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 and optimal network utilization with 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 for operators to accomplish these goals.

Cisco Crosswork Network Controller is an integrated 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 with operator selected manual, or automated, remediation. 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

The information below lists the primary new features and functionality introduced in Cisco Crosswork Network Controller 6.0.x.

Traffic Engineering

- **Local Congestion Mitigation (LCM) feature pack**:
  - Automated Mode—This option allows LCM to automatically deploy TE tunnel recommendations based on thresholds that you configure.

Note Automated mode is accessible through Limited Availability. Engage your account team for further details.

  - Manual Mode (default)—This option, which was available in previous releases, requires a user to view the LCM Operational Dashboard and decide whether to commit TE tunnel recommendations.
  - Pause Mode—This option can pause LCM operations on a particular interface when LCM is in either Automated or Manual mode. Pausing operations in Automated mode are necessary in cases where deployed solutions do not result in the intended resolution, there is uneven ECMP traffic, there are policies that are not carrying traffic, or when an interface is continuously throttling between different solutions.

Note Pausing LCM operations removes all existing TE policies that were deployed for that interface.

- **SR Circuit Style Manager (CSM) feature pack**:
  - Hop count is now available as a metric type when computing SR-TE Circuit Style policies.
  - In response to feedback from customers, we have changed some events to alarms. For example, an alarm is triggered when policy traffic exceeds the reserved bandwidth pool size or threshold.
  - APIs:
    - RESTCONF APIs—Manually re-optimize (single or multiple) SR-TE Circuit Style policies. These APIs can be initiated after network topology changes.
    - CSPolicyPathsOnLinks—Lists Circuit Style SR-TE policies on a specified link and filtered by its operational state (up,down,active, and unknown) of the specified policies.
    - AllCSPolicyPaths—Lists Circuit Style SR-TE policies filtered by its operational state and if it has hops (segment lists).
    - CSPolicyPathsonNode—Lists all Circuit Style SR-TE policies on specified nodes filtered by its operational state (up,down,active, and unknown).

  To view API documentation, see Cisco Devnet.

- **Bandwidth on Demand feature pack**:
  - In previous releases, BWoD required protected adjacency SID constraints. Now user can elect BWoD to prefer to use protected (default option) or unprotected adjacency SIDs.
  - The Policy Violation now has two options: Strict or Loose.

- The process of changing delegation from one PCE to another has been improved to guarantee a clean transfer of PCE roles.
- Enhanced batch processing of queued BWoD policy computations. The queue is initially cleared prior to running a list of new pending delegations/undelegations instead of running each delegation one at a time.

- **Flexible Algorithm**:
  - You can now view Application-Specific Link Attribute ASLA Flexible Algorithm metrics (TE and Delay) in the link details:

Note ASLA is supported on PCC and core routers that are Cisco IOS XR 7.4.1 or later versions.

1. From the Traffic Engineering topology map, click on a participating Flex Algorithm link.
2. From the Links page, click ***Link_Type_entry*** > **Traffic Engineering** tab > **General**. For example:

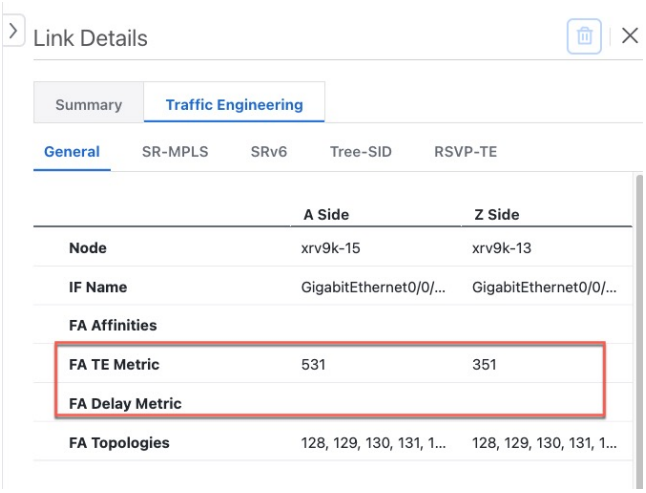

- An overlay on the topology map has been added when Flexible Algorithms are selected. This is to help identify which Flexible Algorithms are selected more easily. For example:

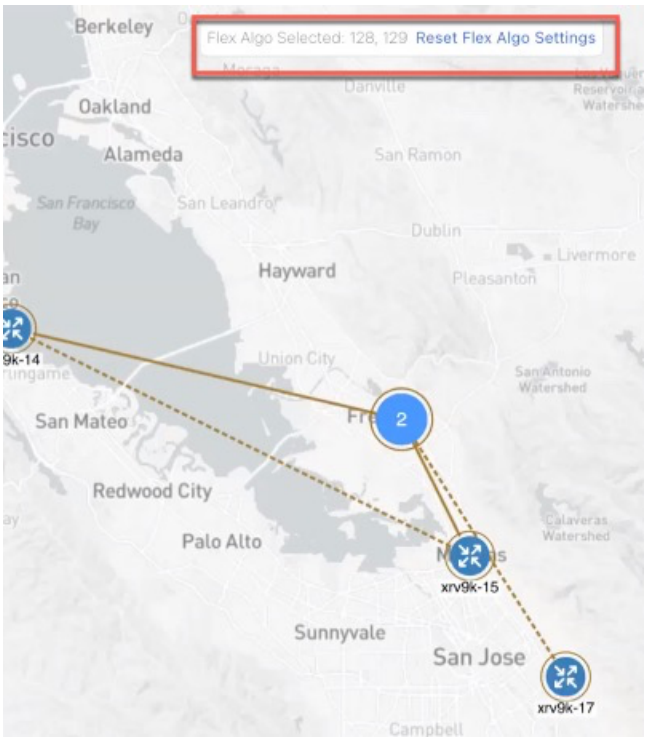

- **Tree-SID**: PCE warnings and path compute elements are displayed in Tree-SID policy details.

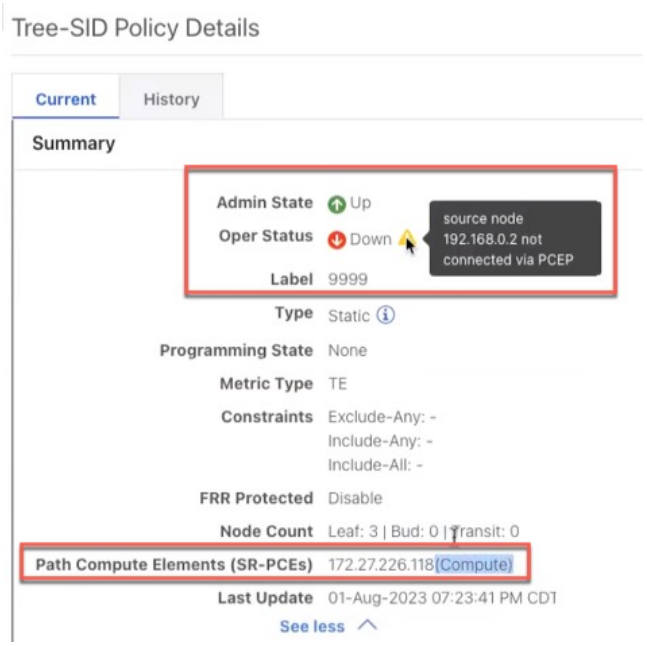

- **Performance Metrics of TE policies**: When Service Health is installed and SR-PM collection is enabled, you can view KPI metrics (Delay, Jitter, and Liveness) from the Traffic Engineering table or from the TE tunnel details.

  You can view the following KPI metrics for the policies:

SR-MPLS and SRv6 policies : Delay, Delay Variance (Jitter) or Liveness (Boolean value) along with traffic utilization.

RSVP-TE policy: Delay and Delay Variance (Jitter) along with Utilization.

- **Asymmetric delay for links**: In previous releases, only one side of the link delay value for an interface was considered during computation. When you configure delays on both remote and local nodes, the calculation of each delay on each interface is now taken into consideration when computing a path.

Note To configure link delay over an interface, refer to the device platform configuration guide. For example, Segment Routing Configuration Guide for Cisco NCS 540 Series Routers.

- **Unique TE tunnel and device detail URLs**: TE tunnel or device details are now assigned unique URLs that can be shared. The URL sends the user to the Policy or Device Details page after logging in.
  - SR-MPLS, SRv6, Tree-SID, and RSVP-TE tunnels —From the Traffic Engineering table, click **Actions > View Details** for a particular row.
  - Devices—From the Traffic Engineering topology map, click on a device to view its details.
- **Increased performance and memory footprint**: Improvements made in topology discovery time, network model building, and processing cache, bandwidth, metric, and TE tunnel type information.
- **Transport Slicing**: Cisco Crosswork Network Controller offers direct support for network slicing at the transport level. This slice “instance” is a unique slice provisioned in the network but with a set of Service Level Requirements chosen from a set of pre-created Network Slice Templates (NST). The SMF in turn communicates with each sub-domain controller, called a Network Slice Subnet Management Function (NSSMF) which in turn provisions the corresponding domain specific slice instance across its own sub-domain boundaries (called a Network Slice Subnet Instance (NSSI) using a similar set of domain specific Slice Subnet Templates (NSST).

  Cisco Crosswork Network Controller also offers:
  - Slice design and deployment from the perspective of two user persona: Slice Designer and Slice Instance Requester.
  - Deploying the Slice Catalog using the new slice template UI.
    - Adding Service Assurance into the catalog using the NSO CLI.
  - Requesting a new Slice Instance by picking intent from the catalog using IETG Slice YANG model, select endpoints and submit.
    - You can deploy a slice instance after providing information, in the UI, after following four easy steps.
  - Automated slice instance deployment:
    - QoS: The Slicing CFP will apply input and output QoS policy maps on all slice endpoint interfaces (policy-maps pre-deployed). Both L2 & L3 QoS supported.

- Path Forwarding: The Slicing CFP will deploy SR-TE ODN templates on all head-ends (metrics= latency, igp, TE, BWoD, FA, etc). Additionally, it will set BGP color community accordingly on all slice advertised VPN prefixes.
- Service Assurance: The Slicing CFP will setup:
  1. Cisco Crosswork Network Controller Heuristic packages for Cisco Crosswork Network Controller Automated Assurance/Service Health.
  2. Configure Y1731 probing for P2P L2 slices.
  3. Configure SR-PM probing for delay and liveness on all slice SR-TE tunnels.
- Connectivity: The Slicing CFP will use the L2/L3VPN IETF NM to setup L3 or L2 connectivity automatically across defined slice endpoints. All VPN parameters inferred and abstracted.
  - Setup eVPN VPWS for P2P L2 slices.
  - Setup eVPN any-to-any or hub-spoke for L2 multipoint or L3 multipoint slices.
  - Setup up “extranet” connectivity between dedicated and shared slice types. (more on this later).
  - Setup PE-CE eBGP for L3 based slices.

Note If you are creating an L2 point-to-point transport slice, prerequisites include the following: (a) The route-policy needs to be configured on the PE nodes (for example: L2-ATTACH). (b) In global-settings on NSO, configure this sample command: `set network-slice-services global-settings parent-rr-route-policy L2-ATTACH`.

- Using the UI, visualize the slice components: VPN, Transport, Health:
  - Display a slice on the map.
  - View Slice and VPN view along with Shared Slices and CE (Neighbor) connected in Logical View.
  - Visualize Shared Slices associated to dedicated slice.
  - From the VPN list, display VPN details including Assurance data if monitoring is enabled.
  - From the Transport list, display SR TE details including SR-PM data if SR-PM is enabled.
  - Using Health details, view symptom details and any failed subexpressions and metrics (which will provide information on any active symptoms and root causes).

Service Health

- **Introduced a new monitoring status - Monitoring Error**: Errors due to a component failures, operational errors or device errors are now displayed as **Monitoring Errors** on the UI. You can filter these errors using the mini-dashboard or the filters.
- **Ability to rate-limit monitoring requests**:

To efficiently manage service monitoring requests, Service Health has implemented a rate-limiting process. This means that there may be a delay in publishing service monitoring requests if the number of requests raised per minute exceeds a specific threshold. The thresholds are defined as follows:

- Basic monitoring requests – 50 services per minute
- Advanced monitoring requests – 5 services per minute
- Delete monitoring requests – 30 services per minute

The rate-limiting process also extends to the monitoring data, that is metrics and Events of Significance (EOS), sent by Crosswork Data Gateways to the Crosswork Tracker component. For example, during a restore process, when all Crosswork Data Gateways send metrics again to the Crosswork Tracker component, the rate at which the Crosswork Tracker processes this data and forwards it to Assurance Graph Manager is regulated. This may lead to a delayed reporting of Events of Significance (EOS) following the restore.

In the event of delays, an event is triggered with a severity level of 'Warning' and a corresponding description to notify you of the delay. The event is cleared once Service Health resumes normal publishing of monitoring requests.

- **Ability to monitor performance metrics of TE policies using SR-PM**: To measure the performance metrics of VPN services using the SR-MPLS or RSVP-TE Traffic Engineering policies, Service Health leverages Segment Routing Performance Measurement (SR-PM). This feature enables measuring metrics on the underlay SR-TE policy to enforce Service Level Agreements in VPN services.
- **Monitor service health with external probes from Accedian Skylight**: Crosswork Network Controller can leverage external probing, provided by Accedian Skylight, to measure metrics of the network services. The metrics are compared with the contracted SLA (defined in the Heuristic package), and the results are made available on the Crosswork Network Controller UI.

  After an L3VPN service is provisioned and service monitoring is enabled, the probe intent and probe topology are learned (from provisioned service) and a probe session to monitor the service starts automatically by invoking relevant RESTConf APIs. Service Health processes the metrics and raises symptoms as needed to be displayed on the UI. You can view historical data for upto 24 hours from the Probe Sessions.

  The maximum number of probe sessions per service are capped at 200 (for all connection types).

Note Accedian Skylight integration is available as a limited-availability feature in this release. Engage with your account team for more information.

Topology

- **Simplified Topology Rebuild Tool**: If the topology is not displaying status as expected, you can now place the system into maintenance mode and then choose to rebuild the topology. This will force the system to create a new topology model and avoid the complicated steps from previous versions.

Note Only users with write permission can Rebuild Topology.

Crosswork Data Gateway

- **Ability to reattempt the import of Controller Certificate file**: When Crosswork Infrastructure and Crosswork Data Gateway are deployed simultaneously, on the first reboot Data Gateway attempts to the download the Controller Certificate file from Crosswork Infrastructure. If the Infrastructure deployment is in-progress, Crosswork Data Gateway may not find the certificate. In the past, you had to wait for the Data Gateway VM to restart before downloading the certificate through the Interactive Console menu.

  With Crosswork Data Gateway's latest release, you can let Data Gateway retry the certificate download multiple times. If the file download fails, the Crosswork Data Gateway will now retry automatically.

  For information on importing the certificate, see the *Import Controller Signing Certificate File* section in Cisco Crosswork Network Controller 6.0 Installation Guide.

- **Dynamic reallocation of the Crosswork Data Gateway vCPU resources**: The Crosswork Data Gateway vCPU resources are now dynamically configured to meet the scaling requirements in response to the number of CPUs assigned to the VM.

- **Parameter to configure the CLI session timeouts for devices**: The **SSH Session Timeout** parameter is implemented to indicate the duration of the CLI connection on a device.

  For information on how to configure the **SSH Session Timeout** parameter, see the *Configure Crosswork Data Gateway Global Parameters* section in Cisco Crosswork Network Controller 6.0 Administration Guide.

- **Changes to the Crosswork Data Gateway APIs**:

  The Crosswork Data Gateway APIs have been altered in the following ways:

  - The new dg-manager APIs are compatible with the OpenAPI v2/v3 specification.
  - The change logs include the deprecated APIs. In the subsequent release, the deprecated APIs are removed.
  - A change log is created for each modified API. The change log includes the APIs that have been deprecated, removed, or updated.

  For information on change logs, see Cisco Devnet.

- **Netconf Collector support is decommissioned**: The NETCONF collector enabled data collection over the NETCONF protocol.

  Support for the NETCONF collector has been discontinued in configurations, such as the base VM, application layer, Docker, and dg-manager.

Infrastructure

- **Device Level RBAC**: This release introduces role-based access control (RBAC) at a device granularity for provisioning and device configuration workflows. Each user must be assigned a role that determines what functions they can access along with a Device Group that determines on which devices they can manage or deploy services. For more information, see the *Manage Device Access Groups* section in the Cisco Crosswork Network Controller 6.0 Administration Guide.

- **Geo Redundancy**: This release introduces the first phase of the geo redundancy solution for Crosswork Network Controller and its components in case of a region or data center failure. For more information, see the *Enable Geo Redundancy* section in the Cisco Crosswork Network Controller 6.0 Installation Guide.

Note Geo Redundancy is accessible through Limited Availability. Engage your account team for further details.

Documentation

- An Information Portal is now available for Crosswork Network Controller 6.0. Information is categorized per functional area, making it easy to find and easy to access.
- Cisco Crosswork Network Controller 6.0 Service Health Monitoring is a new Crosswork Network Controller specific guide that provides information on monitoring the health of L2VPN and L3VPN services. It provides insights into analyzing and troubleshooting degraded services, as well as visualizing service health status and logical dependency trees.
- Cisco Crosswork Network Controller 6.0 Traffic Engineering and Optimization is a new Crosswork Network Controller specific guide that provides information on how to visualize and configure traffic engineering in Crosswork Network Controller.
- Cisco Crosswork Network Controller 6.0 Network Bandwidth Management is a new Crosswork Network Controller specific guide that provides information on how to use Crosswork Network Controller feature packs. Feature packs are tools that tackle congestion mitigation and the management of SR-TE policies to find and maintain intent based bandwidth requirements.

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 supports deployment as either "human in the loop" or fully automated implementations allowing the operator to decide how they want to use the feature. See the Local Congestion Mitigation chapter in the Crosswork Network Controller 6.0 Network Bandwidth Management guide for more information.
- **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.
- **Transport Slice Provisioning:** Cisco Crosswork Network Controller offers direct support for network slicing at the OSI transport layer. Using this solution, network engineering experts can design slice profiles around customer intents and then add them to a catalog. Network line operators can then simply pick the slice that best meets the customer's needs, specify the slice endpoints, and (where needed) set any custom constraints or options built into the chosen slice. Using the UI, you can inspect the slice

details for active symptoms, failures, and root causes. In addition, the slice can be visualized on a geographical map.

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.

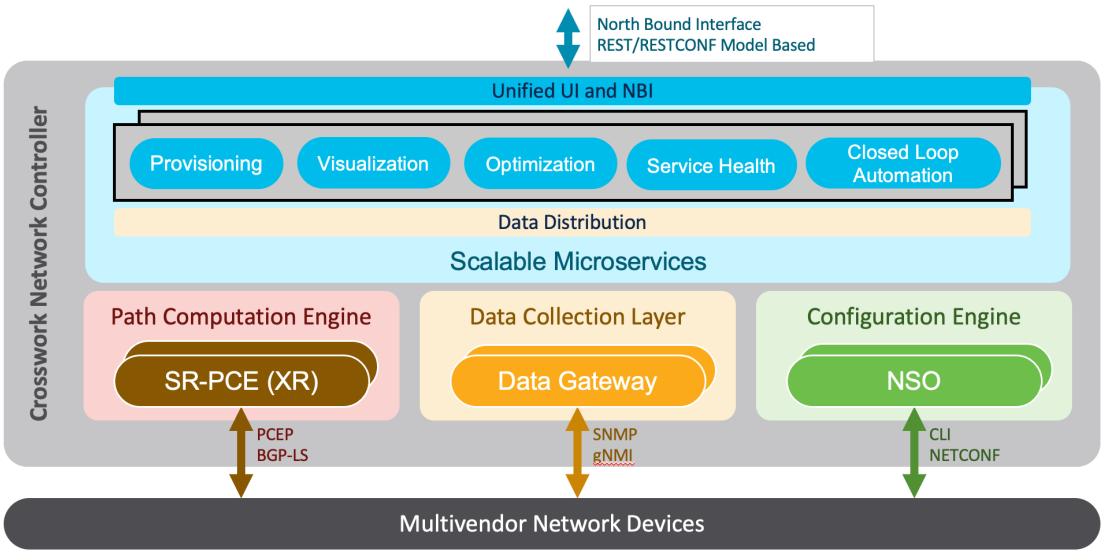

The following components make up the Cisco Crosswork Network Controller 5.0 solution:

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 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.

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.

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.