom Switch pattern is part of the Telecommunications Discovery Patterns application (sn_tsom_patterns), which is part of TSOM Visibility.

Request apps on the Store

Visit the [ServiceNow Store](#) to view all the available apps, and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Prerequisites

- Subscription to TSOM.
- Confirm that your network router devices have SNMP access.
- On the ServiceNow® instance, configure SNMP credentials. For more information, see [SNMP support for Discovery](#).
- For setting up TSOM Patterns, see [Install Horizontal Discovery and set up Discovery Patterns](#).

Impacted CMDB CIs and CI Relationships (Physical Layer)

CIs	CI Relationships
IP Switch CI	<p>IP Switch Device is represented by the IP Switch CI.</p> <p>Table name: cmdb_ci_ip_switch</p> <p>IP Switch CI contain Slots or Network Interfaces.</p>
Slot CI	<p>Slots are represented by the Slot CI.</p> <p>Table name: cmdb_ci_container_slot</p> <p>Slot is contained by the IP Switch.</p> <p>Slot CI contains the Interface Card.</p>
Subslot CI	<p>Subslots are represented by the Subslot CI.</p> <p>Table name: cmdb_ci_container_subslot</p> <p>Subslot is contained by Interface Card CI.</p> <p>Subslot CI contains the Interface Card CI.</p>
Interface Card CI	<p>Different types of cards are represented by the Interface Card CI.</p> <p>Fan and Power Supply Units are also represented by the Card CI.</p> <p>Table name: cmdb_ci_interface_card</p> <p>Interface Cards are contained by Slots or Subslots.</p>

CI	CI Relationships
	Interface Cards can contain Network Interface or Subslots.
Network Interface CI	<p>Any type of Network Interface is represented by the Network Interface CI.</p> <p>Table name: cmdb_ci_ni_interface</p> <p>Network Interface is contained by Interface Card, IP Switch.</p>

Indirect Discovery using Service Graph Connectors

Telecom Discovery using Service Graph Connectors (SGC s) integrates network infrastructure data from external management systems into the Configuration Management Database (CMDB).

This approach helps Communication Service Providers (CSPs) maintain a current and accurate view of their multivendor telecom network resources, services, and configurations.

By using predefined connectors and robust data transformation tools, you can unify your inventory across domains and confirm that the CMDB and Telecom Network Inventory (TNI) reflect real-time network insights aligned with your telecom data model.

Note: Telecom Service Graph Connectors are part of the TSOM Visibility subscription and extend the capabilities of the standard Service Graph Connector framework.

Service Graph Connector overview

With Telecom SGC, you can:

- Ingest data from EMS/NMS/Controllers using northbound REST APIs.
- Automatically populate and update CMDB and TNI records with enriched, telecom-aligned data.
- Reconcile incoming data with existing CI records using the Identification and Reconciliation Engine (IRE).
- Generate TNI entity records automatically when the TNI plugin is installed.
- Support discrepancy detection and remediation as part of Telecom Discrepancy Identification & Reconciliation.
- Use low-code tools to configure, test, and manage your integrations end-to-end.

Architecture overview

Telecom Service Graph Connectors rely on a modular and scalable architecture:

Key components of Service Graph Connector

Component	Role
Service Graph Connector	Defines the integration logic to extract and stage data from EMS/NMS systems (for example, Nokia Altiplano or Nokia NSP).

Key components of Service Graph Connector (continued)

Component	Role
MID Server	Acts as a secure bridge between your ServiceNow instance and the external network system. For more information, see IntegrationHub ETL .
IntegrationHub ETL (3.2)	Provides a guided UI for creating, testing, and managing ETL transform maps. For more information, see MID Server .
Robust Transform Engine (RTE)	Transforms staged source data into CMDB-compliant records using defined ETL logic. For more information, see Create a robust import set transformer .
Identification and Reconciliation Engine (IRE)	Ensures data consistency and prevents duplicates by identifying and reconciling CIs. For more information, see the CMDB Identification and Reconciliation (IRE) .
CMDB/TNI	Stores structured, accurate telecom infrastructure data for visibility and downstream processes.

Supported Service Graph Connectors

- Nokia Altiplano SGC (sn_sgc_altiplano): Integrates with the Nokia Altiplano Access Network SDN Controller via REST APIs.
- Cisco Meraki SGC (sn_sgc_meraki): Integrates with the Cisco Meraki Dashboard via REST APIs.
- Fortinet SGC (sn_sgc_fortinet): Integrates with FortiManager via REST APIs.
- Arista VeloCloud SGC (sn_sgc_vcloud): Integrates with VeloCloud Orchestrator via REST APIs.

Note: The connector can coexist with IT and Cloud Service Graph Connectors (e.g., for servers, monitoring tools, IoT, etc.).

Key benefits

- Fast time to value – Use predefined, supported connectors that require minimal configuration.
- Multivendor support – Integrate with various management platforms across access, core, and transport networks.
- Model-aligned visibility – verify telecom-specific hierarchy and relationships are accurately modeled in the CMDB.

- Discrepancy detection ready – Feed network data directly into Telecom Discrepancy Identification & Reconciliation for CMDB compliance.
- Scalable integration – Use ServiceNow’s proven integration framework built for performance and extensibility.

Telecom Network Inventory (TNI) data model

Telecom SGCs include logic to ensure compatibility with the Telecom Network Inventory (TNI) data model:

- When the TNI plugin is installed, each discovered network element automatically includes a TNI entity record.
- A `tni_entity` is created alongside its corresponding `cmdb_ci` record, using system-generated payload mappings (e.g., `inventory_category`).
- This ensures consistency across operational and planning systems.

If TNI is installed, a payload like the following one will be added to the IRE payload for each item (with `inventory_category` populated based on the `className`):

```
related = [{
  "className": "tni_entity",
  "values": {
    "inventory_category": ""
  }
}];
```

As a result, the discovered CI is in both the `cmdb_ci` and `tni_entity` tables.

Related topics

[Telecom Discovery via Nokia Altiplano](#)

[Telecom Discovery via Cisco Meraki SD-WAN](#)

[Telecom Discovery via Fortinet SD-WAN](#)

[Telecom discovery via Arista VeloCloud SD-WAN](#)

Telecom Discovery via Nokia Altiplano

The Service Graph Connector for Nokia Altiplano offers a telecom-aware integration that brings real-time network inventory from the Nokia Altiplano Access Network SDN Controller into your ServiceNow CMDB.

Designed for service providers and telecom enterprises, this connector enables complete visibility, control, and synchronization of your physical and logical network infrastructure. This integration uses REST APIs and a MID Server to deliver a telecom-model-aligned view of your network, enabling more accurate service modeling, inventory management, and operational efficiency.

Key benefits

- **Accurate inventory synchronization:** Automatically populate the CMDB with up-to-date physical and logical inventory from Nokia Altiplano using secure, REST API-based ingestion. This includes OLTs, ONUs/ONTs, interface cards, ports, slots, and logical connections.
- **Telecom-Aware CI modeling:** Model your network infrastructure in a telecom-aligned format using purpose-built CI classes and relationships. The connector verifies the accurate representation of devices and their dependencies—mirroring your actual network topology.

- **Simplified and Guided Setup:** Reduce time-to-value with a built-in guided setup that walks you through connection configuration, data source management, and import scheduling with ease.
- **Flexible discovery options:** Choose how and when to run discovery based on your needs —whether it's full bulk loads, targeted filtered discovery, or phased OLT-only imports. Apply custom filters by device IP or name.
- **Multi-Instance support:** Scale effortlessly by onboarding multiple Altiplano instances independently. Configure connection aliases and define import schedules per instance for complete operational flexibility.

Note: Supported Nokia Altiplano Controller minimum version is 24.6. For a general overview of Service Graph Connector technology, see [Getting started with Service Graph Connectors](#).

Nokia Altiplano SGC Architecture

The following table illustrates the architecture of Nokia Altiplano service graph connector.

Nokia Altiplano 2.0 Service Graph Connector architecture

#	Stage	Location	Components or sub-flow	Description
1	Source	Network/IT	Altiplano Controller (REST)	External system holding OLT and ONU data. REST is the active protocol for this connector.
2	Altiplano data source (custom script)	MID Server	<ul style="list-style-type: none"> • REST action (get access token) • REST action (get list of OLT devices) • REST sub-flow (get OLT and its related ONUs, physical and logical details) • Bulk data source • Filtering data source • Filter devices list 	REST actions run on the MID Server to get the access token from Altiplano, retrieve the device list for discovery, and get full device details.
3	Import sets	ServiceNow instance	Altiplano import set	<p>Imports data from external sources, triggered by Scheduled Data Import. The bulk data source pulls all Altiplano data; the filtering data source pulls only data for OLTs that match the filter criteria. Use Connection selects the Altiplano connection alias.</p> <p>The import set includes:</p> <ul style="list-style-type: none"> • OLT physical and logical • ONU physical and logical

Nokia Altiplano 2.0 Service Graph Connector architecture (continued)

#	Stage	Location	Components or sub-flow	Description
				<ul style="list-style-type: none"> • Logical ports relationships • Logical connections between OLTs and ONUs
4	IH-ETL RTE	ServiceNow instance	<ul style="list-style-type: none"> • Clean, transform, enrich data • Map to CMDB classes and add relationships • Prepare payloads for IRE 	Determines which classes are populated based on discovered data and models. Handles transform maps, references of models and BW, CI relationships, and CI references.
5	IRE	ServiceNow instance	<ul style="list-style-type: none"> • IRE rules execution • CMDB integration 	RTE output is automatically passed to the IRE as a payload. The IRE matches to existing CIs based on identification rules and creates new CIs if no match is found.
6	CMDB compliance	ServiceNow instance	<ul style="list-style-type: none"> • Audit • Follow-on tasks • Remediation (writes to CMDB and TNI) 	Detects CMDB discrepancies, such as relationship mismatches or CIs present in the CMDB but not discovered. Creates a discrepancy report. Resolves discrepancies automatically or manually through the remediation workflow.

Use cases

The following are examples on how you can use the Nokia Altiplano Service Graph Connector:

- Automatically ingest and structure network data into ServiceNow's CMDB for both physical and logical network elements: This creates a telecom-model-aligned CMDB view that mirrors your live network environment.
 - Physical Components: OLTs, ONUs/ONTs, slots, subslots, cards, and ports
 - Logical Components: Logical ports, VLANs, Link Aggregation Groups (LAGs), and logical paths between devices
- Maintain the integrity of your network data with scheduled, automated reconciliation that:
 - Detects and responds to real-time changes
 - Prevents data drift or stale records
 - Supports operational processes like service assurance, order fulfillment, and network planning
- Discover beyond physical infrastructure by capturing and managing logical connectivity within ServiceNow:

- Identify and model logical ports, LAGs, and logical paths like PON and VLAN connections
- Link logical CIs with physical components using parent-child and member-of relationships
- Enhance diagnostics, impact analysis, and service modeling by visualizing end-to-end logical topologies
- Leverage ServiceNow's built-in Extract, Transform, Load (ETL) framework to simplify and accelerate integration:
 - Predefined transformation maps and CI class definitions reduce development overhead
 - Reuse existing CMDB structures to minimize customization
 - Rapidly onboard new Altiplano instances and scale across your network infrastructure

Key capabilities and components

Key capabilities and components

Capability	Description	Supporting component
Telecom-aware CMDB modeling	Maps Altiplano physical and logical inventory (OLT, ONU/ONT, ports, slots, cards, interfaces) into CMDB with telecom-specific CI classes and relationships.	RTE, IRE, CMDB tables (cmdb_ci_optical_line_terminal, cmdb_ci_optical_network_terminal, etc.)
Automated data ingestion	Retrieves device inventory from Nokia Altiplano via REST APIs with secure, scheduled imports.	MID Server, Data Sources (SGC-Nokia Altiplano Bulk/Filtered Discovery)
Custom discovery control	Choose OLT-only or include ONU data, apply filters by IP or name, schedule jobs per instance.	Import Schedules, System Properties (sn_sgc_altiplano.enable_onu_discovery, etc.)
Multi-instance support	Configure and manage discovery independently for multiple Altiplano controllers.	Connection Aliases, Credential Aliases
Parallel data processing	Improve performance by running concurrent data-source jobs for large datasets.	System Property: sn_sgc_altiplano.parallel_number Enable Parallel Loading
Guided configuration	Simplifies setup with step-by-step interface for creating connections, credentials, and jobs.	Guided Setup UI (Navigation: All > Service Graph Connectors > Nokia Altiplano > Setup)
Model-driven CI classification	Matches discovered devices to models and assigns correct CI class (OLT, ONU, ONT), or falls back to Network Gear.	Model Tables, System Property: sn_sgc_altiplano.onu_ci_class
Relationship	Establishes telecom-specific CI relationships (e.g., member-of, contains, logical path).	IRE, Logical Connections CI (cmdb_ci_ni_logical_path)

Key capabilities and components (continued)

Capability	Description	Supporting component
Dashboards and monitoring	View status, results, and errors of each run; filter by connector or timeframe.	Integration Commons for CMDB Dashboard
Scalable and reusable architecture	Decoupled data sources, transformation, and CI reconciliation for easy scaling and customization.	Import Sets (sn_sgc_altiplano_tsom_inventory), Transformation Maps, System Properties

CMDB Integrations Dashboard

The Integration Commons for CMDB store app provides a dashboard with a central view of the status, processing results, and processing errors of all installed Service Graph Connectors. You can see metrics for all integration runs. You can filter the view to a specific integration, a specific time duration, or a specific integration run. For more details about monitoring integrations in the CMDB Integrations Dashboard, see [Integration Commons for CMDB](#).

Related topics

[Configure Nokia Altiplano service graph connector](#)

[System components installed with Nokia Altiplano](#)

Telecom Discovery via Cisco Meraki SD-WAN

The Service Graph Connector (SGC) for Cisco Meraki provides a cloud-based management platform with a visual representation of network traffic flow between services and applications. It enables centralized configuration, monitoring, and management. Network administrators can easily discover and map services, enforce security policies, and troubleshoot issues in real time.

The Cisco Meraki SGC provides visibility, control, and synchronization of physical and logical network infrastructure for service providers and telecom enterprises. It leverages REST APIs and a MID Server to deliver a unified, telecom-model-aligned view.

Key benefits

- **Automated Device Discovery:** Periodically, as determined by import schedules, scans the Cisco Meraki dashboard to import all configuration item (CI) data.
- **Real-time synchronization:** Ensures the Configuration Management Database (CMDB) is up to date with the latest CI information.
- **Visualization:** Provides a graphical representation of Cisco Meraki networks, device relationships, and dependencies through the Integration Hub ETL.

Note: For a general overview of Service Graph Connector technology, see [Getting started with Service Graph Connectors](#).

Cisco Meraki SGC Architecture

The Cisco Meraki architecture consists of the following key components:

- **Setup:** A guided multi-step process that establishes connection alliances, configures credentials, and sets up REST API endpoints to integrate with the Meraki dashboard.
- **Data sources:** Scripts that fetch information related to various components such as organizations, networks, devices, and uplinks from Meraki APIs.

- **Import schedules:** Determines when data sources are executed to facilitate the discovery and updating of CIs.
- **Connection and credential aliases:** Details of the protocols used for API requests to extract data.
- **Connections:** Contains the necessary connection details for accessing the Meraki environment during the discovery process.
- **Credentials:** Provides the necessary authorization information to obtain data from the Meraki environment.
- **Properties:** System properties that contribute to configuring Cisco Meraki SGC and can be referenced within data sources.

CMDB Integrations Dashboard

The Integration Commons for CMDB application provides a dashboard with a central view of the status, processing results, and processing errors of all installed Service Graph Connectors. The dashboard displays metrics for all integration runs. You can filter the view to a specific integration, time duration, or integration run. For more details about monitoring integrations in the CMDB Integrations Dashboard, see [Integration Commons for CMDB](#).

Related topics

[Configure Cisco Meraki Service Graph Connector](#)

[Run and verify an import schedule for Cisco Meraki SGC](#)

Extensibility in Meraki discovery

The Meraki connector defines its own scripted extension points. Use them to customize API call sequencing, response parsing, field mappings, tag transformation, and CI lifecycle stage status without modifying the connector's core scripts.

The Meraki connector exposes five scripted extension points, each registered in the connector's own *sn_sgc_meraki* scope. Each extension point is a separate customization surface: you can implement one, several, or none, and each is evaluated independently at runtime.

What you can customize

The *MerakiCollectionPlan* extension point controls which API calls the connector makes, in what order, and how each response is routed. Use it to add, remove, or reorder collection steps, or to change how a step's response is fetched and written.

The *MerakiParserHooks* extension point controls how a raw API response is unwrapped, enriched, filtered, and stored before it reaches field mapping. Use it to register additional entity types and to customize the five required handlers: *unwrapResponse*, *beforeMapping*, *afterMapping*, *postProcess*, and *shouldStore*.

Note:

The *MerakiParserHooks* handlers *beforeMapping* and *afterMapping* were previously named *beforeExtract* and *afterExtract*. Rename the handlers in a custom implementation to the current names.

The *MerakiFieldMappings* extension point controls how raw Meraki API data is mapped into CI fields. Use it to add or change entries inside a resource type's *additional_attributes* block. Core CI identity fields such as *key*, *name*, and *serial_number* can't be changed this way.

The *TagTransformationExtensionPoint* extension point controls how an array of Meraki tag strings is converted into key/value pairs for CMDB attribute mapping.

The *MerakiCustomizedLifeCycleStageStatus* extension point controls the life cycle stage status value assigned to a CI, per CI class.

For the full handler contract for each extension point, see the [Meraki extension points reference](#).

How implementations are selected

For *MerakiCollectionPlan*, the connector uses the first registered implementation whose *getCollectionPlan* handler is defined. An implementation missing *getCollectionPlan* is skipped and an error naming the missing handler is logged. If no registered implementation is valid, the connector falls back to its default collection plan.

For *MerakiParserHooks*, the connector uses the first registered implementation that defines all five required handlers: *unwrapResponse*, *beforeMapping*, *afterMapping*, *postProcess*, and *shouldStore*. An implementation missing any of the five is rejected in full rather than partially merged, and an error naming the missing handlers is logged. If no registered implementation is valid, the connector falls back to its default parser hooks.

For *MerakiFieldMappings*, the connector uses the first registered implementation whose *getFieldMappings* handler returns every required resource-type key. An implementation that's missing a required key is skipped and an error naming the missing keys is logged. If no registered implementation is valid, the connector falls back to its default field mappings.

For *TagTransformationExtensionPoint*, the connector uses the first registered implementation whose *handles* handler returns *true* for the Meraki source identifier.

For *MerakiCustomizedLifeCycleStageStatus*, the connector uses the first registered implementation that defines the *getLifeCycleStageStatus* handler, and uses its returned value directly.

Note:

A default implementation of each extension point ships with the connector and is registered out of the box at *order 100*. Adding a custom implementation doesn't replace the default one automatically — give your implementation's extension instance a lower *order* value so it's evaluated first. You don't need to deactivate the default extension instance.

Telecom Discovery via Fortinet SD-WAN

The Fortinet Service Graph Connector provides a telecom-aware integration that imports real-time network inventory data from the Fortinet network management system into the Configuration Management Database (CMDB). This integration enhances visibility and management of network resources.

Key benefits

- **Accurate inventory synchronization:** Automatically populates the CMDB with up-to-date physical and logical inventory from Fortinet using secure, REST API-based ingestion. SD-WAN devices, interface cards, ports, slots, and logical connections are included.
- **Telecom-aware CI modeling:** Models your network infrastructure in a telecom-aligned format using purpose-built CI classes and relationships. The connector promotes an accurate representation of devices and their dependencies that mirrors your network topology.

- Simplified guided setup: Reduces time-to-value with a built-in guided setup that walks you through connection configuration, data source management, and import scheduling.
- Flexible discovery options: Supports custom filters by device IP address, device name, device model, or organization/ADOM. Choose how and when to run discovery.

Note: For a general overview of Service Graph Connector technology, see [Getting started with Service Graph Connectors](#).

Devices with high port counts

Devices with a large number of ports (for example, 600 or more) are processed by splitting the port data across multiple import batches. This approach imports ports incrementally instead of importing all ports for a device at once. This prevents out-of-memory errors during discovery, so devices with high port counts are imported successfully.

Fortinet SD-WAN architecture

The Fortinet SD-WAN architecture is composed of three key components:

1. FortiPortal: A cloud-based service through which you register for Fortinet accounts and manage their licenses.
2. FortiManager: A centralized management solution that oversees and manages SD-WAN networks.
3. FortiAnalyzer: A security analytics and reporting tool that collects and analyzes logs from Fortinet devices to identify potential security threats and conformance issues. Alerts are sent through a preconfigured webhook, and metrics are retrieved through API calls.

Fortinet APIs use JSON-RPC to enable programmatic interaction with FortiManager and FortiAnalyzer for bulk configuration, device management, and inventory collection.

The following diagram illustrates the architecture of the Fortinet Service Graph Connector.

CMDB Integrations Dashboard

The Integration Commons for CMDB application provides a dashboard with a central view of the status, processing results, and processing errors of all installed Service Graph Connectors. The dashboard displays metrics for all integration runs. You can filter the view to a specific integration, time duration, or integration run. For more details about monitoring integrations in the CMDB Integrations Dashboard, see [Integration Commons for CMDB](#) .

Related topics

[Configure a Fortinet SD-WAN Service Graph Connector](#)

[Run and verify an import schedule for Fortinet SGC](#)

Telecom discovery via Arista VeloCloud SD-WAN

The Service Graph Connector for Arista VeloCloud discovers SD-WAN inventory from VeloCloud Orchestrator and imports it into the CMDB.

Key benefits

- Automatically scans VeloCloud Orchestrator to import all configuration item (CI) data based on import schedules.
- Keeps the CMDB up to date with the latest CI information.
- Uses Integration Hub ETL to provide a graphical representation of VeloCloud SD-WAN networks, device relationships, and dependencies.

Note: For a general overview of Service Graph Connector technology, see [Getting started with Service Graph Connectors](#).

Arista VeloCloud SD-WAN architecture

The Arista VeloCloud SD-WAN architecture consists of three key components:

- **VeloCloud SD-WAN Edge:** An enterprise-class appliance deployed at branch sites that provides secure, optimized connectivity to applications. The Edge aggregates multiple links and steers traffic using Dynamic Multipath Optimization (DMPO) and deep application recognition (DAR).
- **VeloCloud SD-WAN Gateway:** A distributed network component that optimizes data paths between branches, data centers, and cloud services. Gateways provide scalability, redundancy, and on-demand flexibility.
- **VeloCloud Orchestrator:** A cloud-hosted or on-premises central management platform for configuring, provisioning, and monitoring the SD-WAN environment. The Service Graph Connector communicates with VeloCloud Orchestrator to discover and import inventory data into the CMDB.

Discovery modes

The Service Graph Connector for Arista VeloCloud supports two VeloCloud account topologies:

- **Partner (MSP) mode:** Applies to an Enterprise Proxy (partner) account that owns and manages many customer enterprises.
- **Operator mode:** Applies to a direct operator account that manages its own enterprises without a partner proxy.

You don't need to specify the mode or provide an organization identifier. When discovery runs, the connector checks whether the supplied credentials belong to a Partner (Enterprise Proxy) account. If they do, discovery proceeds in Partner (MSP) mode; if not, the connector uses Operator mode. The same scheduled data import works for both account types.

Both modes follow the same discovery sequence:

1. **Account and partner lookup:** The connector determines whether the account is a partner proxy.
2. **Gateway pools:** The connector retrieves the gateway pool information available to the account.

3. Enterprise list: The connector retrieves the list of enterprises (organizations). In Partner (MSP) mode, the list includes edge information. In Operator mode, the list is retrieved without edge information.
4. Edge devices: For each enterprise, the connector retrieves its edge devices, including high-availability status, configuration, licenses, links, and site details. Any edge device without a name is skipped during import in both modes.

The main differences between the two modes are described in the following table.

Discovery mode differences

Aspect	Partner (MSP) mode	Operator mode
Account type	Enterprise Proxy (partner)	Direct operator
Enterprise list source	<code>getEnterpriseProxyEnterprises</code>	<code>getNetworkEnterprises</code>
Edge data in the enterprise list	Included	Retrieved separately for each enterprise
Enterprises with no edge devices	Skipped	Retained
Partner association	Each enterprise is linked to its parent partner	No partner parent
Enterprise identification	Identified within the scope of the parent partner	Identified by a globally unique identifier

Regardless of the mode, the connector builds the same inventory — organizations, network sites, edge devices, and network service instances. The connector establishes the relationships between organizations and their network sites.

For the full list of VeloCloud Orchestrator endpoints called during discovery, with sample responses, see [Arista VeloCloud Service Graph Connector API endpoints](#).

CMDB Integrations Dashboard

The Integration Commons for CMDB application provides a dashboard with a central view of the status, processing results, and processing errors of all installed Service Graph Connectors. The dashboard displays metrics for all integration runs. You can filter the view to a specific integration, time duration, or integration run. For more details about monitoring integrations in the CMDB Integrations Dashboard, see [Integration Commons for CMDB](#).

Related topics

[Configure Arista VeloCloud Service Graph Connector](#)

[Run Arista VeloCloud SD-WAN Import](#)

Telecom Discovery Builder framework

The Telecom Discovery Builder framework ETL (Extract, Transform, Load) is a reusable and schema-aligned component delivered with the Telecommunications Service Operations Management (TSOM) Core application. It provides a consistent and extensible method for ingesting telecom inventory data into the ServiceNow configuration management database (CMDB) across multiple Service Graph Connectors (SGCs).

The Telecom Discovery Builder framework ETL serves as a baseline data ingestion utility to handle telecom-specific configuration item (CI) data. It enables connector development teams to avoid building ETLs from scratch by offering a standardized transform logic that can be duplicated and customized for each connector. It provides a standardized, reusable foundation that promotes consistency across implementations.

Once the TSOM Core plugin is activated, the generic ETL is auto-provisioned and becomes available in the Integration Hub ETL Studio. It is ready to be reused and adapted to connector-specific needs.

Why and when to use the Telecom Discovery Builder framework

Use the Telecom Discovery Builder framework when you want to do the following:

- Deploy a new Service Graph Connector in a telecom environment and want to avoid building the ETL from scratch.
- Maintain schema-aligned consistency across multiple connectors and platform instances.
- Discovery payloads and CI relationships compliant with TNI (Telecommunications Network Inventory). For more information, see [Telecommunications Network Inventory](#).
- Work within an SGC (Service Graph Connector) application scope and want to customize ETL behavior without modifying core logic.
- Benefit from predefined mappings, validated JSON schema support, and a UI-driven configuration interface.

Key capabilities and features

- Auto-provisioned with TSOM core: Installs automatically and is ready to use across telecom connectors.
- Generic schema-based ingestion: Supports a unified data schema for telecom CIs.
- Supports duplication and customization: Duplicate the Telecom Discovery Builder ETL framework into your application scope using Integration Hub ETL Studio.
- TNI support: Aligned to the TNI data model and the ETL logic can be extended to generate and link TNI entities.
- Flexible field-mapping interface: Configure import sets, data sources, targets, and transformation logic through a UI-driven experience.

Key user roles

The `tsom_visibility_admin` and `tsom_assurance_admin` roles provide granular access control across TSOM Visibility and Assurance application scopes, supporting stronger security and compliance in line with the Granular admin directive.

- The `tsom_assurance_admin` role provides full operational and administrative access to TSOM Assurance application scopes.
- The `tsom_visibility_admin` role manages the operational tasks for the TSOM visibility application.

These roles control access to the TSOM Visibility and Assurance application scopes listed in the following table.

TSOM Visibility and Assurance application scopes

Scope name	Description
sn_tsom_core (Telecom Service Operations Core)	Includes the audit application and various system properties within that app scope.
sn_tsom_em_core(Event Management Core)	Includes the monitoring and analysis of telecommunications network alerts and performance metrics from SD-WAN infrastructure.
sn_tsom_em_conns(Event Management Connectors)	Include push connectors and connector definitions used to ingest alerts from external telecommunications network sources.
sn_tsom_patterns(Telecom Discovery Patterns)	Include the MID Server and system IP Service properties.
sn_sgc_aliplano(Service Graph Connector for Nokia Altiplano)	Includes the application properties.
sn_sgc_meraki(Service Graph Conenctor for Cisco Meraki Telco SD-WAN)	Includes the application properties.
sn_sgc_fortinet(Service Graph Connector for Fortinet Telco SD-WAN)	Includes the application properties.

The following table lists the roles contained in the tsom_visibility_admin role.

Visibility granular admin roles

Role name	Description
import_admin	<ul style="list-style-type: none"> Monitors scheduled data import executions, including viewing the hierarchy and flow of chained imports, accessing execution contexts, and diagnosing import issues. Manages all aspects of import set records and imports.
connection_admin	<ul style="list-style-type: none"> Responsible for managing connections and credentials, especially for Integration Hub. Create and configure HTTP and HTTPS connections, which are used by custom actions or activities to connect to external endpoints. These actions include setting up connection details such as credentials, connection aliases, URLs, MID Server options, and additional attributes.
sn_cmdb_admin	<ul style="list-style-type: none"> Highest-level administrative role for CMDB privileges. Grants full access to CMDB data, tools, and user interfaces. Enables administrators

Visibility granular admin roles (continued)

Role name	Description
	to configure policies such as class management and application service requirements that aren't available to editors.
discovery_admin	<ul style="list-style-type: none"> Grants access to the tables in the Discovery application. Essential for configuring and executing discovery in your network. Includes setting up MID Servers, managing credentials, defining discovery schedules, and validating results.
usage_admin	<ul style="list-style-type: none"> For instances using Discovery, this role enables users to track the number of active users and hardware configuration items (CIs) discovered. Monitors instance usage through dashboards, providing information to help your organization understand application adoption and usage patterns.
cmdb_inst_admin	<ul style="list-style-type: none"> Enables users to create and customize integrations using Integration Hub ETL, including integrating third-party data into CMDB or supported non-CMDB tables. Necessary to use the Integration Hub ETL store app and access CMDB Integrations.
certification_admin	<ul style="list-style-type: none"> Enables users to manage certification filters for auditing subsets of records. Create, update, delete, and run certification audits, compare records against expected values, and assign follow-on tasks to remediate discrepancies. Manage and reassign follow-on tasks from compliance audits and view all such tasks. Users with the certification role can access only tasks assigned to them.

Visibility granular admin roles (continued)

Role name	Description
agent_admin	<ul style="list-style-type: none"> Responsible for downloading and administering the system's built-in MID Server agent. Enables users to manage scripts and properties related to the MID Server.

Key benefits

Benefits of the Telecom Discovery Builder framework ETL include:

- Consistency in how telecom configuration items (CIs) are structured and loaded into the CMDB.
- Reusability through duplication and customization across connectors.
- Compliance with TNI schemas and discovery model requirements.
- Flexibility to extend and tailor data transformations without compromising the base schema.

Related topics

[Configure the Telecom Discovery Builder ETL](#)

[Extend TNI entity support for duplicated ETLs](#)

Telecom Discrepancy Identification and Reconciliation

Use the Telecom Discrepancy Identification and Reconciliation capability to keep your CMDB or Telecommunications Network Inventory (TNI) in sync with your live telecom network. This solution continuously audits and compares discovered data with inventory records. It helps you detect, classify, and automatically remediate inconsistencies before they impact service quality, assurance, or fulfillment processes.

Telecom Discrepancy Identification and Reconciliation is a telecom-specific capability included with the TSOM Visibility plugin. It confirms the integrity of your network inventory by identifying mismatches between real-time network data and inventory records. The comparison includes data from Discovery or external systems against records stored in the CMDB or TNI.

This solution uses:

- CMDB Compliance Certification audits to detect anomalies.
- Follow-on tasks to log and track issues.
- Automated remediation subflows to resolve discrepancies.

i Note: Keeping inventory accurate enables automation, reduces service errors, and supports regulatory compliance. It is also a foundational component of TM Forum's Autonomous Network Operations (ANO) framework.

Telecom Discrepancy Identification & Reconciliation architecture

Key features

Key features of Telecom Discrepancy Identification and Reconciliation

Feature	Description
Continuous audit checks	Compares discovered network data with inventory records using Certification Audits.
Discrepancy detection	Identifies CIs that are missing, misconfigured, or misaligned in terms of hierarchy or attribute values.
Follow-On task generation	Automatically creates tasks for each discrepancy to enable tracking and remediation.
Automated remediation	Uses Flow Designer subflows to resolve issues such as outdated CIs, invalid relationships, and missing discovery data.

How Telecom Discrepancy Identification & Reconciliation works

Once the discovery process is complete, the CMDB Compliance framework identifies mismatches between discovered data and inventory records using scheduled or on-demand compliance audits.

- **Compliance Audits:** Compare discovered network data with CMDB inventory records to detect inconsistencies.
- **Follow-On Tasks:** Automatically generated for each failed audit. These tasks document and categorize the identified discrepancies.
- **Remediation Subflows:** Launched from Follow-On Tasks to resolve discrepancies by updating, decommissioning, or realigning Configuration Items (CIs).

Note: For more information, see [Discrepancy identification – types of discrepancies](#).

CMDB Compliance and Telecom Discrepancy Identification & Reconciliation

CMDB Compliance is a toolset that enables administrators to certify CMDB data for accuracy and identify discrepancies detected during compliance audits. It can also automatically generate and assign Follow-on Tasks for failed audit records, which serve as tasks to trigger an appropriate remediation subflow to correct discrepancies. CMDB Compliance Audits form the foundation of the Telecom Discrepancy Identification & Reconciliation.

- CMDB Compliance runs audits as a post-processing rule, identifying anomalies (discrepancies) in the CMDB.
- CMDB Compliance creates a Follow-On Task for each Audit Record in a failed state. The failed state results from an audit finding an anomaly or discrepancy in the CMDB. A remediation flow can be designed and triggered for each Follow-On Task to address and resolve the discrepancy.

The logic for Telecom Discrepancy Identification & Reconciliation, and the example remediation subflows, are automatically with the TSOM Visibility plugin. For more information on the general CMDB Compliance toolset, see [CMDB Compliance](#).

Discrepancy Identification Scenarios (using Compliance Audits)

Discrepancy identification in TSOM Visibility relies on using CMDB Compliance (Certification Audits). It extends this capability by adding specific logic that uses model relationships and information to identify mismatches. To support remediation, the system generates specific reconciliation task types for each issue found, such as:

- Slots occupied discrepancy
- Most recent discovery date not set
- Most recent discovery date not within configured threshold
- Model relationships not defined
- CI model not found
- Incorrect number of relationships
- Reference to Logical Interface not found

You can use the following audits to identify the discrepancies in the discovered physical and logical entities:

- Telecom Discrepancy Audit
- Telecom Logical Connections Discrepancy Audit
- Telecom Network Topology Discrepancy Audit

Note: For more information on the general compliance audits, see [Certification audits](#).

Automation and UI Integration

- Every failed audit automatically creates a Follow-On Task.
- Tasks can trigger prebuilt or custom subflows using Flow Designer.
- Manual remediation can be initiated using a "Remediate" UI action button in the task form.
- Remediation steps are recorded in work notes for visibility and auditing.

Note: For more information on how to build a subflow, see [Building subflows](#).

Follow-On Task types created for failed Audit Result Records

The following discrepancy types (Audit Results) can be found for Parent CI and child CIs for each relationship record in the CI Relationship table (cmdb_rel_ci) that matches the conditions. You can create Follow-On Tasks for each failed Audit Result:

1. The most recent discovery date not set - generated in case the Most recent discovery date field in CI is missing.
2. The most recent discovery date not within configured threshold - generated when the field value difference between a Parent CI and child CI exceeds the configured threshold. The default threshold is 2.5 days. To change this value, use the `sn_tsom_core.discovered_date.diff.threshold.in.days` system property.
3. CI model not found (the 'Model ID' field isn't set or data is invalid) - generated in case a corresponding CI model isn't found. If a CI model isn't found, the next validations (4-6) are irrelevant because they rely on CI models. In case a CI model is found, the audit will continue to the next validations (4-6).
4. Slots occupied discrepancy - generated in case a card occupies an incorrect number of Slots.

5. Model relationships not defined - relevant only if TNI is installed. Generated if the audit is unable to find a relationship between Parent and child CI models in the Network Model Relationships table.
6. Incorrect number of relationships - relevant only if TNI is installed. Generated if the audit finds that the number of discovered child CI records exceeds the maximum number of its corresponding Parent CI record in the model relationship Count field in the Network Model Relationship table.
7. Incorrect number of relationships - generated during the Logical Connection Discrepancy Audit when a logical interface is associated with more than one logical connection, violating the expected one-to-one mapping.
8. Incorrect number of relationships - generated during the Network Topology Discrepancy Audit when a network topology record does not meet the required relationship criteria. Specifically:
 - The record must have at least one "Contains:Contained By" relationship with an equipment CI.
 - The record must also have at least one "Members:Member Of" relationship with a logical connection CI.
9. Reference to Logical Interface not found - generated during the Logical Connection Discrepancy Audit when a logical connection is missing one of the end points (Port A or Port z, or both).

Telecom Reconciliation

Automate reconciliation of network inventory discrepancies and enhance operational efficiency. Telecom Reconciliation helps to:

- Identify and resolve discrepancies between live network inventory and CMDB inventory to ensure alignment and boost productivity
- Empower users with auto generated discrepancy reports detailing the type of discrepancy
- Provide users with recommended corrective actions and the flexibility and control to choose between manual or automated methods to resolve discrepancies
- Improve operations by automatically aligning the operational status of network resources with the same status in the inventory CMDB

Related topics

[Discrepancy identification – types of discrepancies](#)

[Activate Telecom Discrepancy and Reconciliation](#)

[Run Telecom Discrepancy audit](#)

[Telecom discrepancy identification and reconciliation](#)

Discrepancy identification – types of discrepancies

The Telecom Discrepancy Identification and Reconciliation capability identifies and classifies mismatches between the discovered network state and inventory data stored in the CMDB or TNI.

Discrepancy identification is powered by the CMDB Compliance Certification Audit, which:

- Runs on CI and relationship data.
- Compares discovered and inventoried records.
- Generates Follow-on tasks when mismatches are detected.

Discrepancy types

The audit process identifies discrepancies by comparing discovered network data with the existing inventory in the CMDB/TNI. Discrepancies fall into four categories:

Missing in network - entities existing in inventory but missing in network

Definition: A CI exists in the CMDB/TNI inventory but is no longer detected on the network by Discovery. This indicates that the equipment has been removed, decommissioned, or is otherwise no longer present, but the inventory record remains.

Impact:

- A discrepancy task is generated to highlight the missing network entity.
- Optional remediation subflows may be triggered to retire or update the inventory record.

Mismatched configuration items (CIs) - entities that exist in inventory and network but differ in attribute values and hierarchical relationships

Definition: The CI exists in both Discovery and CMDB/TNI, but discrepancies exist in relationships, hierarchy, or attribute values. The following are the sub types:

- Hierarchy Mismatches - Occur when the structural relationships between CIs (e.g., parent-child associations) are inconsistent. The following are the examples:
 - A chassis contains more child cards than allowed by the model definition.
 - A card is incorrectly associated with a slot in the CMDB that does not align with Discovery data. The validation sources are:
 - `cmdb_rel_ci` records for `Contains::Contained` by relationships.
 - `sn_ni_core_network_model_relationship` table for enforcing model-specific constraints.
- Attribute Value Mismatches - Involve discrepancies in CI field-level properties. The common issues are:
 - Outdated or incorrect Discovery Dates.
 - Inaccurate Model Configurations.
 - Invalid Slot Assignments that breach model rules.

The impact is the affected records are marked as failed in audit reports. Follow-On Tasks may invoke context-specific remediation subflows to realign inventory data with the actual network state.

Note: For more information, see [Configure attribute value discrepancy in CMDB 360](#).

Missing in CMDB – entities discovered on the network but not present in inventory

Definition: A CI is discovered on the network through TSOM Discovery or a Service Graph Connector, but no corresponding record exists in the CMDB. This indicates that the equipment is physically present and operational but has not yet been inventoried. For example, Discovery detects a Fortinet FortiGate-80F device on the network, but no matching CI exists in the CMDB.

Impact:

- The audit identifies the CI as new when its creation date matches its discovery date.
- A single follow-on task is created at the equipment level to prompt review. Individual tasks are not generated for each underlying child CI (such as IP addresses) associated with the equipment, keeping the task list focused and actionable.
- The follow-on task includes the short description **New CI has been discovered**, along with the CI reference and associated product model.
- You can enable or disable follow-on task creation for this discrepancy type using a system property.

SD-WAN discrepancy identification

In addition to telecom-specific discrepancies, the Telecom SD-WAN Network Discrepancy Audit identifies mismatches in your SD-WAN inventory. Standard telecom discrepancy detection focuses on CI classes such as slot, card, and interface. The SD-WAN audit targets a different set of CI classes and validates them against rules specific to the SD-WAN data model.

By default, the audit scopes to CIs whose discovery source contains TSOM. The audit compares discovered SD-WAN data with the inventory in the CMDB and generates follow-on tasks when mismatches are detected.

SD-WAN CI classes and validation requirements

CI class	Validation requirements	Relationship type
Group (cmdb_ci_group)	<ul style="list-style-type: none"> • Company reference • At least 1 Service child • At least 1 Site child 	Members::Member of
Site (cmdb_ci_ni_site)	<ul style="list-style-type: none"> • Company reference • Location reference • Exactly 1 Group parent 	Members::Member of
Service (cmdb_ci_network_service_instance)	<ul style="list-style-type: none"> • Company reference • Exactly 1 Group parent 	Members::Member of
Equipment (sn_tsom_core.audit.equipment_tables)	<ul style="list-style-type: none"> • Company reference • Max 1 Service parent • Max 1 Site parent • At least 1 IP Address child • Each IP Address child has max 1 Equipment parent 	<ul style="list-style-type: none"> • Service/Site: • Contains::Contained by • IP Address: Owns::Owned by

The SD-WAN audit validates the following CI classes. For each CI, the audit checks that required references are populated and that relationships match the expected structure. A CI that meets all requirements is marked as Certified. If one or more requirements are not met, the CI is marked as Failed and a follow-on task is created.

SD-WAN discrepancy types

The SD-WAN audit detects the same categories of discrepancy as the telecom audit, applied to the SD-WAN CI classes listed above.

- **Missing reference:** A required reference (such as Company or Location) is not populated on the CI. A follow-on task is created with the description identifying which reference is missing.
- **Incorrect relationships:** A required relationship is missing or exceeds the allowed count. For example, a Site CI without a Group parent, or an Equipment CI with more than one Service parent. A follow-on task is created with the description identifying the relationship issue.
- **Missing in CMDB:** A CI is discovered on the network but no corresponding record exists in the CMDB. The audit identifies the CI as new when its creation date matches its discovery date, and creates a single follow-on task at the equipment level. For details on this discrepancy type, see the section **Missing in CMDB** above.

Related topics

[Activate Telecom Discrepancy and Reconciliation](#)

[Run Telecom Discrepancy audit](#)

Certification Audit Logics

Audit Results are created for each audit executed on records that matched the selection (see matching conditions in Initial Certification Audit Run).

The result's state can be certified or failed. A Follow-On Task is created for each 'failed' Audit Result record.

Initial Certification Audit Run

Validating specific CMDB tables for anomalies.

The Service Operation CMDB Compliance Audit starts to run on the CI Relationship table (cmdb_rel_ci), but only on specific records with matching conditions as follows:

- Parent AND child CI classes are supported classes, including extended tables as follows:

slot (cmdb_ci_container_slot), subslot (cmdb_ci_container_subslot), card (cmdb_ci_interface_card), interface (cmdb_ci_ni_interface), telco equipment (cmdb_ci_ni_telco_equipment), IP switch (cmdb_ci_ip_switch), and IP router (cmdb_ci_ip_router).

Note: These properties can be configured via sn_tsom_core.audit.* system properties.

- Parent OR child is created or updated by Discovery (discovery_source = SG-Altiplano, ServiceNow). Audit uses its own filter mechanism that applies on equipment CIs. Default filter is "discovery_source CONTAINS TSOM"
- Parent AND child life-cycle Stage is Operational.

- The CI Relationship Type is Contains::Contained By.

Note: This property can be configured in the `sn_tsom_core.audit.relationship_types` system property.

Subsequent Certification Audit Runs

Follows the same logic as the Initial Certification Audit Run, but with the following additional matching selection criteria:

The audit runs when the Updated field timestamp is later than the Last run date in the Telecom Discrepancy Audit. The timestamp can be from the CI Relationship table, Parent CI, or child CIs. This indicates a change since the last audit.

Telecom Assurance

Telecom assurance monitors network performance, detects faults, and maintains service quality. ServiceNow Telecommunications Service Operations Management software is integrated with existing monitoring tools to consolidate alerts into a single platform, delivering AI-driven insights and automated workflows from fault detection to resolution.

Some key components of ServiceNow AI Platform[®] Telecommunications Service Operations Management telecom assurance are:

- Alert management: Managing, escalating, and resolving alerts
- Performance management: Analyzing network performance data to optimize service delivery

Related topics

[Configure Telecom Assurance](#)

[Now Assist for IT Operations Management \(ITOM\)](#) 

[Metric Intelligence](#) 

[Learning Enhanced Automation Platform \(LEAP\)](#) 

External event management via Telecommunications API notifications

Use the Telecommunications API notification to receive the external events that occurring in the customer network system so that you can promptly respond to them in the ServiceNow.

Introduction to Telecommunications API notification

Telecommunications API notification is a feature available in the Telecommunications Alarm Management Open API application. The Telecommunications API notification enables ServiceNow to receive the incoming notifications that occurring in the external network system and responds to them in a timely manner. It enables the broadcasting of events to the external systems through platform capabilities by eliminating the need for point connections.

Telecommunications API notification receives incoming notifications from the external systems that are subscribed on your network. When the notifications are received from the external system, you can create the events for the responses by using the Event Management application. Based on the collected information, the Event Management provides dashboards showing a consolidated view of all service-impact events.

Telecommunications API notification data model

The following diagram shows the components in the data model for the API

notification.

The Telecommunications API notification enables ServiceNow to receive incoming notifications through the event-driven architectures such as the Publisher/Subscriber (Pub/Sub) subscription model, Hermes, and Kafka Stream Connect. While cloud customers have the flexibility to select between both architectures, on-premise customers are limited to using their own Kafka or Pub/Sub subscription model.

- To learn more about Stream connect for Apache Kafka Stream, see [Using Stream Connect for Apache Kafka](#) .
- To learn more about Hermes Messaging Service, see [Hermes Messaging Service](#) .

In the Pub/Sub model, incoming notifications are categorized into topics. You use ServiceNow to publish the incoming notifications to these topics, and subscribers (customers) have the flexibility to select the topics to which they want to subscribe. This process enables subscribers to select only those messages that align with their interests. For example, if there are 10 topics for incoming messages from the external system, a customer can opt to subscribe to two of them based on their requirement. Consequently, when notifications are received from the external system, events are generated specifically for the two topics to which the customer has subscribed.

Related topics

[Configuring Telecommunications API notifications](#)

[System components installed with Telecommunications API notifications](#)

Fault Management: Events and alerts

Fault Management supports the monitoring, detection, and resolution of configuration and performance issues across SD-WAN-managed network devices. It integrates with Event Management to generate alerts and track events automatically.

Fault Management helps you monitor network devices by collecting alerts related to configuration and performance issues. Key features include alerting, customizable event rules, webhook-based notifications, and integration through the Integrations Launchpad.

Alerts are central to detecting and responding to configuration issues. Each alert is generated based on three core attributes:

- **Device name:** The affected device (for example, IP router, IP switch, access point)
- **Alert type:** The category of the event
- **Alert severity:** One of Critical, Major, Minor, Warning, or Clear. The source system's severity or alert level is mapped to one of these values.

When a monitored event occurs, you can create an alert with these attributes and begin tracking it. You can configure event rules to control what happens next: grouping related alerts, escalating critical ones, or suppressing noise so your team only sees issues that require action. The Integrations Launchpad provides a central place to manage integrations with external systems.

Managing alerts with event rules

Event rules and alert suppression give you control over how alerts are generated and managed. Configure the following options to control alert generation:

- **Event rules:** Set criteria such as the device name, alert type, severity, and frequency to determine the events and create the appropriate alert.
- **Alert suppression:** Consolidate multiple alerts from the same device into a single record.
- **Alert grouping:** Combine related events into a single alert for easier management.
- **Correlation IDs:** Link related events and alerts with unique identifiers, improving traceability across the event life cycle. These IDs create connections between relevant objects and alerts, enabling you to navigate directly between them.

For instructions on how to configure Event Management solutions, see the following documentation:

- [Configure a webhook](#)
- [Configure a metric pull connector](#)
- [Cisco Meraki installed integrations](#)
- [Fortinet installed integrations](#)

SD-WAN Alerts Dashboard


Monitor and analyze SD-WAN network alerts across customers and devices, with views for today's alerts, 30-day alert history, and alert trends over time.

Access the SD-WAN Alerts Dashboard

Access the SD-WAN Alerts Dashboard by navigating to **All > Platform Analytics > Dashboards** and searching for or selecting SD-WAN Alerts Dashboard. By default, the **Alert Summary** tab is selected.

Alert Summary


The **Alert Summary** tab provides the capability to filter SD-WAN alert activity by severity, state, customer, or CI.



Title	Type	Description
All alerts created today	Bar chart	The number of alerts created today, grouped by CI
All alerts created the last 30 days	Bar chart	The number of alerts created in the last 30 days, grouped by CI
Alert Trends Over Time	Line graph	The trend of alert activity over time

Alert Categorization

The **Alert Categorization** tab displays SD-WAN alerts grouped by customer and severity.




Title	Type	Description
Alerts per Customer	Donut chart	Distribution of alerts grouped by customer
Alert Severity HeatMap	Heat map	Distribution of alerts by severity (for example, Critical, Major, Minor, Warning, and Clear) over time

SD-WAN Inventory Dashboard

The SD-WAN Inventory Dashboard provides a visual summary of your SD-WAN network inventory, including total customers, sites, devices, and models, with breakdowns by device model and class.

Accessing the SD-WAN Inventory Dashboard

If you're a user with the `tsom_assurance_admin` role, you can access the SD-WAN Inventory Dashboard by navigating to **All > Platform Analytics > Dashboards** and searching for and selecting SD-WAN Inventory Dashboard.



Network Overview

The **Network Overview** tab provides a high-level summary of all discovered SD-WAN network inventory in the CMDB.

The following breakdowns are available in the SD-WAN Inventory Dashboard:

- Total devices by model
- Total devices by class

Title	Type	Description
Total customers	Single score	The total number of customers with SD-WAN network inventory discovered in the CMDB.

Title	Type	Description
Total Sites	Single score	The total number of sites across all customers with discovered SD-WAN inventory.
Total Devices	Single score	The total number of SD-WAN devices discovered in the CMDB.
Total Models	Single score	The total number of distinct device models discovered across the SD-WAN network.
Total Devices by Model	Donut chart	The distribution of discovered devices by hardware model.
Total Devices by Class	Donut chart	The distribution of discovered devices by device class, such as IP Router, Wireless Access Point, and IP Camera.


Filtered View

The **Filtered View** tab provides the ability to filter inventory data by customer, manufacturer, model, or discovery source.

SD-WAN Inventory Dashboard ▾

🔍 ⌂ Edit ⓘ

Network Overview **Filtered View**



Title	Type	Description
Customers	Single score	The number of customers matching the applied filter criteria.
Sites	Single score	The number of sites matching the applied filter criteria.
Sites by Customer	Bar Chart	The number of sites per customer matching the applied filter criteria.
Devices by Model	Donut Chart	The distribution of filtered devices by hardware model.

Title	Type	Description
Devices by Class	Donut Chart	The distribution of filtered devices by device class.

MPN latency metrics

Monitor UE-to-switch latency and related KPIs across MPN-connected devices, with combined and per-KPI trend views, an instance summary table, and a last-recorded-value panel.

Average Latency for All KPIs

The **Average Latency for All KPIs** panel is a combined time-series chart showing all three latency KPIs together, in 30-second buckets. Filter the chart by Assurance ID, Instance, or DNN.

Average Latency for All KPIs

Title	Type	Description
Average Latency for All KPIs	Combined time-series chart	The average value of all three latency KPIs (<code>collector2ue</code> , <code>collector2sw</code> , <code>ue2sw</code>) over time, in 30-second buckets, filterable by Assurance ID, Instance, or DNN

Per-KPI trend views

Three separate panels, one per latency KPI, each showing the trend in 2-minute buckets with legends by DNN and Instance.

Per-KPI trend views

Title	Type	Description
Collector-to-UE latency	Time-series chart	Trend of the <code>collector2ue</code> KPI, in 2-minute buckets, with legends by DNN and Instance
Collector-to-switch latency	Time-series chart	Trend of the <code>collector2sw</code> KPI, in 2-minute buckets, with legends by DNN and Instance
UE-to-switch latency	Time-series chart	Trend of the derived <code>ue2sw</code> KPI, in 2-minute buckets, with legends by DNN and Instance

Additional Info

The **Additional Info** table lists latency data by instance, sorted by Instance.

Additional Info columns

Column	Description
Instance	The instance identifier

Additional Info columns (continued)

Column	Description
Customer	The customer associated with the instance
opCo	The operating company associated with the instance
Assurance ID	The assurance identifier for the instance
DNN	The data network name associated with the instance
Product	The product associated with the instance
Vendor	The vendor associated with the instance
Last value of Latency	The most recently recorded latency value for the instance
KPI	The latency KPI the row's value belongs to

Instances with Latency (Last Recorded Value)

The **Instances with Latency (Last Recorded Value)** panel shows the most recently recorded latency value for each instance.

Understanding data gaps

TSOM flags or logs a latency metric as stale when the latest published value for a CI is older than expected, rather than silently displaying it as current. If a switch-side record is missing for a DNN during a collection cycle, the ue2sw calculation is skipped for that cycle. This may appear as a gap in the dashboard panels before rather than a repeated or estimated value.

Related topics

[MPN latency KPIs](#)

[SD-WAN Alerts Dashboard](#)

Mobile private network health status

The health status feature provides a quick indicator of radio and core server component health collected from a mobile private network (MPN). The system classifies each component using a configurable rule set and assigns error codes that identify specific detected conditions.

How health status is calculated

Each collected component is checked against a set of rules and assigned one of the following statuses:

- **Red:** The component reports a condition that indicates a critical issue, such as an alarm status of major or critical, or the component is not reachable.
- **Amber:** The component reports a condition that indicates a degraded but non-critical issue, such as a minor alarm or an indeterminate status.
- **Green:** The component reports normal operating conditions across all evaluated fields.
- **Unknown:** The component's status can't be determined from the available data.

The system checks the rules in a fixed priority order: red conditions first, then amber, then green. A component is marked green only when every green condition matches; if no rule matches at all, the status is unknown.

Rules are currently defined for two component types: radio and core server.

Rule configuration

The active rule set is stored as a versioned system property, so you can trace which rule version produced a given status. The system provides a default rule set and supports an administrator-editable override.

The system uses the override when one is provided. If the override is empty or contains invalid data, the system automatically falls back to the default rule set so that health status calculation does not stop.

The rule set supports multiple vendors, each with its own component types and evaluation logic. Currently, only the Nokia rule set is implemented.

Performance management: Metric collection

Performance management enables you to capture and analyze operational metrics to identify anomalies. Detected anomalies can generate alerts that surface in Service Operations Workspace, helping you avoid service outages.

Use provided solutions to collect metrics and monitor your devices for different device platforms:

- [Cisco Meraki installed integrations](#)
- [Fortinet installed integrations](#)
- [Metric type IDs](#)
- [Arista VeloCloud installed integrations](#)

Pagination in Meraki pull connector metrics requests

The Meraki pull connector handles pagination automatically for metrics API requests, ensuring complete data retrieval for large organizations where API responses are returned in multiple pages.

The Meraki API limits the number of items returned in a single response. In large Meraki organizations with many networks or devices, a single metrics request may not return all available data.

The pull connector now follows pagination links in API responses and issues subsequent requests until all pages are retrieved. This applies to metrics HTTP requests and verifies that data collection is complete regardless of organization size.

No configuration is required to enable this behavior. The connector handles pagination transparently as part of each scheduled metrics collection cycle.

For the page size used by each endpoint, see [Page limits for Meraki pull connector endpoints](#).

Pull connector granularity

Granularity controls the time resolution of metrics samples collected by a pull connector from the source system API. Aligning granularity with the polling interval reduces API calls to the source system.

Pull connectors for Cisco Meraki and Fortinet use two related timing settings:

Polling interval

How often the connector calls the source system API. The default polling interval is 15 minutes.

Granularity

The time resolution of each sample returned by the API. A granularity of 5 minutes means each API call returns data in 5-minute increments.

When the polling interval is 15 minutes and the granularity is 5 minutes, each polling cycle retrieves 3 samples, covering the full 15-minute window. Polling every 5 minutes for a single sample results in three times as many API calls for the same data.

Each connector type and API endpoint enforces minimum and maximum values for the granularity and metrics collection schedule fields. The system rejects values outside the supported range when you save the connector instance. For the specific constraints per connector and API type, see [Pull connector granularity constraints](#).

Related topics

[Configure a metric pull connector](#)

[Configure elastic event pull connectors for MPN](#)

MPN Formula Engine processing

The Formula Engine processes raw KPI formulas and stores the result in the **Formatted KPI Formula** field when a record is inserted or updated in the MPN Formulas [sn_tsom_em_conns_kpi_definitions] table.

When a record is inserted or updated in the MPN Formulas table, a business rule named Transform KPI Formula to Script fires automatically. It invokes the `NokiaMpnFormulaEngineSEP` extension point to process the raw formula before it is stored in the **Formatted KPI Formula** field.

The processing runs in two steps:

1. Cleanup: The raw value in **KPI Formula** is normalized into a JavaScript-evaluable expression:
 - Whitespace and line breaks are collapsed to a single space.
 - `Sum()` wrappers are removed; `Max()` and `Min()` are converted to `Math.max()` and `Math.min()`.
 - `**` exponentiation notation is converted to `Math.pow()` for compatibility with the MID Server JavaScript engine.
2. Validation — The processed formula is checked for balanced parentheses. If the parentheses are unbalanced, the raw value is still stored but a warning is logged. Records with unbalanced parentheses are skipped during metric collection at runtime.

The resulting value is written to the **Formatted KPI Formula** field and is the expression that the metric calculation engine references at runtime.

Note: If no `NokiaMpnFormulaEngineSEP` implementation is registered, the raw value in the **KPI Formula** value is stored unchanged and a warning is logged.

Related topics

[MPN Formulas table](#)

Metric-to-CI binding

Metric-to-CI binding associates metrics collected by a pull connector with configuration items (CIs) in the Configuration Management Database (CMDB). Default binding logic is used for each supported connector, which you can override with a custom implementation.

How metric-to-CI binding works

When a metric is collected, the platform runs the following sequence to associate each metric with a CMDB CI:

1. The connector pulls metric data from the vendor and stores it in the metric base (Clotho).
2. If the metric is not yet bound to a CI, the metric base framework automatically creates an event to record the request for a binding.
3. An event field mapping rule fires on the new event. Telecommunications Service Operations Management ships pre-configured event field mapping rules for each supported connector source.
4. The rule delegates to a scripted extension that resolves the appropriate CI for the event. The script updates the `cmdb_ci` field on the event with the resolved CI.
5. The binding appears in the Metric to CI mappings table. From that point on, all subsequent metric data for that source-resource pair is associated with the bound CI directly, without creating another event.

If the scripted extension can't resolve a CI, the event's `cmdb_ci` field is left empty and the metric data is stored in the metric base unbound.

Default behavior for MPN metrics

No additional configuration is required for CI binding. A scripted extension named `NokiaMPNMetricCIMapper` handles metric-to-CI binding for the MPN pull connector by firing on matching events through an event field mapping rule.

The extension tries the following event fields in order, stopping at the first match. The first lookup uses the CMDB table named in `additional_info.ciClass` on the event (defaulting to `cmdb_ci`); the remaining lookups always run against `cmdb_ci`:

MPN metric event CI lookup order

Order	Event field	Lookup behavior
1	<code>additional_info.name</code>	Matched against name on the class specified by <code>additional_info.ciClass</code>
2	<code>additional_info.pmDataSource</code>	Distinguished name traversal in <code>cmdb_ci</code> . The deepest matching component wins
3	<code>additional_info.pmDataSource</code>	Matched against <code>serial_number</code> in <code>cmdb_ci</code>
4	<code>additional_info.pmDataSource</code>	Matched against name in <code>cmdb_ci</code>
5	<code>additional_info.pmDataSource</code>	Matched against name in <code>cmdb_ci</code>

The `additional_info.name` and `additional_info.pmDataSource.dn` fields can each hold a comma-separated list of values. The extension tries the values from right to left, and the first value that resolves to a CI wins.

On a match, the extension sets the `cmdb_ci` and `ci_type` fields on the event.

Override the default binding behavior

If the shipped lookup logic does not match how your CMDB models the source data, you can override it by providing your own implementation of the `EventFieldMapping` extension point. The platform resolves the implementation at runtime by source name and rule name, so you can substitute your own logic without modifying product code.

For the procedure, see [Override default metric-to-CI binding](#).

Related topics

- [View metric to CI and resource binding](#) 
- [Event field mapping configuration](#) 
- [Create an event rule to bind metric events to host CIs](#) 
- [Configure elastic event pull connectors for MPN](#)
- [View metric values in the Insights Explorer](#) 
- [Create an Insights Explorer view](#) 
- [Resource binding](#) 

Metric aggregation modes

Metric aggregation rolls up raw metrics collected from individual network resources into calculated metrics on parent or anchor configuration items (CIs). The aggregation mode determines how source resources are selected and where the aggregated result is published.

By default, a pull connector collects raw metrics and publishes one value per network resource, such as a single metric value for each mobile cell. There is no built-in way to combine those per-resource values into a higher-level metric.

Metric aggregation addresses this gap. It reads the raw metric values for a set of related resources and applies an aggregate function such as average, maximum, or minimum. It then publishes a calculated metric so that the result can be viewed at a higher level of the network. An aggregation runs as a scheduled job that calls the metric aggregation scripted extension point.

Key benefits

Metric aggregation provides the following benefits:

- View metrics at higher levels of the network, according to Baseband Unit (BBU) or per Network Site, instead of only per cell.
- Aggregate metrics across a defined range of resources, such as a range of SIM cards.
- Define new aggregations without changing product code by configuring scheduled jobs that call a scripted extension point.

How it works

Each aggregation runs in one of two modes. The mode determines how source resources are selected and where the calculated metric is published.

Hierarchy mode

The aggregation traverses configuration management database (CMDB) relationships downward from each target CI to collect the matching source CIs, then writes one calculated metric onto each target CI. Use hierarchy mode to roll per-cell metrics up to each parent resource, according to BBU or per Network Site.

Flat mode

The aggregation selects source CIs directly, optionally filtered by a name range or by field values, and writes a single calculated metric onto one anchor CI. Use flat mode for range-based or fleet-wide metrics where no single parent CI exists to hold the result.

A range identifies a contiguous set of resources by name, such as a range of SIM cards. You can also apply field filters so that only resources meeting a condition are included in the aggregation. For example, you can filter to include only resources whose operational status indicates that they are active.

An aggregation runs one of the following functions, set through the `aggregate` parameter: average (`avg`), sum (`sum`), maximum (`max`), minimum (`min`), count (`count`), and the 95th percentile (`p95`).

After an aggregation run produces a calculated metric, the metric base framework raises an event to bind the metric to a CI. After the metric-to-CI binding completes, the calculated metric is stored in the metric base (Clotho) against the bound CI and becomes available to view.

Considerations

Consider the following when you plan an aggregation:

- In hierarchy mode, a depth limit bounds how far the traversal descends through the CMDB. The depth limit helps prevent the traversal from following relationships beyond the resources you intend to aggregate.
- In flat mode, the calculated metric is not tied to an existing parent CI, so you specify an anchor CI to hold the result.
- An aggregation runs on the schedule defined for its scheduled job. Calculated metrics appear after the next successful run and metric-to-CI binding.

Related topics

[Configure a metric aggregation job](#)

[Create a custom metric aggregation implementation](#)

[Metric aggregation configuration reference](#)

[Metric aggregation scripted extension](#)

[Metric-to-CI binding](#)

[MPN data model](#)

Metric aggregation scripted extension

You can define custom KPI calculations on top of raw metrics collected from source systems. Configurable formulas combine, aggregate, and label metrics without requiring code changes to the platform.

By default, TSOM Assurance collects raw metrics from pull connectors and publishes them as-is. The metric aggregation scripted extension provides an extension point on the aggregation script, allowing you to implement custom logic that derives KPIs from those raw metrics. One or more KPIs can be calculated from a single metric value, and KPI definitions can be added, updated, or deleted on demand without redeployment.

Key capabilities

Configurable formulas

Define KPIs using mathematical expressions that combine one or more raw or already-transformed metrics. Formulas are vendor-agnostic and can be conditioned on resource attributes such as vendor or software version.

Temporal aggregation

Apply aggregation operators over a configured time period to produce a single derived value per metric. Supported operators are `avg`, `max`, `min`, and `p95`. Supported periods are 5 minutes, 15 minutes, and 1 hour.

Spatial aggregation

Aggregate the same KPI or metric across multiple resources that match a filter criterion based on metric or resource attributes. The supported aggregate functions are `avg`, `sum`, `max`, `min`, `count`, and `p95`.

User-defined labels

Attach custom labels to calculated KPIs to support downstream categorization, routing, or display logic.

Formula error handling

When a formula can't be evaluated (for example, due to a division by zero or missing input metric), the system generates an error code for that KPI. This prevents values from being discarded without notification.

Implementation

The extension point is implemented using the `sn_tsom_em_conns.TSOMMetricAggregatorSEP` scripted extension point. To create a custom implementation, navigate to **All > System Scripted Extension Points > Scripted Extension Points** and search for `sn_tsom_em_conns.TSOMMetricAggregatorSEP`. Select **Create implementation** and define your KPI formulas in the script body.

KPI definitions, including formulas and resource attribute filters, can be configured and managed independently. Add, update, or delete definitions to adapt KPI coverage as your network environment changes.

Configuring Telecommunications Service Operations Management

Configure Telecommunications Service Operations Management (TSOM) to enable real-time event ingestion, correlation, and automated remediation by integrating with external network monitoring and discovery systems.

Set up TSOM to enable end-to-end telecom service operations, including alarm ingestion, CMDB population, discrepancy detection, and service impact visibility. This configuration involves activating Telecommunications API notifications, setting up discovery and service graph connectors, enabling visibility, and configuring audit and reconciliation frameworks.

Configuration overview

Telecommunications Service Operations Management (TSOM) requires configuring multiple components across alarm ingestion, discovery, data normalization, and reconciliation to support end-to-end telecom operations. The configuration flow typically includes the next plugins (The plugin dependency will be automatically resolved during the installation).

1. Install the following plugins:

TSOM plugins

Name	ID	Product/Service
TSOM Event Management Connectors	sn_tsom_em_connectors	Telecom Assurance
Service Graph Connector for Fortinet Telco SD-WAN	sn_tsom_fortinet_connector	Telecom Visibility
Service Graph Connector for Meraki Telco SD-WAN	sn_tsom_meraki_connector	Telecom Visibility
Service Graph Connector for VeloCloud Telco SD-WAN	sn_tsom_vcloud_connector	Telecom Visibility

2. Enable alarm ingestion: Activate the Telecommunications API notifications and set up topic subscriptions to receive alarms from external systems. For more information, see [Configuring Telecommunications API notifications](#).
3. Set up visibility or Event Management by following the appropriate setup wizards. See [Configure Telecom Visibility](#) and [Fault Management: Events and alerts](#).

Each step is modular and can be configured based on your environment and available integrations.

Configure Telecom Visibility

Telecom Visibility provides foundational capabilities that support both Telecom Discovery and Telecom Discrepancy Identification & Reconciliation. It includes shared logic, enhanced CI class models, and Identification and Reconciliation Engine (IRE) updates tailored for telecom network elements.

Before you begin

Role required: tsom_visibility_admin

Confirm that:

- Your ServiceNow AI Platform instance is licensed for Telecommunications Service Operations Management.
- You review any custom IRE rules applied to telecom CIs to help prevent overrides during upgrade.

About this task

To configure Telecom Visibility, you must install the required plugins and update the CMDB CI Class Models to version 1.82 or later. This version introduces telecom-specific IRE identification rules. This ensures accurate CI identification and reconciliation across telecom domains.

Note: If you have customized IRE identification rules for any of the affected telecom CIs, upgrading to version 1.82 or later may override or impact those rules.

The following capabilities are included with the Telecom Visibility subscription:

Plugin Name	Plugin ID	Plugin Description	Store App
Service Graph Connector for Nokia Altiplano	sn_sgc_altiplano	Service Graph Connector for Nokia Altiplano	Yes
Service Graph Connector for Meraki Telco SD-WAN	sn_sgc_meraki	Service Graph Connector for Cisco Meraki	Yes
Service Graph Connector for Fortinet Telco SD-WAN	sn_sgc_fortinet	Service Graph Connector for Fortinet	Yes
Service Graph Connector for VeloCloud Telco SD-WAN	sn_sgc_vcloud	Service Graph Connector for Arista VeloCloud	Yes
Telecommunications Discovery Patterns	sn_tsom_patterns	Telecommunication Discovery Patterns	Yes
Telecom Visibility Plugin (Core logic)	sn_tsom_core	Core logic for Telecom Visibility and Discrepancy Reconciliation	No

Procedure

1. Navigate to **All > Plugins**.

2. Search for and install the following:

- a. Telecommunication Discovery Patterns (sn_tsom_patterns)
- b. Telecom Visibility (sn_tsom_core)
- c. Service Graph Connector for your vendor.

3. **Optional:** Update the CMDB CI Class Models.

- a. Navigate to **All > Applications** list in the ServiceNow Store.
- b. Search for CMDB CI Class Models (sn_cmdb_ci_class).
- c. Install or upgrade to version 1.82 or later.

Note: Installing any of the TSOM plugins listed above will automatically update or install version 1.82 or later of the CMDB CI Class Models app. If your instance doesn't include TSOM plugins or you're on a pre-Yokohama release (for example, Washington DC or Xanadu), you can manually install or upgrade the store app.

4. Navigate to **CMDB > Identification Rules**.

a. Review rules related to:

- cmdb_ci_interface_card
- cmdb_ci_slot
- cmdb_ci_subslot
- cmdb_ci_network_adapter

b. Validate that your custom logic remains functional.

c. Update or merge custom rules as needed.

Result

Confirm that the following are listed as active:

- Telecom Visibility (sn_tsom_core)
- Telecommunication Discovery Patterns (sn_tsom_patterns)
- Service Graph Connector for your vendor
- CMDB CI Class Models version 1.82 or later

Related topics

[Telecom Visibility](#)

Install Horizontal Discovery and set up Discovery Patterns

Install Horizontal Discovery patterns understanding the dependencies and requirements.

Before you begin

Role required: admin

Ensure you have a subscription to TSOM.

About this task

Visit the [ServiceNow Store](#) website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

Dependencies and Requirements:

- Telecom Visibility Core (sn_tsom_core)
- Discovery Core plugin (com.snc.discovery.core), which is automatically installed by Discovery.
- ITOM Discovery License plugin (com.snc.itom.discovery.license). You must activate this plugin.
- ITOM Licensing plugin (com.snc.itom.license). For more information, see [Request Discovery](#).

Pattern execution logic: By default, when a TSOM pattern is run, it executes both the TSOM-specific and the corresponding ITOM pattern (e.g., Telco Router runs Router). This ensures shared ITOM libraries are reused when needed. To override this behavior:

- Use the system property sn_tsom_patterns.itom_pattern_enabled.
- Setting this property to false ensures only the TSOM-specific pattern is executed.

Procedure

1. Install the Horizontal Discovery application.
See [Discovery setup](#), as it is foundational for running Telecommunications Discovery Patterns.
2. Obtain and install Telecommunications Discovery Patterns:
 - a. Install the Telecommunications Discovery Patterns (sn_tsom_patterns) from the ServiceNow® Store.
3. Set up a MID Server and synchronization Patterns:

- a. Synchronization the installed patterns with the appropriate MID Servers to confirm they're ready for use:

- i. Navigate to **Discovery > MID Servers**.
- ii. Select **Pattern Sync to Mid**.

 **Note:** This action synchronizes both TSOM and ITOM patterns.

For more information on how to configure a MID Server, see [Configuring MID Server](#) .

4. Configure TSOM System Properties:

- a. Set the system property `sn_tsom_patterns.itom_pattern_enabled` to define the logic for whether to use only the TSOM Pattern or a combination of ITOM and TSOM patterns.
 - i. Navigate to **All > System Properties > All Properties**.
 - ii. Select **`sn_tsom_patterns.itom_pattern_enabled`**.
 - iii. Check that the Value is set to **true** (default).

If you want TSOM to run only TSOM patterns and exclude ITOM patterns, set the Value to **false**.

 **Note:** The default setting is configured to use both TSOM and ITOM patterns.

5. Enable the replacement of various ITOM patterns with TSOM patterns on a specific MID Server:

For example: The Telecom Router pattern replaces the Network Router pattern for a specific MID Server when **`mid.telecom.discovery.patterns.enabled`** is set to true for that MID Server.

- a. Go to the **Filter Navigator** and type **`ecc_agent_config.list`**.
- b. Select **`mid.telecom.discovery.patterns.enabled`** (each MID Server has this parameter).
- c. Check that the Value is set to **true**.

Repeat this configuration for each MID Server that you want to use for running TSOM patterns.

Related topics

[Direct Discovery using Discovery Patterns](#)

Activate Telecom Discrepancy and Reconciliation

Activate the Telecom Discrepancy Identification and Reconciliation feature to ensure consistency between telecom network resources and CMDB or Telecom Network Inventory (TNI) data. This feature helps detect and remediate mismatches automatically, supporting service accuracy and data integrity across your telecom environment.

Before you begin

Role required: admin

To use this feature, your organization must have an active subscription to TSOM. Telecom Discovery, Telecom Visibility, and this feature are licensed together.

About this task

Plugin dependencies: verify the following plugins are installed and activated:

Plugin	ID / App ID	Type
Telecom Service Operation Core	sn_tsom_core	Store
CMDB CI Class Models	sn_cmdb_ci_class	Store
Expanded Model and Asset Classes	sn_ent	Store
Visibility Content	sn_pattern_design	Store
Integration Commons for CMDB	sn_cmdb_int_util	Store
Discovery Core	com.snc.discovery.core	Family
ITOM Discovery License	com.snc.itom.discovery.license	Family
ITOM Licensing	com.snc.itom.license	Family

Procedure

1. Ensure the sn_tsom_core plugin is automatically installed when you install Telecommunications Discovery Patterns or the Nokia Altiplano Service Graph Connector.
2. Ensure the mentioned dependent plugins are activated.

Related topics

- [Discrepancy identification – types of discrepancies](#)
- [Telecom discrepancy identification and reconciliation](#)
- [Telecom Discrepancy Identification and Reconciliation](#)
- [Run Telecom Discrepancy audit](#)
- [Control CI attribute updates using Reconciliation rules](#)

Configure filter for audit

Define filtering conditions to control the scope of Telecom Discrepancy audits. These filters ensure that audits run only on the desired set of CIs, improving performance and targeting accuracy.

Before you begin

Role required: admin

About this task

Understanding Filtering Conditions in Telecom Discrepancy Audits

- Audits use filtering conditions to narrow the scope of CIs being evaluated.
- Filtering conditions are essential for narrowing down audit targets.
- You can define and modify conditions per audit instance to meet specific audit requirements.
- Filtering by `discovery_source` is a common condition to filter records associated with particular integration sources (e.g., TSOM and Altiplano). Additionally, you can customize the filter according to the audit requirements.
- Each version of the audit can have a different filtering condition.

The following are the audit filters:

1. **Telecom Logical Connections Discrepancy Audit:** This audit runs at the logical connection table. The default filter condition is defined as `discovery_source` like TSOM. This condition is applied on each CI in the logical connection table.
2. **Telecom Network Topology Discrepancy Audit:** This audit runs at the network topology table. The default filter condition is defined as `discovery_source` like TSOM. This condition is applied on each CI in the network topology table.
3. **Telecom Discrepancy Audit:** This audit runs at the relationships table. The default filter condition is defined as `discovery_source` like TSOM. This condition is applied on each CI at the equipment level.

Procedure

1. Navigate to **All > Certification > Filters**.
2. Select **New**.
The certification filter interface appears.
3. Fill in the fields.
For more guidance, see [Create a filter](#) .
4. Add the filter conditions.
For example, Discovery source contains TSOM.
5. Select **Submit**.
The filter is created.

Result

You can use the filter in the logical connection audit.

Related topics

[Run Telecom Discrepancy audit](#)
[Example for Telecom Discrepancy Audit and Remediation](#)
[Telecom discrepancy identification and reconciliation](#)
[Configure attribute value discrepancy in CMDB 360](#)
[Control CI attribute updates using Reconciliation rules](#)

Configure attribute value discrepancy in CMDB 360

Configure attribute comparison settings in CMDB 360 to detect data inconsistencies across multiple discovery sources.


Before you begin

Role required: admin

Confirm that CMDB 360 is enabled and configured in your instance.

About this task

The following screenshot can help you to set up the fields to identify the



The screenshot displays the 'CMDB MultiSource Column Metadata' configuration page for the table 'cmdb_ci_ni_interface'. It includes the following fields and options:

- * MultiSource Column:** A dropdown menu with 'Col 12' selected.
- Application:** A dropdown menu with 'Telecom Service Operations Core' selected.
- * Table:** A dropdown menu with 'Network Interface [cmdb_ci_ni_interface]' selected.
- * Table field:** A dropdown menu with 'Port bandwidth' selected.

At the bottom left, there are 'Update' and 'Delete' buttons. At the bottom right, there is a link to '[SN Utils] Versions (1)'.

discrepancy. [\[SN Utils\] Versions \(1\)](#)

Procedure

1. Navigate to the table **All > cmdb_multisource_column_metadata.list**.
2. Select **New** to create a new record.
3. In the **MultiSource Column** field, select the attribute to compare.
4. Select the table to compare the attributes to the fields.

Note: If the table you need isn't listed, select and hold (or right-click) on the Table label and select **Configure Dictionary**.

5. **Optional:** To configure a dictionary entry, do the following.
 - a. In the dictionary entry, select **View** and then select **Advanced**.
 - b. In the **Attributes** field, set the attributes as `base_start=true` and `allow_public=true`.

Result

You can use the configured attributes in the CMDB 360 query.

Related topics

[Control CI attribute updates using Reconciliation rules](#)

[Generate reports for attribute value discrepancies](#)

Control CI attribute updates using Reconciliation rules

To prevent specific attributes of a configuration item (CI) from being overwritten by Discovery or other data sources, use Reconciliation Rules. These rules define which data source is trusted to update a particular attribute when multiple sources provide values.

Before you begin

Role required: admin

About this task

Reconciliation Rules are processed by the Identification and Reconciliation Engine (IRE) and are essential for maintaining data integrity in the CMDB.

Use Case: You can prevent Discovery from updating specific attributes while allowing other sources to update them. Define a rule that excludes the Discovery source and permits updates from sources like SCCM or manual entry.

Note: Only one Reconciliation Rule should be active for a specific CI class and attribute combination to avoid conflicts.

Procedure

1. Navigate to **All > CI Class Manager**.
2. Select the target CI class (e.g., `cmdb_ci_computer`).
3. In the **Reconciliation Rules** tab, click **New** to create a rule.
4. Provide the rule name and select the Discovery source to apply the rule to.
5. Specify the attributes that this rule will govern.
6. Set the source precedence to allow only the selected sources to update the specified attributes.

Result

Certain CI attributes will no longer be updated by untrusted or lower-priority data sources like Discovery. Only the sources defined in the rule (e.g., SCCM, manual updates) will be allowed to update those attributes.

Related topics

[Configure attribute value discrepancy in CMDB 360](#)

Configure the Telecom Discovery Builder ETL

Duplicate the prebuilt Telecom Discovery Builder framework ETL into your connector scope. Assign the appropriate data source and deploy it as part of a new service graph connector.

The Telecom Discovery Builder framework ETL is provided with the Telecom Service Operations Management (TSOM) Core. This ready-to-use framework streamlines data ingestion into the CMDB across telecom connectors. Rather than configuring it directly, duplicate the ETL into a connector's application scope. Then update the data source to align with the connector's discovery payload.

Steps to use the Telecom Discovery Builder framework ETL in a connector

1. Create a temporary data source: Create a data source in your connector's application scope based on the Generic Schema multi source data source from TSOM Core. This enables you to run and test the duplicated ETL with schema-aligned placeholder data. For more information, see [Create a data source similar to Telecom core data source](#)
2. Duplicate the Telecom Discovery Builder framework ETL: Access ETL Studio, locate the Telco Generic Schema ETL, and duplicate it into the connector's scope. During duplication: Provide a new name. Assign the temporary data source. Use importSet as the discovery source. For more information, see [Duplicate the Telecom Discovery Builder ETL](#)
3. After duplication, open the new ETL and replace the temporary data source with the connector's actual discovery data source. For more information, see [Update the data source of the connector](#).
4. Test or deploy the service graph connector. For more information, see [Deploy service graph connector with existing ETL](#)

Create a data source similar to Telecom core data source

Set up a schema-compliant data source in the connector's application scope to simulate telecom data. Test the Telecom Discovery Builder framework ETL before integrating with live device data.

Before you begin

Role required: admin

Ensure the following:

- Access to the TSOM Core application and its data sources.
- Switch to the connector's application scope.
- Have a sample payload that conforms to the Telco Generic Schema (optional, but recommended for testing).

About this task

When duplicating the Telecom Discovery Builder framework ETL into a Service Graph Connector (SGC) application scope, you must first create a data source. This data source must replicate the exact structure of the TSOM Core data source. It provides the required schema and structure for testing and validating the ETL before connecting it to live telecom data.

The Telecom Discovery Builder framework expects data that conforms to the Telco Generic Schema, as defined in the Telecom Core. Duplicating the original TSOM data source ensures:

- Schema alignment during ETL duplication.
- A valid import set structure for testing.
- Separation of core and connector scopes for customization and upgrade safety.

When to create a data source similar to the TSOM Core data source:

- Before duplicating the Telecom Discovery Builder framework ETL.
- When you want to run a test load using simulated or placeholder telecom data.
- When preparing the connector's application scope for ingestion configuration.

The following screenshot can help you understand to fill the field values while you create the data

Data Source
New record

Application: Service Graph Connector for NOKIA NSP

* Name: NSP Generic Data Source

Import set table label: NSP Generic Data Source

* Import set table name: sn_sgc_nsp_generic_data_source

* Type: Custom (Load by Script)

Data in single column: ☒

Use Batch Import: ☐

Enable parallel loading: ☐

Properties:

Data Loader: Turn on ECMAScript 2021 (ES11) mode

```
(function loadData(import_set_table, data_source, import_log, last_success_import_time) {  
  
    const LOGICAL_COMPOSITES = '{ "logical_composites": [{ "name": "generated_SR_7s_core1_LogicalComposite",  
        "class": "logical_composite", "key": "NSD08E595F50:logicalComposite", "description": null,  
        "power_units": [{ "name": "Power Shelf-1(ps-b10-shelf-ac/hv)", "class": "pdpu", "unit_position": 1, "serial_number": "NSA47C6E8507",  
            "model_name": "ps-b10-shelf-ac/hv", "model_number": null, "key": "NSD08E595F50:NSD08E595F50:powerShelf1/powerShelf1",  
            "description": null, "slots": [{ "name": "Power Module Slot-1/1(ps-b-ac/hv-6000)", "class": "slot", "description": null,  
                "model_name": "pm-s1-slot", "model_number": null, "manufacturer": "Nokia", "unit_position": 1, "key": "NSD08E595F50:powerShelf1slot1/  
powerShelf1/powerModuleSlot1"}, {"card": { "name": "Power Module-1(ps-b-ac/hv-6000)", "class": "card", "unit_position": 1,  
                    "key": "NSD08E595F50:powerShelf1slot1/powerModuleSlot1/powerModule1", "model_name": "ps-b-ac/hv-6000",  
                    "model_number": null, "manufacturer": "Nokia", "firmware_version": null, "mac_address": null, "description": null,  
                    "serial_number": "NS15E40AD9A1"} } ] } ] }';  
  
    const DEVICES = '{ "devices": [{ "description": "Modular device, 1 card in slot 1, 1 MDA card in subplot 1, 36 connector ports  
with 2 transceivers in port 1/1/C1 and 1/1/C2, 2 connector ports with 1 breakout ports each", "class": "device",  
        "model_name": "T750 SR-7s", "model_number": null, "manufacturer": "Nokia", "mac_address": "80:B0:90:46:F1:11:31", "ip_address": "11.128.  
28.27", "name": "generated_SR_7s_core1", "serial_number": "NSD08E595F50X", "key": "NSD08E595F50X", "firmware_version": "T1M0S-C-23.  
10.R4-1", "ports": [{ "name": "device_port_1", "class": "port", "model_name": "Ethernet", "description": null, "ip_address": "11.128.28.  
101", "unit_position": 1, "key": "NSD08E595F50X:device_port_1", "virtual": false, "bandwidth_name": null,  
            "bandwidth_group": "Ethernet", "bandwidth": 1400000, "mtu_size": 1400}, {"name": "device_port_1-BNET", "class": "port",  
                "model_name": "BNET", "description": null, "ip_address": "11.128.28.261", "unit_position": 1,  
                "key": "NSD08E595F50X:device_port_1-BNET", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD08E595F50X:lag_1-VLAN", "virtual": true, "bandwidth_name": null, "bandwidth_group": null,  
                "mtu_size": null, "name": "lag_1-VLAN", "class": "port", "model_name": "VLAN", "description": null, "ip_address": "11.128.28.251",  
                "unit_position": 1, "key": "NSD0
```

source.

Procedure

1. Navigate to **All > System Import Sets > Administration > Data Sources**.
2. Find the **Generic Schema Multi Source v2** or a similar baseline source provided with TSOM Core.
3. Open the TSOM Core data source record and copy the script from the **Data Loader** field.
4. Create a data source by clicking **New**.
5. On the form, fill in the fields
For more information, see [Create a Custom \(Load by Script\) type data source](#) .
6. In the **Type** field, select **Custom (Load by Script)**.
7. Select the **Data in single column** field.
8. In the **Data Loader** field, paste the copied script.
9. Select **Submit**.
The data source is created.
10. **Optional:** To test load the data source:

- a. Click Test Load 20 Records (or similar) to generate an import set.
- b. Ensure that records are created without errors.
- c. Confirm that classes like Logical Composite, Network Gear, or Port appear in the staging table.

What to do next

After the data source is created and tested:

- Use it as the import source when duplicating the Telco Generic ETL.
- Replace it with the actual connector-specific data source once simulation and validation are complete.

Related topics

[Standardized JSON data set for service graph connectors](#)

[Duplicate the Telecom Discovery Builder ETL](#)

Duplicate the Telecom Discovery Builder ETL

The Telecom Discovery Builder framework ETL enables Service Graph Connector (SGC) teams to rapidly adopt a standardized, schema-compliant data ingestion pipeline without building ETL logic from scratch.

Before you begin

Role required: admin

Ensure that:


- The Telecom Core plugin is activated and the Telco Generic ETL v2 (automatically installed) is available.
- You have admin access to IntegrationHub ETL Studio.
- You have created a temporary data source in the connector's application scope. After creating a temporary data source, you can duplicate the Generic ETL into your connector's application scope to customize and extend it for your specific use case.

About this task

Duplicating the Telco Generic ETL allows you to:

- Reuse standardized ETL mappings across multiple connectors.
- Customize ETL behavior without altering the original baseline provided by the Telecom Service Operations Management (TSOM) Core.
- Align with the Telco Generic Schema for consistency and TNI compliance.
- Save time and reduce errors by working with a tested and proven ETL framework.

The following screenshot can help you duplicate the ETL Transform



Procedure

1. Navigate to **All > IntegrationHub ETL**.

2. On the **IntegrationHub ETL** page, select the connector's application scope.

Note: For example, **Service Graph Connector for Nokia NSP**.

3. Expand the CMDB Application Telco Generic Schema and select **Telco Generic Schema** record.

4. Click **Duplicate**.

The Duplicate ETL Transform Map page appears.

5. Configure the duplicated ETL transform map

- In the **Duplicate From** field, ensure Telco Generic Schema is selected.
- Select an existing CMDB Application or click **Add new...** to add a new CMDB Application.
- For a new CMDB Application, enter the CMDB application name (for example, Nokia NSP Generic ETL v2)
- In the **Discovery Source** field, select the import set.
- In the **Name** field, enter a name for the duplicate ETL.
- In the **Description** field, enter a description for the duplicate ETL.
- Select the data source that is used for the duplicate ETL transform map.

Note: This needs to be different than the existing default Data Source that is attached to the Telco Generic Schema ETL. For more information, see [Create a data source similar to Telecom core data source](#).

h. Enable **Auto-pull a new import set** option to automatically pull data into a new import set.

i. In the **Preview Size Override** field, set a custom preview size for testing and validation.

6. Select **Create Duplicate.**

A duplicate ETL of Telco Generic ETL is created.

7. Optional: Open the duplicated ETL record and review the mappings and settings.**What to do next****1. Test the duplicated ETL (optional but recommended):**

- Run a test load or a simulation using the temporary data source.
- Verify that import sets are processed successfully, CIs are created according to the Telco Generic Schema, and relationships are established properly.

2. After successful duplication and testing:

- Update the ETL's data source configuration to point to the actual production data source (for real device data).
- Deploy the connector integration into a test or production environment.
- Monitor import runs to validate that inventory data is ingested into the CMDB correctly.

Related topics

[Update the data source of the connector](#)

Example - Duplicate the Telco Generic ETL Schema

This example walks you through how to duplicate the Telco Generic Schema ETL to set up a customized Service Graph Connector (SGC) ETL for your telecom integration. Use this procedure to create a baseline ETL in your connector's application scope based on the standardized Telco Generic Schema.

Scenario

You are deploying a new Service Graph Connector and need to duplicate the Telco Generic Schema ETL into your connector's application scope.

Create a temporary data source, duplicate the ETL, and configure it to work with your connector's device payloads.

Steps to duplicate the Telco Generic Schema ETL

- 1.** Switch to the connector's application scope where you want to duplicate the ETL (e.g., Nokia NSP Connector).
- 2.** Create the data source:
 - Navigate to System Import Sets > Administration > Data Sources.
 - Locate the Generic Schema Multi-Source Data Source provided by the TSOM Core application.
 - Copy this data source into your connector's application scope.
 - Test the copied data source by running Test Load 20 Records to create a sample import set.
- 3.** Open the Duplicate ETL Transform Map Dialog: In ETL Studio, select Duplicate ETL to start the duplication process.
- 4.** In the **Duplicate from** list, choose Telco Generic Schema.
- 5.** Select **Add new...** and enter a name for the duplicated ETL.
- 6.** Select importSet as the discovery source for your duplicated ETL.

7. Enter a new name for the duplicated transform map.
8. Specify the newly created temporary data source that you copied from the Generic Schema Multi Source.
9. Optionally, enable **Auto-pull a new import set** to automatically load new records after duplication.
10. Click **Create Duplicate** to complete the duplication.
11. Update **Basic Details**:
 - After duplication, open the newly created ETL.
 - In step 1: specify **Basic Details**, replace the temporary data source with your connector's production data source.
12. Save the ETL Configuration: click **Save** to finalize your changes.
13. Run the **Data Source**: From Import Schedules, run your connector's data source to ingest real device data.

Result: The system creates a new ETL based on the Telco Generic Schema settings and processes the payloads from your specified connector data source. The imported data is transformed into Configuration Items (CIs) and inserted into the ServiceNow CMDB with the expected relationships and structures.

Update the data source of the connector


Link the duplicated ETL to a valid data source for your specific connector.

Before you begin

Role required: admin

About this task

The following screenshot can help you to replace the default source with your connector's data



Procedure

1. Navigate to **All > IntegrationHub ETL**.
2. In **Specify Basic Details**, replace the default data source with your connector's data source.
3. Select **Save**.
4. Open **Import Schedules** and run your data source.
The system creates a new transform map based on the selected settings and create CIs generated from your payload.

Related topics

[Deploy service graph connector with existing ETL](#)

Deploy service graph connector with existing ETL

Associate a new Service Graph Connector (SGC) with an existing ETL.

Before you begin

Role required: admin

Ensure the required Service Graph Connector plugin is activated.

Procedure

1. Navigate to **All > IntegrationHub ETL**.
2. Open the duplicated ETL configuration.
3. Select the data source of the service graph connector and click **Save**.
4. Select **Import Schedules** to execute the ETL to transform and load the data from the new SGC.
5. Monitor execution status.

Result

Verify that CI records and relationships are created in the CMDB using CMDB Maps or list views.

Extend TNI entity support for duplicated ETLs

When Telecom Network Inventory (TNI) is enabled, every Configuration Item (CI) created by a duplicated ETL must have a corresponding TNI Entity.

Before you begin

Role required: admin

- Ensure TNI is installed and active in your instance.
- Complete the duplication of the Telecom Discovery Builder framework ETL into the target application scope.

Procedure

1. Navigate to **All > IntegrationHub ETL**.
2. Select the CMDB Application associated with the duplicated ETL.
The CMDB Integration Studio application in a new page
3. In the ETL configuration page, check the **Execute Before Script** option.
4. Replace the default script with the following:

```
(function(input, runId) {
    new
    sn_tsom_core.TelcoGenericMappingHelper().checkAndUpdateIrePayloadForTni(input);
})(input, runId);
```

5. Click **Update** to save the changes.
The duplicated ETL links TNI entities for discovered CIs.

Standardized JSON data set for service graph connectors

Use the TSOM architecture to support standardized Service Graph Connectors using a unified schema and reusable ETL logic. This reduces onboarding time for new connectors and simplifies integration with CMDB.

1. Define the common JSON Schema

Standardize the output of all service graph connectors to align with a single JSON format that conforms to the TNI schema.

Ensure the following points:

1. Implement conversion logic in the collector or adaptor to output data in the common schema.
2. Ensure:
 - Slot-on-slot or card-on-card hierarchies are excluded.
 - Logical interfaces are clearly marked with `virtual=true`.
 - Equipment types align to model-class mappings.

Note: The schema should support runtime adaptability to TNI changes if available.

The following is the JSON schema common data set that supports all TSOM service graph connectors:

```
{
  "$schema": "http://json-schema.org/draft-07/schema#",
  "title": "Telco Generic Schema",
  "version": "2.0.1",
  "oneOf": [
    {
      "type": "object",
      "properties": {
        "logical_composites": {
          "type": "array",
          "items": {
            "$ref": "#/$defs/logical_composite"
          }
        }
      },
      "required": [
        "logical_composites"
      ],
      "additionalProperties": false
    },
    {
      "type": "object",
      "properties": {
        "devices": {
          "type": "array",
          "items": {
            "$ref": "#/$defs/device"
          }
        }
      },
      "required": [
        "devices"
      ]
    }
  ]
}
```

```

    ],
    "additionalProperties": false
  },
  {
    "type": "object",
    "properties": {
      "logical_connections": {
        "type": "array",
        "items": {
          "$ref": "#/$defs/logical_connection"
        }
      }
    },
    "required": [
      "logical_connections"
    ],
    "additionalProperties": false
  },
  {
    "type": "object",
    "properties": {
      "port_relations": {
        "type": "array",
        "items": {
          "$ref": "#/$defs/port_relation"
        }
      }
    },
    "required": [
      "port_relations"
    ],
    "additionalProperties": false
  },
  {
    "type": "object",
    "properties": {
      "logical_connection_relations": {
        "type": "array",
        "items": {
          "$ref": "#/$defs/logical_connection_relation"
        }
      }
    },
    "required": [
      "logical_connection_relations"
    ],
    "additionalProperties": false
  },
  {
    "type": "object",
    "properties": {
      "numbers": {
        "type": "array",
        "items": {
          "$ref": "#/$defs/number"
        }
      }
    }
  }
}

```

```

    },
    "required": [
        "numbers"
    ],
    "additionalProperties": false
},
{
    "type": "object",
    "properties": {
        "topologies": {
            "type": "array",
            "items": {
                "$ref": "#/$defs/network_topology"
            }
        }
    },
    "required": [
        "topologies"
    ],
    "additionalProperties": false
},
{
    "type": "object",
    "properties": {
        "topology_relations": {
            "type": "array",
            "items": {
                "$ref": "#/$defs/network_topology_relation"
            }
        }
    },
    "required": [
        "topology_relations"
    ],
    "additionalProperties": false
}
],
"$defs": {
    "keyRef": {
        "type": "object",
        "properties": { "key": { "type": "string" } },
        "required": [ "key" ],
        "additionalProperties": false
    },
    "optionalKeyRef": { "type": [ "object", "null" ],
        "properties": {
            "key": { "type": "string" }
        },
        "additionalProperties": false
    },
    "value": {
        "type": "object",
        "properties": {
            "from": { "type": [ "integer" ], "default": 0,
            "minimum": 0 },

```

```

    "to": { "type": [ "integer" ], "default": 0, "minimum":
0 }
  },
  "required": [ "from", "to" ],
  "additionalProperties": false
},
"logical_composite": {
  "type": "object",
  "properties": {
    "class": { "type": "string", "enum":
[ "logical_composite" ] },
    "key": { "type": "string" },
    "name": { "type": [ "string", "null" ] },
    "description": { "type": [ "string", "null" ] },
    "devices": { "type": "array", "items": { "$ref":
"#/$defs/keyRef" } },
    "power_units": { "type": "array", "items": { "$ref":
"#/$defs/pdu" } },
    "fan_shelves": { "type": "array", "items": { "$ref":
"#/$defs/fan_shelf" } } },
    "required": [ "key", "name" ]
  },
  "pdu": {
    "type": "object",
    "properties": {
      "class": { "type": "string", "enum": [ "pdu" ] },
      "key": { "type": "string" },
      "name": { "type": [ "string", "null" ] },
      "description": { "type": [ "string", "null" ] },
      "model_name": { "type": [ "string", "null" ] },
      "model_number": { "type": [ "string", "null" ] },
      "unit_position": { "type": [ "integer", "null" ] },
      "minimum": 1 },
      "slots": { "type": "array", "items": { "$ref":
"#/$defs/slot" } }
    },
    "required": [ "key", "name" ]
  },
  "fan_shelf": {
    "type": "object",
    "properties": {
      "class": { "type": "string", "enum": [ "fan_shelf" ] },
      "key": { "type": "string" },
      "name": { "type": [ "string", "null" ] },
      "description": { "type": [ "string", "null" ] },
      "slots": { "type": "array", "items": { "$ref":
"#/$defs/slot" } }
    },
    "required": [ "key", "name" ]
  },
  "device": {
    "type": "object",
    "properties": {
      "class": { "type": "string", "enum": [ "device" ] },
      "key": { "type": "string" },
      "name": { "type": [ "string", "null" ] },
      "description": { "type": [ "string", "null" ] },

```

```

        "ip_address": { "type": [ "string", "null" ] },
        "mac_address": { "type": [ "string", "null" ] },
        "serial_number": { "type": [ "string", "null" ] },
        "model_name": { "type": [ "string", "null" ] },
        "model_number": { "type": [ "string", "null" ] },
        "manufacturer": { "type": [ "string", "null" ] },
        "firmware_version": { "type": [ "string", "null" ] },
        "slots": { "type": "array", "items": { "$ref":
"#/$defs/slot" } },
        "ports": { "type": "array", "items": { "$ref":
"#/$defs/port" } } },
        "required": [ "key", "name", "serial_number" ]
    },
    "slot": {
        "type": "object",
        "properties": {
            "class": { "type": "string", "enum": [ "slot" ] },
            "key": { "type": "string" },
            "name": { "type": [ "string", "null" ] },
            "description": { "type": [ "string", "null" ] },
            "model_name": { "type": [ "string", "null" ] },
            "model_number": { "type": [ "string", "null" ] },
            "manufacturer": { "type": [ "string", "null" ] },
            "unit_position": { "type": [ "integer", "null" ],
"minimum": 1 },
            "cards": { "type": "array", "items": { "$ref":
"#/$defs/card" } }
        },
        "required": [ "key", "name" ]
    },
    "card": {
        "type": "object",
        "properties": {
            "class": { "type": "string", "enum": [ "card" ] },
            "key": { "type": "string" },
            "name": { "type": [ "string", "null" ] },
            "description": { "type": [ "string", "null" ] },
            "mac_address": { "type": [ "string", "null" ] },
            "serial_number": { "type": [ "string", "null" ] },
            "firmware_version": { "type": [ "string", "null" ] },
            "model_name": { "type": [ "string", "null" ] },
            "model_number": { "type": [ "string", "null" ] },
            "manufacturer": { "type": [ "string", "null" ] },
            "unit_position": { "type": [ "integer", "null" ],
"minimum": 1 },
            "slots": { "type": "array", "items": { "$ref":
"#/$defs/slot" } },
            "ports": { "type": "array", "items": { "$ref":
"#/$defs/port" } }
        },
        "required": [ "key", "name" ]
    },
    "port": {
        "type": "object",
        "properties": {
            "class": { "type": "string", "enum": [ "port" ] },
            "key": { "type": "string" },

```

```

    "name": { "type": [ "string", "null" ] },
    "description": { "type": [ "string", "null" ] },
    "model_name": { "type": [ "string", "null" ] },
    "ip_address": { "type": [ "string", "null" ] },
    "virtual": { "type": "boolean", "default": false },
    "unit_position": { "type": [ "integer", "null" ],
"minimum": 1 },
    "bandwidth_name": { "type": [ "string", "null" ] },
    "bandwidth_group": { "type": [ "string", "null" ] },
    "bandwidth": { "type": [ "integer", "null" ], "minimum":
0 },
    "mtu_size": { "type": [ "integer", "null" ], "minimum":
0 }
  },
  "required": [ "key", "name" ]
},
"number": {
  "type": "object",
  "properties": {
    "class": { "type": "string", "enum": [ "number" ] },
    "key": { "type": "string" },
    "name": { "type": [ "string", "null" ] },
    "related_ci_type": { "type": "string", "enum":
[ "Network Interface", "Physical Connection", "Logical
Connection", "Equipment", "Topology" ] },
    "related_ci": { "$ref": "#/$defs/keyRef" },
    "type": { "type": "string", "enum": [ "vlan_range",
"vlan_subrange", "vlan", "lag_range", "lag" ] },
    "vlan_type": { "type": "string", "enum": [ "inner",
"outer" ] },
    "value": { "$ref": "#/$defs/value" }
  },
  "required": [ "key", "name", "type", "related_ci_type",
"related_ci", "value" ] },
  "logical_connection": {
    "type": "object",
    "properties": {
      "class": { "type": "string", "enum":
[ "logical_connection" ] },
      "key": { "type": "string" },
      "name": { "type": [ "string", "null" ] },
      "description": { "type": [ "string", "null" ] },
      "model_name": { "type": [ "string", "null" ] },
      "bandwidth_group": { "type": [ "string", "null" ] },
      "bandwidth_name_a_to_z": { "type": [ "string",
"null" ] },
      "bandwidth_name_z_to_a": { "type": [ "string",
"null" ] },
      "bandwidth_a_to_z": { "type": [ "integer", "null" ],
"minimum": 0 },
      "bandwidth_z_to_a": { "type": [ "integer", "null" ],
"minimum": 0 },
      "equipment_a": { "$ref": "#/$defs/optionalKeyRef" },
      "equipment_z": { "$ref": "#/$defs/optionalKeyRef" },
      "port_a": { "$ref": "#/$defs/keyRef" },
      "port_z": { "$ref": "#/$defs/keyRef" }
    },
  },

```

```

      "required": [ "key", "name", "equipment_a", "equipment_z",
"port_a", "port_z" ]
    },
    "port_relation": {
      "type": "object",
      "properties": {
        "class": { "type": "string", "enum":
[ "port_relation" ] },
        "parent": { "$ref": "#/$defs/keyRef" },
        "child": { "$ref": "#/$defs/keyRef" }
      },
      "required": [ "parent", "child" ]
    },
    "logical_connection_relation": {
      "type": "object",
      "properties": {
        "class": { "type": "string", "enum":
[ "logical_connection_relation" ] },
        "sequence": { "type": [ "integer", "null" ], "minimum":
1 },
        "route": { "type": [ "integer", "null" ], "minimum":
1 },
        "parent": { "$ref": "#/$defs/keyRef" },
        "child": { "$ref": "#/$defs/keyRef" }
      },
      "required": [ "parent", "child" ]
    },
    "network_topology": {
      "type": "object",
      "properties": {
        "class": { "type": "string", "enum":
[ "network_topology" ] },
        "key": { "type": "string" },
        "name": { "type": "string" },
        "model_name": { "type": [ "string", "null" ] },
        "devices": { "type": "array", "items": { "key":
"#/$defs/optionalKeyRef" } },
        "logical_connections": { "type": "array", "items":
{ "key": "#/$defs/optionalKeyRef" } }
      },
      "required": [ "key", "name", "devices",
"logical_connections" ]
    }
  }
}

```

2. Decouple Connectivity from ETL

Ensure flexibility by separating device interaction logic from transformation logic (ETL ingestion). Users can develop their own collectors independently of the ETL logic.

Ensure the following points:

1. Design collectors to focus only on connectivity and data conversion to the unified schema.
2. Use or reuse any EMS/NMS adaptor (including third-party adaptors like Atrinet).
3. Push the standardized data into ServiceNow import sets.

3. Configure Generic ETL for CMDB Updates

Use a single reusable ETL to process all standardized import sets and update the CMDB accurately.

Ensure the following points:

1. Validate the import set contains data in the standardized JSON format.
2. Use the TSOM Generic ETL to:
 - Detect the entity type (e.g., Slot, Card, Logical Interface).
 - Map each entity to the correct CI class based on model mapping logic.
 - Populate fields including Inventory Category where applicable.

Special Handling:

- Logical Interfaces → Port table (`virtual=true`)
- Logical Connections → Logical Connection table
- PDU Cards → PDU table
- Equipment → Model-based CI class (most recent mapping if multiple exist)

4. Configure support for Multi-Chassis and Composite Devices

Model complex devices such as multi-chassis routers with correct CMDB relationships.

Ensure the following points:

1. Use the Logical Composite construct to represent grouped entities like Router + PDU.
2. Define individual components (Fan, Management, Slot) under their respective hierarchy.
3. Map each entity as per the TNI modeling guidance.

Example: For a 7750-2s multi-chassis device:

- Logical Composite → contains Router and PDU
- PDU → contains Slots → Cards → Sub-slots

5. Enable TNI Entity Creation Based on Installation

Ensure consistency with TNI standards without unnecessary record creation.

Ensure the following points:

- If TNI is installed:
 - Automatically create a TNI Entity for each discovered CI.
 - Set Inventory Category appropriately (e.g., "Logical Connection", "Interface").
- If TNI is not installed: Skip TNI entity creation to avoid orphaned or invalid records.

Configure service graph connectors

Configure Nokia Altiplano service graph connector

Configure the Nokia Altiplano service graph connector to import physical and logical inventory data from the Nokia Altiplano Access Network SDN Controller into your CMDB.

This integration uses REST APIs (via a MID Server) to ensure the CMDB reflects accurate, up-to-date telecom inventory aligned with the TM Forum-based data model.

Note: A valid Telecommunications Service Operations Management subscription is required to use this connector.

Required plugins

Plugin	ID
Telecom Service Operations Core	sn_tsom_core
Service Graph Connector for Nokia Altiplano	sn_tsom_altiplano_connector

Note: External requirements:

- A running Nokia Altiplano Controller instance with access to its REST northbound API.
- A MID Server with secure connectivity to the Altiplano instance.

Configuration tasks overview

The following sections are available under the Nokia Altiplano navigation pane. Use the following table for post-guided setup or to perform manual configurations.

Section	Description
Setup	Configure MID Server, define Altiplano connections, and schedule imports.
Data Sources	Predefined data sources for bulk and filtered discovery (SGC-Nokia Altiplano Bulk Discovery, SGC-Nokia Altiplano Filtering Discovery). Enable parallel loading if needed. For more information on parallel loading, see Configure concurrent import and parallel loading .
Import Schedules	Manage scheduling for each Altiplano connection alias. Run jobs manually or at defined intervals.
Connections & Credential Aliases	Define aliases for each Altiplano instance. Store connection metadata and credentials.
Connections	Define Altiplano instance details, such as URL, selected MID Server, credential reference, and connection alias reference.
Credentials	Create Altiplano credentials using Basic Auth.
Filters	Configure filtering parameters used in filtered discovery (for example, by device IP or name).
Properties	Modify system behavior using connector-specific properties. For more information, see System components installed with Nokia Altiplano .

Supported Nokia Altiplano versions

- Minimum supported version: 24.6
- Supported OLT: Lightspan MF-2

Note: For more information, see [Mapping Nokia Altiplano CIs and Relationships in CMDB](#).

Access the Guided Setup

Use the guided setup to simplify the configuration process. This setup provides an organized sequence of steps to help you complete integration quickly and correctly. To access the guided setup:

1. Navigate to **All > Service Graph Connectors > Nokia Altiplano > Setup**.
2. Follow the task sequence to configure MID Server settings, Nokia Altiplano connections, filters, and import schedules. For more information, see [Set up Nokia Altiplano](#).

Related topics

[Set up Nokia Altiplano](#)
[Set up multiple Nokia Altiplano instances](#)
[Configure concurrent import and parallel loading](#)
[Mapping Nokia Altiplano CIs and Relationships in CMDB](#)
[Run and verify a Nokia Altiplano import](#)
[System components installed with Nokia Altiplano](#)
[Telecom Discovery via Nokia Altiplano](#)

Set up Nokia Altiplano

Learn how to install and configure the Service Graph Connector for Nokia Altiplano, including optional demo data, connectivity setup, and data collection schedules.

Before you begin

Role required: TSOM Visibility admin

About this task

Verify the active application scope is Service Graph Connector (SGC) for Nokia Altiplano.

1. Select the application picker () in the header.
2. Search for **Nokia Altiplano**.
3. Select Service Graph Connector for Nokia Altiplano from the list.

Procedure

1. Navigate to **All > Service Graph Connectors > Nokia Altiplano > Setup**.
2. On the Getting started page, select **Get Started**.
The Guided Setup home page opens in a new tab.
3. **Optional:** Create and configure MID Server or skip based on your environment.
 - a. Select **Configure** to complete the process.
 - b. Follow the on-screen instructions to download and install.
 - c. Select **Mark as Complete** when done.

Note: For more information on how to install and configure MID Server, see [Configuring MID Server](#).

- d. Configure or skip the validation step.
 - e. Select **Mark as Complete** when validated.
Once all MID Server steps are complete, proceed to **Configure Connectivity**.
4. Create connection aliases, credentials, and HTTP connections to your Nokia Altiplano instance with **Configure Connectivity**
- a. Select **Get Started**.
 - b. Create and configure aliases for the connections and credentials:
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. Leave the rest of the fields as default, select **Submit** and then **Mark as Complete**.

Note: This enables using the connection by name rather than directly, enabling the collector to extract all active aliases from the CMDB and start performing data collection on the HTTP connection bound to it.
 - c. Create the credentials to access the Nokia Altiplano Controller by selecting **Configure**.
 - i. In the **Name** field, specify the alias name.
 - ii. In the **User name** field, specify your Nokia Altiplano instance user name.
 - iii. In the **Password** field, specify your Nokia Altiplano instance password.

Note: Other authentication fields may be required depending on the authentication methods used in your Nokia Altiplano instance. By default, we use [Basic authentication credentials](#) (as part of the Guided setup).
 - iv. Leave the rest of the fields as default, select **Submit** and then select **Mark as Complete**.
 - d. Create HTTP Connection by selecting **Configure**
 - i. In the **Name** field, specify the connection name.
 - ii. Choose the **Credentials** and **Connection Alias** created earlier.
 - iii. Specify the **Connection URL** for Nokia Altiplano.
 - iv. Check **Use MID Server** and choose either:
 - Auto-select
 - Specific MID Server
 - Specific MID Cluster
 - v. Leave the rest of the fields as default, select **Submit** and then select **Mark as Complete**.
5. Configure Data Collection Schedule either for bulk discovery or filtered discovery by clicking **Get Started** and then click **Configure**.

To	Do the following
<p>Schedule data collection for bulk discovery (Use bulk discovery data source for all devices in connection alias)</p>	<ul style="list-style-type: none"> ○ Provide a Name for the scheduler. ○ Ensure the bulk data source SGC-Nokia Alti plano Bulk Discovery is selected. ○ Set Active the scheduler automatically runs as mentioned in the Run and Time fields. If inactive, you must manually execute. ○ Select Use Connection. ○ In the Connection field, select connection alias. For example, sn_sgc_altiplano.Ati plano_Alias. ○ In the Run field, select the frequency. Specify when you want this schedule to run: Daily, Weekly, Monthly, Periodically, After Parent Runs, or Once. ○ In the Time field, enter the time in hours, minutes, and seconds.
<p>Schedule data collection for filtered discovery (Use Filtering Discovery data source for specific OLT filter criteria for each connection alias)</p>	<ol style="list-style-type: none"> Add Filtering Parameters <ol style="list-style-type: none"> In the Connection Alias field, choose a connection alias. Optionally, in the Filtered IPs field, add IP filters in various formats defined as following: <ul style="list-style-type: none"> ▪ Single IP: 10 . 10 . 10 . 10 ▪ List: 10 . 10 . 10 . 10 , 10 . 10 . 10 . 20 ▪ Ranges: 10 . 10 . 10 . 30 - 10 . 10 . 10 . 40 ▪ CIDR: 10 . 10 . 10 . 0 / 24 ▪ Mask: 10 . 10 . 10 . 0 : 255 . 255 . 255 . 0 Optionally, add name filters to search by device name or part of the name. You can specify a single string or a list of names to filter the results. Schedule filtered Collection <ul style="list-style-type: none"> ▪ Provide a Name for the scheduler. ▪ Ensure the data source SGC-Nokia Alti plano Filtering Discovery is selected. ▪ Set Active the scheduler automatically runs as mentioned in the Run and Time fields. If inactive, you must manually execute.

To	Do the following
	<ul style="list-style-type: none"> In the Run field, select the frequency. Specify when you want this schedule to run: Daily, Weekly, Monthly, Periodically, After Parent Runs, or Once. In the Time field, enter the time in hours, minutes, and seconds. Select Use Connection.. In the Connection field, select connection alias. For example, sn_sgc_altiplano.Atiplano_Alias.

6. Select **Submit** and then select **Mark as Complete**.

Result

The configured connections, aliases, credentials, and import schedules can also be accessed from the navigation Navigate to **All > Service Graph Connectors > Nokia Atiplano**.

The following snapshot helps you confirm the configuration set up of Nokia Atiplano service graph connector.

Nokia Atiplano Integration with CMDB

100% Complete

MID Server Skip **3 / 3 Tasks completed**

- ✓ Create MID User
- ✓ Download & Install MID
- ✓ Validate MID

Configure Connectivity **3 / 3 Tasks completed**

- ✓ Create Connection Aliases
- ✓ Create Credentials
- ✓ Create HTTP Connection

Configure Data Collection schedule Skip **2 / 2 Tasks completed**

- ✓ Schedule Data Collection for Bulk Discovery
- ✓ Schedule Data Collection for Filtered Discovery

Related topics

[Set up multiple Nokia Atiplano instances](#)

[Configure concurrent import and parallel loading](#)

[Mapping Nokia Atiplano CIs and Relationships in CMDB](#)

[System components installed with Nokia Atiplano](#)

Set up multiple Nokia Atiplano instances

Learn how to configure and manage multiple Nokia Atiplano instances within a single ServiceNow environment. This enables administrators and integrators to create independent connection aliases and import schedules customized to specific filtering or frequency needs.

Before you begin

Role required: TSOM Visibility admin

About this task

Confirm you have:

- You're operating in the Service Graph Connector for NOKIA Altiplano application scope.
- Installed the Service Graph Connector for Nokia Altiplano.
- Completed the guided setup for the initial Altiplano instance.
- MID Server set up and validated.

You can configure additional Nokia Altiplano instances or reuse the same Altiplano instance with different connection aliases and import schedules. To add a new Altiplano instance, rerun the Guided setup to configure a new connection alias. Go through all connectivity stages in the setup for each new alias.

Procedure

1. Navigate to **All > Service Graph Connectors > Nokia Altiplano > Setup**.
2. On the Getting Started page, select **Get Started**.
Repeat all the steps under the Configured Connectivity section. It creates configuration entries for the new instance of Nokia Altiplano. For more information, see [Set up Nokia Altiplano](#).
3. Configure the connectivity by creating a connection alias, credentials, and HTTP connection.
 - a. In the Configure Connectivity section, select Get Started.
 - b. Select Configure to create a new connection alias by entering a unique alias name (for example, `Altiplano_Prod_01`).
 - c. Select **Submit** and mark the step as complete.
4. Configure credentials for the new Altiplano instance.
 - a. Specify a user name and password for the Altiplano Controller.
 - b. Submit and mark the step as complete.
5. Configure an HTTP connection.
 - a. Provide the connection name.
 - b. Select the newly created credentials and connection alias.
 - c. Enter the connection URL of the Nokia Altiplano instance.
 - d. Enable use MID Server and choose the appropriate MID option.
 - e. Select **Submit** and mark the step as complete.
6. Schedule data imports using either bulk or filtered discovery by configuring the import schedule.
 - a. Fill in the fields.
For more information, see [Set up Nokia Altiplano](#).
 - b. In the **Use connection** field, choose the new Altiplano instance.
7. Select **Submit**.

Result

To confirm that your new instance setup is successful, you can verify the configuration.

- Navigate to **All > Service Graph Connectors > Nokia Altiplano > Connection & Credential Aliases** and confirm that the new alias is listed.
- Navigate to **All > Service Graph Connectors > Nokia Altiplano > Import Schedule** and confirm that the new import schedule is listed and confirm that corresponding jobs exist for each alias.

You can configure multiple connection aliases over the same Nokia Altiplano instance. This flexibility enables you to run imports at different frequencies and apply different filters to each alias.

Example: `Altiplano_Weekly_OLT_10.10.10.*`: Runs weekly, filters OLTs with IPs in the 10.10.10.* range.
`Altiplano_Daily_Prod`: Runs daily, filters OLTs whose name contains "prod_olt".

Related topics

[Configure concurrent import and parallel loading](#)

[Mapping Nokia Altiplano CIs and Relationships in CMDB](#)

[System components installed with Nokia Altiplano](#)

Configure concurrent import and parallel loading

Improve the efficiency of large data imports from Nokia Altiplano by configuring concurrent imports and enabling parallel loading. This configuration allows the platform to run multiple data import and transformation jobs simultaneously.

Before you begin

Role required: TSOM Visibility admin

About this task

Concurrent import and parallel loading improve performance by dividing large datasets into smaller partitions, allowing multiple transformation and collection jobs to run in parallel. This setup is useful when importing large volumes of equipment, logical connections, and network topology data during bulk or filtered discoveries.

Use the following guidance to optimize performance:

- If transformation is slow: Enable Concurrent import in the import schedule.
- If data collection is slow: In addition to enabling concurrent import, configure Parallel loading in the data source and update the system property.

Note: Parallel loading works in addition to concurrent import. It does not replace it.

Procedure

1. Navigate to **All > Service Graph Connectors > Nokia Altiplano > Import Schedule**.
2. Configure concurrent import or concurrent import

To	Do the following
To enable concurrent import	<ol style="list-style-type: none"> a. Navigate to the Import Schedule for the discovery job. b. Select the Concurrent import checkbox. c. Set the Partition method to Custom size. d. In the Partition size field, enter the number of records per partition (for example, 1000).

To	Do the following
	<p>Note: The system splits the dataset into import sets based on the partition size. Each import set is processed in parallel, improving the speed of data transformation.</p>
Configure parallel loading in the data source	<ol style="list-style-type: none"> Navigate to the corresponding Data Source record. Select the Parallel loading checkbox. Navigate to Nokia Altiplano > Properties. Open the system property <code>sn_sgc_altiplano.parallel_number_of_data_</code> Set the value to the desired number of parallel jobs (for example, 3). <p>Note: By default, the number of data source jobs is set to 1. Increasing this value allows the platform to execute multiple data collection jobs concurrently.</p>

Related topics

[Set up Nokia Altiplano](#)

[Mapping Nokia Altiplano CIs and Relationships in CMDB](#)

[System components installed with Nokia Altiplano](#)

Mapping Nokia Altiplano CIs and Relationships in CMDB

Use the Service Graph Connector for Nokia Altiplano to map discovered physical and logical network resources to telecom-aligned configuration item (CI) classes in the CMDB. The connector supports consistent service modeling, visibility into chassis-level components, and automation of logical and physical relationships.

To confirm accurate CI classification and insertion, the connector uses the Robust Transform Engine (RTE) and Identification and Reconciliation Engine (IRE).

The connector classifies and relates discovered CIs using telecom-specific models based on device type, function, and chassis structure. This helps maintain a clean and normalized CMDB across vendors. Discovered model names from Nokia Altiplano are automatically transformed into ServiceNow standard model identifiers and categories for slot and subslot components.

CI mapping and relationships

The following tables describe how Altiplano CIs are represented in the CMDB and how they relate to each other across physical and logical layers.

CMDB CI Mapping and Relationships (Physical Layer)

CI Type	CMDB Table	Description and Relationships
OLT CI	<code>cmdb_ci_optical_line_terminal</code>	Represents the OLT device. Contains Slot CIs and Logical Network Interface CIs.
ONU/ONT CI	<code>cmdb_ci_optical_network_terminal</code> or <code>cmdb_ci_optical_network_unit</code>	Represents ONU or ONT devices. The class is determined by the system property <code>sn_sgc_altiplano.onu_ci_class</code> . Contains Network Interface CIs.
Slot CI	<code>cmdb_ci_container_slot</code>	Represents main chassis slots. Contained by OLT CIs. Contains Interface Card CIs (for example, LT/NT, PSU, fan). Model transformations are applied in the data source. For more information, see the above- mentioned Model transformation for slot and subslot CIs table.
Subslot CI	<code>cmdb_ci_container_subslot</code>	Represents subcomponents within interface cards (for example, cages for SFPs). Contained by LT/NT cards. Contains transceiver card CIs. For more information, see the above- mentioned Model transformation for slot and subslot CIs table.
Interface Card CI	<code>cmdb_ci_interface_card</code>	Represents LT/NT cards, transceivers, and control units. Can contain subslots and network interfaces.
Network Interface CI	<code>cmdb_ci_ni_interface</code>	Represents both physical (for example, PON, Ethernet) and logical (for example, VLAN) ports. Contained by interface cards or ONU/ONT CIs. Logical ports are related to physical ports using <code>Members::Member of</code> .
Logical Connection CI	<code>cmdb_ci_ni_logical_path</code>	Represents logical paths such as PON or VLAN between OLT and ONU. Defined with Port A and Port Z attributes referencing terminating Network Interface CIs. VLAN paths consume PON paths.
IP Address CI	<code>cmdb_ci_ip_address</code>	Represents discovered IP addresses for OLTs. Owned by the corresponding OLT CI.

Key Relationship Examples

- Containment relationships
 - OLT CI → contains Slot CI
 - Slot CI → contains Interface Card CI
 - Interface Card CI → contains Subslot CI

- Subslot CI → contains transceiver Interface Card CI
- ONU/ONT CI → contains Network Interface CIs
- Interface relationships:
 - Logical Connections -> terminated by Network Interfaces
 - Network Interfaces -> members of Network Interfaces
- Logical path relationships:
 - VLAN path (parent) → consumes → PON path (child)
 - Logical paths → terminate at → Network Interface CIs via Port A and Port Z
- Ownership: IP Address CI → owned by OLT device

Supported models

1. Network equipment models (sn_ent_nw_equipment_model)

- The supported OLT is Nokia Lightspan MF-2, by default the model name is "Nokia MF-2"
- ONU/ONT models are manufacturer + ONU/ONT. the system property sn_sgc_altiplano.onu_ci_class defines if ONU or ONT will be used.
- If the model wasn't found in the model table, a new model is created in the CI. and the CI will be created as "Network gear"

2. Equipment holder models: (sn_ent_nw_holder_model)

- Slots models: "Traffic Slot", "FAN Slot", "Power Slot"
- Subslots models: "SFP Subslot"
- The used model name can be customized by the customer through the Altiplano extension point (sn_sgc_altiplano.AltiplanoCustomizedModels)
- If the model wasn't found in the model table, a new model is created in the CI.

3. Network card models (sn_ent_nw_card_model)

- Card models are found by the model name, manufacturer, and model number discovered from the Altiplano API
- If the model wasn't found in the model table, a new model is created in the CI.

4. Network Interface Models: (sn_ent_nw_interface_model).

- Ethernet ports models are found by the "port bandwidth" column in the Network Interface table (sn_ent_nw_interface_model). the port bandwidth of the port CI is located by the discovered port speed in the Bandwidth table (bandwidth)
- PON physical ports models: "PON Access Interface", "PON Network Interfaces"
- Logical ports models: "ENET Interface", "VLAN Interface", "LAG Interface", "PON Logical Interface"
- If the model wasn't found in the model table, the reference to the "Model ID" will remain empty.

5. Logical network connection models (sn_ent_logical_nw_connection_model)

- PON Access Path
- VLAN Path
- If the model wasn't found in the model table, the reference to the "Model ID" will remain empty.

Note:

- If the connector can't match a discovered equipment to an existing model in the product model table the CI is created as `Network Gear` by default.
- If demo data is installed, default models are created for OLT, ONU, ONT, Slots, Subslots, Cards, Network Interfaces, and Logical connections.
- Equipment and Equipment holder model names can be customized using an extension point

Model transformation for slot and subslot CIs

During ingestion, specific discovered model names are mapped to predefined CMDB model identifiers to confirm consistent slot categorization. The transformation logic is embedded in the SGC data source script and applies to the Nokia Altiplano source.

Slot components such as fan, power, and traffic slots are mapped to the `Slot` category in the CMDB. Subslot components such as SFP cages or synthetic subslots are mapped to the `Subslot` category in the CMDB.

Example: Nokia Altiplano slot and sub slot model mappings


Source	Discovered Model Name	Target CMDB Model ID	Model Category
Altiplano	slot-fan	Fan Slot	Slot
Altiplano	slot-It	Traffic Slot	Slot
Altiplano	cage	SFP Subslot	Subslot
Altiplano	slot-nt	Traffic Slot	Slot
Altiplano	synthetic nt-slot	SFP Subslot	Subslot
Altiplano	slot-psu	Power Slot	Slot

CI relationship structure

The following infographics describes the CI relationships.

CIs Relationships

Contains::Contained by:



Owns::Owned by (Dependent Relationship)



Members::Member of



Consumes::Consumed by



Terminates::Terminated by



LAG Interface

Virtual Interface

Physical Interface

Device

Example: OLT to ONU Structure

The following structure enables end-to-end traceability from OLT through its hardware layers to the connected ONU and associated logical paths.

```

OLT (cmdb_ci_optical_line_terminal)
├── Slot
│   ├── Interface Card (LT/NT)
│   │   ├── Subslot
│   │   │   ├── Interface Card (Transceiver)
│   │   │   │   └── Network Interface (PON port)
│   └── Logical Connections (PON/VLAN) → Connected to →
│       ├── ONU/ONT (cmdb_ci_optical_network_terminal)
│       └── Network Interface (ONT port)
  
```

Configure Cisco Meraki Service Graph Connector

Configure the Cisco Meraki Service Graph Connector (SGC) to import physical and logical inventory data from the Cisco Meraki (SGC) into your Configuration Management Database (CMDB).

Authentication

You must authenticate before initiating discovery. During the authentication process, the discovery service receives an access token, which it then uses for bulk or specific discovery

operations. The integration uses Cisco Meraki REST APIs to discover managed elements such as network equipment, interfaces, and services.

This integration uses REST APIs (through a MID Server) so that the CMDB reflects accurate, up-to-date telecom inventory aligned with the TM Forum-based data model. For a list of API references, see [Cisco Meraki Service Graph Connector API endpoints](#).

Note: A valid Telecommunications Service Operations Management subscription is required to use this connector.

Required plugins

Plugin	Plugin ID
Telecom Service Operations Core	sn_tsom_core
Service Graph Connector for Meraki Telco SD-WAN	sn_tsom_meraki_connector

Note: External requirements:

- A running Cisco Meraki instance with access to its REST API.
- A MID Server with secure connectivity to the Cisco Meraki instance.

Configuration tasks overview

The following sections are available under the Cisco Meraki navigation pane. Use the following table for post-guided setup or to perform manual configurations.

Section	Description
Setup	Configure MID Server, define Cisco Meraki connections, and schedule imports.
Data Sources	Predefined data sources for bulk and filtered discovery (Cisco Meraki Full Discovery, Cisco Meraki Filtered Discovery). Enable parallel loading if needed. For more information on parallel loading.
Import Schedules	Manage scheduling for each Cisco Meraki connection alias. Run jobs manually or at defined intervals.
Connections & Credential Aliases	Define aliases for each Cisco Meraki instance. Store connection metadata and credentials.
Connections	Define Cisco Meraki instance details, such as URL, selected MID Server, credential reference, and connection alias reference.
Credentials	Create Cisco Meraki credentials using Basic Auth.
Filters	Configure filtering parameters used in filtered discovery (for example, by device IP or name).
Properties	Modify system behavior using connector-specific properties. For more information, see.

Access the guided setup

Use the guided setup to simplify the configuration process. This setup provides an organized sequence of steps to help you complete integration quickly and correctly. For more information, see [Set up the Service Graph Connector for Cisco Meraki schedule](#).

Set up the Service Graph Connector for Cisco Meraki schedule

Set up and configure the Service Graph Connector for Cisco Meraki. The process includes installation, optional demo data integration, connectivity configuration, and scheduling data collection for Configuration Management Database (CMDB) integration.

Before you begin

Verify the active application scope is Service Graph Connector (SGC) for Cisco Meraki.

1. Select the application picker () in the header.
2. Search for **Meraki**.
3. Select Service Graph Connector for Meraki from the list.

Role required: TSOM Visibility admin

Procedure

1. Navigate to **All > Service Graph Connectors > Meraki > Setup**.
2. On the Getting started page, select **Get Started**.
3. **Optional:** On the Guided setup home page, create and configure the MID Server or skip based on your environment.

 **Note:** For more information about how to activate and configure a MID Server, see [Configuring MID Server](#) .

- a. Select **Get Started**.
- b. Select **Configure**.
- c. Follow the on-screen instructions to download and install the appropriate MID Server installer archive for the operating system.
- d. Select **Mark as Complete** when done.
After all MID Server steps are complete, proceed to **Configure Connectivity**.
4. Create connection aliases, credentials, and HTTP connections.
This enables the Service Graph Connector to reference the Meraki connection by name and extract active aliases from the CMDB.
 - a. Select **Get Started**.
 - b. Create and configure aliases for the connections and credentials:
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. Retain the default values in the rest of the fields.
 - iv. Select **Submit** and then select **Mark as Complete**.
 - c. Create the basic credentials to access the Meraki Cloud Controller.
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.

iii. In the **API Key** field, enter the API key created in the dashboard.

Note: Other authentication fields might be required depending on the authentication methods used in your Cisco Meraki instance. By default, use basic authentication credentials as part of the Guided setup. For more information, see [Basic authentication credentials](#).

iv. Retain the default values in the rest of the fields.

v. Select **Submit** and then select **Mark as Complete**.

d. Create the HTTP Connection.

i. Select **Configure**.

ii. In the **Name** field, specify the connection name.

iii. Choose the **Credentials** and **Connection Alias** created earlier.

iv. Specify the **Connection URL** for Meraki.

v. Select the **Use MID Server** check box and indicate how the MID Server should be selected:

- Auto-select: Selects from MID Servers based on MID Server criteria, regardless of whether they are members of a cluster.
- Specific MID Server: Automatic reassignment to another MID Server is not supported.
- Specific MID Cluster: Automatic reassignment to another MID Server only selects from members of the specified cluster.

vi. Retain the default values in the rest of the fields.

vii. Select **Submit** and then select **Mark as Complete**.

5. Configure the data collection schedule.

a. Select **Get Started**.

b. Select **Configure**.

c. Configure the schedule.

i. In the **Name** field, provide a name for the scheduler.

ii. Select the data source based on whether the schedule should be for bulk discovery or filtered discovery.

- To use the bulk discovery data source for all devices in the connection alias, select the data source Cisco Meraki Full Discovery.
- To set specific filter criteria for each connection alias, select the data source Cisco Meraki Filtered Discovery.

iii. Determine whether the schedule should run automatically or manually.

iv. ▪ Select the **Active** check box to run the schedule automatically based on the values in the **Run** and **Time** fields.

- Clear the **Active** check box to run the schedule manually.

v. From the **Run** drop-down list, select the frequency at which the scheduler runs automatically. The available choices are Daily, Weekly, Monthly, Periodically, after Parent Runs, or once.

vi. In the **Time** field, set the time of day at which the scheduler should run in hours, minutes, and seconds.

vii. Provide the connection information.

1. Select the **Use Connection** check box.

2. From the **Connection** drop-down list, select the connection alias. For example, sn_sgc_meraki.Meraki_Alias.

d. **Optional:** Set specific filter criteria for each connection alias.

i. Verify that the [FIELD NAME FOR DATA SOURCE] field is set to the data source Cisco Meraki Filtered Discovery.

ii. In the **Connection Alias** field, choose a connection alias.

iii. Set the filtering parameters.

- To filter by IP address, in the **Filtered IPs** field, add IP filters in the desired format.
 - Single IP: 10 . 10 . 10 . 10
 - List: 10 . 10 . 10 . 10 , 10 . 10 . 10 . 20
 - Ranges: 10 . 10 . 10 . 30 - 10 . 10 . 10 . 40
 - CIDR: 10 . 10 . 10 . 0 / 24
 - Mask: 10 . 10 . 10 . 0 : 255 . 255 . 255 . 0
- To filter by name, add name filters to search by device name or part of the name.

(Optional) You can specify a single string or a list of names.

6. Select **Submit** and then select **Mark as Complete**.

Result

After the setup is complete, a confirmation screen appears indicating all tasks have been completed successfully.

You can access the configured connections, aliases, credentials, and import schedules by navigating to **All > Service Graph Connectors > Meraki**.

Manage multiple Cisco Meraki SD-WAN instances

Configure and manage multiple Cisco Meraki instances within a single ServiceNow AI Platform[®] environment. Create distinct connection aliases and independent import schedules. Customize data filtering and frequency requirements.

Before you begin

Verify the following:

- The active application scope is Service Graph Connector (SGC) for Meraki.
- The SGC for Cisco Meraki has been installed.
- The initial Cisco Meraki instance has been set up. For more information, see [Set up the Service Graph Connector for Cisco Meraki schedule](#).
- The associated MID Server has been set up and validated. For more information, see [Configuring MID Server](#) .

Role required: TSOM Visibility admin

About this task

You can configure additional Cisco Meraki instances or reuse the same Cisco Meraki instance with different connection aliases and import schedules. To add a new Cisco Meraki instance, run the guided setup to configure a new connection alias. All connectivity stages within the setup must be completed for each new alias.

Procedure

1. Navigate to **All > Service Graph Connectors > Meraki > Setup**.
2. On the Getting Started page, select **Get Started**.
3. Create and configure aliases for the connections and credentials:
 - a. Select **Configure**.
 - b. In the **Name** field, specify the alias name.
 - c. Retain the default values in the rest of the fields.
 - d. Select **Submit** and then select **Mark as Complete**.
4. Create the basic credentials to access the Meraki Cloud Controller.
 - a. Select **Configure**.
 - b. In the **API Key** field, enter the API key created in the dashboard.

Note: Other authentication fields might be required depending on the authentication methods used in your Cisco Meraki instance. By default, use basic authentication credentials as part of the Guided setup. For more information, see [Basic authentication credentials](#).

- c. In the **User name** field, specify your Cisco Meraki instance user name.
 - d. Retain the default values in the rest of the fields.
 - e. Select **Submit** and then select **Mark as Complete**.
5. Create the HTTP Connection.
 - a. Select **Configure**.
 - b. In the **Name** field, specify the connection name.
 - c. In the **Credentials** and **Connection Alias** fields, choose the items created earlier.
 - d. In the **Connection URL** field, select the URL for Cisco Meraki.
 - e. Select the **Use MID Server** check box and indicate how the MID Server should be selected:
 - Auto-select: Automatically chooses the most appropriate MID Server.
 - Specific MID Server: Select the name of the MID Server from the **MID Server** field that is displayed.
 - Specific MID Cluster: Select the name of the MID cluster from the **MID Cluster** field that is displayed.
 - f. Retain the default values in the rest of the fields.
 - g. Select **Submit** and then select **Mark as Complete**.
6. Configure the connectivity by creating a connection alias, credentials, and HTTP connection.

- a. In the Configure Connectivity section, select **Get Started**.
- b. Select **Configure** to create a new connection alias by entering a unique alias name (for example, `Meraki_Prod_01`).
- c. Select **Submit** and then select **Mark as Complete**.
7. Configure credentials for the new Meraki instance.
 - a. Specify a user name and password for the Meraki instance.
 - b. Select **Submit** and then select **Mark as Complete**.
8. Configure an HTTP connection.
 - a. Provide the connection name.
 - b. Select the newly created credentials and connection alias.
 - c. Enter the connection URL of the Meraki instance.
 - d. Enable use MID Server and choose the appropriate MID option.
 - e. Select **Submit** and then select **Mark as Complete**.
9. Schedule data imports using either bulk or filtered discovery by configuring the import schedule.
 - a. Select **Configure** and fill in the fields.
For more information, see [Set up the Service Graph Connector for Cisco Meraki schedule](#).
 - b. In the **Use connection** field, choose the new Meraki instance.
10. Select **Submit**.
11. Confirm that your new instance setup is successful by verifying the configuration.
 - a. Navigate to **> All > Service Graph Connectors > Meraki > Connection & Credential Aliases**.
 - b. Confirm that the new alias is listed.
You can configure multiple connection aliases over the same Meraki instance. This flexibility enables you to run imports at different frequencies and apply different filters to each alias.

Configure Cisco Meraki allowlists

Configure connector definitions to limit polling to specific Cisco Meraki organizations.

Before you begin

Role required: `tsom_visibility_admin`

About this task

Use multiple connector instances with disjoint allowlists to distribute collection across MID Servers and avoid polling interval overruns in large-scale deployments. Only listed entities are included in the polling scope; all others are silently excluded.

Procedure

1. Navigate to **All > Event Management > Integrations > Connector Definitions**.
2. Search for `Meraki` in the search field.
3. Select the connector from the list, then the connector instance.
4. In the **organizationsToWhitelist** field in the **Connector Instance Values** section, enter the organization names to be included in the polling scope, separated by commas.

For example, CompanyG, Corp, SoutheastRegion.

5. Select **Update**.

Result

The connector applies these values at each polling cycle.

Mapping Cisco Meraki CIs and relationships

Use the Service Graph Connector (SGC) for Cisco Meraki to map discovered physical and logical network resources to telecom-aligned configuration item (CI) classes in the Configuration Management Database (CMDB). Service Graph Connectors support consistent service modeling, provide visibility into chassis-level components, and automate the creation of logical and physical relationships.

To confirm accurate CI classification and insertion, the connector uses the Robust Transform Engine (RTE) and Identification and Reconciliation Engine (IRE).

The connector classifies and relates discovered CIs using telecom-specific models based on device type, function, and chassis structure. This organization helps maintain a clean and normalized CMDB across vendors. Discovered model names from Cisco Meraki are automatically transformed into ServiceNow AI Platform standard model identifiers and categories for slot and subslot components.

CI mapping and relationships

The following table lists the CI object types in the CMDB that can be discovered, along with their representations in the CMDB and how they relate to one another.

CMDB CI mapping and relationships (physical layer)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
Network site	cmdb_ci_ni_site	CI → Site → Network Site	Organization network	<ul style="list-style-type: none"> Represents the physical location of IP routers according to their longitude, latitude, and address. Network site contains IP routers and network interfaces. Network site is a member of a group.
IP Camera	cmdb_ci_ip_camera		IP Camera	<ul style="list-style-type: none"> Represents an IP camera

CMDB CI mapping and relationships (physical layer) (continued)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
				<p>device discovered from a Meraki network.</p> <ul style="list-style-type: none"> • Related to the network site. • Related to the network service instance.
OT industrial sensor	cldb_ci_ot_industrial_sensor		Industrial Sensor	<ul style="list-style-type: none"> • Represents an industrial sensor device discovered from a Meraki network. • Related to the network site. • Related to the network service instance.
IP router	cldb_ci_ip_router	CI → HW → NG → IP router	SD-WAN Edge / network or service router is represented by the IP router	<ul style="list-style-type: none"> • Represents the Meraki device. • Contains network interface CIs. • Contained by network sites and network

CMDB CI mapping and relationships (physical layer) (continued)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
				<p>service instances.</p> <ul style="list-style-type: none"> IP router is a member of a group.
Slot	cmdb_ci_container_slot	HW → Equipment → holder → Slot	Slot	<ul style="list-style-type: none"> Slot is the main device in the network hierarchy. Contains slots for IP routers, IP switches, power supply units, and fans.
Subslot	cmdb_ci_container_subslot	HW → Equipment → holder → Slot	Subslot	<ul style="list-style-type: none"> Network card is the main device in the network hierarchy. Contains IP router, IP switch subslots (small form-factor pluggable or child cards).
Network interface CI	cmdb_ci_ni_interface	Port → Network Port → Network interface	The list of support port models is defined in the vendor-specific network physical information.	<ul style="list-style-type: none"> IP router, IP switch, or wireless access point is the main device in the network hierarchy. Network card within

CMDB CI mapping and relationships (physical layer) (continued)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
				<p>the IP router or IP switch is the primary component.</p> <ul style="list-style-type: none"> Represents the physical ports contained within the device (IP router).
Network service instance	cmdb_ci_network_service_instance	CI → Service instance → Network service instance	Network service instance	<ul style="list-style-type: none"> Network service instance includes IP routers and network interfaces. Network site is a member of a group.
Group	cmdb_ci_group	CI → Group	Represents organization	Network sites and network service instance are members.
IP address CI	cmdb_ci_ip_address	CI → IP address	Represents discovered IP addresses for CIs.	Owned by the corresponding CI.

Configure a Fortinet SD-WAN Service Graph Connector

Configuring the Fortinet SD-WAN Service Graph Connector (SGC) enables you to import physical inventory data from FortiManager into the Configuration Management Database (CMDB) of your ServiceNow instance.

FortiManager is the Fortinet centralized management platform that enables you to configure, monitor, and manage multiple Fortinet security devices, including firewalls and SD-WAN appliances, from a single interface.

The FortiManager JSON API is used to perform configuration and monitoring tasks on a FortiManager device. The SGC for Fortinet SD-WAN uses JSON to gather information and populate the CMDB.

By using structured JSON requests over HTTPS, you can efficiently interact with FortiManager to streamline operations and scale network management tasks across multiple Fortinet devices. These APIs enable automated tasks within the Fortinet ecosystem, such as bulk configuration changes, device management, status monitoring, and inventory collection. To access the Fortinet APIs, create a user and key credentials from FortiManager. For API reference examples, see [Fortinet Service Graph Connector API endpoints](#).

Note: A valid Telecommunications Service Operations Management subscription is required to use this connector.

Required plugins

Plugin	Plugin ID
Telecom Service Operations Core	sn_tsom_core
Service Graph Connector for Fortinet Telco SD-WAN	sn_tsom_fortinet_connector

Note: External requirements:

- A running FortiManager instance with access to its JSON-RPC based northbound API.
- A configured API key in the FortiPortal to enable access to JSON-RPC requests.
- A MID Server with secure connectivity to the FortiManager instance.

Configuration tasks overview

The following sections are available under the Fortinet navigation pane. Use the following table for post-guided setup or to perform manual configurations.

Section	Description
Setup	Configure the MID Server, define Fortinet connections, and schedule imports.
Data Sources	Predefined data sources for bulk (SGC-Fortinet Bulk Discovery) and filtered discovery (SGC-Fortinet Filtering Discovery).
Import Schedules	Manage scheduling for each Fortinet connection alias. Run jobs manually or at defined intervals.
Connections & Credential Aliases	Define aliases for each Fortinet instance. Store connection metadata and credentials.
Connections	Define Fortinet instance details, such as the URL, the selected MID Server, credential reference, and connection alias reference.
Credentials	Create Fortinet credentials using API Key Credentials.
Filters	Configure filtering parameters used in discovery filtered by device IP or name.
Properties	Modify system behavior using connector-specific properties.

Access the Guided Setup

Use the guided setup to simplify the configuration process. This setup provides an organized sequence of steps to help you complete integration quickly and correctly. For more information, see [Set up the Service Graph Connector for Fortinet schedule](#).

Set up the Service Graph Connector for Fortinet schedule

Set up and configure the Service Graph Connector for Fortinet SD-WAN. The process includes installation, optional demo data integration, connectivity configuration, and scheduling data collection for seamless Configuration Management Database (CMDB) integration.

Before you begin

Verify the active application scope is Service Graph Connector (SGC) for Fortinet.

1. Select the application picker () in the header.
2. Search for **Fortinet**.
3. Select Service Graph Connector for Fortinet from the list.

Role required: TSOM Visibility admin

Procedure

1. Navigate to **All > Service Graph Connectors > Fortinet > Setup**.
2. On the Getting started page, select **Get Started**.
3. **Optional:** On the Guided setup home page, create and configure the MID Server or skip based on your environment.

 **Note:** For more information on how to activate and configure a MID Server, see [Configuring MID Server](#) .

- a. Select **Get Started**.
- b. Select **Configure**.
- c. Follow the on-screen instructions to download and install the appropriate MID Server installer archive for the operating system.
- d. Select **Mark as Complete** when done.
Once all MID Server steps are complete, proceed to **Configure Connectivity**.
4. Create connection aliases, credentials, and HTTP connections to enable data collection.
 - a. Select **Get Started**.
 - b. Create and configure aliases for the connections and credentials:
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. Retain the default values in the rest of the fields.
 - iv. Select **Submit** and then select **Mark as Complete**.
 - c. Create the basic credentials to access the FortiManager.
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. In the **API Key** field, enter the API key created in the FortiPortal.

Note: Other authentication fields might be required depending on the authentication methods used in your FortiManager instance. By default, use basic authentication credentials as part of the Guided setup. For more information, see [Basic authentication credentials](#).

iv. Retain the default values in the rest of the fields.

v. Select **Submit** and then select **Mark as Complete**.

d. Create the HTTP Connection.

i. Select **Configure**.

ii. In the **Name** field, specify the connection name.

iii. Choose the **Credentials** and **Connection Alias** created earlier.

iv. Specify the **Connection URL** for FortiManager.

v. Select the **Use MID Server** check box and indicate how the MID Server should be selected:

- Auto-select: Selects from MID Servers based on MID Server criteria, regardless of whether they are members of a cluster.
- Specific MID Server: Automatic reassignment to another MID Server is not supported.
- Specific MID Cluster: Automatic reassignment to another MID Server only selects from members of the specified cluster.

vi. If **Specific MID Server** or **Specific MID Cluster** is selected, select a name in the appropriate field.

vii. Retain the default values in the rest of the fields.

viii. Select **Submit** and then select **Mark as Complete**.

5. Configure the data collection schedule.

a. Select **Get Started**.

b. Select **Configure**.

c. Configure the schedule.

i. In the **Name** field, provide a name for the scheduler.

ii. Select the data source based on whether the schedule should be for bulk discovery or filtered discovery.

- To use the bulk discovery data source for all devices in connection alias, select the data source SGC-Fortinet Bulk Discovery.
- To set specific filter criteria for each connection alias, select the data source SGC-Fortinet Filtering Discovery.

iii. Determine whether the schedule should automatically run or you want to have it run it manually.

iv. ▪ Select the **Active** check box, to run the schedule automatically based on the values in the **Run** and **Time** fields.

- Clear the **Active** check box to run the schedule manually.

v. From the **Run** drop-down list, select the frequency at which you want the scheduler to run automatically. The available choices are Daily, Weekly, Monthly, Periodically, after Parent Runs, or once.

vi. In the **Time** field, set the time of day at which the scheduler should run in hours, minutes, and seconds.

vii. Provide the connection information.

1. Select the **Use Connection** check box.

2. From the **Connection** drop-down list, select the connection alias. For example, sn_sgc_fortinet.Fortinet_Alias.

d. **Optional:** Set specific filter criteria for each connection alias.

i. Verify that the [FIELD NAME FOR DATA SOURCE] field is set to the data source SGC-Fortinet Filtering Discovery.

ii. In the **Connection Alias** field, choose a connection alias.

iii. Set the filtering parameters.

- To filter by IP address, in the **Filtered IPs** field, add IP filters in the desired format.
 - Single IP: 10 . 10 . 10 . 10
 - List: 10 . 10 . 10 . 10 , 10 . 10 . 10 . 20
 - Ranges: 10 . 10 . 10 . 30 - 10 . 10 . 10 . 40
 - CIDR: 10 . 10 . 10 . 0 / 24
 - Mask: 10 . 10 . 10 . 0 : 255 . 255 . 255 . 0
- To filter by name, add name filters to search by device name or part of the name.

(Optional) You can specify a single string or a list of names.

6. Select **Submit** and then select **Mark as Complete**.

Result

Once the setup is complete, a confirmation screen appears indicating all tasks have been completed successfully.

You can access the configured connections, aliases, credentials, and import schedules from the navigation. Navigate to **All > Service Graph Connectors > Fortinet**.

Manage multiple Fortinet SD-WAN instances

Configure and manage multiple Fortinet instances within a single ServiceNow AI Platform[®] environment. Create distinct connection aliases and independent import schedules customized for specific data filtering and frequency requirements.

Before you begin

Verify the following:

- The active application scope is Service Graph Connector (SGC) for Fortinet.
- The SGC for Fortinet has been installed.
- The initial Fortinet instance has been set up. For more information, see [Set up the Service Graph Connector for Fortinet schedule](#).
- The associated MID Server has been set up and validated. For more information, see [Configuring MID Server](#) .

Role required: TSOM Visibility admin

About this task

You can configure additional Fortinet instances or reuse the same Fortinet instance with different connection aliases and import schedules. To add a new Fortinet instance, run the guided setup to configure a new connection alias. All connectivity stages within the setup must be completed for each new alias.

Procedure

1. Navigate to **All > Service Graph Connectors > Fortinet > Setup**.

2. On the Getting Started page, select **Get Started**.

3. Create and configure aliases for the connections and credentials:

- a. Select **Configure**.
- b. In the **Name** field, specify the alias name.
- c. Retain the default values in the rest of the fields.
- d. Select **Submit** and then select **Mark as Complete**.

4. Create the basic credentials to access the FortiManager.

- a. Select **Configure**.
- b. In the **API Key** field, enter the API key created in the FortiPortal.

 **Note:** Other authentication fields might be required depending on the authentication methods used in your Fortinet instance. By default, use basic authentication credentials as part of the Guided setup. For more information, see [Basic authentication credentials](#) .

- c. In the **User name** field, specify your FortiManager instance user name.
- d. Retain the default values in the rest of the fields.
- e. Select **Submit** and then select **Mark as Complete**.

5. Create the HTTP Connection.

- a. Select **Configure**.
- b. In the **Name** field, specify the connection name.
- c. In the **Credentials** and **Connection Alias** fields, choose the items created earlier.
- d. In the **Connection URL** field, select the URL for Fortinet.
- e. Select the **Use MID Server** check box and indicate how the MID Server should be selected:
 - Auto-select: Automatically chooses the most appropriate MID Server.
 - Specific MID Server: Select the name of the MID Server from the **MID Server** field that is displayed.
 - Specific MID Cluster: Select the name of the MID cluster from the **MID Cluster** field that is displayed.
- f. Retain the default values in the rest of the fields.
- g. Select **Submit** and then select **Mark as Complete**.

6. Configure the connectivity by creating a connection alias, credentials, and HTTP connection.

- a. In the Configure Connectivity section, select **Get Started**.
- b. Select **Configure** to create a new connection alias by entering a unique alias name.
- c. Select **Submit** and then select **Mark as Complete**.
7. Configure credentials for the new Fortinet instance.
 - a. Specify a user name and password for the Fortinet instance.
 - b. Select **Submit** and then select **Mark as Complete**.
8. Configure an HTTP connection.
 - a. Provide the connection name.
 - b. Select the newly created credentials and connection alias.
 - c. Enter the connection URL of the Fortinet instance.
 - d. Enable use MID Server and choose the appropriate MID option.
 - e. Select **Submit** and then select **Mark as Complete**.
9. Schedule data imports using either bulk or filtered discovery by configuring the import schedule.
 - a. Select **Configure** and fill in the fields.
For more information, see [Set up the Service Graph Connector for Fortinet schedule](#).
 - b. In the **Use connection** field, choose the new Fortinet instance.
10. Select **Submit**.
11. Confirm that your new instance setup is successful by verifying the configuration.
 - a. Navigate to > **All** > **Service Graph Connectors** > **Fortinet** > **Connection & Credential Aliases**.
 - b. Confirm that the new alias is listed.
You can configure multiple connection aliases over the same Fortinet instance. This flexibility enables you to run imports at different frequencies and apply different filters to each alias.

Configure Fortinet ADOM for connector instance

Configure a connector instance to poll a specific Fortinet ADOM for metric data.

Before you begin

Role required: tsom_visibility_admin

About this task

Each connector instance polls a single ADOM. To collect data from multiple ADOMs, create multiple connector instances, each configured with a different ADOM value. Distribute connector instances across MID Servers to avoid polling interval overruns in large-scale deployments.

Procedure

1. Navigate to **All** > **Event Management** > **Integrations** > **Connector Definitions**.
2. Search for **Fortinet** in the search field.
3. Select the connector from the list, then the Connector Instance.
4. In the **Connector Instance Values** section, specify the ADOM from which you want to receive metric data.
The connector retrieves logs only for the specified ADOM.
5. Select **Update**.

Result

The connector applies these values at each polling cycle.

Filter ADOMs for the Fortinet pull connector

Additional filter conditions required in Fortinet pull connector.

Before you begin

Role required: `tsom_visibility_admin`

About this task

By default, the Fortinet pull connector collects metric data from all ADOMs. To limit collection to ADOMs that meet specific conditions, configure the following parameters on the connector instance (for example, **Fortinet SD-WAN SLA Log**):

ADOM filter parameters

Parameter	Default	Description
<code>adom</code>		
<code>adomFilterEnabled</code>	<code>false</code>	Turns ADOM filtering on or off. When <code>false</code> , the connector collects all ADOMs and the filter is bypassed. When <code>true</code> , the connector collects only the ADOMs that the filter class accepts.
<code>adomFilterClass</code>	<code>FortinetAdomFilterVF</code>	Name of the MID Server script include that implements the filter. Evaluated only when <code>adomFilterEnabled</code> is <code>true</code> . The class must be active, deployed to the MID Server, and expose a <code>shouldIncludeAdom(adomData)</code> handler that returns a Boolean.

Two filter classes are provided:

- `FortinetAdomFilterVF` (default) collects only ADOMs where the `restricted_prds` property equals `fabric` and the `VF SDWAN Service Provided` meta field equals `yes`.
- `FortinetAdomFilter` collects all ADOMs unconditionally.

Warning: Filtering fails closed. If you set `adomFilterEnabled` to `true` and `adomFilterClass` is not a valid script include that is deployed to the MID Server, all ADOMs are blocked and metric collection stops. Enable filtering only when the named class populates the conditions your FortiManager deployment supports.

Procedure

1. Navigate to **All > Event Management > Connectors**.
2. Open the Fortinet pull connector instance.
3. On the **Additional Parameters** related list, set `adomFilterEnabled` to `true`.
4. Set `adomFilterClass` to the name of the filter class you want to apply.
Use `FortinetAdomFilterVF` for the Vodafone filter, `FortinetAdomFilter` to collect all ADOMs, or the name of a custom filter class.

5. Save the connector instance.

The connector applies the filter at the next polling cycle and collects metric data only from the ADOMs that the filter class accepts.

What to do next

To apply your own filtering logic, create a MID Server script include (`ecc_agent_script_include`) that exposes a `shouldIncludeAdom(adomData)` handler. The handler returns `true` to collect the ADOM or `false` to skip it. Deploy the script include to the MID Server and set `adomFilterClass` to its name.

Filter ADOM discovery for the Fortinet Service Graph Connector

Restrict which Fortinet ADOMs the SGC discovers by activating the predefined Vodafone filter or by providing your own implementation of the ADOM filter extension point.

Before you begin

Role required: `tsom_visibility_admin`

About this task

By default, the SGC for Fortinet discovers all ADOMs returned by FortiManager.

To restrict discovery to a subset of ADOMs based on their properties, use the `sn_sgc_fortinet.FortinetCustomizedAdomFilter` extension point. An implementation defines a `shouldIncludeAdom(adom)` handler that returns `true` to discover the ADOM or `false` to skip it, where `adom` is the ADOM object returned by the Fortinet API.

A predefined implementation, `FortinetDefaultAdomFilter`, is provided for Vodafone deployments. It discovers only ADOMs where the **VF SDWAN Service Provided** meta field equals `yes` and the `restricted_prds` property equals `fabric`. This implementation is inactive by default, so it doesn't affect discovery unless you activate it.

Warning: Activate the predefined filter only if your FortiManager deployment populates the required meta fields. If you activate it without them, no ADOM matches the filter and discovery returns no data.

Procedure

1. Navigate to **All > System Definition > Extension Instances**.
2. Filter the list by **Point** equals `sn_sgc_fortinet.FortinetCustomizedAdomFilter`.
3. Open the **FortinetDefaultAdomFilter** record.
4. Set **Active** to true and save the record.
The SGC applies the filter at the next polling cycle and discovers only the ADOMs that satisfy it.

What to do next

To apply your own filtering logic instead, create an implementation of the `sn_sgc_fortinet.FortinetCustomizedAdomFilter` extension point in your own scope, define the `shouldIncludeAdom(adom)` handler, and activate its extension instance. Leave the predefined `FortinetDefaultAdomFilter` instance inactive.

Customize scripted extension points

Customize specific parts of the Fortinet connector's behavior by implementing scripted extension points.

Before you begin

Service Graph Connector for Fortinet must be installed. For instructions, see [Configure a Fortinet SD-WAN Service Graph Connector](#).

Role required: tsom_visibility_admin

About this task

The Service Graph Connector for Fortinet exposes nine scripted extension points. Each extension point controls a specific part of the connector's behavior, such as request payloads, ADOM filtering, or how a CIs Life Cycle Stage Status is set. Every extension point follows the same pattern: an extension point record defines the contract (the functions your implementation must provide), and an extension instance record wires your implementation to that extension point. See [Fortinet extension points reference](#) to find the extension point that controls the behavior you want to change.

The following steps apply to any of the nine extension points. This example customizes `FortinetCustomizedContractParsing` to change how license expiration dates are stored on CIs.

Procedure

1. Navigate to **All > System Scripted Extension Points > Scripted Extension Points**.
2. Search for Fortinet in the **Extension Points** search field.
3. From the **API Name** search results list, select the extension point you want to customize, for example **sn_sgc_fortinet.FortinetCustomizedContractParsing**.
4. Select the **Create implementation** related link.
5. In the **Script** field, implement the function stubs for the extension point's contract. The docstring on each stub describes what your function must return.
6. Run your implementation before the default by setting the **Order** field to a value less than 100. The default implementation has an order of 100; implementations with a lower order number execute first and take precedence.
7. Select **Update**.

Result


Your custom implementation is saved and active. The connector picks it up automatically on its next scheduled run — no restart or cache clear needed.

Example: Example: Customize license expiration date storage

By default, Fortinet SGC stores license expiration dates as separate CI key-value pairs for each device, using the naming convention `license_expiration_date_SERVICE_CODE`. This example customizes `FortinetCustomizedContractParsing` to store only the earliest expiration date across all devices instead.

In the **Script** field, modify the `buildLicenseAttributes` function to return an array of `{key, value}` objects containing the CI key-value pairs you need to store. Loop through `contractItems` to find the minimum date value, then return

it as a single object: { key: "license_expiration_date", value:



"<earliest_date>" }.

The next time Fortinet SGC discovers devices, your script runs first. It stores the license key-value pairs you defined on the relevant CIs, overriding the default per-device key-value pairs behavior.

Related topics

[Fortinet extension points reference](#)

[Extension points](#)

[Create a scripted extension point](#)

Mapping Fortinet CIs and relationships

The Service Graph Connector (SGC) for Fortinet SD-WAN maps discovered physical and logical network resources to telecom-aligned configuration item (CI) classes in the Configuration Management Database (CMDB). Service Graph Connectors support consistent service modeling, provide visibility into chassis-level components, and automate the creation of logical and physical relationships.

The connector uses the Robust Transform Engine (RTE) and Identification and Reconciliation Engine (IRE) to confirm accurate CI classification and insertion.

The connector classifies and relates discovered CIs using telecom-specific models based on device type, function, and chassis structure. This organization helps maintain a clean and normalized CMDB across vendors. Discovered model names from Fortinet are automatically transformed into ServiceNow AI Platform standard model identifiers and categories for slot and subslot components.

CI mapping and relationships

The following table lists the CI object types in the CMDB that can be discovered, along with their representations in the CMDB and how they relate to one another.

CMDB CI Mapping and Relationships (Physical Layer)

CMDB CI Class	CMDB CI Table	CMDB Hierarchy	Object types/ models	Description and Relationships
Network site	cmdb_ci_ni_site	CI → Site → Network Site	Organization network	<ul style="list-style-type: none"> Represents the physical location of IP routers by longitude, latitude, and address. Network site contains IP routers and network interfaces. Network site is a member of a group.
IP router	cmdb_ci_ip_router	CI → HW → NG → IP router	SD-WAN Edge/ network or service router is represented by the IP router	<ul style="list-style-type: none"> Represents the FortiGate device. Contains network interface CIs. Contained by network sites and network service instances. IP router is a member of a group.
Slot	cmdb_ci_container_slot	HW → Equipment → holder → Slot	Slot	<ul style="list-style-type: none"> Slot is the main device in the network hierarchy. Contains slots for IP routers, IP switches, power

CMDB CI Mapping and Relationships (Physical Layer) (continued)

CMDB CI Class	CMDB CI Table	CMDB Hierarchy	Object types/ models	Description and Relationships
				supply units, and fans.
Subslot	cmdb_ci_container_subslot	HW → Equipment → holder → Slot	Subslot	<ul style="list-style-type: none"> • Network card is the main device in the network hierarchy. • Contains IP router, IP switch sublots (small form-factor pluggable or child cards).
Network Interface CI	cmdb_ci_ni_interface	Port → Network Port → Network interface	The list of support port models is defined in the vendor-specific network physical information.	<ul style="list-style-type: none"> • IP router, IP switch, or wireless access point is the main device in the network hierarchy. • Network card within the IP router or IP switch is the primary component. • Represents the physical ports contained within the device (IP router).
Network service instance	cmdb_ci_network_service_instance	CI → Service instance → Network	Network service instance	<ul style="list-style-type: none"> • Network service instance includes

CMDB CI Mapping and Relationships (Physical Layer) (continued)

CMDB CI Class	CMDB CI Table	CMDB Hierarchy	Object types/ models	Description and Relationships
		service instance		IP routers and network interfaces. <ul style="list-style-type: none">• Network site is a member of a group.
Group	<code>cmdb_ci_group</code>	CI → Group	Represents organization	Network sites and network service instance are members.
IP Address CI	<code>cmdb_ci_ip_address</code>	CI → IP address	Represents discovered IP addresses for CIs.	Owned by the corresponding CI.

Discovered logical interfaces

In addition to physical ports, the SGC discovers logical interfaces, such as tunnel and VPN interfaces (for example, `Hub1 - inetVPN`, `inetVPN`, and `mplsVPN`), from the Fortinet inventory. Each logical interface is created as a network interface CI (`cmdb_ci_ni_interface`) and related to its parent device (IP router) CI.

A CMDB CI exists for each logical interface. Metrics that report against a logical interface name are mapped to the corresponding logical interface CI rather than to the parent device CI. This mapping enables interface-level metric correlation and reporting for tunnel and VPN interfaces.

Firmware version calculation

When the SGC discovers an IP router, it sets the firmware version on the `cmdb_ci_ip_router` CI by joining the device `os_ver`, `mr`, and `patch` fields from the Fortinet inventory with periods. For example, a device with `os_ver=7`, `mr=4`, and `patch=11` resolves to `7.4.11`. This matches the value shown in the Fortinet GUI. If any of these fields is missing, the firmware version is left empty. The firmware value is also left empty if the major version is not a whole number.

Override the default with the `sn_sgc_fortinet.FortinetCustomizedFirmwareVersion` extension point to calculate the firmware version differently. Create an implementation that defines a `formatFirmwareVersion()` handler, where `device` is the Fortinet device object returned by the API. Return your own version string from the handler. A custom firmware handler's return value is used exactly as returned, including an empty result; the built-in calculation is used only when no custom handler is active.

Configure Arista VeloCloud Service Graph Connector

Configure the Arista VeloCloud Service Graph Connector (SGC) to discover and import physical and logical inventory data from your Arista VeloCloud environment into the Configuration Management Database (CMDB).

Arista VeloCloud relies on VeloCloud Orchestrator (VCO) as its centralized management platform, available as either a self-hosted or SaaS deployment. VCO configures, monitors, and manages SD-WAN inventory across the network.

Authentication

VeloCloud Orchestrator supports two types of user roles: MSP admins and organization admins. To authenticate API requests, users must generate an API token from their VCO account.

There are two types of API tokens:

- User token: Inherits the same access permissions as the user who created it.
- Organization token: Provides access scoped to a single organization.

You must authenticate before initiating discovery. During the authentication process, the discovery service receives an access token, which it then uses for bulk or specific discovery operations. The integration uses VeloCloud REST APIs to discover managed elements such as network equipment, interfaces, and services.

This integration uses REST APIs (through a MID Server) so that the CMDB reflects accurate, up-to-date telecom inventory aligned with the TM Forum-based data model. For a list of API references, see [Arista VeloCloud Service Graph Connector API endpoints](#).

Note: A valid Telecommunications Service Operations Management subscription is required to use this connector.

Required plugins

Plugin	Plugin ID
Telecom Service Operations Core	sn_tsom_core
Service Graph Connector for Velocloud Telco SD-WAN	sn_tsom_vcloud_connector

Note: External requirements:

- A running Arista VeloCloud instance with access to its REST API.
- A MID Server with secure connectivity to the Arista VeloCloud instance.

Configuration tasks overview

The following sections are available under the Arista VeloCloud navigation pane. Use the following table for post-guided setup or to perform manual configurations.

Section	Description
Setup	Configure MID Server, define Arista VeloCloud connections, and schedule imports.
Data Sources	Predefined data sources for bulk and filtered discovery (VeloCloud Full Discovery, VeloCloud Filtered Discovery). Enable parallel loading if needed.
Import Schedules	Manage scheduling for each Arista VeloCloud connection alias. Run jobs manually or at defined intervals.
Connections & Credential Aliases	Define aliases for each Arista VeloCloud instance. Store connection metadata and credentials.

Section	Description
Connections	Define Arista VeloCloud instance details, such as URL, selected MID Server, credential reference, and connection alias reference.
Credentials	Create Arista VeloCloud credentials using Basic Auth.
Filters	Configure filtering parameters used in filtered discovery (for example, by device IP or name).
Properties	Modify system behavior using connector-specific properties.

Access the guided setup

Use the guided setup to simplify the configuration process. This setup provides an organized sequence of steps to help you complete integration quickly and correctly. For more information, see [Set up the Service Graph Connector for Arista VeloCloud schedule](#).

Set up the Service Graph Connector for Arista VeloCloud schedule

Set up and configure the Service Graph Connector for Arista VeloCloud SD-WAN. The process includes installation, optional demo data integration, connectivity configuration, and scheduling data collection for seamless Configuration Management Database (CMDB) integration.

Before you begin

Verify the active application scope is Service Graph Connector (SGC) for VeloCloud.

1. Select the application picker () in the header.
2. Search for **VeloCloud**.
3. Select Service Graph Connector for VeloCloud from the list.

Role required: TSOM Visibility admin

Procedure

1. Navigate to **All > Service Graph Connectors > VeloCloud > Setup**.
2. On the Getting started page, select **Get Started**.
3. **Optional:** On the Guided setup home page, create and configure the MID Server or skip based on your environment.

 **Note:** For more information on how to activate and configure a MID Server, see [Configuring MID Server](#) .

- a. Select **Get Started**.
 - b. Select **Configure**.
 - c. Follow the on-screen instructions to download and install the appropriate MID Server installer archive for the operating system.
 - d. Select **Mark as Complete** when done.
Once all MID Server steps are complete, proceed to **Configure Connectivity**.
4. Create connection aliases, credentials, and HTTP connections to enable data collection.
This configuration allows the Service Graph Connector to reference the VeloCloud connection by name and extract active aliases from the CMDB.

- a. Select **Get Started**.
 - b. Create and configure aliases for the connections and credentials:
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. Retain the default values in the rest of the fields.
 - iv. Select **Submit** and then select **Mark as Complete**.
 - c. Create the basic credentials to access the VeloCloud Orchestrator.
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the alias name.
 - iii. In the **API Key** field, enter the API key created in the VeloCloud Orchestrator.

Note: Other authentication fields might be required depending on the authentication methods used in your instance. By default, use basic authentication credentials as part of the Guided setup. For more information, see [Basic authentication credentials](#).
 - iv. Retain the default values in the rest of the fields.
 - v. Select **Submit** and then select **Mark as Complete**.
 - d. Create the HTTP Connection.
 - i. Select **Configure**.
 - ii. In the **Name** field, specify the connection name.
 - iii. Choose the **Credentials** and **Connection Alias** created earlier.
 - iv. Specify the **Connection URL** for VeloCloud Orchestrator.
 - v. Select the **Use MID Server** check box and indicate how the MID Server should be selected:
 - Auto-Select: Selects from MID Servers based on MID Server criteria, regardless of whether they are members of a cluster.
 - Specific MID Server: Automatic reassignment to another MID Server is not supported.
 - Specific MID Cluster: Automatic reassignment to another MID Server only selects from members of the specified cluster.
 - vi. If **Specific MID Server** or **Specific MID Cluster** is selected, select a name in the appropriate field.
 - vii. Retain the default values in the rest of the fields.
 - viii. Select **Submit** and then select **Mark as Complete**.
5. Configure the data collection schedule.
- a. Select **Get Started**.
 - b. Select **Configure**.

c. Configure the schedule.

- i.** In the **Name** field, provide a name for the scheduler.
- ii.** Select the data source based on whether the schedule should be for bulk discovery or filtered discovery.
 - To use the bulk discovery data source for all devices in connection alias, select the data source VeloCloud Full Discovery.
 - To set specific filter criteria for each connection alias, select the data source VeloCloud Filtered Discovery.
- iii.** Determine whether the schedule should automatically run or you want to have it run it manually.
- iv.**
 - Select the **Active** check box, to run the schedule automatically based on the values in the **Run** and **Time** fields.
 - Clear the **Active** check box to run the schedule manually.
- v.** From the **Run** drop-down list, select the frequency at which you want the scheduler to run automatically. The available choices are Daily, Weekly, Monthly, Periodically, After Parent Runs, or Once.
- vi.** In the **Time** field, set the time of day at which the scheduler should run in hours, minutes, and seconds.
- vii.** Provide the connection information.
 - 1.** Select the **Use Connection** check box.
 - 2.** From the **Connection** drop-down list, select the connection alias. For example, sn_sgc_vcloud.Velocloud_Alias.

d. Optional: Set specific filter criteria for each connection alias.

- i.** Ensure that the [FIELD NAME FOR DATA SOURCE] field is set to the data source VeloCloud Filtered Discovery.
- ii.** In the **Connection Alias** field, choose a connection alias.
- iii.** Set the filtering parameters.
 - To filter by IP address, in the **Filtered IPs** field, add IP filters in the desired format.
 - Single IP: 10 . 10 . 10 . 10
 - List: 10 . 10 . 10 . 10 , 10 . 10 . 10 . 20
 - Ranges: 10 . 10 . 10 . 30 - 10 . 10 . 10 . 40
 - CIDR: 10 . 10 . 10 . 0 / 24
 - Mask: 10 . 10 . 10 . 0 : 255 . 255 . 255 . 0
 - To filter by name, add name filters to search by device name or part of the name.

(Optional) You can specify a single string or a list of names.

6. Select **Submit and then select **Mark as Complete**.**

Result

Once the setup is complete, a confirmation screen appears indicating all tasks have been completed successfully.

You can access the configured connections, aliases, credentials, and import schedules from the navigation. Navigate to **All > Service Graph Connectors > VeloCloud**.

Manage multiple Arista VeloCloud instances

Configure and manage multiple Arista VeloCloud instances within a single ServiceNow AI Platform[®] environment. Create distinct connection aliases and establish independent import schedules. Customize data filtering and frequency requirements for each instance.

Before you begin

Verify the following:

- The active application scope is Service Graph Connector (SGC) for VeloCloud.
- The SGC for Arista VeloCloud has been installed.
- The initial Arista VeloCloud instance has been set up. For more information, see [Set up the Service Graph Connector for Arista VeloCloud schedule](#).
- The associated MID Server has been set up and validated. For more information, see [Configuring MID Server](#).

Role required: TSOM Visibility admin

About this task

You can configure additional VeloCloud instances or reuse the same VeloCloud instance with different connection aliases and import schedules. To add an VeloCloud instance, run the guided setup to configure a new connection alias. All connectivity stages within the setup must be completed for each new alias.

Procedure

1. Navigate to **All > Service Graph Connectors > VeloCloud > Setup**.
2. On the Getting Started page, select **Get Started**.
3. Create and configure aliases for the connections and credentials:
 - a. Select **Configure**.
 - b. In the **Name** field, specify the alias name.
 - c. Retain the default values in the rest of the fields.
 - d. Select **Submit** and then select **Mark as Complete**.
4. Create the basic credentials to access the VeloCloud Orchestrator.
 - a. Select **Get Started**.
 - b. Select **Configure**.
 - c. In the **API Key** field, enter the API key created in the Orchestrator.

Note: Other authentication fields might be required depending on the authentication methods used in your Arista VeloCloud instance. By default, use basic authentication credentials as part of the guided setup. For more information, see [Basic authentication credentials](#).

- d. In the **User name** field, specify your Arista VeloCloud instance user name.
 - e. Retain the default values in the rest of the fields.
 - f. Select **Submit** and then select **Mark as Complete**.
5. Create the HTTP Connection.

- a. Select **Get Started**.
 - b. Select **Configure**.
 - c. In the **Name** field, specify the connection name.
 - d. In the **Credentials** and **Connection Alias** fields, choose the items created earlier.
 - e. In the **Connection URL** field, select the URL for Arista VeloCloud.
 - f. Select the **Use MID Server** check box and indicate how the MID Server should be selected:
 - Auto-select: Automatically selects the most appropriate MID Server.
 - Specific MID Server: Select the name of the MID Server from the **MID Server** field that is displayed.
 - Specific MID Cluster: Select the name of the MID cluster from the **MID Cluster** field that is displayed.
 - g. Retain the default values in the rest of the fields.
 - h. Select **Submit** and then select **Mark as Complete**.
6. Configure the connectivity by creating a connection alias, credentials, and HTTP connection.
- a. In the Configure Connectivity section, select **Get Started**.
 - b. Select **Configure** to create a connection alias by entering a unique alias name (for example, `Velocloud_Prod_01`).
 - c. Select **Submit** and then select **Mark as Complete**.
7. Configure credentials for the new VeloCloud instance.
- a. Specify a user name and password for the VeloCloud instance.
 - b. Select **Submit** and then select **Mark as Complete**.
8. Configure an HTTP connection.
- a. Provide the connection name.
 - b. Select the newly created credentials and connection alias.
 - c. Enter the connection URL of the VeloCloud instance.
 - d. Enable use MID Server and choose the appropriate MID option.
 - e. Select **Submit** and then select **Mark as Complete**.
9. Schedule data imports using either bulk or filtered discovery by configuring the import schedule.
- a. Select **Configure** and fill in the fields.
For more information, see [Set up the Service Graph Connector for Arista VeloCloud schedule](#).
 - b. In the **Use connection** field, choose the new VeloCloud instance.
10. Select **Submit**.
11. Confirm that your new instance setup is successful by verifying the configuration.

a. Navigate to **All > Service Graph Connectors > VeloCloud > Connection & Credential Aliases**.

b. Confirm that the new alias is listed.

You can configure multiple connection aliases over the same VeloCloud instance. This flexibility enables you to run imports at different frequencies and apply different filters to each alias.

Mapping Arista VeloCloud CIs and relationships

Use the Service Graph Connector (SGC) for Arista VeloCloud to map discovered physical and logical SD-WAN resources to telecom-aligned configuration item (CI) classes in the Configuration Management Database (CMDB). SGC support consistent service modeling, provide visibility into network components, and automate the creation of logical and physical relationships.

To confirm accurate CI classification and insertion, the connector uses the Robust Transform Engine (RTE) and Identification and Reconciliation Engine (IRE).

The connector classifies and relates discovered CIs using telecom-specific models based on device type, function, and chassis structure. This organization helps maintain a clean and normalized CMDB across vendors. Discovered model names from Arista VeloCloud are automatically transformed into ServiceNow AI Platform standard model identifiers and categories for slot and subslot components.

CI mapping and relationships

The following table lists the CI object types in the CMDB that can be discovered, along with their representations in the CMDB and how they relate to one another.

CMDB CI mapping and relationships (physical layer)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
Network site	cmbd_ci_ni_site	CI → Site → Network Site	Organization network	<ul style="list-style-type: none"> Represents the physical location of IP routers according to their longitude, latitude, and address. Network site contains IP routers and network interfaces. Network site is a member of a group.

CMDB CI mapping and relationships (physical layer) (continued)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
IP router	cmdb_ci_ip_router	CI → HW → NG → IP router	SD-WAN Edge / network or service router is represented by the IP router	<ul style="list-style-type: none"> Represents the VeloCloud device. Contains network interface CIs. Contained by network sites and network service instances. IP router is a member of a group.
Network interface CI	cmdb_ci_ni_interface	Port → Network Port → Network interface	The list of support port models is defined in the vendor-specific network physical information.	<ul style="list-style-type: none"> IP router is the main device in the network hierarchy. Network card within the IP router is the primary component. Represents the physical ports contained within the device (IP router).
Network service instance	cmdb_ci_network_service_instance	CI → Service instance → Network service instance	Network service instance	<ul style="list-style-type: none"> Network service instance includes IP routers and network interfaces.

CMDB CI mapping and relationships (physical layer) (continued)

CMDB CI class	CMDB CI table	CMDB hierarchy	Object types / models	Description and relationships
				<ul style="list-style-type: none"> The network service instance contains the SD-WAN edge devices and is a member of the organization group.
Group	cmdb_ci_group	CI → Group	Represents organization	Network sites and network service instance are members.
IP address CI	cmdb_ci_ip_address	CI → IP address	Represents discovered IP addresses for CIs.	Owned by the corresponding CI.

Configure Telecom Assurance

Setting up and configuring fault management enables you to monitor network devices for configuration and performance issues.

Configure a webhook to monitor SD-WAN network devices using Event Management. This application collects event data from your device and generates alerts.

All events are received in the ServiceNow AI Platform[®] dashboard and automatically mapped to alerts. Event rules evaluate each incoming event and determine whether to create an alert or link it to an existing one. You can define custom event rules, receive notifications using webhook mechanism, and integrate with external systems through the Integrations Launchpad.

Configuring Telecommunications API notifications

Configure the Telecommunications API notification in the ServiceNow instance.

Modeling the Telecommunications API notification workflow

The following steps help to configure the Telecommunications API notification in the ServiceNow instance.

1. **Create a topic:** You can create topics either by manually typing the external message details or automatically collecting the available topics from the external system.
2. **Create a topic subscription:** You subscribe to the available topics for incoming notifications from the external system, based on the customer preference. Additionally, you generate the callback URL and register the subscription.
3. **Activate the Telecommunications Alarm Management Open API endpoint:** To receive responses from the external system, activate the subscribed endpoints of the Telecommunications Alarm Management Open API connection in the Workflow Studio.
4. Provide the callback URL to the external system for receiving notifications. Customer can also reuse the callback URL. When requests from TMF 688 hit the Callback URL, it initiates the *Default Alarm Event Notification Trigger* flow to create an event.

External trigger definitions trigger Event Notification Management Open API requests to create, update, and delete events. To learn more about the functions to handle these requests, see [Event Notification Management Open API](#) and [TMFTopicEventAPIUtilOOB - Scoped](#).

This workflow creates an event in the Event Management application. To learn more about using Event Management, see [Event Management](#).

Related topics

[System components installed with Telecommunications API notifications](#)

[External event management via Telecommunications API notifications](#)

Create a topic

Create a topic and publish the incoming notifications from the external system to the topic. By creating the topics, subscribers can select the topics to which they want to subscribe.

Before you begin

Make sure that the Telecommunications Alarm Management Open API (sn_ind_tmf642) application is installed with the ServiceNow AI Platform.

Role required: admin, sn_api_notif_mgmt.topic_creator

About this task

You can create topics either by manually typing the external message details or automatically collecting the available topics from the external system. When you create a topic, it creates a record in the Topic [sn_api_notif_mgmt_topic] table.

Procedure

1. **All > Telecom API Notification > Topics.**
2. Select **New**.
If you've integrated with an external system, you can select **Get Topics** to get the available topics automatically. This action triggers the *Event Alarm Notification API* subflow. To learn more about the functions that enable you to query and manipulate records in the topic, see [TopicUtilOOB - Scoped](#).
3. On the form, fill in the fields.

Topic form

Field	Description
Topic id	Unique topic id.

Field	Description
Topic name	Name of the topic.
Type	Type of topic. Select one from the following: <ul style="list-style-type: none"> ○ Ingress: Option for inbound notification. ○ Egress: Option for outbound notification.
Header query	Encoded header query parameters. To learn more about the query parameters that follow the TMF 688 standards, see the TM Forum .
Content query	Encoded content query parameters. To learn more about the query parameters that follow the TMF 688 standards, see TM Forum .
Description	A brief description about the topic.

4. Select **Submit**.

Result

A topic is created.

What to do next

You can create the topic subscription according to the customer requirement. For more information, see .

Create a topic subscription

Subscribe to the topic in the ServiceNow AI Platform that you want respond to the incoming notification from the external system. By subscribing to the topic, the subscriber receives the notifications based on the topics that you subscribe to.

Before you begin

- Make sure that the Telecommunications Alarm Management Open API (sn_ind_tmf642) application is installed with the ServiceNow AI Platform.
- Create topics for the incoming notifications.

Role required: admin, sn_api_notif_mgmt.subscription_creator

About this task

You subscribe to the available topics for the incoming notifications from the external system, based on the customer preference. You generate the callback URL to share with the customers. When a request from an external system hits the callback URL, it initiates the creation of an event in the Event Management application.

Additionally, you register the topic subscription to start receiving the incoming notifications. When you create a topic subscription, it creates a record in the Topic Subscription [sn_api_notif_mgmt_subscription] table. To learn more about the methods to query and manipulate records in the Topic Subscription, see [TopicSubscriptionUtilOOB - Scoped](#).

Procedure

1. **All > Telecom API Notification > Subscription.**
2. Select **New**.
3. On the form, fill in the fields.

Topics Subscription form

Field	Description
Topic	Topic that you want to subscribe.
CallbackURL	The callback URL that you're sharing with the external system to capture the incoming notification. The URL is generated automatically when you select Generate CallbackURL .
Filter query	Encoded content query parameters from the topic. You can also modify the filter query. To learn more about the query parameters that follow the TMF 688 standards, see TM Forum .
Registration status	Status of the Topic registration with the external system. By default, it's Unregistered . If the process is successful, the field value changes to Registered . Otherwise it's Error .
Registration message	Registration status message from the external system.
Subscription id	Unique subscription id from the external system.

4. Get the callback URL by selecting **Generate CallbackURL**.
5. Register the subscription by selecting **Register**.

Result

A trigger definition is created for the callback URL and the topic is registered to the external system.

What to do next

In the Workflow Studio, you activate the endpoints of the Telecommunications Alarm Management Open API connection. For more information, see [Activate the Telecommunications Alarm Management Open API endpoint](#).

Activate the Telecommunications Alarm Management Open API endpoint

Activate the endpoint of the Telecommunications Alarm Management Open API connection. By activating the endpoint, you receive the incoming notifications from the external system for the topic that you registered.

Before you begin

- Create the topic and subscribe to it to receive the incoming notifications.
- Generate a callback URL and register the topic subscription.

Role required: admin

About this task

You activate the subscribed endpoints of the Telecommunications Alarm Management Open API connection in the Workflow Studio to receive responses from the external system.

Procedure

1. Navigate to **All > Process Automation > Flow Designer**.
2. Select the **Integrations** tab, then select **Connections**.
3. Switch the toggle to **Inbound**.
4. Select **Telecommunications Alarm Management Open API**.
5. Open the endpoint record that you want to activate.
6. Select **Activate**.

Configure event and metric connectors

Configure a webhook

Integrate with a webhook to connect to an external event source and push event information to your ServiceNow instance.

Before you begin

The plugins listed in the following table must be installed.

Plugin	Plugin ID	Plugin description
TSOM Event Management Connectors	sn_tsom_em_conns	TSOM Assurance Connectors for events and metrics
TSOM Event Management Core	sn_tsom_em_core	TSOM Assurance Core features

Role required: TSOM Assurance admin

Procedure

1. Navigate to **> Workspaces > Service Operations Workspace**.
2. From the navigation pane, select the Integrations Launchpad icon .
3. Select the **Browse Integrations** tab
4. Search for the desired integration (for example, Meraki, Fortinet, or VeloCloud).
5. Select the integration tile.

For most integrations, select the **Events** tile.

Note: For Fortinet, two integration tiles are available: **Events** and **Metrics (Push)**. Select the tile that matches the type of data you want to collect from your Fortinet system.

6. On the Provide details page, provide details about the connector type.

- a. In the **Connector name** field, enter a unique name for the connector type.
- b. **Optional:** In the **Tags** field, enter tags to help locate and identify connectors of this type on the Express List.
 - c. **Optional:** Add additional tags.
 - i. Select the plus sign next to the **Tags** field.
 - ii. In the Add tags window, provide the tag key and tag value.
 - iii. Select **Add**.
 - d. **Optional:** In the **Description** field, enter information to help identify this connector type.
 - e. Verify that the vendor name in the **Source** field is correct.
 - f. **Optional:** From the **Assignment group** drop-down list, select the group or team that is responsible for managing and maintaining the Push connector.
 - g. Select **Next**.
7. On the Set-up push connector page, copy the auto-generated URL parameter value to the clipboard by selecting **Copy to clipboard** in the **URL parameter value** field.

This URL is automatically generated.
8. Follow the instructions on the page to integrate your third-party connector with Event Management.
9. Select **Activate**.
10. In the confirmation message, select **OK**.
11. **Optional:** In the **Details** tab, view the connector details.
The connector is saved and a notification is sent after the connector is active.
12. **Optional:** After you receive a notification that the connector is active, select the **Events** tab to confirm successful configuration and event flow.
Incoming events might take a few minutes to appear in the dashboard.

Related topics

[Exploring Event Management](#) 

[Arista VeloCloud installed integrations](#)

[Cisco Meraki installed integrations](#)

[Fortinet installed integrations](#)

Configure a metric pull connector

Configure metric pull connectors to retrieve performance data from external management systems such as Meraki, Fortinet, or VeloCloud.

Before you begin

Role required: TSOM Assurance admin

Procedure

1. Navigate to **Workspaces > Service Operations Workspace**.
2. From the navigation pane, select the Integrations Launchpad icon .
3. Select the **Browse Integrations** tab
4. Search for the external management system integration (for example, Meraki, Fortinet, or VeloCloud).
5. Select the integration tile labeled **Metrics**.
6. On the Connector Instance New record page, in the **Name** field, enter a unique name for the connector type.
7. In the **Connector definition** field, select the type of integration you're setting up.
8. **Optional:** In the **Description** field, enter brief information about this connector.
9. **Optional:** In the **Assignment group** field, select the group or team responsible for managing and maintaining the connector.
10. In the **Host IP** field, enter the IP address or host name used to select the appropriate MID Server for communicating with the event source host.
11. In the **Credential** field, select the valid credentials to access the event source host.
12. In the **Metrics collection schedule (seconds)** field, enter the polling interval for metric collection.

i Important: For a device performance (CPU usage) connector instance, if you enter a value below 1800 seconds (30 minutes), the system automatically resets the schedule to 1800 seconds and displays a message. For a Fortinet interface logs connector instance, if you enter a value above 600 seconds (10 minutes), the system automatically resets the schedule to 600 seconds and displays a message. In both cases, the record saves successfully. This enforcement applies only when a numeric value is entered; an empty or blank field bypasses it.

13. **Optional:** Validate the connectivity of the connector before activating it by selecting **Test Connector**.
14. Select **Update**.
15. **Optional:** Verify that the pull connector is configured correctly and events are flowing into the system by returning to the Integrations Launchpad.
The tiles appear under the **Installed Integrations** tab.

Related topics

[Exploring Metric Intelligence](#) 

[Arista VeloCloud installed integrations](#)

[Cisco Meraki installed integrations](#)

[Fortinet installed integrations](#)

Configure elastic event pull connectors for MPN

Configure a mobile private network (MPN) pull connector instance to collect metrics, including latency KPIs, from specified KPI domains in Elastic data store. The connector publishes metrics for assurance monitoring and analytics.

Before you begin

The MID Server MID Server must have Access Control (ACL) records for each of the following tables:

- `sn_tsom_em_conns_kpi_definitions` (KPI definitions table)
- `cmdb_ci_network_service_instance`
- `cmdb_ci_app`

For more information, see [Set up a MID Server role](#) .

Roles required: `tsom_assurance_admin` and assigned roles for MID Server.

Procedure

1. Navigate to **All > Event Management > Integrations > Connector Definitions**.
2. Select **Nokia MPN Pull Connector**.
3. Open an existing connector instance or create one:
 - To use an existing instance, select it from the **Connector Instances** list and edit only the fields you need to change.
 - To create an instance, select **New** and complete all required fields as described in the following steps.
4. On the form, fill in the fields.

Field	Description
Host IP	URL of the MPN host
Credential	Username and password for basic authentication
kpi_domain identifier	Domains to collect metrics from. Leaving the field empty collects metrics from all domains. Entering one or more domain identifiers, with values separated by commas (for example, <code>RAN, Core</code>) collect metrics only from specific domains.
Metrics collection schedule	How long, in seconds, each metrics collection should run.

Note: Active connector instances also collect the latency KPIs, from a separate Elasticsearch index, on the same metrics collection schedule. For a list of the latency KPIs collected and how the derived KPI is calculated, see the MPN latency KPIs reference.

5. Add a MID Server.
 - a. In the **MID Servers for Connectors** section, select the plus icon next to **Insert a new row**.
 - b. In the **MID Server** field, enter the name of the MID Server.
 - c. Save your selection by selecting the green check mark.
6. Enable the connector instance by selecting the **Active** check box.
The connector begins collecting metrics from Elastic data store.
7. Select **Update**.

Related topics

[MPN latency KPIs](#)

[View metric to CI and resource binding](#) ↗

[Event field mapping configuration](#) ↗

[Create an event rule to bind metric events to host CIs](#) ↗

[View metric values in the Insights Explorer](#) ↗

[Create an Insights Explorer view](#) ↗

Configure elastic connectors for MPN alarm collection

Configure a connector instance to collect fault management alarm data from a Mobile Private Network (MPN) Elastic index and forward events to Event Management.

Before you begin

1. [Create Basic Auth Credentials](#) ↗
2. [Create HTTP\(S\) Connection](#) ↗
3. [Activate the Telecommunications Alarm Management Open API endpoint](#)

Role required: tsom_assurance_admin

About this task

The MPN connector collects alarm and event data from managed network infrastructure by querying Elastic fault management indices at a configured polling interval. Each connector instance collects from one or more of these indices, each corresponding to a specific network domain. The MPN pull connector definition supports three connector instances: one for alarm collection and two for metrics collection. This procedure covers configuring the alarm collection instance.

You can create multiple connector instances, one per index, and assign each to a dedicated MID Server to distribute the polling workload across your deployment. A default index value is provided for each connector instance.

The alarm collection instance also clears stale alarms: alarms that were cleared at the source but were missed during a normal collection cycle and would otherwise remain open in Event Management. At a regular interval, the connector rechecks the source for active alarms that have not been updated within a set time window. It then clears any that are no longer active. Default values are provided, so no action is required unless you want to tune this behavior.

Procedure

1. Navigate to **All > Event Management > Integrations > Connector Definitions**.
2. Select **Nokia MPN Pull Connector**.
3. In the **Default schedule** field, enter the polling time (in seconds) to be used for all connector instances.
The default value is 120 seconds.
4. In the Connector Instances related list, select **Nokia MPN Alarms**.
5. On the Connector Instances page, in the **Event collection schedule** field, enter how long each event collection should run, in seconds.
This value overrides the default schedule set on the connector definition. The MPN Alarms instance defaults to 60 seconds.
6. In the **Index** field, provide the index connector instance values.

Enter the index pattern for the network domain this connector instance collects from. To collect from more than one domain, enter the patterns as a comma-separated list. The following table lists the standard index patterns:

Connector Instance Values	
Name ▲	Value
index	radio-fm*, core-fm*
last_stale_check	
log_payload_for_debug	false
stale_check_interval_minutes	30
stale_threshold_minutes	60

Index connector instance values

Description	Name
Radio network elements	radio-fm*
Core network elements	core-fm*
Application-layer network elements	application-fm*
IXR (Interconnect Router) devices	ixr-fm*
DAC (Digital Access Controller) devices	dac-fm*

7. Add the MID Server.

- In the **MID Servers for Connectors** section, select the plus icon next to **Insert a new row**.
- Enter the name of the MID Server in the **MID Server** field.
- Select the green check mark to save your selection.

8. Enable the connector instance by selecting the **Active** check box.

9. **Optional:** Tune how stale alarms are detected and cleared.

Set the following connector instance values on the alarm collection instance. If you don't set them, the default values apply.

Stale alarm detection parameters

Parameter	Default	Description
stale_threshold_minutes	60	How long, in minutes, an active alarm can go without an update at the source before the connector treats it as stale and clears it in Event Management.
stale_check_interval_minutes	30	How often, in minutes, the connector runs the stale alarm check. This check runs independently of the event collection schedule.

10. Select **Update**.

11. **Optional:** View your configured alarms by navigating to **em_event.LIST** and searching for API Alarm Notification.

Configure security log collection for MPN

Configure the out-of-box Health Log Analytics security log data input, an Elasticsearch collector, to collect Mobile Private Network (MPN) security logs and convert them into structured log records.

Before you begin

Roles required:

- `tsom_assurance_admin`
- Assigned role for MID Server

For instructions on setting up a MID Server, see [Set up a MID Server role](#) .

About this task

MPN security log collection uses the out-of-box security log data input under Health Log Analytics, configured with values specific to the MPN security log source. The data input is an Elasticsearch collector that queries a security log index and passes each raw log document through a parser. The parser extracts structured fields, including a customer identifier, from the raw message text. For a list of the fields produced, see the MPN security log fields reference.

Procedure

1. Navigate to **All > Health Log Analytics > Data inputs** and open the out-of-box security log data input.

This data input is an Elasticsearch collector provided predefined for MPN security logs.

2. Set **Execute on** to **Specific MID Server** and select the **MID** to use for the connection. Use the lookup to select a MID Server that is installed on your instance.

3. Select the **Service instance** for this data input. The service instance is required. It is used to correlate configuration items (CIs) and service instances in the CMDB for root cause analysis.

4.  **Note:** Credentials aren't provided predefined.

On the **Transport** tab, configure the connection to your security log store.

Configure the following fields:

- **Server URL:** The endpoint of your security log store. A sample URL is provided predefined; replace it with the URL for your environment.
- **Authentication method:** The authentication mechanism for your store. The out-of-box configuration uses basic authentication.
- **Basic auth credentials:** The credentials for your store. Shown when **Authentication method** is set to **Basic authentication**.

5. On the **Query Settings** tab, provide the following MPN-specific values.

MPN security log query settings

Field	Value
Index	Index pattern for your security log store, for example <code>app-security-logging-staging</code> . Use your staging index pattern in non-production environments, and your production index pattern once you move to production.
Timestamp field	<code>@timestamp</code>

Field	Value
Timestamp format	yyyy-MM-dd'T'HH:mm:ss.SSS'Z'

Warning: Don't modify the **Timestamp field** or **Timestamp format** values. Changing them can break log ingestion.

6. Attach the MPN security log source type parser to the data input.

The parser is provided predefined, automatically detects the log source subtype—NDAC, COVE, AWS CloudWatch, or Management-Plane—and extracts structured fields from each raw log document. If a structured `raw_data` field is present, the parser extracts from it directly; otherwise it falls back to regex parsing of the `message` field.

The parser also maps log severity: INFO logs are treated as OK (normal/informational activity), WARN logs as Minor (informational alerts), and ERROR logs as Major/Critical (for example, unauthorized access, denied operations, off-hours sensitive access, or failed logins).

All fields except the customer identifier are extracted automatically and require no configuration. The customer identifier is extracted using a configurable pattern in the parser's JavaScript function. To change how the customer identifier is derived, edit this pattern.

Note: Properties longer than 256 characters aren't extracted and aren't indexed as keywords in Elastic.

7. On the **Advanced** tab, set the collection **Frequency**.

The frequency controls how often the data input pulls data. The default is every minute; adjust it as needed, for example every 5 or 10 minutes.

8. Save and test the data input connection, then select **Publish**.

When you publish the data input, its dependencies, including the streaming source, are created automatically by the Health Log Analytics infrastructure.

Result

The data input starts streaming security log data from the configured Elasticsearch index and produces structured log records. For example, a raw log entry such as:

```
SECURITY: Wrong password **** customer: "C1234567", site:
"<site-id>", user: "<user>", client: "<client-ip>", server:
"<server>", request: "POST /auth **** HTTP/1.1", host: "<host>"
```

produces a structured record containing fields such as customer, site, user, client IP, and host, shown concatenated in the log message in the log viewer.

Related topics

#unique_116

Configure an anomaly alert for an MPN data input

Configure a log anomaly alert for a Mobile Private Network (MPN) data input that generated its own source type.

Before you begin

You created a data input and, if applicable, a source type, as described in [Configure security log collection for MPN](#).

Role required: `tsom_assurance_admin`

About this task

An anomaly alert is configured for the out-of-box MPN source type and fires on Major-severity ERROR events with no aggregation window. If your data input generated its own source type, configure a matching alert for it.

Procedure

1. Navigate to **Health Log Analytics > Log Anomaly Detection > Define Alert** and select **Create Alert**.
2. Select the index corresponding to the data input and source type pairing you configured.
3. Define the filter or query for the condition to detect.
For example, specify a severity or field value. Use the **Lexical Keywords**, **Metrics Threshold Raises**, or **Advanced Log Alert Filter** related lists if a simple filter is not sufficient.
4. Set the trigger thresholds, then set the alert severity, name, and description.
Trigger thresholds include the number of matches, time period, and consecutive periods.
5. Activate the alert and validate it with a test event.
Use **Muted Metrics** to temporarily suppress noisy conditions while you tune the alert.

Configure elastic connectors for MPN health status collection

Configure a connector instance to collect health status data from an MPN Elastic index, and optionally customize the rules used to calculate that status.

Before you begin

1. [Create Basic Auth Credentials](#) 
2. [Create HTTP\(S\) Connection](#) 

Role required: `tsom_assurance_admin`

About this task

Health status collection lets you monitor the health of radio and core server components in your MPN so you can identify and respond to issues.

The connector instance for health status collection is created automatically but is inactive by default. Complete this procedure to activate it.

Procedure

1. Navigate to **All > Event Management > Integrations > Connector Definitions**.
2. Select **Nokia MPN Pull Connector**.
3. In the Connector Instances related list, select **Nokia MPN RAG Status**.
4. Review the **Event collection schedule** field.

This value defaults to 60 seconds. Change it if you need a different polling interval for this instance.

5. **Optional:** Customize the classification rules.

(Optional) By default, health status is calculated using a predefined rule set stored in the `nokia.rag.rules.default` system property. To override it, enter your own rule set, in the same JSON structure, in the `nokia.rag.rules.override` system property.

Note:

If the override property is empty or contains invalid data, the system uses the default rule set instead.

6. Enable the connector instance by selecting the **Active** check box.

7. Select **Update**.

Result

The connector instance begins polling the configured MPN Elastic index and calculates a health status for each radio and core server component, using the active rule set.

Configure the Equinix pull connector

Configure a connector instance to collect power and environmental metrics from Equinix IBX data center facilities and forward alerts to Event Management.

Before you begin

- The TSOM EM Connectors plugin (*sn_tsom_em_conns*) must be installed.
- A credential record must exist for the Equinix OAuth2 client secret. Do not store the client secret as a connector instance parameter.

Role required: *tsom_assurance_admin*

About this task

The Equinix connector collects power and environmental readings and system alerts from the Equinix Colo Portal Open API for each configured account. On each run, the connector authenticates with OAuth2 and retrieves the facility hierarchy for the account (IBX, zone, cage, and cabinet). The connector then requests power and environment readings for that hierarchy and polls for system alerts on a separate, paginated request.

You can configure multiple accounts on a single connector instance by entering a comma-separated list of account numbers.

Procedure

1. Navigate to **All > Event Management > Integrations > Connector Definitions**.

- Note:** You can also discover and open this connector from Service Operations Workspace. From the navigation pane, select the AIOps Configuration Center icon, select **Add integration**, and search for **Equinix IBX SmartView** in the Integration Launchpad.

2. Select **Equinix IBX SmartView**.

3. Create a connector instance or select an existing instance to edit by selecting **New** or selecting an existing instance.

4. Complete the following fields.

Equinix connector instance parameters

Field	Description
<i>host</i>	Base URL of the Equinix Colo Portal Open API. Required.
<i>client_id</i>	OAuth2 client ID for your Equinix account.

Field	Description
<i>accounts</i>	Comma-separated list of Equinix account numbers to poll. Example: 578555 , 548686 , 111850.
<i>auth_url</i>	OAuth2 token endpoint path.
<i>last_event</i>	ISO 8601 timestamp used as the starting point for incremental alert collection. Example: 2026 - 06 - 19T00 : 00 : 00 . 000Z.
<i>debug</i>	Enables additional logging for troubleshooting.

5. Link the credential record that contains the OAuth2 client secret.
The connector retrieves the client secret from this credential record at runtime. The connector never reads the secret from a plain connector instance parameter.
6. Add the MID Server.
 - a. In the **MID Servers for Connectors** section, select the plus icon next to **Insert a new row**.
 - b. In the **MID Server** field, enter the name of the MID Server.
 - c. Select the green check mark to save your selection.
7. Enable the connector instance by selecting the **Active** check box.
8. Select **Update**.
9. **Optional:** Verify the connection by running `testConnection()` on the connector instance.
A successful test returns a structured result with individual pass or fail status for the authentication, hierarchy, and power checks.

Result

The connector begins polling on the configured schedule. Power and environmental readings appear in MetricBase for each cabinet and zone in the account's facility hierarchy, and system alerts are forwarded to Event Management.

Override default metric-to-CI binding

Replace the shipped logic that binds collected metrics to configuration items (CIs) for a Telecommunications Service Operations Management metric source. Create your own implementation of the `EventFieldMapping` extension point and wire it into an event field mapping rule.

Before you begin

- Set your application scope to **Telecommunications Service Operations Management Event Management Connectors** (`sn_tsom_em_conns`) using the scope switcher. Working in the Global scope generates the implementation in the wrong scope, preventing the resolver from finding it at run time.
- Identify the source name as it appears on the metric and on the corresponding event. The source name is the value used to resolve which scripted extension handles the event.
- Decide which rule name your custom implementation will own. Each implementation is identified by a unique combination of source name and rule name. You will use this same rule name in the script include's `getRuleName()` method and in the rule's **Name** field, and the two must match exactly.
- Determine which CMDB class and field your implementation should match against, and which event fields contain the values to match.

Role required: `tsom_assurance_admin`

For an overview of how metric-to-CI binding works in Telecommunications Service Operations Management service graph connectors, see [Metric-to-CI binding](#).

About this task

When a connector collects metric data, the metric framework stores it and creates a metric-to-CI mapping record. Because the record does not yet have a configuration item (CI), the framework automatically generates an event for it. That event triggers the event field mapping rules. A rule whose source matches the event runs a script that resolves the CI and sets the `cmdb_ci` field on the event. The resolved CI is then reflected in the CI column of the Metric To CI Mappings table:

metric-to-CI mapping record → event created → event field mapping rule → `cmdb_ci` set on the event → configuration item shown in the Metric To CI Mappings table

The shipped behavior is provided by an `EventFieldMapping` implementation that the rule's script retrieves through a resolver. The resolver, `EventFieldMappingResolver.getByRule()`, returns the implementation whose `getVendor()` matches the source name you pass and whose `getRuleName()` matches the rule name. That implementation exposes a `mapCi()` method that performs the CI lookup. To override the shipped behavior for a connector source, create your own implementation under the same source name with a new rule name. Set its `getVendor()`, `getRuleName()`, and `mapCi()` methods. Wire it into an event field mapping rule whose script delegates to the resolver.

i Important: The resolver matches on names, and a mismatch fails silently: no error is thrown, the CI is simply never set. The source name returned by `getVendor()` must match the literal you pass to `getByRule()` in the rule script, and the rule name returned by `getRuleName()` must match the rule's **Name**. A typo, a case difference, or the wrong application scope all produce the same result: the resolver returns nothing and the CI is not set. The resolver itself doesn't log the mismatch; the "No EventFieldMapping implementation found" message is written by the fallback in the rule script you create in step 3.

Procedure

1. Create an implementation of the `sn_tsom_em_conns.EventFieldMapping` extension point.

- a. Navigate to **All > System Extension Points > Scripted Extension Points**.

- b. Search for the API name `sn_tsom_em_conns.EventFieldMapping` and open it.

- c. In the **Related Links** section, select **Create implementation**.

The platform generates a script include and registers it as an extension instance automatically, so the resolver can find it. Creating the implementation from this link is what registers the instance; you don't register it as a separate step.

- d. Set the script include name (for example, `NokiaMPNMappingRule`) and save it.

2. Implement the extension point methods in the generated script include.

- a. In `getVendor()`, return your source name (for example, `return 'Nokia MPN';`). The resolver matches this value, so it must be exact.

- b. In `getRuleName()`, return your rule name (for example, `return 'NokiaMPNMappingRule';`).

This value must match the **Name** of the rule you create in step 3.

- c. In `mapCi(eventGr, origEventSysId, fieldMappingRuleName)`, read values from the event, look up the matching configuration item, and set the `cmdb_ci` and `ci_type` fields.
For example, read fields under `additional_info`.

- d. Return `true` from `mapCi()` to continue event processing, or `false` to drop the event.

- e. Save the script include.

i Note: For example, the shipped Nokia MPN implementation resolves the CI on the table named by `additional_info.ciClass` (default `cmdb_ci`), trying these event fields in order: `additional_info.name` (direct name match), `additional_info.pmDataSource.dn` (bottom-up DN traversal, deepest matching component wins), `serial_no`, `hw_id`, and `nhg_alias`. On a match, it sets `cmdb_ci` and `ci_type` on the event.

3. Create an event field mapping rule that delegates to your implementation.

- a. Navigate to **All > Event Management > Rules > Event Field Mapping**, and then select **New**.

- b. In **Name**, enter your rule name.

The value must match what `getRuleName()` returns. The platform passes this name to the rule script as `fieldMappingRuleName`, and the resolver uses it to find your implementation.

- c. Set **Source** to your source name.

This field filters which events the rule applies to.

- d. Set **Mapping type** to **Advanced mapping using script**.

- e. Set **Active** to true.

- f. Set **Order** to a value lower than any existing rule for the same source that you want this rule to take priority over.

Lower values run first. The default rule runs at order 100. To make your rule run first, set a value below 100, for example, 50.

- g. In **Filter**, add the condition `[Source] [is] [your source name]`.

- h. In the **Script** field, retrieve your implementation through the resolver and call `mapCi()` on it:

```
(function eventFieldMappingScript(eventGr, origEventSysId,
    fieldMappingRuleName) {
    var impl = new
    sn_tsom_em_conns.EventFieldMappingResolver()
        .getByRule('<your source name>', fieldMappingRuleName);
    if (impl) {
        return impl.mapCi(eventGr, origEventSysId,
            fieldMappingRuleName);
    }
    gs.warn('No EventFieldMapping implementation found for rule: ' +
        fieldMappingRuleName);
}
```

```

    return true;
  })(eventGr, origEventSysId, fieldMappingRuleName);

```

Replace `<your source name>` with the same value that `getVendor()` returns. If you reuse this script for another source, update it in both places.

i. Submit the rule.

4. Confirm the override.

a. Trigger the connector to collect metrics for the source.

The metric framework creates the metric record and its event automatically.

b. Navigate to **All > Event Management > Metrics > Metric to CI**.

In the **Configuration item** column, verify that new metrics for your source show the CI resolved by your implementation. To trace a binding, open the linked event and verify that the `cmdb_ci` field is set.

c. If the CI is not set, check the system logs for the message "No EventFieldMapping implementation found".

That message means the resolver could not match your `getVendor()` and `getRuleName()` values, or your implementation was created in the wrong application scope. Confirm the names match across the script include and the rule, and that the script include is in the `sn_tsom_em_conns` scope.

Result

New metrics collected from the source are bound to configuration items by your custom implementation.

Related topics

[Event field mapping configuration](#) 

[Create event field mappings](#) 

[MPN Formulas table](#)

Configure a metric aggregation job

Configure a scheduled job to aggregate raw metrics into a calculated metric on a parent or anchor configuration item (CI). Aggregated metrics let you view performance data at a higher level of the network.

Before you begin

A pull connector must be configured and collecting metrics for the source resources to aggregate.

Role required: `tsom_assurance_admin`

About this task

Each scheduled job handles one aggregation. The job script defines the source data, aggregation method, and result destination.

Procedure

1. Navigate to **All > System Definition > Scheduled Jobs**.

2. Select **New**.

3. Select **Automatically run a script of your choosing**.
4. In the **Name** field, enter a meaningful name for the aggregation.
5. Set the run schedule for the job.
For example, to run the aggregation periodically, set **Run** to **Periodically**, set the repeat interval, and set the start date and time.
6. In the **Script** field, enter the script that runs the aggregation.

The script calls the metric aggregation scripted extension point and passes a configuration object that defines the aggregation. The following example aggregates a per-cell metric into an average value on each Baseband Unit.

```
var impls = new GlideScriptedExtensionPoint()
    .getExtensions('sn_tsom_em_conns.TSOMMetricAggregatorSEP');

if (!impls || !impls.length) {
    gs.warn('TSOMMetricAggregator: No implementations
    registered for TSOMMetricAggregatorSEP');
} else {
    impls[0].aggregate({
        mode: 'hierarchy',
        targetCIClass: 'cldb_ci_baseband_unit',
        sourceCIClass: 'cldb_ci_mobile_cell',
        sourceMetric: 'NR_5076b',
        aggregate: 'avg',
        calculatedMetricName: 'NR_5076b_avg_bbu',
        unit: 'Mbps',
        level: 4
    });
}
```

For a flat, range-based aggregation, set the mode to flat, identify the source resources by range, and optionally apply field filters. Provide the anchor CI that holds the result, as shown in the following example.

```
var impls = new GlideScriptedExtensionPoint()
    .getExtensions('sn_tsom_em_conns.TSOMMetricAggregatorSEP');

if (!impls || !impls.length) {
    gs.warn('TSOMMetricAggregator: No implementations
    registered for TSOMMetricAggregatorSEP');
} else {
    impls[0].aggregate({
        mode: 'flat',
        sourceCIClass: 'cldb_ci_sim',
        rangeStart: 'SIM010',
        rangeEnd: 'SIM050',
        filters: [
            { field: 'operational_status', operator: '=',
value: '1' }
        ],
        sourceMetric: 'sim_throughput',
        aggregate: 'avg',
        calculatedMetricName: 'sim_throughput_avg_global',
        anchorCiSysId: '<sys_id>',
        unit: 'Mbps'
    });
}
```

```
});
}
```

To roll up only a subset of source CIs into a parent CI, use hierarchy mode. Narrow the source resources by range and field filters, as shown in the following example. This example averages a metric across a range of SIM cards with an active operational status. The result is published to the parent mobile private network (MPN) service.

```
var impls = new GlideScriptedExtensionPoint()
    .getExtensions('sn_tsom_em_conns.TSOMMetricAggregatorSEP');

if (impls && impls.length) {
    impls[0].aggregate({
        mode: 'hierarchy',
        targetCIClass: 'cmdb_ci_network_service_instance',
        sourceCIClass: 'cmdb_ci_sim_card',
        rangeStart: 'PLMN-PLMN/SST-1/SD-000011',
        rangeEnd: 'PLMN-PLMN/SST-1/SD-000015',
        filters: [
            { field: 'operational_status', operator: '=',
value: '1' }
        ],
        sourceMetric: 'NR_5076b',
        aggregate: 'avg',
        calculatedMetricName:
'NR_5076b_avg_11_to_15_per_MPN_Service',
        unit: 'Mbps',
        level: 4
    });
}
```

To reference the scripted extension point by its fully qualified scoped name and log a warning when no implementation is registered, use the following pattern.

```
var impls = new GlideScriptedExtensionPoint()
    .getExtensions('sn_tsom_em_conns.TSOMMetricAggregatorSEP');

if (!impls || !impls.length) {
    gs.warn('TSOMMetricAggregator: No implementations
registered for TSOMMetricAggregatorSEP');
} else {
    impls[0].aggregate({
        mode: 'hierarchy',
        targetCIClass: 'cmdb_ci_network_service_instance',
        sourceCIClass: 'cmdb_ci_sim_card',
        rangeStart: 'PLMN-PLMN/SST-1/SD-000011',
        rangeEnd: 'PLMN-PLMN/SST-1/SD-000015',
        filters: [
            { field: 'operational_status', operator: '=',
value: '1' }
        ],
        sourceMetric: 'NR_5076b',
        aggregate: 'avg',
        calculatedMetricName:
'NR_5076b_avg_11_to_15_per_MPN_Service',
        unit: 'Mbps',
        level: 4
    });
}
```

```
});  
}
```

7. Select **Submit**.

Note: For a description of each configuration value, see [Metric aggregation configuration reference](#).

Result

The scheduled job runs on the defined schedule. Each run produces the calculated metric and raises an event to bind the metric to a CI. After the metric-to-CI binding completes, the calculated metric is stored in the metric base and is available to view in the Insight Explorer.

Related topics

[Metric aggregation modes](#)

[Create a custom metric aggregation implementation](#)

[Metric-to-CI binding](#)

Create a custom metric aggregation implementation

Create a custom implementation of the metric aggregation scripted extension point when the default aggregation modes don't cover a calculation you need.

Before you begin

Role required: tsom_visibility_admin

About this task

The metric aggregation scripted extension point includes a default implementation that covers the standard aggregation modes. If the default implementation doesn't cover the aggregation you need, create your own implementation of the extension point. Your implementation runs without requiring a product change.

Procedure

1. Navigate to **All > System Scripted Extension Points > Scripted Extension Points**.
2. Search for **TSOMMetricAggregator**.
3. Select **Create implementation**.
4. In the **Script** field, define your aggregation logic.

Provide the aggregation configuration that your implementation uses, following the same configuration structure as the default aggregation modes.

5. In the **Order** field, set the execution priority for your implementation.
6. Select **Submit**.

Result

Your custom implementation is active. A scheduled job that calls the metric aggregation scripted extension point runs your implementation along with the other registered implementations.

Related topics

[Metric aggregation modes](#)

[Configure a metric aggregation job](#)

Using Telecommunications Service Operations Management

Use Telecommunications Service Operations Management (TSOM) to proactively monitor telecom services, validate data integrity, and reconcile discrepancies across network inventory and discovery sources. TSOM enables operations teams to maintain an accurate telecom-aware CMDB and act on real-time network insights.

Common operational tasks of TSOM

Once TSOM is configured, you can use its capabilities to:

- **Run Telecom Discrepancy audits:** Identify inconsistencies in network relationships and configurations by running certification audits based on discovered data and reconciliation rules.
- **Validate attribute value mismatches in CMDB 360:** Use the Attribute Value Discrepancy feature in CMDB 360 to compare attribute values across different discovery sources and detect conflicts or outdated records.
- **Use Telecom Discovery Patterns:** Apply Telecom Discovery Patterns to accurately discover and populate telecom-specific Configuration Items (CIs) and their relationships.
- **Run and verify import jobs from service graph connectors:** Manually trigger import schedules (for example, for the Nokia Altiplano connector). Verify that topology and device data is successfully imported into the CMDB and mapped to the correct CI classes.

These tasks help verify that your telecom network is continuously synchronized, accurate, and operationally visible within ServiceNow.

Use Telecom Discovery patterns

Use Telecom Discovery patterns to identify and classify network functions (xNFs) from vendor devices such as Cisco, Juniper, routers, or switches. Use pattern-based discovery to map telecom resources into the CMDB.

Telecom Discovery Patterns automatically classify some discovered network attributes into standardized CMDB values, so that patterns from different vendors produce consistent CI data.

Port bandwidth classification

When a Telecom Discovery pattern discovers a network interface (port), it uses the reported bandwidth to automatically assign an interface hardware model to the **Model ID** field. The assignment is based on the following fixed set of bandwidth values:

Discovered bandwidth	Assigned interface model
100000	100 GigabitEthernet
40000	40 GigabitEthernet
25000	25 GigabitEthernet
10000	10 GigabitEthernet

Discovered bandwidth	Assigned interface model
1000	GigabitEthernet
10–100	FastEthernet
0	1/10/25G Ethernet

A bandwidth value of **0** is not treated as "no bandwidth" or an error—it resolves to the **1/10/25G Ethernet** model. A discovered bandwidth that doesn't match any value in this table is left without a model assignment.

Note: The interface model is only assigned if a matching Product Model (cmdb_model) record already exists in your instance. If the corresponding model record is missing, the Model ID field is left blank even for a recognized bandwidth value.

Classify xNFs for Telecom Router Pattern

To access a full list of OIDs that will be classified.

Before you begin

Role required: admin

Classifier name: **Standard Network Router**.

Procedure

1. Navigate to **All > Discovery Definition > CI Classification > SNMP**.
2. From the list, select **Standard Network Router**.
3. Open the tab **SNMP OID Classifications** and see the list of OIDs.

Note:

For more information on how to add additional OIDs to the classifier, see [Direct Discovery using Discovery Patterns](#).

MiB Tables Used on an xNF:

- SystemMIB
- EntityPhysicalMIB
- IfMIB
- IfXMIB
- IpMIB

Classify xNFs for Telecom Cisco 7613 Router Pattern

To access a full list of OIDs that will be classified.

Before you begin

Role required: admin

Classifier name: **Standard Network Router**.

Procedure

1. Navigate to **All > Discovery Definition > CI Classification > SNMP**.
2. From the list, select **Standard Network Router**.

3. Open the tab **SNMP OID Classifications** and see the list of OIDs.

Note:

For more information on how to add additional OIDs to the classifier, see [Direct Discovery using Discovery Patterns](#).

The List of Specific OIDs to call this Pattern:

Vendor	Model	OID	Pattern
Cisco	7613	1.3.6.1.4.1.9.1.528	Telecom Cisco 7613 Router

MiB Tables Used on an xNF:

- SystemMIB
- EntityPhysicalMIB
- IfMIB
- IfXMIB
- IpMIB

Classify xNFs for Telecom Juniper MX SSH Router Pattern

To access a full list of OIDs that will be classified.

Before you begin

Role required: admin

Classifier name: **Standard Network Router**.

Procedure

1. Navigate to **All > Discovery Definition > CI Classification > SNMP**.
2. From the list, select **Standard Network Router**.
3. Open the tab **SNMP OID Classifications** and see the list of OIDs.

Note:

For more information on how to add additional OIDs to the classifier, see [Direct Discovery using Discovery Patterns](#).

The List of Specific OIDs to call this Pattern:

Vendor	Model	OID	Pattern
Juniper	MX80	1.3.6.1.4.1.2636.1.1.1.2.57	Telecom Juniper MX SSH Router
Juniper	MX104	1.3.6.1.4.1.2636.1.1.1.2.97	Telecom Juniper MX SSH Router
Juniper	MX240	1.3.6.1.4.1.2636.1.1.1.2.29	Telecom Juniper MX SSH Router
Juniper	MX480	1.3.6.1.4.1.2636.1.1.1.2.25	Telecom Juniper MX SSH Router

MiB Tables Used on an xNF: SystemMIB.

CLI Commands Used.

- show chassis hardware | no-more | display xml
- show interface media | no-more | display xml

Classify xNFs for Telecom Cisco Switch Pattern

To access a full list of OIDs that will be classified.

Before you begin

Role required: admin

Classifier name: **Standard Network Switch.**

Procedure

1. Navigate to **All > Discovery Definition > CI Classification > SNMP.**
2. From the list, select **Standard Network Switch.**
3. Open the tab **SNMP OID Classifications** and see the list of OIDs.

i Note:

For more information on how to add additional OIDs to the classifier, see [Direct Discovery using Discovery Patterns](#).

The List of Specific OIDs to call this Pattern:

Vendor	Model	OID	Pattern
Cisco	Nexus 9000	1.3.6.1.4.1.9.12.3.1.3.1954	Telecom Cisco Switch
Cisco	Nexus 3548	1.3.6.1.4.1.9.12.3.1.3.1666	Telecom Cisco Switch

MiB Tables Used on an xNF:

- SystemMIB
- EntityPhysicalMIB
- IfMIB
- IfXMIB
- IpMIB

Classify xNFs for Telecom Switch Pattern

To access a full list of OIDs that will be classified.

Before you begin

Role required: admin

Classifier name: **Standard Network Switch.**

Procedure

1. Navigate to **All > Discovery Definition > CI Classification > SNMP.**
2. From the list, select **Standard Network Switch.**
3. Open the tab **SNMP OID Classifications** and see the list of OIDs.

Note:

For more information on how to add additional OIDs to the classifier, see [Direct Discovery using Discovery Patterns](#).

MiB Tables Used on an xNF:

- SystemMIB
- EntityPhysicalMIB
- IfMIB
- IfXMIB
- IpMIB

Run SGC imports

Run and verify a Nokia Altiplano import

Manually execute a configured import schedule for the Nokia Altiplano Service Graph Connector. You can also execute the import to validate the connector setup, run ad-hoc imports, or test newly configured connection aliases. This task helps ensure that data from Altiplano is successfully imported or updated in the CMDB.

Before you begin

Role required: TSOM Visibility admin

About this task

After configuring the Nokia Altiplano Service Graph Connector and setting up one or more import schedules, you can execute an import job. This task helps you verify that the integration is working as expected and that CIs are successfully imported or updated in the CMDB. You can either enable the import schedules to run automatically based on their defined frequency, or trigger them manually for immediate execution and validation.

Note: If you have configured multiple Altiplano connection aliases (for different environments or filters), confirm that the import schedule is set up correctly for each alias.

The following screenshot helps you understand the scheduled data import process and displays the executions in the form of import sets.

Scheduled Data Import
Altiplano bulk

Name:

Data source:

Run as:

Active: ☒

Use connection: ☒

Connection:

Concurrent Import: ☐

Execute pre-import script: ☐

Execute post-import script: ☐

Application:

Run:

Time: Hours Minutes Seconds

Conditional: ☐

[Execute Now](#) [Update](#) [Delete](#)

Related Links
[View Dashboard](#)
[Run Point Scan](#)
[\[SN Utils\] Versions \(4\)](#)

Executions Number

Number	Execution Context	Import Set	Created
EXE00001043	EC00001043	Import Set: ISET0010043	2025-07-13 00:00:00
EXE00001042	EC00001042	Import Set: ISET0010042	2025-07-12 23:36:01
EXE00001040	EC00001040	Import Set: ISET0010040	2025-07-12 00:00:00

Procedure

1. Navigate to **All > Nokia Altiplano > Import Schedules**.
2. Select the import schedule that you want to run.
The list of scheduled data imports appear.
3. Select the scheduled data import.
4. Do one of the following:
 - If the schedule is **Active**, the job runs automatically at the defined interval. So wait until it's executed.
 - Select **Execute Now** to run the import immediately.
5. To monitor the execution and verify results:
 - a. Scroll down to the **Executions** list or the table linked to the import schedule.
 - b. Open the most recent **Import Set** record created by the execution.
6. Review the import set log to verify the numbers of rows read, number of rows inserted or updated in the CMDB, or transformation success status.
7. **Optional:** Navigate to **All > CMDB > CI Classes** or your custom telecom tables (such as TNI base item) to confirm the CIs were created or updated as expected.

Result

The import schedule runs immediately, and you can verify the import set execution and the corresponding CI updates in the CMDB. If configured correctly, the connector brings in the network inventory data from Nokia Altiplano into your ServiceNow instance.

Example: After executing the import schedule:

- The log shows: "12 rows read, 0 inserted/updated." This indicates that CIs exist.
- If you delete existing CIs and rerun the import, the log might show multiple inserts and updates, validating end-to-end functionality.

Note:

- If **Concurrent Import** option is enabled in the import schedule, you'll see records in the **Concurrent Import Set** table instead of the standard Import Set table.
- Each concurrent job creates its own import set and log entry. The structure of the execution records remains the same.

Related topics

[Configure Nokia Altiplano service graph connector](#)

[Telecom Discovery via Nokia Altiplano](#)

Run and verify an import schedule for Fortinet SGC

Manually run a configured import schedule for the Fortinet Service Graph Connector (SGC) to verify that data from Fortinet was successfully imported or updated in the Configuration Management Database (CMDB). You can also perform the import to validate the connector setup, run one-off imports, or test newly configured connection aliases.

Before you begin

Role required: TSOM Visibility admin

If you have configured multiple Fortinet connection aliases for different environments or filters, confirm that the import schedule is set up correctly for each alias.

About this task**Note:**

- If the Concurrent Import option is enabled in the import schedule, records are imported into the Concurrent Import Set table instead of the standard Import Set table.
- Each concurrent job creates its own import set and log entry. The structure of the execution records remains the same.

Procedure

1. Navigate to **All > Service Graph Connectors > Fortinet > Import Schedules**.
2. Select the import schedule to run.
3. From the list of scheduled data imports, select the scheduled data import.
4. If the import needs to be run manually rather than having run automatically on a specified schedule, select **Execute now**.
5. In the Executions related list, verify the import results by selecting the most recent Import Set record created by the execution of the import.
6. Review the import set log to verify the number of rows read, the number of rows inserted or updated, or the transformation success status.
These values in the CMDB indicate that the CIs exist. If you delete existing CIs and rerun the import, the log might show multiple inserts and updates, validating end-to-end functionality.

Result

If configured correctly, the connector brings in the network inventory data from Fortinet into your ServiceNow AI Platform[®] instance.

Run and verify an import schedule for Cisco Meraki SGC

Manually run a configured import schedule for the Cisco Meraki Service Graph Connector (SGC) to verify that data from Cisco Meraki was successfully imported or updated in the Configuration

Management Database (CMDB). You can also perform the import to validate the connector setup, run one-off imports, or test newly configured connection aliases.

Before you begin

Role required: TSOM Visibility admin

If you have configured multiple Cisco Meraki connection aliases for different environments or filters, confirm that the import schedule is set up correctly for each alias.

About this task

Note:

- If Concurrent Import option is enabled in the import schedule, records are imported into the Concurrent Import Set table instead of the standard Import Set table.
- Each concurrent job creates its own import set and log entry. The structure of the execution records remains the same.

Procedure

1. Navigate to **All > Meraki > Import Schedules**.
2. Select the import schedule that you want to run.
3. From the list of scheduled data imports, select the scheduled data import.
4. If the import needs to be run manually rather than having run automatically on a specified schedule, select **Execute now**.
5. In the Executions related list, verify the import results by selecting the most recent Import Set record created by the execution of the import.
6. Review the import set log to verify the numbers of rows read, rows inserted or updated, and transformation success status in the CMDB.
If you delete existing CIs and rerun the import, the log might show multiple inserts and updates, validating end-to-end functionality.

Result

If configured correctly, the connector brings in the network inventory data from Cisco Meraki into your ServiceNow AI Platform[®] instance.

Run Arista VeloCloud SD-WAN Import

Manually run a configured import schedule for the Arista VeloCloud Service Graph Connector (SGC) to verify that data from Arista VeloCloud was successfully imported or updated in the Configuration Management Database (CMDB). You can also perform the import to validate the connector setup, run one-off imports, or test newly configured connection aliases.

Before you begin

Role required: TSOM Visibility admin

If you have configured multiple Arista VeloCloud connection aliases for different environments or filters, confirm that the import schedule is set up correctly for each alias.

About this task

Note:

- If Concurrent Import option is enabled in the import schedule, records are imported into the Concurrent Import Set table instead of the standard Import Set table.
- Each concurrent job creates its own import set and log entry. The structure of the execution records remains the same.

Procedure

1. Navigate to **All > VeloCloud > Import Schedules**.
2. Select the import schedule that you want to run.
3. From the list of scheduled data imports, select the scheduled data import.
4. If the import needs to be run manually rather than having run automatically on a specified schedule, select **Execute now**.
5. In the Executions related list, verify the import results by selecting the most recent Import Set record created by the execution of the import.
6. Review the import set log to verify the number of rows read, number of rows inserted or updated, and transformation success status.
These values indicate that the CIs exist in the CMDB. If you delete existing CIs and rerun the import, the log might show multiple inserts and updates, validating end-to-end functionality.

Result

If configured correctly, the connector brings in the network inventory data from Arista VeloCloud into your ServiceNow AI Platform[®] instance.

Validating JSON payloads using TSOM Schema Validator

Use the `TsomSchemaValidator` utility class to validate JSON payloads against TSOM schemas before importing data. This helps identify errors early, reduce ETL failures, and confirm data quality.

Use this validator to check whether your JSON payloads conform to the expected schema for telco objects like Devices, Logical Connections, and Topologies before creating import sets. This pre-validation step helps to prevent schema mismatch errors and improves debugging.

Supported schema types

The validator supports multiple schema types for different telco data structures:

- Logical Composites - Representing groupings of components: Equipment, PDUs, Fan Shelves
- Devices - Equipment and their contained components
- Logical Connections - Connections between network interfaces
- Port Relations - Relationships between network interfaces: physical, logical, lags
- Logical Connection Relations - Relationships between logical connections
- Topologies - Network topology

Class structure

```
let TsomSchemaValidator = Class.create();
TsomSchemaValidator.prototype = {
```

```

initialize: function() {
this.schemas = new TsomGenericSchema();
},
isValidJson: function(payload) {
// Validation logic that determines if the JSON structure is
valid
// Returns boolean (true/false)
},
checkJsonValidation: function(payload) {
// Validation logic that determines if the JSON structure is
valid
// Returns a JSON object containing errors (if exist)
},
type: 'TsomSchemaValidator'
};

```

Steps

1. Instantiate the Schema Validator

```

var TsomSchemaValidator = new
  sn_tsom_core.TsomSchemaValidator();

```

2. Run a Boolean Validation Check

```

if (!TsomSchemaValidator.isValidJson(target_json)) {
  gs.error('Invalid JSON: ' + JSON.stringify(target_json));
  return;
}

```

3. Run a Detailed Validation Check

```

let result =
  TsomSchemaValidator.checkJsonValidation(target_json);
if (!result.valid) {
  gs.error('Invalid JSON: ' + JSON.stringify(result, null,
  2));
  return;
}

```

Example output

Example Output

```

{
  "schemaName": "devices",
  "errors": [
    {
      "message": "Missing required property: model_name",
      "params": { "key": "model_name" },
      "code": 302,
      "dataPath": "/devices/0/ports/0",
      "schemaPath":
"/properties/devices/items/properties/ports/items/required/4"
    }
  ],
  "valid": false
}

```

Related topics

[Configure the Telecom Discovery Builder ETL](#)

[Telecom Discovery Builder framework](#)

Run Telecom Discrepancy audit

The Telecom Discrepancy Audit validates the integrity of configuration items (CIs) and relationships across your telecom inventory using the CMDB Compliance framework.

Before you begin

Role required: admin

About this task

The Telecom Discrepancy Audit is part of the CMDB Compliance framework and supports phased auditing. It detects and addresses discrepancies in CI relationships and attributes discovered through ServiceNow Discovery or Service Graph Connectors (for example, Nokia Altiplano).

You can run the audit manually or set it to run at regular intervals. During manual execution, you can select from existing filters to limit the audit scope to specific CIs—helpful for testing or targeted reconciliation.

The Telecom Discrepancy Audit works in two phases:

1. Initial compliance run: Validates CI relationships and selected CI tables using the following default criteria:
 - Supported CI classes: Slot, Subslot, Card, Interface, Network Gear (includes all extension tables).
 - Source: CI was discovered via Discovery (For example `discovery_source = SG-TSOM-Altiplano`).
 - Relationship type: `Contains::Contained By` (customizable via `sn_tsom_core.audit.relationship_types` property).
2. Subsequent compliance runs: In addition to initial checks, evaluates whether the Updated timestamp of the relationship or related CIs is more recent than the Last run date in the audit record.

Note: Each failed audit creates a Follow-On Task for manual or automated remediation, confirming telecom CMDB data stays aligned with the network state.

The following screenshot helps you understand the Telecom Discrepancy

You are editing a record in the Telecom Service Operations Core application (cancel)

Name: Telecom Discrepancy Audit

Filter: Filter for TSOM Discrepancy Audit - 1

Table: Configuration Item [cmdb_ci]

Application: Telecom Service Operations Core

Active: ☒

Run: On Demand

Last run date: 2025-06-30 04:07:00

Create tasks: ☒

Assignment type: User Field

Assign to: Assigned to

Assign to empty: Create Unassigned Task

Audit type: Scripted

Short description: Telecom CMDb Discrepancy Audit

Task description: Remedy the discrepancies listed in the Audit Results below

Run this script: Turn on ECMAScript 2021 (ES12) mode

```

1  /* global gs, current, SNC, CertificationProcessing, TsomDiscrepancyAudit, GlideDateTime */
2
3  var certProcessor;
4  if (gs.getCurrentScopeName().includes("global")) {
5    certProcessor = new SNC.CertificationProcessing();
6  } else {
7    certProcessor = new CertificationProcessing();
8  }
9
10 //
11 let tsomDiscrepancyAudit = new TsomDiscrepancyAudit();
12 tsomDiscrepancyAudit.printInfoLog("Audit Started.");
13 let auditStartTime = new GlideDateTime();
14 auditStartTime.addSeconds(-1); // change the audit start time to 1 second before the current time, so that
15 // relationships added at the current time will be picked up later.
16 let lastAuditRunDate = current.last_run_date.toString();
17 let auditId = current.sys_id;
18 tsomDiscrepancyAudit.initAudit(auditId, certProcessor);

```

Audit.

Procedure

1. Navigate to **All > Compliance > Audits**

2. Open the Telecom Discrepancy Audit (cert_audit table) record.

3. Optional: Select a filter.

- If multiple audit filters have been defined during configuration, you can choose one before running the audit.
- Filters enable you to restrict the audit scope based on criteria such as discovery source, CI class, or specific CI attributes.
- This step is especially useful when troubleshooting a subset of records or validating specific discovery results.

4. Do either of the following to run the audit.

- Click **Run Audit** to trigger a manual execution.
- Configure the audit to run on a scheduled basis by setting up a recurring job.

The audit uses CMDB Compliance to validate relationship and CI integrity across telecom inventory.

Result

Review audit results

- The audit identifies compliance failures such as missing relationships or misaligned attributes.
- For each failed audit record, a Follow-On Task is automatically generated.
- Tasks include recommended or automated remediation actions such as updating or decommissioning CIs.

Example use case: The configured filters can be used for different discovery sources. For example, the Nokia Altiplano. When running the audit manually, you can select the appropriate filter to validate only those CIs discovered by that specific source—confirming targeted and efficient auditing.

Related topics[Configure filter for audit](#)[Example for Telecom Discrepancy Audit and Remediation](#)**Example for Telecom Discrepancy Audit and Remediation**

The following example illustrates how the audit works in a scenario.

Before you begin

Role required: admin

About this task

A card (Card04) was initially discovered in Slot04. Later, Card04 was replaced by Card05 in the physical network, but the CMDB still shows Card04 in Slot04. When a new discovery run executes, Card05 is discovered and added to the same slot, creating a data conflict in CMDB.

Audit Behavior:

- The Telecom Discrepancy Audit detects this inconsistency and creates a failed audit record (e.g., AUDR0001283).
- A Follow-On Task is created automatically (e.g., TASK0020215) with a detailed description of the discrepancy.

Task Description Example (Incorrect Number of Relationships):

Card04 was last discovered more than 2.5 days ago.

Relationships between the following CIs

CI	Model
Slot04	DEMO 20532Tree
Card04	Nokia 7360 FANT-F CARD MODULE
Card05	<model not identified>

Procedure

1. Navigate **All > Compliance > Audits > Telecom Discrepancy Audit > Run Audits**.

A Follow-on Task is automatically created for each failed audit record (for example, TASK0020215).

2. Select **TASK0020215** and review the discrepancy description.

Note: This is an example of the TASK0020215 description created for the "Incorrect number of relationships" scenario. Other scenarios and environments might have different descriptions.

The Follow-On Task contains a detailed description of the discrepancy. As you can see in the description, the Card04 CI is in discrepancy.

3. To remediate, perform the following steps:

- a. Navigate **All > System Definition > UI Actions**.

- b. Select **Remediate**.

The Remediate button is a UI action that triggers the **TSOM CI Decommission** subflow. This subflow:

- Retires the outdated Card04 CI.
- Removes the incorrect Slot04 → Card04 relationship.
- Synchronizes CMDB records with the network state.

Result

After remediation:

- The Follow-On Task (TASK0020215) updates with work notes summarizing the resolution steps.
- The Card04 CI is marked as retired, and incorrect relationships are removed.
- CMDB is now aligned with the latest discovered state from the network.

Note: Customize Remediation - The Telecom Discrepancy Audit ships with example remediation subflows. You can create and attach custom subflows using Flow Designer to suit your operational requirements.

Related topics

[Telecom Discrepancy Identification and Reconciliation](#)

[Discrepancy identification – types of discrepancies](#)

[Telecom discrepancy identification and reconciliation](#)

[Activate Telecom Discrepancy and Reconciliation](#)

[Defining UI actions](#) 

Generate reports for attribute value discrepancies

Use CMDB 360 to generate reports that highlight discrepancies in attribute values between different discovery sources or between a discovery source and the CMDB baseline.

Before you begin

Role required: admin

About this task

Attribute value discrepancy reports help identify conflicting data from multiple sources updating the same configuration item (CI). This enables better data quality and integrity across your CMDB.

Procedure

1. Navigate to **All > CMDB Workspace > CMDB 360**.
2. Select **Create Query**.
3. In the query type window, select **Compare attribute values**.
4. Define the query parameters:
 - a. Select **CI Class** you want to analyze (for example, `cmdb_ci_computer`).
 - b. Apply filters to narrow the scope (for example, Discovery source contains "TSOM").
 - c. Select attributes to compare values across different discovery sources.
 - d. Choose the sources to compare (for example, Discovery, SCCM).

e. **Optional:** Select **Compare to CMDB** to compare discovery source data to the current Baseline CMDB.

f. Select **Save**, provide a name for the query, and then select **Run** to generate the results.

5. **Optional:** Select **Schedule** to set up the query to run at defined intervals.

Result

The report displays a comparison of attribute values across sources:

- Source 1: Current CMDB record.
- Source 1 Value: Value updated by the most recent discovery source.
- Source 2: Previous discovery source.
- Source 2 Value: Value provided by the earlier discovery source.

Each discrepancy is listed as a separate record when a CI has been updated by multiple sources.

Related topics

[Configure attribute value discrepancy in CMDB 360](#)

[Control CI attribute updates using Reconciliation rules](#)

Telecommunications Service Operations Management reference

Several types of components are installed with Telecommunications Service Operations Management applications and plugins.

Arista VeloCloud installed integrations

Predefined system integrations use VeloCloud REST APIs to pull events into your ServiceNow instance to monitor your Arista VeloCloud devices.

VeloCloud installed integrations

VeloCloud integration	Description
VeloCloud Events (sn_tsom_em_conns)	Pulls event data from VeloCloud into your ServiceNow instance for monitoring and alerting.

Fortinet installed integrations

Predefined system integrations use Fortinet REST APIs to pull metric data into your ServiceNow instance to monitor your Fortinet devices.

Fortinet installed integrations

Fortinet integration	Description
Fortinet Performance Status	Monitors resource utilization and system health for FortiGate devices, including CPU and memory usage.

Fortinet installed integrations (continued)

Fortinet integration	Description
Fortinet SD-WAN SLA Log	Collects SD-WAN link performance data, including latency, jitter, and packet loss against configured SLA thresholds.
Fortinet SD-WAN Interface Log	Monitors interface transmit and receive bandwidth and bytes metrics.

Related topics

[Metric type IDs](#)

Fortinet extension points reference

The Service Graph Connector for Fortinet exposes nine scripted extension points that let you customize specific parts of the connector's behavior without modifying the shipped application. Use this reference to find which extension point controls the behavior you want to change.

Each extension point defines a contract — a set of functions — at a specific point in the connector's data-collection pipeline. The connector ships a default implementation for every extension point and picks up any active custom implementation automatically on its next scheduled run. For the steps to create a custom implementation, see [Customize scripted extension points](#).

Available extension points

Extension point	Lets you customize	Enabled by default
ActionRequestBody	FortiManager request payloads	Yes
FortinetCollectionPlan	Which APIs are called, and how responses are processed	Yes
FortinetFieldMappings	Extra CI attributes populated from raw API fields	Yes
FortinetCustomAttributes	Extra additional_attributes computed in script, on devices, organizations, network sites, network service instances, and ports	Yes
FortinetParserHooks	Custom enrichment logic during data extraction	Yes
FortinetCustomizedFirmwareVersion	How a device's firmware version string is built	Yes
FortinetCustomizedContractData	How a contract data is parsed	Yes
FortinetCustomizedAdom	Which ADOMs are collected at all	No
FortinetCustomizedLifecycleStage	The Lifecycle Stage Status value set on new CIs, per CI class	Yes

How the platform selects an implementation

At each extension seam, the connector retrieves every active implementation registered for that extension point, sorted by the extension instance's **Order** field in ascending order. The first implementation, in that order, that provides the required function is used.

Every default implementation shipped with the connector registers its extension instance with an **Order** of 100. Give your custom implementation a lower **Order** value (for example, 20) so it's evaluated first. If no active implementation is found, or a custom implementation throws an error, the connector's default behavior applies instead.

Request payloads (`ActionRequestBody`)

The connector talks to FortiManager over a single shared JSON-RPC endpoint. This extension point lets you customize the request body sent for any of the connector's four API calls (get ADOMs, get device contracts, get devices per ADOM, get interfaces per device). For example, you can add or remove fields, change query parameters, or point at a different FortiManager URL pattern.

Default implementation: `ActionRequestBodyFortinetDefault`, matched to the `fortinet_alias` connection alias.

Note: The resolver matches by exact connection alias name. Your implementation's `getConnectionAlias()` function must exactly match the name field on the `sys_alias` record configured on the instance. If nothing matches, the connector has no request bodies and every FortiManager call fails.

Collection steps (`FortinetCollectionPlan`)

Lets you add a FortiManager API call to the collection run — for example, a firmware-upgrade-status endpoint that doesn't exist in the out-of-box connector. You can also control the dependency order of collection steps and how each response is turned into stored data.

Default implementation: `FortinetCollectionPlan`, which declares the `contracts`, `adoms`, and `devices` steps.

Note: Your implementation must return all three existing steps plus any new one you add. Omitting `contracts`, `adoms`, or `devices` removes that built-in collection entirely; the connector doesn't merge your implementation with the default plan.

CI attributes (`FortinetFieldMappings`)

Lets you expose additional fields from the FortiManager response as CI attributes on devices, ports, organizations, network sites, and network service instances, without writing custom enrichment logic.

Default implementation: `FortinetFieldMappings`.

Note: Add fields only to the `additional_attributes` list. Don't change core identity fields such as `key`, `name`, `serial_number`, `ip_address`, `class`, or `company` — doing so can break how the connector recognizes the same CI across runs.

Additional attributes (`FortinetCustomAttributes`)

Lets you compute extra `additional_attributes` in script for devices, organizations, network sites, network service instances, and ports, for logic too involved to express as a `FortinetFieldMappings` entry. Your custom attributes are merged with the attributes the connector already computed for that entity, and win on any key conflict.

Default implementation: `FortinetDefaultCustomAttributes`, which returns no additional attributes.

Note: All five methods — `device`, `organization`, `network_site`, `network_service_instance`, and `port` — are required. If your implementation is missing any of them, the connector skips it entirely and logs a warning naming the missing methods. Return an empty object, `{ }`, from a method you don't need to customize.

Enrichment hooks (FortinetParserHooks)

Provides hooks into the extraction pipeline that run as each raw API item is turned into a CI or import-set payload. Use these hooks for customizations that a field mapping alone can't express, such as pulling in data collected earlier in the same run. Combined with a custom collection plan step, you can use this extension point to support an entirely new CI type.

Default implementation: `FortinetParserHooks`, which fetches each device's ports and license data before mapping, and creates network sites and service instances after.

Required handlers

Handler	Runs
<code>unwrapResponse</code>	Once per fetched API response, before any item is mapped. Normalizes the response into the array to iterate over.
<code>beforeExtract</code>	Once per raw item, before field mapping runs. The default implementation uses this to attach a device's ports and license data.
<code>afterExtract</code>	Once per item, after field mapping runs. The default implementation uses this to create network sites and service instances.
<code>postProcess</code>	Once per response, with the full array of items already mapped for that type.
<code>shouldStore</code>	Once per mapped item. Return <code>false</code> to exclude that item from the shared store entirely.

Note: All five required handlers must be present in your implementation, or the connector rejects it entirely, logs an error naming the missing handlers, and falls back to the default implementation. Start your implementation from a copy of the default's device-handling logic rather than writing it from scratch — dropping that logic silently removes port attachment and site or service-instance creation for every device.

Three further handlers are optional, and only needed if a custom collection plan step introduces a new entity type. `getEntityTypes` registers the new type's key in the shared store. `getWriteKeyMap` controls the top-level key it's written under. `getDerivedWriteTypes` marks it as written automatically after collection finishes, for a type built entirely in `beforeExtract/afterExtract` rather than fetched from its own API call.

Firmware version string (FortinetCustomizedFirmwareVersion)

Lets you change how the firmware version string stored on a device CI is built, for customers who display firmware versions differently than FortiManager's raw `os_ver`, `mr`, and `patch` fields.

Default implementation: `FortinetDefaultFirmwareVersion`, which joins the three fields with periods (for example, `7.4.11`).

License contract parsing (FortinetCustomizedContractParsing)

Lets you change how a license contract item string is turned into a named license-expiration attribute. FortiManager returns license contract data as an opaque string (for example, COMP - 1 - 20 - 20290503 : 0 : 1 : 1 : 0); customers with a different license format need a different parser.

Default implementation: `FortinetDefaultContractParsing`, which extracts the expiration date and service code from the standard Fortinet format.

For the steps to customize this extension point, see [Customize scripted extension points](#).

ADOM filtering (FortinetCustomizedAdomFilter)

Lets you exclude entire ADOMs from collection based on business criteria beyond the standard IP, name, device model, or organization filters. For example, you can collect only ADOMs flagged for a specific service.

Default implementation: `FortinetDefaultAdomFilter`. This implementation is included but disabled by default; every ADOM is collected.

Note: This extension point fails closed, not open. If no implementation is active, all ADOMs pass through. But once an implementation is active and it throws an error, the connector collects zero ADOMs rather than falling back to collecting everything. Test any custom filter carefully before activating it.

Life Cycle Stage Status (FortinetCustomizedLifeCycleStageStatus)

Lets you set a CI's **Life Cycle Stage Status** value differently per CI class — for example, a router might default to `In Use` while a decommissioned-looking interface should read `In Maintenance`.

Default implementation: returns `In Use` for every CI class.

Note: Unlike the other seven extension points, this one doesn't run during data collection. A business rule on `cmdb_ci` calls it when a CI record is created or updated, after the import set has already been transformed into a CI.

Related topics

[Customize scripted extension points](#)

[Configure a Fortinet SD-WAN Service Graph Connector](#)

Cisco Meraki installed integrations

Predefined system integrations use Meraki REST APIs to pull metric data into your ServiceNow instance to monitor your Cisco Meraki devices. Each integration is driven by a connector instance whose metric parameter determines the data that the instance collects.

Cisco Meraki installed integrations

Cisco Meraki integration	Description	Metric parameter value
Device Performance	Collects performance metrics from devices, including utilization, throughput, and availability data.	<code>getDevicePerformance</code>

Cisco Meraki installed integrations (continued)

Cisco Meraki integration	Description	Metric parameter value
Uplink Statuses	Monitors the connectivity status of WAN uplinks on MX and Z series devices.	getUplinkStatuses
Configuration Changes	Tracks configuration changes made to devices and networks in the dashboard.	getConfigurationChanges
Network Shaping Uplink Bandwidth	Collects configured bandwidth limits for WAN uplinks used in traffic shaping.	getNetworkShapingUplinkBandwidth
Switch Port Statuses by Switch	Monitors port status and connectivity for MS switches.	getSwitchPortsStatusesBySwitch
Uplink Usage by Network	Collects uplink bandwidth usage data aggregated by network for MX and Z series devices.	getUplinksUsageByNetwork
Uplinks Loss and Latency	Monitors packet loss and latency for WAN uplinks on MX appliances.	getUplinksLossAndLatency
VPN Stats	Monitors site-to-site VPN tunnel status for MX appliances.	getVpnStats

Related topics

[Metric type IDs](#)

System components installed with Telecommunications API notifications

Administrators can assign user roles to grant access to the API notification database tables. The following standard roles for the Topic [sn_api_notif_mgmt_topic] and Topic Subscription [sn_api_notif_mgmt_subscription] tables are included in the ServiceNow system.

Telecommunications API notification roles

Role	Description
sn_api_notif_mgmt.topic_subscription_viewer	Role that enables with read access to the Topic and Topic Subscription tables.
sn_api_notif_mgmt.topic_creator	Role that enables with create, read, and edit access to the Topic table.
sn_api_notif_mgmt.subscription_creator	Role that enables with create and read access to the Topic Subscription table.
sn_api_notif_mgmt.subscription_admin	Role that enables with the following permissions: <ul style="list-style-type: none"> Create and read access to the Topic and Topic Subscription tables. Change the status of registration to deregister a topic subscription.

Related topics

[External event management via Telecommunications API notifications](#)

[Configuring Telecommunications API notifications](#)

System components installed with Nokia Altiplano

System properties control how the connector operates, including discovery options and performance settings.

Properties installed with Nokia Altiplano

Property	Description
<code>sn_sgc_altiplano.enable_onu_discovery</code>	<p>Enable or disable discovery of ONU devices and logical connections between OLT and ONU.</p> <ul style="list-style-type: none"> • Default value: True • Location: All>Service Graph Connectors>Nokia Altiplano>Properties or System Properties [sys_properties] table filtering by the name <code>'*altiplano'</code>
<code>sn_sgc_altiplano.devices_list_batch_size</code>	<p>Controls the batch size for Altiplano REST API calls.</p> <ul style="list-style-type: none"> • Default value: 1000 • Location: All>Service Graph Connectors>Nokia Altiplano>Properties or System Properties [sys_properties] table filtering by the name <code>'*altiplano'</code>

Properties installed with Nokia Altiplano (continued)

Property	Description
<code>sn_sgc_altiplano.parallel_number_of_data_source_jobs</code>	<p>Number of parallel jobs for collecting Altiplano data (requires "Enable parallel loading" setting).</p> <ul style="list-style-type: none"> • Default value: 2 • Location: All>Service Graph Connectors>Nokia Altiplano>Properties or System Properties [sys_properties] table filtering by the name <code>'*altiplano'</code>
<code>sn_sgc_altiplano.onu_ci_class</code>	<p>Defines whether ONUs are stored as ONU or ONT CI class.</p> <ul style="list-style-type: none"> • Default value: ONU • Location: All>Service Graph Connectors>Nokia Altiplano>Properties or System Properties [sys_properties] table filtering by the name <code>'*altiplano'</code>

Related topics

[Telecom Discovery via Nokia Altiplano](#)

[Configure Nokia Altiplano service graph connector](#)

Telecom discrepancy identification and reconciliation

The system properties are part of the TSOM Visibility plugin (sn_tsom_core) and control the Telecom Discrepancy Identification & Reconciliation log (TSOM CMDB Audit). The TSOM Visibility plugin enables the TSOM Visibility applications. It contains logic shared across the Telecom Discovery and Telecom Discrepancy Identification & Reconciliation solution.

Telecom Discrepancy Identification & Reconciliation System Properties (Impacts CMDB Audit)

Property Name	Recommended Default Value	Description
sn_tsom_core.audit.log.level	Telecom Discrepancy Audits' Logging.	debug,info,warn Default value: info
sn_tsom_core.audit.relationship_types	Displayed Names of Relationship Types that will be handled by Telecom Discrepancy Audit.	Comma separated Default value: Com
sn_tsom_core.audit.discovered_date.diff.threshold.in.days	Most recent discovered date threshold in days. Used for Telecom Discrepancy Audit.	String Default value: 2.5
sn_tsom_core.audit.suppress_CI_Model_missing_discrepancy_task_creation	Suppress creation of CI Model missing discrepancy task for the tables.Used for Telecom Discrepancy Audit.	cmdb_ci_ni_telco cmdb_ci_containe
sn_tsom_core.audit.equipment_tables	Equipment tables that will be handled by Telecom Discrepancy Audit.	Comma separated Default value: cm
sn_tsom_core.audit.interface_card_tables	Interface Cards tables that will be handled by Telecom Discrepancy Audit.	Comma separated Default value: cm
sn_tsom_core.audit.interface_tables	Network Interface	Comma separated

Telecom Discrepancy Identification & Reconciliation System Properties (Impacts CMDB Audit) (continued)

Property Name	Recommended Default Value	Description
	tables that will be handled by Telecom Discrepancy Audit.	Default value: cm
sn_tsom_core.audit.slot_tables	Slot tables names that will be handled by Telecom Discrepancy Audit.	Comma separated Default value: cm
sn_tsom_core.audit.subslot_tables	Subslot tables names that will be handled by Telecom Discrepancy Audit.	Comma separated Default value is cr

Telecom Discrepancy Identification & Reconciliation Audits

Audit Name	Description
Telecom Discrepancy Audit	<p>Audits records in the <code>cmdb_rel_ci</code> table. For each relationship record discovered by Nokia Altiplano or NSP, it validates the parent and child CIs based on: - Most recent discovered relationship (if TNI Core is installed).</p> <p>Note: All discrepancy audits use a filtering mechanism that applies to the filter condition is: <code>discovery_source CONTAINS TSOM</code>.</p>
Telecom Logical Connections Discrepancy Audit	Audits records in the <code>cmdb_ci_ni_logical_path</code> table. For each logical connection discovered by Nokia Altiplano or NSP, it checks: - Empty endpoint (Port A or Port Z) - Each endpoint only one unique logical connection.
Telecom Network Topology Discrepancy Audit	Audits records in the <code>cmdb_ci_network_topology</code> table. For each network topology discovered by Nokia Altiplano or NSP, it verifies: - At least one Contains:Contained By relationship - At least one Members:Member Of relationship with a logical connection.

Related topics

[Telecom Discrepancy Identification and Reconciliation](#)

[Activate Telecom Discrepancy and Reconciliation](#)

[Run Telecom Discrepancy audit](#)

Fortinet Service Graph Connector API endpoints

The Service Graph Connector for Fortinet integrates Fortinet Dashboard API data into ServiceNow AI Platform Configuration Management Database (CMDB). This document details the API endpoints used and how data flows through the system.

Note: The `meta_fields` values in the ADOM and device requests are configurable. Define the meta field names that match your Fortinet deployment. If you don't configure any meta fields, the connector omits the `meta_fields` parameter from the request, and Fortinet returns standard data. The values shown in the following examples are sample meta fields and are required for discovery.

API reference examples

Description	API response
Authenticate JSON-RPC session	<pre>{ "method": "exec", "params": [{ "url": "/sys/login/user", "data": { "user": "user", "passwd": "abc1123" } }], "verbose": 1 }</pre>
Logout JSON-RPC session	<pre>{ "method": "exec", "params": [{ "url": "/sys/logout" }], "verbose": 1, "session": "TxMw/zbjw+tVZ/JL3bhmYg0vTUpVguYHIQesJXpye4j2vvsqtZaSgKwA+0iLqug+3/074jq8Qqml/KOx4GHkwQ==" }</pre>
GET ADOMs	<pre>{ "method": "get", "params": [{ "url": "/dvmdb/adom", "meta_fields": ["VF SDWAN Service Provided", "VF OpCo Name", "VF Customer ID", "SERVICE_ID", "SDN_MS"] }], "verbose": 1, "session": "TxMw/zbjw+tVZ/JL3bhmYg0vTUpVguYHIQesJXpye4j2vvsqtZaSgKwA+0iLqug+3/074jq8Qqml/KOx4GHkwQ==" }</pre>
GET ADOM devices	<pre>{ "method": "get", "params": [{ "url": "/dvmdb/adom/{ADOM_NAME}/device", "meta_fields": ["Company/Organization", "VF 3C Ref", "Address", "Contact Email", "Contact Phone Number", "VF Order Ref"] }], "verbose": 1, "session": "TxMw/zbjw+tVZ/JL3bhmYg0vTUpVguYHIQesJXpye4j2vvsqtZaSgKwA+0iLqug+3/074jq8Qqml/KOx4GHkwQ==" }</pre>
GET interfaces by device name	<pre>{ "method": "get", "params": [{ "url": "/dbcache/system/interface"</pre>

API reference examples (continued)

Description	API response
	<pre>, "option": ["scope member"], "scope member": [{ "name": "{DEVICE_NAME}", "vdo m": "global" }], "fields": ["name", "alias", "comments", "spe ed", "mediatype", "ip", "mode", "type" , "mtu", "interface", "vlanid", "member ", "vdom", "role", "allowaccess", "zon e", "mac", "status"], "current_adom": "{ADOM_NAME}" }], "verbose": 1, "session": "TxMw/zbjw+tVZ/JL3bhmYg0 vTUpVguYHIQesJXpye4j2vvsqtZaSgKWa+0i Lqug+3/074jq8Qqml/KOx4GHkwQ==" }</pre>

Cisco Meraki Service Graph Connector API endpoints

The Service Graph Connector for Meraki integrates Cisco Meraki Dashboard API data into ServiceNow AI Platform® Configuration Management Database (CMDB). This document details the API endpoints used and how data flows through the system.

Organizations endpoint

Description	API response
Organizations API response URL: /organizations	<pre>{ "id": "123456", "name": "My Organization", "management": { "details": { "MSP ID": "...", "customer number": "...", "IP restriction mode for API": "..." } } }</pre>

Organizations endpoint (continued)

Description	API response
	<pre>{ }</pre>

Licenses overview endpoint

Description	API response
Licenses overview API responseURL : / organizations/{orgId}/licenses/overview	Called for organizations with a CO-TERM licensing model (org.licensing.model === "co-term") to collect the license expiration date, which is not available on a per-device basis for these organizations.

Networks Endpoint

Description	Endpoint
Networks API response URL : /organizations/{orgId}/networks	<pre>{ "id": "N_123456", "organizationId": "123456", "name": "My Network", "notes": "Network description" }</pre>

Devices Endpoint

Description	Endpoint
Devices API response URL : /organizations/{orgId}/devices	<pre>{ "serial": "Q2XX-XXXX-XXXX", "name": "My Device", }</pre>

Devices Endpoint (continued)

Description	Endpoint
	<pre>{ "networkId": "N_123456", "model": "MX68", "mac": "00:11:22:33:44:55", "productType": "appliance", "firmware": "mx-18.1", "lat": 37.7749, "lng": -122.4194, "address": "123 Main St", "notes": "Device description" }</pre>

Device statuses endpoint

Description	Endpoint
Devices statuses API response URL: /organizations/{orgId}/devices/statuses	<pre>{ "serial": "Q2XX-XXXX-XXXX", "status": "online", "publicIp": "203.0.113.1", "wan1Ip": "192.168.1.1", "wan2Ip": "192.168.2.1", }</pre>

Uplink statuses endpoint (continued)

Description	Endpoint
	<pre> "status": "active", "ip": "192.168.1.1", "gateway": "192.168.1.254", "publicIp": "203.0.113.1", "primaryDns": "8.8.8.8", "secondaryDns": "8.8.4.4", "ipAssignedBy": "dhcp" }, { "interface": "cellular", "status": "ready", "apn": "broadband", "iccid": "...", "imsi": "...", "msisdn": "..." }] }</pre>

Switch ports endpoint

Description	Endpoint
<div>Switch ports endpoint</div> <div>URL: /organizations/{orgId}/switch/ports/statuses/bySwitch</div>	<pre>{ "items": [{ "serial": "Q2XX-XXXX-XXXX", "network": { "id": "N_123456" }, "ports": [{ "portId": "1", "status": "connected" }, { "portId": "2", "status": "disconnected" }] }] }</pre>

Switch ports endpoint (continued)

Description	Endpoint
	<pre>[{ }]</pre>

Device inventory endpoint

Description	Endpoint
Device inventory endpoint URL: /organizations/{orgId}/inventory/devices	<pre>{ "serial": "Q2XX-XXXX-XXXX", "orderNumber": "ORD-12345", "licenseExpirationDate": "2025-12-31" }</pre>

Cisco Meraki extension points

Each Meraki extension point controls one part of the connector's data collection. Use these tables to find the handler that governs the behavior you want to change and the arguments it receives.

Extension points

The Meraki connector exposes five scripted extension points, each registered in the connector's own `sn_sgc_meraki` scope: `MerakiCollectionPlan`, `MerakiParserHooks`, `MerakiFieldMappings`, `TagTransformationExtensionPoint`, and `MerakiCustomizedLifeCycleStageStatus`. A default implementation of each ships with the connector and is registered out of the box.

MerakiCollectionPlan

Controls which API calls the connector makes, in what order, and how each response is routed. Before any data is fetched, the connector reads the active implementation's collection plan and sorts its steps by their `dependsOn` declarations into execution levels. A circular dependency stops collection.

MerakiCollectionPlan handlers

Handler	Required	Controls
<code>getCollectionPlan(constants)</code>	Yes	Returns the collection plan object, keyed by step name. Each step's <code>url</code> function returns the endpoint path.

MerakiCollectionPlan handlers (continued)

Handler	Required	Controls
		<i>fetch.scope</i> and <i>fetch.strategy</i> control whether the step is fetched once or once per parent entity. <i>dependsOn</i> lists step names that must finish first. <i>write</i> is deferred or immediate with <i>writeKeys</i> . A step that defines <i>onResponse</i> handles its own response directly, bypassing <i>MerakiParserHooks</i> and <i>MerakiFieldMappings</i> for that step.
<i>getConstants()</i>	No	Returns additional named constants available to step definitions, merged over the connector's built-in constants.

⚠ Warning:

An implementation must return the `organization`, `networks`, and `devices` steps in addition to any step you add. Omitting one of these from a custom plan silently drops that collection — the connector doesn't merge your implementation with the default plan.

MerakiParserHooks

Controls how a raw API response is unwrapped, enriched, filtered, and stored. An implementation can also register additional entity types for a collection plan step that introduces a brand-new type. In every handler, `type` is the step name from the active collection plan.

Required MerakiParserHooks handlers

Handler	Arguments	Controls
<i>unwrapResponse</i>	<code>responseJson</code> , <code>type</code> , <code>parser</code>	Once per fetched response, before any item is mapped. Returns the value to iterate over, usually an array.
<i>beforeMapping</i>	<code>rawItem</code> , <code>type</code> , <code>store</code> , <code>parser</code>	Once per raw item, before field mapping. Mutates <code>rawItem</code> in place so a mapping can reference the added data as a normal source path. Returns nothing.
<i>afterMapping</i>	<code>mappedItem</code> , <code>type</code> , <code>store</code> , <code>rawItem</code> , <code>parser</code>	Once per item, after field mapping. Mutates <code>mappedItem</code> in place to add attributes or relationships that depend on the mapped shape. Returns nothing.
<i>postProcess</i>	<code>results</code> , <code>type</code> , <code>parser</code>	Once per response, with every item already mapped for that type. Returns the items to store or write, optionally filtered or transformed.
<i>shouldStore</i>	<code>mappedItem</code> , <code>type</code> , <code>parser</code>	Per item. Returns <i>false</i> to exclude the item from the shared store.

Note:

An implementation missing any of the five required handlers is rejected in full, not partially merged. The connector logs an error naming the missing handlers and falls back to the default implementation.

beforeMapping and *afterMapping* were previously named *beforeExtract* and *afterExtract*. Rename the handlers in a custom implementation to the current names.

Optional MerakiParserHooks handlers

Handler	Controls
<i>getEntityTypes(constants)</i>	Registers extra entity type keys in the shared store. Required if a custom MerakiCollectionPlan step introduces a brand-new type.
<i>getWriteKeyMap(constants)</i>	Returns extra { entityType: outputKey } entries merged over the out-of-box write key map. Use this when a deferred-write entity type should be emitted under a different top-level key than its store key.
<i>getDerivedWriteTypes()</i>	Returns extra entity type keys whose entities are built during parsing, such as inside <i>afterMapping</i> , rather than fetched from their own collection plan step. The connector writes these types automatically after deferred writes complete.

MerakiFieldMappings

Controls how raw Meraki API data is mapped into CI fields. The connector uses the first registered implementation whose *getFieldMappings* handler returns every required resource-type key. An implementation that's missing a required key is skipped, an error naming the missing keys is logged, and the connector tries the next registered implementation. If no implementation is valid, the connector falls back to its default field mappings.

MerakiFieldMappings handlers

Handler	Required	Controls
<i>getFieldMappings</i>	Yes	Returns the field mappings object, keyed by resource type. The returned object must include an entry for every required resource type: networks, organization, devices, device_statuses, networkSites, uplink_status, switch_ports, and device_inventory. Add or change fields only inside a resource type's additional_attributes block — changing a core CI attribute such as <i>key</i> , <i>name</i> , <i>serial_number</i> , or <i>company</i> can break CI identity resolution and deduplication.
<i>getTransforms</i>	No	Returns an object mapping custom transform names to functions, for use with the transform property in a field mapping entry. Each transform function runs with the active parser instance as <i>this</i> , giving access to <i>this.store</i> and <i>this.log</i> . Return null if you have no custom transforms. Custom transform names can't reuse a built-in transform name.

Field mapping entry properties

Property	Description
<i>source</i>	Dot-path into the raw API object, for example <code>deviceStatus.ip_address</code> . Pass an array of paths to forward multiple values to a <i>transform</i> .
<i>value</i>	A hardcoded value, used when no source field exists.
<i>transform</i>	Name of a transform function, built-in or from <i>getTransforms</i> , that post-processes the resolved source value.
<i>defaultValue</i>	Fallback value used when the source resolves to <code>null</code> or <code>undefined</code> .
<i>isArray</i>	Marks the field as an array placeholder with no source mapping.
<i>expandArray</i>	Flattens an array value into individual entries.
<i>mapArray</i>	A nested mapping applied to each element of a source array.
<i>\$parent</i>	In a <i>mapArray</i> context, a prefix used to access the parent record's fields.

Note:

The following transform names are built in and can't be overridden from *getTransforms*: *lookupOrgName*, *lookupNetworkCompany*, *transformTags*, *defaultIfNull*, *defaultIfNullOrEmpty*, *roundNumberToTwoDecimalPlaces*, *getNetworkSiteKey*, *getSiteName*, *getSiteAddress*, *getOrgIdFromNetwork*, *getInterfaceKey*, *getInterfaceModel*, *getInterfacePosition*, *getInterfaceBandwidth*, *getInterfaceBandwidthName*, *getInterfaceBandwidthGroup*, *getPortName*, and *getPortPosition*.

TagTransformationExtensionPoint

Controls how an array of Meraki tag strings is converted into key/value pairs for CMDB attribute mapping. The connector uses the first registered implementation whose *handles* handler returns *true* for the Meraki source identifier.

TagTransformationExtensionPoint handlers

Handler	Arguments	Controls
<i>handles</i>	<code>sourceSystem</code>	Whether this implementation applies to the given source system. The connector calls this handler with the source identifier <code>meraki - SDWAN</code> ; return <i>true</i> to have this implementation used.
<i>transform</i>	<code>tags</code>	Converts an array of tag strings into an array of { <code>key</code> , <code>value</code> } pairs used for CMDB attribute mapping. Only called on the implementation selected by <i>handles</i> .

MerakiCustomizedLifeCycleStageStatus

Controls the life cycle stage status value assigned to a CI, per CI class. The connector uses the first registered implementation that defines the *getLifeCycleStageStatus* handler, and uses its returned value directly. This handler only runs when a CI doesn't already have a life cycle stage status set.

MerakiCustomizedLifeCycleStageStatus handlers

Handler	Arguments	Controls
<code>getLifeCycleStageStatus</code>	<code>entityClass</code>	Returns the life cycle stage status string for a given CI class, such as <code>cmdb_ci_ip_router</code> . The default implementation returns <code>In Use</code> for every CI class.

Arista VeloCloud Service Graph Connector API endpoints

The Service Graph Connector for Arista VeloCloud integrates VeloCloud Orchestrator API data into ServiceNow AI Platform Configuration Management Database (CMDB). This document details the API endpoints used and how data flows through the system.

Arista VeloCloud Orchestrator endpoints called during discovery

Description	API response
<p>MSP details</p> <p>URL: <code>/portal/rest/enterpriseProxy/getEnterpriseProxy</code></p>	<pre>{ "id": 1, "created": "2020-02-25T14:13:13.000Z", "networkId": 1, "proxyType": "MSP", "operateGateways": 1, "logicalId": "084265a7-96dd-4dd4-9e9f-d8e924850663", "name": "Vodafone Group POC", "domain": "VFGroupPoc", "contactName": "Samsul Arefin", "contactPhone": "+1", "contactMobile": "+44", "contactEmail": "samsul.arefin@vodafone.com", "streetAddress": "", "streetAddress2": "The Connection", "city": "Newbury", "state": "Berkshire", "postalCode": "RG14 2FN", "country": "United Kingdom", "lat": 37.402866, "lon": -122.117332, "modified": "2024-11-21T14:32:30.000Z" }</pre>
<p>Gateway pools and details</p> <p>URL: <code>/portal/rest/enterpriseProxy/getEnterpriseProxyGatewayPools</code></p>	<pre>[{ "id": 2, "networkId": 1, "logicalId": "4f19ad54-06b7-11ec-b96c-026628323786",</pre>

Arista VeloCloud Orchestrator endpoints called during discovery (continued)

Description	API response
	<pre> "name": "Vodafone VCC Pool 1", "gateways": [], "enterprises": [] }] </pre>
<p>Enterprises</p> <p>URL: /portal/rest/enterpriseProxy/getEnterpriseProxyEnterprises</p>	<pre> [{ "id": 1, "created": "2020-03-05T14:56:45.000Z", "networkId": 1, "gatewayPoolId": 2, "alertsEnabled": 1, "operatorAlertsEnabled": 1, "endpointPkiMode": "CERTIFICATE_DISABLED", "name": "UseCase1", "domain": "domain-test", "logicalId": "135c30c3-90a7-499c-b778-942794c205e5", "accountNumber": "WAT-4N6RVGN", "modified": "2025-05-29T14:51:03.000Z" }] </pre>
<p>Network enterprises (Operator mode)</p> <p>URL: /portal/rest/network/getNetworkEnterprises</p>	<pre> [{ "id": 43, "created": "2019-03-11T13:38:01.000Z", "networkId": 1, "gatewayPoolId": 3, "alertsEnabled": 1, "operatorAlertsEnabled": 1, "endpointPkiMode": "CERTIFICATE_REQUIRED", "name": "DAWN WING/TIME FREIGHT", "domain": null, "prefix": null, "logicalId": "550b0770-23ce-4884-889c-fadb21d72acb", </pre>

Arista VeloCloud Orchestrator endpoints called during discovery (continued)

Description	API response
	<pre> "accountNumber": "DAW-WIN-5EQ", "description": null, "contactName": null, "contactPhone": null, "contactMobile": null, "contactEmail": null, "streetAddress": null, "streetAddress2": null, "city": null, "state": null, "postalCode": null, "country": null, "lat": 37.402866, "lon": -122.117332, "timezone": "America/Los_Angeles", "locale": "en-US", "bastionState": "UNCONFIGURED", "modified": "2022-08-18T18:14:09.000 Z", "enterpriseProxyId": 11, "enterpriseProxyName": "DAWN WING" }] </pre>
<p>Edges</p> <p>URL: /portal/rest/enterprise/getEnterpriseEdges</p>	<pre> [{ " id": 4523, "enterpriseId": 1552, "logicalId": "2b16c0a6-1d22-4a2f-861d- 3994e43f7982", "name": "610-05-08", "serialNumber": "2937V43", "modelNumber": "edge610-lte", "edgeState": "OFFLINE" }] </pre>

MPN Formulas table

Reference for the MPN Formulas [sn_tsom_em_conns_kpi_definitions] table and the automatic processing that occurs when a record is inserted or updated.

MPN Formulas table columns

MPN Formulas table columns

Column	Description
Release	Release version
Vendor	Vendor name
Technology	Radio technology type (for example, NR, LTE)
KPI Formula	Raw formula string
Index	Elasticsearch index pattern to query (for example, <code>radio - pm *</code>)
Labels Up Down	Indicates the direction of the KPI metric. The available values are up and down.
Labels Vendor KPI ID	Unique KPI identifier assigned by the vendor
Labels KPI Standard Name	Standardized name for the KPI
Labels KPI Node	Node-level label associated with the KPI. The available values are RAN, Core, or Latency Collector.
Unit	Unit of measurement for the KPI value, for example, %, Mbps
Source Time Interval	Data collection time interval from the source system
Aggregation Type	How the KPI value is aggregated, for example, AVG, SUM
CI Label	Path in the Elasticsearch document used to extract the CI identifier.
CI Class	ServiceNow CI class the KPI is mapped to
Component Aggregation	Aggregation method applied at the component level. The available values are flat, hierarchy, or flat+hierarchy.
MPC Description	Full description of the KPI from the MPC specification
MPC UX Name	Display name for the KPI as shown in the MPC UI
MPC KPI Group	Top-level KPI grouping as defined in the MPC
MPC KPI Sub Group	Sub-grouping of the KPI within the MPC group
Formatted KPI Formula	Auto-populated. Don't enter a value manually.

Related topics

[MPN Formula Engine processing](#)

MPN latency KPIs

Reference for the latency KPIs collected by the MPN Pull Connector and the derived UE-to-switch latency metric.

Source KPIs

The MPN Pull Connector collects the following latency KPIs from the `e2e - latency - collector - {env}` Elasticsearch index, filtered to round-trip-time samples (`phase=rtt`).

Source KPIs

KPI	Description
collector2ue	Latency between the collector and the user equipment (UE).
collector2sw	Latency between the collector and the switch.

Derived KPI

TSOM calculates one additional KPI from the source KPIs. This value is not reported directly by the vendor; it is computed after both source KPIs are collected.

Derived KPI

KPI	Description
ue2sw	User-equipment-to-switch latency, calculated as collector2sw minus collector2ue. If a switch-side record is missing for a given DNN in a collection cycle, the calculation is skipped for that cycle.

Related topics

[MPN Formulas table](#)

[MPN Formula Engine processing](#)

[MPN latency metrics](#)

MPN health status rule reference

Default rule set (version 0.0.11) for calculating health status for MPN radio and core server components, including error codes.

Rule structure

Rules are organized by component type. Health status uses a RAG (Red, Amber, Green) indicator system. For radio and core server components, each condition is a simple field-and-value match tied to an error code. A green result requires all conditions in the component's green rule to match simultaneously. If no rule matches, the status is unknown.

Default rules for radio components

Default health rules for radio components

Health status	Field	Value	Error code
Red	alarm_status	MAJOR	EC2002
Red	alarm_status	CRITICAL	EC2003
Red	operational_status	DISABLED	EC2004
Red	operational_status	NOT_AVAILABLE	EC2001
Red	cell_status	INACTIVE	EC2005
Red	cell_status	NOT_AVAILABLE	EC2006

Default health rules for radio components (continued)

Health status	Field	Value	Error code
Red	mme_status	DISCONNECTED	EC2001
Red	mme_status	NOT_AVAILABLE	EC2001
Red	ap_connection_status	DISCONNECTED	EC2001
Red	ap_connection_status	NOT_AVAILABLE	EC2001
Amber	alarm_status	MINOR	EC2101
Amber	alarm_status	INDETERMINATE	EC2102
Amber	alarm_status	WARNING	EC2103
Amber	operational_status	UNKNOWN	EC2104
Amber	cell_status	UNKNOWN	EC2105
Amber	cell_status	PARTIALLY_ACTIVE	EC2106
Amber	mme_status	UNKNOWN	EC2107
Amber	ap_connection_status	UNKNOWN	EC2107
Green	alarm_status AND operational_status AND cell_status	NORMAL, ENABLED, ACTIVE (all three required)	—
Unknown	(no rule matched)	—	EC2108

Default rules for core server components

Default health rules for core server components

Health status	Field	Value	Error code
Red	edge_connection	EDGE_CONN_STATUS_UNAVAILABLE	EC3001
Red	edge_connection	OFFLINE	EC3002
Green	edge_connection	ONLINE	—
Unknown	(no rule matched)	—	EC3101

Default health rules for network-level aggregation

Default health rules for network-level aggregation

Health status	Condition	Error code
Red	All radio components are red	EC1001
Red	All core server components are red	EC1002
Amber	Not all radio components are green	EC1101
Amber	Not all core server components are green	EC1102
Green	All radio and core server components are green	—

Default health rules for network-level aggregation (continued)

Health status	Condition	Error code
Unknown	(no rule matched)	EC1103

Related topics

[Mobile private network health status](#)

[MPN Formula Engine processing](#)

Equinix connector metrics

Reference for the power metrics collected at the cabinet level and the environmental metrics collected at the zone level by the Equinix pull connector.

Cabinet metrics

The connector collects the following power metrics for each cabinet in the account's facility hierarchy. The cabinet is identified in MetricBase by its Equinix level value, for example SG3 : 03 : 070210 : 0308.

Cabinet metrics

Metric	Description
kw	Real power draw, in kilowatts.
kva	Apparent power draw, in kilovolt-amperes.
amps	Current, in amperes.
powerFactor	The cabinet's power factor.
percentageKva	Current utilization as a percentage of the cabinet's contracted capacity. Values can exceed 100% for an oversubscribed cabinet.
peakKvaLastSevenDays	The highest apparent power reading recorded over the previous seven days.
peakKvaLastSevenDaysPercentage	The seven-day peak apparent power as a percentage of the cabinet's contracted capacity.
primaryKva	Apparent power drawn from the primary feed.
redundantKva	Apparent power drawn from the redundant feed.
cabinetRating	The cabinet's contracted power capacity, in kilovolt-amperes.

Zone metrics

The connector collects the following environmental metrics for each zone in the account's facility hierarchy. The zone is identified in MetricBase by its Equinix zone label, for example SG3 : 3 : 03 : HALL7 : Z1.

Zone metrics

Metric	Description
Temperature	Ambient temperature for the zone, in degrees Celsius.
Humidity	Relative humidity for the zone.

System alerts

The connector also retrieves system alerts for each account and forwards them as events. Alert retrieval is paginated, so accounts with many active alerts are collected across multiple requests rather than truncated at the first page.

Related topics

[Configure the Equinix pull connector](#)

Retrieve data from Nokia Altiplano via REST API

Examples of Retrieving Data from Nokia Altiplano via REST API.

URL format

Versioned URL: POST: `altiplano-indexsearch/latestcompleted-inv/_search`

Example: For OLT

```
{
  "_source": [
    "deviceAVmetadata",
    "inventorymetadata",
    "inventorydata.ietf-hardware:hardware",
    "inventorydata.ietf-hardware:hardware-state",
    "inventorydata.nokia-state:state"
  ],
  "sort": [{"_id": {"order": "asc"}}],
  "from": 0,
  "size": 300
}
```

Example: For ONU

```
{
  "query": {
    "bool": {
      "should": [
        {
          "exists": {
            "field":
              "inventorydata.ietf-interfaces:interfaces-state.interface.bbf-x
              ponvni:v-ani.onu-present-on-this-olt.detected-serial-number"
          }
        }
      ]
    }
  }
}
```

```

    },
    "_source": [

      "inventorydata.ietf-interfaces:interfaces-state.interface.bbf-x
      ponvani:v-ani.onu-present-on-this-olt.detected-serial-number",

      "inventorydata.bbf-fiber-onu-emulated-mount:onus.onu.root.ie
      tf-hardware-mounted:hardware-state",

      "inventorydata.bbf-fiber-onu-emulated-mount:onus.onu.name"
    ],
    "sort": [{"_id": {"order": "asc"}}],
    "from": 0,
    "size": 3
  }

```

Page limits for Meraki pull connector endpoints

The Meraki pull connector paginates API requests so that all records are retrieved for large organizations. The page size is set per endpoint to match each endpoint's maximum supported value in the Meraki API v1.

How pagination works

The connector follows the pagination links returned in the `Link` header (`rel=next`) until all pages are consumed, then aggregates the data from every page. The page size is controlled by the `perPage` parameter, which is configured per endpoint as shown in the following table.

Endpoint page sizes

Endpoint	Items per page (<code>perPage</code>)
<code>/organizations</code>	1000
<code>/organizations/{orgId}/networks</code>	1000
<code>/organizations/{orgId}/devices</code>	1000
<code>/organizations/{orgId}/appliance/uplink/statuses</code>	1000
<code>/organizations/{orgId}/appliance/uplinks/usage/byNetwork</code>	20
<code>/organizations/{orgId}/appliance/vpn/stats</code>	300
<code>/organizations/{orgId}/switch/ports/statuses/bySwitch</code>	20
<code>/organizations/{orgId}/configurationChanges</code>	1000

Endpoints without a page limit

Endpoints that do not have a configured `perPage` limit (for example, `/devices/{serial}/appliance/performance` and `/networks/{networkId}/appliance/trafficShaping/uplinkBandwidth`) are treated as single requests. These endpoints return all of their data in one response, so the `Link` header is absent and the pagination loop exits after the first page.

Pull connector granularity constraints

Each connector type enforces minimum and maximum values for the granularity and metrics collection schedule fields on the connector instance form. The connector instance form rejects values outside the supported range.

Granularity and schedule constraints by connector and API type

Connector	API / metric type	Field	Constraint
Meraki	All metrics	Granularity	Granularity must be a positive multiple of 60 seconds and must divide evenly into the metrics collection schedule.
	Uplink loss and latency	Granularity	Maximum 5 minutes
	Device performance	Metrics collection schedule	30 minutes (1800 seconds) is the default collection window. If a lower value is set, an error is shown and the value is reset to the default of 1800 seconds.
Fortinet	Interface logs	Metrics collection schedule	10 minutes (600 seconds) is the default collection window. If a higher value is set, an error is shown and the value is reset to the default of 600 seconds.

Field descriptions

Field descriptions

Field	Description
Granularity	Time resolution of each sample returned per API call. Controls how finely the source system data is sliced within a polling cycle.
Metrics collection schedule	Frequency at which the connector polls the source system API. Also referred to as the poller iteration window.

Metric aggregation configuration reference

The configuration object passed to the metric aggregation scripted extension point defines what to aggregate, how to aggregate it, and where to publish the result. The parameters that apply depend on the aggregation mode.

Configuration parameters

Metric aggregation configuration parameters

Parameter	Type	Applies to mode	Required	Description
mode	String	Both	Required	Aggregation mode. One of <code>hierarchy</code> or <code>flat</code> . The mode determines how source resources are selected and where the calculated metric

Metric aggregation configuration parameters (continued)

Parameter	Type	Applies to mode	Required	Description
				is published. See Metric aggregation modes .
targetCIClass	String	Hierarchy	Required (hierarchy)	CMDB class of the parent or target CI that receives the calculated metric, for example <code>cmdb_ci_baseband_unit</code> . The aggregation traverses CMDB relationships downward from each target CI to collect the matching source CIs.
sourceCIClass	String	Both	Required	CMDB class of the source CIs whose raw metric is aggregated, for example <code>cmdb_ci_mobile_cell</code> .
sourceMetric	String	Both	Required	Name of the raw metric collected on the source CIs that is read as input to the aggregation.
aggregate	String	Both	Required	Aggregate function applied to the source metric values. One of <code>avg</code> , <code>sum</code> , <code>max</code> , <code>min</code> , <code>count</code> , or <code>p95</code> . The calculated value is stored unrounded, as a raw floating-point number.
calculatedMetricName	String	Both	Required	Name of the calculated metric that the aggregation produces and publishes.
unit	String	Both	Optional	Unit label applied to the calculated metric, for example <code>Mbps</code> .
level	Integer	Hierarchy	Optional	Depth limit that bounds how far the CMDB traversal descends from each target CI. Prevents the traversal from following relationships beyond the resources you intend to aggregate. Defaults to 4 and is capped at a maximum of 20; larger values are silently capped. A value of 0 is not treated as "no traversal" — it falls back to the default of 4.
rangeStart	String	Both	Optional	Start of a contiguous range of source resources identified by name, for example <code>SIM010</code> . Use with <code>rangeEnd</code> to select source CIs by range.

Metric aggregation configuration parameters (continued)

Parameter	Type	Applies to mode	Required	Description
rangeEnd	String	Both	Optional	End of the contiguous range of source resources identified by name, for example SIM050.
filters	Array	Both	Optional	Field filters that restrict the source CIs to those meeting a condition. Each element is an object with <code>field</code> , <code>operator</code> , and <code>value</code> . See Filter object .
anchorCiSysId	String	Flat	Required (flat)	The <code>sys_id</code> of the anchor CI that holds the calculated metric. In flat mode the result is not tied to an existing parent CI, so you specify an anchor CI to hold it.

Filter object

Each element of the `filters` array selects source CIs that meet a single field condition.

Filter object properties

Property	Type	Description
field	String	Name of the field on the source CI to evaluate, for example <code>operational_status</code> .
operator	String	Comparison operator applied to the field. One of <code>=</code> , <code>!=</code> , <code>></code> , <code>>=</code> , <code><</code> , <code><=</code> , <code>CONTAINS</code> , <code>STARTSWITH</code> , or <code>IN</code> .
value	String	Value the field is compared against, for example <code>1</code> .

Related topics

[Metric aggregation modes](#)

[Configure a metric aggregation job](#)

[Metric type IDs](#)

Metric type IDs

Metric type IDs identify each time-series metric that a connector writes to the metrics database. Use these IDs to locate metrics in Insights Explorer or when you build custom dashboards and indicators.

Cisco Meraki metric type IDs

Cisco Meraki metric type IDs

Metric type ID	Unit	Description
<code>meraki.uplink.status</code>	Boolean	Uplink up or down.
<code>uplink.usage.sent</code>	Bytes	Bytes sent on the uplink.

Cisco Meraki metric type IDs (continued)

Metric type ID	Unit	Description
uplink.usage.received	Bytes	Bytes received on the uplink.
uplink.throughput.sent	bps	Throughput sent on the uplink.
uplink.throughput.received	bps	Throughput received on the uplink.
uplink.latency	ms	Uplink latency.
uplink.loss	Percent	Uplink packet loss.
uplink.bandwidth.limit.upload	bps	Configured upload bandwidth limit for the uplink.
uplink.bandwidth.limit.download	bps	Configured download bandwidth limit for the uplink.
vpn.usage.sent	KB	Data sent over the VPN tunnel.
vpn.usage.received	KB	Data received over the VPN tunnel.
vpn.minJitter, vpn.maxJitter, vpn.avgJitter	ms	Minimum, maximum, and average VPN jitter.
vpn.minLatency, vpn.maxLatency, vpn.avgLatency	ms	Minimum, maximum, and average VPN latency.
vpn.minPercentageLoss, vpn.maxPercentageLoss, vpn.avgPercentageLoss	Percent	Minimum, maximum, and average VPN packet loss.
vpn.minMOS, vpn.maxMOS, vpn.avgMOS	String	Minimum, maximum, and average VPN Mean Opinion Score.
device.performance.score	String	Device performance score.
switch.port.status	Boolean	Switch port status.

Fortinet metric type IDs**Fortinet metric type IDs**

Metric type ID	Unit	Description
fortinet.sdwan.sla.latency	ms	SD-WAN SLA latency.
fortinet.sdwan.sla.jitter	ms	SD-WAN SLA jitter.
fortinet.sdwan.sla.packetloss	Percent	SD-WAN SLA packet loss.
fortinet.sdwan.sla.link_up	Boolean	SD-WAN link up or down status.
fortinet.sdwan.sla_log.count	String	SLA log entry count.
fortinet.performance.cpu.usage	Percent	FortiGate CPU usage.
fortinet.performance.memory.usage	Percent	FortiGate memory usage.

VERIFY: Confirm the metric type IDs written by the Fortinet SD-WAN Interface Log integration. This integration is listed in the installed integrations reference but its metric type IDs aren't yet documented here.