Solution Overview

This section explains the following topics:

Description

The exponential growth of network traffic and the pressures of efficiently running network operations pose huge challenges for network operators. Providing quick intent-based service delivery with optimal network utilization and the ability to react to bandwidth and latency demand fluctuations in real time is vital to success. Migration to Software-Defined Networks (SDNs) and automation of operational tasks is the optimal way to become more efficient and competitive.

Cisco Crosswork Network Controller is a turnkey network automation solution for deploying and operating IP transport networks that delivers increased service agility, cost efficiency, and optimization for faster time-to-customer value and lower operating costs. The solution combines intent-based network automation to deliver critical capabilities for service orchestration and fulfillment, network optimization, service path computation, device deployment and management, and anomaly detection and automatic remediation. Using telemetry gathering and automated responses, Cisco Crosswork Network Controller delivers network optimization capabilities that are nearly impossible to replicate even with a highly skilled and dedicated staff operating the network.

The fully integrated solution combines functionality from multiple Crosswork components installed upon a common Crosswork infrastructure, as well as industry-leading capabilities from Cisco® Network Services Orchestrator (NSO), Cisco Segment Routing Path Computation Element (SR-PCE), and Cisco WAN Automation Engine (WAE). Its unified user interface provides a single pane of glass for real-time visualization of the network topology and services, provisioning, monitoring, and optimization.

What's New in This Release

This release of Cisco Crosswork Network Controller supports the following new features and capabilities:

- **Circuit Style Segment Routing Traffic Engineering (CS SR-TE)**:

  The Circuit Style Manager (CSM) feature pack provides a bandwidth-aware Path Computation Element (PCE) to compute CS SR-TE policy paths, provision them, and visualize them on geographic or logical maps. CS SR-TE policies guarantee allocated bandwidth services with predictable latency and persistent bidirectional path protection for critical traffic. Operators can provision CS SR-TE policies based on the operator's intent. Unlike Bandwidth on Demand, where SR policies with requested bandwidth are created on a best-effort basis, CS SR-TE reserves a percentage of bandwidth in the network and computes CS SR policy bidirectional failover paths with the requested bandwidth, metric type, and constraints. CS SR-TE also maintains a running account of all CS SR reserved bandwidth in the network. CS SR policies are typically used for high-priority services, such as crucial monetary transactions or important live video feeds, for which the operator can now offer a service level that is guaranteed to be better than simple best-effort performance.

  Crosswork Network Controller enables you to provision CS SR-TE policy configurations and easily edit policies, as needed. In addition, the ability to visualize CS SR policies in your network topology allows you to easily verify CS SR policy configurations, details, and path states. With a few clicks you can view Active and Protected paths, operational status, reserved bandwidth pool size, and monitor path failover behavior for individual CS SR policies.

- **Tree Segment Identifier (Tree-SID) policy provisioning and L3VPN service model association**:

  Operators use Tree-SID to implement multicast trees in segment-routed transport networks. Using Crosswork Network Controller, operators can:

  - Create, provision and visualize static Tree-SID policies using the UI, each representing a leaf or node along the path.
  - Create, provision and visualize dynamic Tree-SID policies directly on devices using an API.
  - Visualize dynamic Tree-SID policies using the UI Traffic Engineering page. However, there will be no mapping on the Transport tab if it is attached to an L3 point-to-multipoint VPN service.
  - Associate static mVPN Tree-SID policies associated with existing, or newly created, L3VPN service models (SR MPLS point-to-multipoint) that can also be visualized and analyzed to assist in efficient management and troubleshooting of your multicast network.
  - Modify existing static Tree-SID policies and mVPN Tree-SID policies associated with an L3VPN service model using the UI.
  - Configure link affinities used to specify the link attributes that determine which links are suitable to form a path for the Tree-SID policy and maps each bit position or attribute with a color (making it easier to refer to specific link attributes).

- **Crosswork UI Improvements**:

  Improvements to the Crosswork UI include:

  - Traffic Engineering Dashboard, which includes the TE Dashlet that provides:
    - A high-level summary of RSVP-TE tunnel, SR-MPLS, SRv6, and Tree-SID policy information (total policy count, policy state, metric types for all TE services, and specific data that is filtered upon a one-click selection).
    - Policies and Tunnels under traffic threshold for historical data by displaying RSVP-TE tunnels and SR-MPLS policies that have traffic below the defined threshold in the selected time period. This information may be used to find and filter the unused policies or tunnels.

  - Filtering the data on the dashlet based on the time range you want to view (date, 1 month, 1 week, 1 day, and 1 hour).
  - Policy and Tunnel Change Events: Displays all the policies and tunnels that have had a path or state change event ordered by the event count, within the selected time range. This information helps identify the unstable policies and tunnels.

- Traffic Engineering event and historical data information associated with a device when viewing details for a policy or tunnel. For example, the traffic rate and event history for an SR-MPLS policy can be viewed by selecting the Historical Data tab and clicking on an event. By doing so, you can view the state of the policy or tunnel at that point in time and view additional details, such as Admin and Operational state, Segment type, accumulated metric, delegated PCE, and more so to drill down on the event details.
- Enable/Disable Alarm Status Badge slider: The system allows you to enable or disable the Alarm Status Badge slider for devices and links across various topology views. By disabling the Alarm Status Badge, you can better focus the overlay on an area of interest when troubleshooting.
- Configuration of the Traffic Engineering Data Dashboard Settings (and historical data) for the collection of policy and tunnel metrics, state and path changes, data retention intervals, and the utilization threshold for underutilized LSPs.
- Global search in UI topology: You can now search within the Crosswork Network Controller topology map in the UI. This feature allows you to quickly locate devices based on the following criteria:
  - Civic Location (for example, San Jose)
  - Host/Device name (for example, NAT-01)
  - IP address (for example, 121.10.10.1.1)
- Import and Export geographical objects using Keyhole Markup Language (KML) format:
  - Using the Crosswork Network Controller UI, you can import and export KML files to exam, change, or add device geographical information and see the updates in the UI map. For example, you may use the export function to download your device's data in KML format to your system, exam and/or change the device details, and upload it into a map generator (such as Google Maps) to view your updated device information and coordinates outside of Crosswork Network Controller. You can then use the import function to upload the updated, or browse for a new, KML file back into Crosswork Network Controller. If changes were made, they will now appear in the geographical map after it refreshes. When using the import function, Crosswork Network Controller also provides a sample KML template. The sample KML template provides information on where to identify devices and their coordinates, an optional device name, and the IP address (IPv4 or IPv6) of a device with corresponding coordinates. This template can be used on your system before importing back into Crosswork Network Controller.
- Traffic Engineering device details improvements that will provide options, after selecting a device from the topology map, to select different TE tabs (such as Links, Alarms, SR-MPLS, SRv6, RSVP-TE and others) that provide associated data for the selected device's policies and prefixes.

Note For more information on Crosswork UI improvements, see the Cisco Optimization Engine guide section, Visualize Traffic Engineering Services.

- **Crosswork Provisioning UI Improvements**
  - Dry Run for a deleted service: When decommissioning a service, only the configuration related to a service is deleted on the device. By implementing Dry Run, it shows the user the configuration that is deleted from multiple devices.
  - Edit in json configuration editor: Using the json configuration editor, you can highlight different details that make up the service configuration and edit them directly in the json editor before committing the configuration. For instance, go to **Services & Traffic Engineering > Provisioning (NSO)**. From the Services/Policies panel, select a service (for instance, **L2VPN > L2vpn-Service**). From the list of services available, select the Actions column for the service configuration you want to edit and click **Edit in Json Editor**. The json Configuration editor popup appears. Click within the editor to make changes on select entries or click on the icons on the left to either: **Drag to move this field**, or, **Click to open the actions menu**. If you select **Click to open the actions menu**, a drop-down list appears. Select one of the options (for example: **Insert**, **Duplicate**, or **Remove**). After you complete the configuration edits, click **Commit**.
  - Clone existing services and policies and utilize the json configuration editor to make changes to your cloned configuration. By cloning existing services and policies, you save time and ensure consistency across configurations while maintaining the ability to make specialized modifications. For instance, go to **Services & Traffic Engineering > Provisioning (NSO)**. From the Services/Policies panel, select a service (for instance, **L2VPN > L2vpn-Service**). From the list of services available, select the Actions column for the service configuration you want to copy and click **Clone**. A L2vpn-Service popup appears requiring a new vpn-id for the cloned service. Add a new name and click **Continue**. The json Configuration editor popup appears. After you complete the cloned configuration edits, click **Commit**.
  - Due to NSO Core Function Pack (CFP) model version upgrade, L2VPN, L3VPN or RSVP-TE upgrade is not supported from 4.x to 5.0. SR-TE upgrade from 4.x to 5.0 is supported. Direct upgrade from 3.x to 5.0 is not supported.
  - **Show all fields** toggle option: When editing a service configuration, you can either hide multiple fields that do not pertain to the editable service configuration or you can view all fields by using the **Show all fields** toggle option.

- **Crosswork Infrastructure and Shared Services**:
  - Support for offline licenses, solution-based licenses, and lab licenses
  - Support for user authentication via single sign-on (SSO)
  - Ability to log the user’s source IP address for auditing and accounting
  - High Availability support for Common Licensing Management Service (CLMS)
  - High Availability support for Engineering Management Functions (Inventory, Notification, Fault, and SWIM)
  - Support for visualization of device alarms and events

    - API and Notification support for alarms/events OSS integration
    - Ability to enable SMU installation using a single playbook

- **Services Overlay Visualization Enhancements**:

  Ability to select Basic View or Extended View when visualizing a service overlay. The Basic View is a minimalistic view with no additional details, edge directions, router targets, or EVI/PW IDs. The Extended View includes all details, including edge directions, router targets, and EVI/PW IDs. The services overlay visualization enhancements apply to:

    - Point-to-Point Service Visualization
    - Any-to-Any Service Visualization (L2VPN and L3VPN)
    - Hub and Spoke Service Visualization (L2VPN and L3VPN)
    - Custom Service Visualization (L2VPN and L3VPN)

- **Cisco Service Health**:

  Service Health monitoring is available for both Basic Monitoring and Advanced Monitoring.

Note For help selecting the appropriate monitoring option for your needs, see the section Basic and Advanced Monitoring Rules. In total, Basic + Advanced Monitoring provides up to 52,000 services that can be monitored.

    - Heuristic Package improvements include:
        - IPv6 support that enhances and extends the SRv6 feature support
        - New Basic and Advanced rules for L2VPN (E-LAN and E-Tree) for monitoring (including multi-point feature for E-LAN and E-Tree)
    - High Availability for all Service Health containers.
    - Assurance Graph improvements that include node aggregation and expand/collapse capabilities to view subservice summary information and associated subservices.
    - New subservices, such as:
        - Dynamic subservices implementation (also includes SR-ODN policy)
        - Reservation Protocol for Traffic Engineering (RSVP-TE) Tunnel
        - Bridge Domain
        - Mac Learning
    - Device badge feature displays an orange badge on a healthy device when viewing devices in the topological view and indicate there are down and/or degraded subservices underneath that should be identified and symptoms inspected.
    - Summary node feature summarizes the aggregated health status of child subservices and reports one consolidated health status to a service node. The Summary node feature is available in both L2VPN multipoint Basic and Advanced monitoring models.

- Basic monitoring subservices:
  - Device – Summarizes the health status of all underlying Devices participating in the given L2VPN service.
  - Bridge Domain – Summarizes the L2VPN Service's Bridge Domain health status across all participating devices.
- Advanced monitoring subservices (in addition to what is also available with Basic monitoring)
  - EVPN – Summarizes the health status of all underlying subservices – BGP Neighbor Health & MacLearning Health across all participating PE endpoints and provides a consolidated overall EVPN health summary status.
  - Transport – Summarizes the health status of all underlying subservices – SR-ODN (dynamic), SR Policy (statically configured) and RSVP TE Tunnel, across all participating PE endpoints and provides a consolidated overall Transport health summary status.
  - SR-PCEP – Summarizes the health status of all the underlying subservices that are monitoring the PCEP sessions. Each underlying subservice monitors the PCEP session health on a particular device participating in the given VPN service.

- Dynamic subservices functionality: In contrast to other subservices, dynamic subservices will be added to or removed from the Assurance Graph according to a service's intent and/or SR polices present on the devices. Each Summary node (Transport) has either *dynamic.l3vpn.sr.policy* or *l2vpn.sr.odn.policy.dynamic* child subservices for each device with a defined SR intent. And each dynamic subservice will have several *sr.policy.pcc.pm* subservices: one for each relevant SR policy on that device. Dynamic subservices are only for SR-policies on supported l2vpn/l3vpn services.
- Extended CLI support using new Service Health system device packages, that can derive exact sensor paths for metric health calculation, that can now be installed as a bundle when the Service Health application is deployed.

Supported Use Cases

Crosswork Network Controller supports a wide range of use cases allowing operators to manage many aspects of the network. The following describes specific use cases, with details about the Crosswork applications needed to deliver each capability.

- **Orchestrated service provisioning:** Provisioning of layer 2 VPN (L2VPN) and layer 3 VPN (L3VPN) services with underlay transport policies to define, meet, and maintain service-level agreements (SLA), using the UI or APIs. Using Segment Routing Flexible Algorithm (Flex-Algo) provisioning and visualizing to customize and compute IGP shortest paths over a network according to specified constraints.

Note An SLA defines the expectations set between a service provider and a customer. The SLA details the products or services that are to be delivered, the point of contact for end-user issues, and metrics by which the effectiveness of the process is both monitored and approved.

- **Real-time network and bandwidth optimization:** Intent-based closed-loop automation, congestion mitigation, and dynamic bandwidth management based on Segment Routing and RSVP-TE. Optimization of bandwidth resource utilization by setting utilization thresholds on links and calculating tactical alternate paths when thresholds are exceeded. The ability to provision SR-Circuit Style policies and visualize them in your network topology provides:
  - Straightforward verification of SR-Circuit Style policy configurations
  - Visualization of SR-Circuit Style details, bi-directional active and candidate paths
  - Operational status details
  - Failover behavior monitoring for individual SR-Circuit Style policies
  - A percentage of bandwidth reservation for each link in the network
  - Manually triggered recalculations of existing SR-Circuit Style policy paths that may no longer be optimized due to network topology changes

- **Local Congestion Management:** Local Congestion Mitigation (LCM) provides localized mitigation recommendations within surrounding interfaces, with the use of standard protocols. Data is gathered in real-time and when congestion is detected, solutions are suggested. LCM has a "human-in-the-loop" aspect which ensures that the control of making changes in the network is in the hands of the operator.

- **Visualization of network and service topology and inventory:** Visibility into device and service inventory and visualization of devices, links, and transport or VPN services and their health status on maps with logical or geographical contexts.

- **Performance-based closed-loop automation:** Automated discovery and remediation of problems in the network by allowing Key Performance Indicator (KPI) customization and monitoring of pre-defined remediation tasks when a KPI threshold is breached. For this use case, Cisco Crosswork Health Insights and Cisco Crosswork Change Automation must be installed.

- **Planning, scheduling, and automating network maintenance tasks:** Scheduling an appropriate maintenance window for a maintenance task after evaluating the potential impact of the task (using WAE Design). Automating the execution of maintenance tasks (such as throughput checks, software upgrades, SMU installs) using playbooks. For this use case, Cisco Crosswork Health Insights and Change Automation must be installed.

- **Secured zero-touch onboarding and provisioning of devices:** Onboarding new IOS-XR devices and automatically provisioning Day0 configuration resulting in faster deployment of new hardware at lower operating costs. For this use case, Cisco Crosswork Zero Touch Provisioning must be installed.

- **Visualization of native SR paths:** Visualizing the native path using the traceroute SR-MPLS multipath command to get the actual paths between the source and the destination can be achieved using Path Query. With Cisco Crosswork Network Controller, a traceroute command runs on the source device for the destination TE-Router ID and assists in retrieving the paths.

- **Provision, Visualize, and Analyze Tree Segment Identifier Policies in Multipath Networks:** Creating and visualizing static Tree-SID policies using the UI. Static mVPN Tree-SID policies associated with existing, or newly created, L3VPN service models (SR MPLS point-to-multi-point) using the Crosswork Network Controller that can also be visualized and analyzed to assist in efficient management and troubleshooting of your multicast network. Configuring link affinities used to specify the link attributes that determine which links are suitable to form a path for the Tree-SID policy and maps each bit position or attribute with a color (making it easier to refer to specific link attributes). Modifying existing static

Tree-SID policies and mVPN Tree-SID policies associated with an L3VPN service model – both edit and delete – using the UI.

Solution Components Overview and Integrated Architecture

The following diagram provides a high-level illustration of how the solution's components work together within a single pane of glass to execute the primary supported use cases.

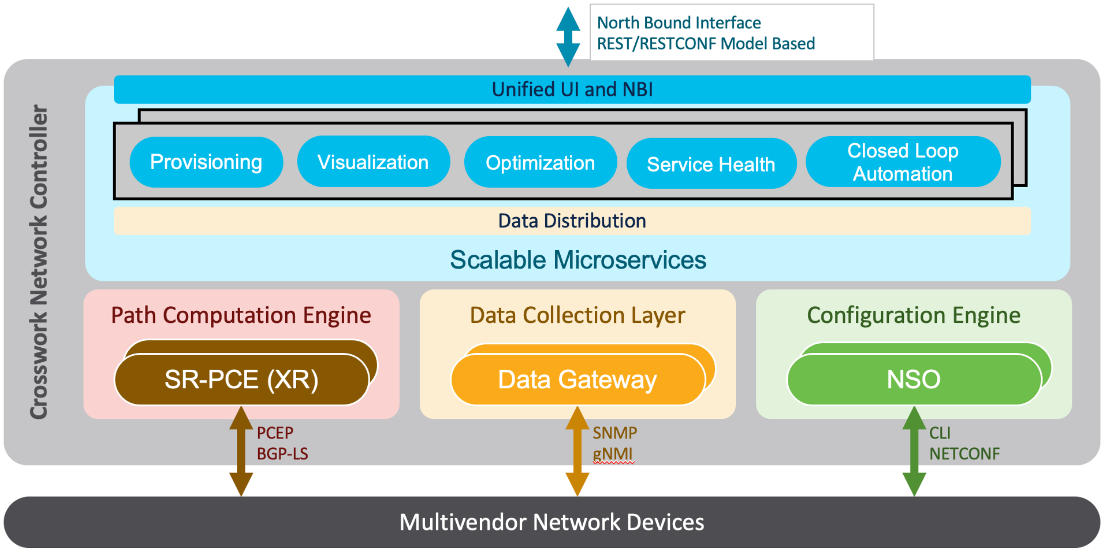

The following components make up the Cisco Crosswork Network Controller 5.0 solution:

Cisco Crosswork Active Topology

Cisco Crosswork Active Topology's logical and geographical maps provide real-time visibility into the physical and logical network topology, service inventory, and SR-TE policies and RSVP-TE tunnels, all within a single pane of glass. They enable operators to see, at-a-glance, the status and health of the devices, services, and policies. Services and transport policies can be visualized end-to-end as an overlay within the context of the topology map. Cisco Crosswork Active Topology provides device grouping functionality so that operators can set up their maps to monitor exactly the set of devices, services, and locations for which they are responsible. In addition, operators can save custom views for quick and easy access to the views and functionality they use on an ongoing basis.

Cisco Crosswork Optimization Engine

Cisco Crosswork Optimization Engine provides real-time network optimization allowing operators to effectively maximize network capacity utilization, as well as increase service velocity. Leveraging real-time protocols, such as BGP-LS and Path Computation Element Communication Protocol (PCEP), SR-PCE and Crosswork Optimization Engine enables closed-loop tracking of the network state, reacting quickly to changes in network conditions to support a self-healing network.

Cisco Crosswork Data Gateway

Cisco Crosswork Data Gateway is a secure, common collection platform for gathering network data from multi-vendor devices. It is an on-premise application deployed close to network devices. Crosswork Data Gateway supports multiple data collection protocols including MDT, SNMP, CLI, standards-based gNMI (dial-in), and syslog. Any type of data can be collected by Crosswork Data Gateway as long as it can be delivered over one of the supported protocols. In this way, it can provide support for a growing set of use cases and customizations.

To address scale challenges, Cisco Crosswork Data Gateway is implemented as a number of VMs and designed with a distributed architecture in mind. Each lightweight VM manages a subset of the overall network and as the network grows, additional VMs can be added horizontally to address the new demands on the compute resources. It also supports a flexible redundancy configuration based on the operator's needs. After the initial setup, Cisco Crosswork Network Controller automatically orchestrates the collection across the multiple Cisco Crosswork Data Gateway VMs.

APIs and configuration examples are available to illustrate how to add new collection jobs (outside of those built for you by Cisco Crosswork Network Controller) to gather additional information from your network. The collected data can be published to approved destinations. Supported destinations are Kafka and gRPC messaging bus.

Crosswork Common UI and API

All Cisco Crosswork Network Controller's functionality are provided within a single, common graphical user interface. This common UI brings together the features of all Crosswork Network Controller's components, including common inventory, network topology and service visualization, service and transport provisioning, and system administration and management functions. When optional add-on Crosswork components are installed, their functionalities are also fully integrated into the common UI. Having all functionality within a common UI, instead of having to separately navigate individual application UIs, enhances the operational experience and increases productivity.

A common API enables Crosswork Network Controller's programmability. The common APIs provides a single access point for all APIs exposed by various built-in components. The API provides a REST-based Northbound Interface to external systems (e.g., OSS systems) to integrate with Cisco Crosswork Network Controller. RESTCONF and YANG data models are made available for optimization use cases. For details about the APIs and examples of their usage, see the Cisco Crosswork Network Automation API Documentation on Cisco DevNet.

Crosswork Infrastructure and Shared Services

The Cisco Crosswork Infrastructure provides a resilient and scalable platform on which all Cisco Crosswork components can be deployed. This infrastructure and shared services provide:

- A single API endpoint for accessing all APIs of Crosswork applications deployed
- A shared Kafka bus to pass data between applications
- Shared database(s) (such as relational and graph) for applications to store data
- A singe shared database to store all gathered time-series data from the network
- A robust Kubernetes-based orchestration layer to provide for process-level resiliency
- Tools for monitoring the health of the infrastructure and the cluster of virtual machines (VMs) on which it resides

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation

Cisco Crosswork Health Insights and Cisco Crosswork Change Automation are components that can optionally be installed with Cisco Crosswork Network Controller.

Cisco Crosswork Health Insights performs real-time Key Performance Indicator (KPI) monitoring, alerting, and troubleshooting. Cisco Crosswork Health Insights enables programmable monitoring and analytics. It provides a platform dynamically for addressing changes to the network infrastructure. Cisco Crosswork Health Insights builds dynamic detection and analytics modules that allow operators to monitor and alert about network events based on user-defined logic.

Cisco Crosswork Change Automation automates the process of deploying changes to the network. Orchestration is defined via an embedded Ansible Playbook and then configuration changes are pushed to Cisco Network Services Orchestrator (NSO) to be deployed to the network.

These components within Cisco Crosswork Network Controller enable closed-loop discovery and remediation of problems in the network. Operators can match alarms to pre-defined remediation tasks, which are performed when a defined Key Performance Indicator (KPI) threshold is breached. This reduces the time it takes to discover and repair a problem while minimizing the risk of human error resulting from manual network operator intervention.

Cisco Crosswork Zero-Touch Provisioning (ZTP)

Cisco Crosswork ZTP can optionally be installed with Cisco Crosswork Network Controller.

Cisco Crosswork ZTP is an integrated turnkey solution for automatically onboarding and provisioning new IOS-XR devices, resulting in faster deployment of new hardware at lower operating costs. Operators can quickly and easily bring up devices using a Cisco-certified software image and a day-zero software configuration. After it is provisioned in this way, the new device is onboarded to the Crosswork device inventory where it can be monitored and managed along with other devices.

Cisco Crosswork ZTP offers Secure ZTP functionality in addition to the Classic ZTP functionality. Secure ZTP is based on RFC 8572 standards and uses secure transport protocols and certificates to verify devices and perform downloads. Secure ZTP is useful when public Internet resources must be traversed to reach remote network devices, or when the devices are from third-party manufacturers. With Secure ZTP, the device and the Cisco Crosswork ZTP bootstrap server authenticate each other using the device's Secure Unique Device Identifier (SUDI) and Crosswork server certificates over TLS/HTTPS. After a secure HTTPS channel is established, the Crosswork bootstrap server allows the device to request to download and apply a set of signed image and configuration artifacts adhering to the RFC 8572 YANG schema. After the image (if any) is downloaded and installed, and the device reloads with the new image, the device downloads configuration scripts and executes them.

Cisco Network Services Orchestrator

Cisco Network Services Orchestrator (NSO) is an orchestration platform that makes use of pluggable function packs to translate network-wide service intent into device-specific configuration. Cisco NSO provides flexible service orchestration and lifecycle management across physical network elements and cloud-based virtual network functions (VNFs), fulfilling the role of the Network Orchestrator (NFVO) within the ETSI (European Telecommunications Standards Institute) architecture. It provides complete support for physical and virtual network elements, with a consistent operational model across both. It can orchestrate across multi-vendor environments and support multiple technology stacks, enabling the extension of end-to-end automation to virtually any use case or device.

Cisco NSO has a rich set of APIs designed to allow developers to implement service applications. It provides the infrastructure for defining and executing the YANG data models that are needed to realize customer services. It is also responsible for providing the overall lifecycle management at the network service level.

Service and device models, written using YANG modelling language, enable Cisco NSO to efficiently 'map' service intent to device capabilities and automatically generate the minimum required configuration to be deployed in the network. This feature, facilitated by Cisco NSO's FASTMAP algorithm, is capable of comparing current configuration states with a service's intent and then generating the minimum set of changes required to instantiate the service in the network.

All Crosswork components that are included in Cisco Crosswork Network Controller or are optional add-ons, with the exception of Cisco Crosswork ZTP, require integration with Cisco NSO.

Cisco Crosswork Network Controller requires the following Cisco NSO function packs:

- SR-TE core function pack (CFP) enables provisioning of explicit and dynamic segment routing policies, including SRv6, and on-demand SR-TE policy instantiation for prefixes with a specific color.
- Sample function packs for IETF-compliant L2VPN and L3VPN provisioning. These function packs provide baseline L2VPN and L3VPN provisioning capabilities, based on IETF NM models. Prior to customization, these sample function packs enable provisioning of the following VPN services:
    - L2VPN:
        - Point-to-point VPWS using Targeted LDP
        - Point-to-point VPWS using EVPN
        - Multipoint VPLS using EVPN (with service topologies ELAN, ETREE, and Custom)
    - L3VPN
    - Sample IETF-compliant RSVP-TE function pack intended as a reference implementation for RSVP-TE tunnel provisioning, to be customized as required.

Note

- By default, the IETF-compliant NM models are used. If your organization wishes to continue to use the Flat models that were provided with the previous version, a manual setup process is required.
- The Cisco NSO sample function packs are provided as a starting point for VPN service provisioning functionality in Cisco Crosswork Network Controller. While the samples can be used “as is” in some limited network configurations, they are intended to demonstrate the extensible design of Cisco Crosswork Network Controller. Answers to common questions can be found on Cisco Devnet and Cisco Customer Experience representatives can provide answers to general questions about the samples. Support for customization of the samples for your specific use cases can be arranged through your Cisco account team.
- Cisco NSO currently does not support bundle ethernet (BE), route distinguisher (RD), or BGP route-target (RT) functions with L2VPN EVPN. Although it does support multihoming and L2VPN route policy, there is no option to specify an RD value in L2VPN for an EVPN ELAN/ETREE, nor is there an option to specify load balancing type. To perform these functions, contact your Cisco account team for a set of custom configuration templates and advice on configuring bundles manually.

Cisco Segment Routing Path Computation Element (SR-PCE)

Cisco SR-PCE is an IOS-XR multi-domain stateful PCE supporting both segment routing (SR) and Resource Reservation Protocol (RSVP). Cisco SR-PCE builds on the native Path Computation Engine (PCE) abilities within IOS-XR devices, and provides the ability to collect topology and segment routing IDs through BGP-LS, calculate paths that adhere to service SLAs, and program them into the source router as an ordered list of segments. A Path Computation Client (PCC) reports and delegates control of head-end tunnels sourced from the PCC to a PCE peer. The PCC and PCE establish a Path Computation Element Communication Protocol (PCEP) connection that SR-PCE uses to push updates to the network and re-optimize paths where necessary.

Cisco SR-PCE can either reside on server resources using virtualized XRv9000 , or as a converged application running within IOS-XR Routers.

Note Adding static routes for auto-discovering the scale nodes from SR-PCE after 2,000 nodes is not supported.

Cisco Service Health

- Service Health substantially reduces the time required to detect and troubleshoot service quality issues. It monitors the health status of provisioned L2/L3 VPN services and enables operators to pinpoint why and where a service is degraded. It can also provide service-specific monitoring, troubleshooting, assurance, and proactive causality through a heuristic model that visualizes the:
  - Health status of sub-services (device, tunnel) to a map when a single service is selected

- Service logical dependency tree and help the operator in troubleshooting in case of degradation by locating where the problem resides, an indication of possible symptoms, and impacting metrics in case of degradation
- Historical view of service health status up to 60 days

Service Health also provides the following:

- Service Health monitoring is available for both Basic Monitoring and Advanced Monitoring options. For help selecting the appropriate monitoring option for your needs, see the section **Basic and Advanced Monitoring Rules**.
- Service Health provides **Internal Storage** of monitoring data up to a maximum limit of 50 GB. This data is stored on your system. If you exceed the limit of the internal storage, historical data will be lost. If you choose to extend Service Health storage capacity, you can optionally configure **External Storage** in the cloud using an Amazon Web Services (AWS) cloud account. By leveraging External Storage, all existing internal storage data will be automatically moved to the external cloud storage (see **Configuring Service Health External Storage Settings** appendix for more details) and your internal storage will act locally as cache storage. Configuring External Storage for Service Health ensures you will not lose historical data for services that continue to monitor a service's health, and will retain service health data for any service you choose to stop monitoring when you select the option to retain historical monitoring service for the data. For more information on Internal and External Storage, and how to retain historical monitoring service data when stopped, see the Appendix sections **Configuring Service Health External Storage Settings** and **Stopping Service Health monitoring.**

Note If you anticipate monitoring a large amount of Service Health services, Cisco recommends you configure External Storage after you install Service Health and before you begin monitoring services so to avoid exceeding the Internal Storage and losing historical data.

- To view subservices supported by Service Health L2VPN/L3VPN, see the **Service Health Supported Subservices** appendix section. Details are provided that define which subservices are supported by each VPN service flavor.
- Service Health supports point-to-point L2VPN.

Note Currently, Service Health does not support multipoint L2VPN.

- Service Health supports integration with standalone Network Services Orchestrator (NSO) or NSO Layered Service Architecture (LSA).
- NSO LSA support is limited to one CFS node and two RFS nodes. These additional NSO types serve as a high availability feature. By distributing your devices across the different types, the LSA feature in Service Health allows for dynamic configurations for assurance.

  To manage the Service Health provider Access, select **Administration** > **Manage Provider Access**. The Providers screen appears. See the Crosswork Administration guide and NSO documentation for additional, detailed information.

- The Service Health Collection Jobs administrative option provides the capability to view Parameterized Jobs (template-based collection jobs) that supports a greater number of jobs, adding the ability to view CLI collection jobs. This is useful when troubleshooting collection job issues by examining details of individual devices using Parameterized Jobs. Devices are identified by their Context ID (protocol) to determine if they are GMNI, SNMP, or CLI-based jobs. Additionally, you may export the collection job information to review. The information is collected at the time the export is initiated and stored in a .csv file.

Note When exporting the collection status, you must fill in the information each time an export is executed. In addition, make sure to review the Steps to Decrypt Exported File content available on the Export Collection Status pop up to ensure you can access and view the exported information.

- Service Health provides expanded redundancy/High Availability (HA) for Assurance Graph Manager, Expression Orchestrator, and Crosswork Expression Tracker microservices (two instances are now available). To view, select **Administration > Crosswork Manager**. In the Crosswork Summary tab, select Crosswork Service Health to view the Application Details screen and Microservices.
  - For example, if you click the Assurance Graph Manager, two redundant/high availability instances appear. In certain situations, one of the instances will be in the active-active mode while the other is in the active-standby mode. This ensures that if one instance goes down, the second acts as a redundant, HA, backup.
- Heuristic Packages: Three additional Rules have been added to assist in Basic monitoring level rules, where a rule to generate Assurance Graph information, for example Basic L2VPN NM P2P services, can be used along with two sub services:
  - Rule-L2VPN-NM- Basic
  - Rule-L2VPN-NM-P2P-Basic
  - Rule-L3VPN-NM-Basic
- Heuristic Package Metrics now has the capability for CLI based metrics and GMNI filtering customizations of packages.

Multi-Vendor Capabilities

Today's networks have typically been built up over time and incorporate multiple vendors and multiple generations of hardware and software. Furthermore, there is a lack of industry standardization, making support for these networks using a single tool challenging.

Service providers require an integrated solution to manage third-party devices that will reduce operational expenses and maintenance overhead, as well as eliminate the need to build custom applications to deploy and maintain different vendor products for a single network.

Because it uses standards-based protocols, Cisco Crosswork Network Controller has multi-vendor capabilities for:

- Network service orchestration via Cisco Network Services Orchestrator using CLI and Netconf/YANG. Cisco Network Services Orchestrator is a YANG model-driven platform for automating provisioning, monitoring, and managing applications and services across multi-vendor networks.
- Telemetry data collection using SNMP with standards-based MIBs, syslog, and gNMI with standard OpenConfig models. Cisco Crosswork Data Gateway also supports Native YANG data models for external destinations and proprietary SNMP MIBs with custom packages.
- Topology and transport discovery via SR-PCE, using IGP and BGP-LS, with link utilization and throughput collected via SNMP using standard MIBs.
- Transport path computation using PCEP.

Note For third-party network device support, use cases must be validated by Cisco Customer Experience representatives in the customer's multi-vendor environment, especially if legacy platforms and non-standard devices or services are involved.

Extensibility

The Cisco Crosswork Network Controller provisioning functionality can be extended using the application programming interfaces (APIs). For more information about the APIs, see the Cisco Crosswork Network Automation API Documentation on Cisco DevNet.

The provisioning UI is extensible as it is rendered based on the YANG model. When new services are introduced, they can be easily incorporated.