APPLICATION NOTE

Infinera

Accelerating Services with Multi-domain, Multi-layer Path Computation Element and Service Provisioning

End-to-End Service Routing and Activation with Transcend Software Suite

ENSURING END-TO-END SERVICE FULFILLMENT

Key priorities for network operator investments are:

- Accelerating the deployment of new services to speed up revenue and improve customer experience
- Reducing costs to improve margin and offer lower end-user prices
- Using network resources in an efficient manner to maximize the return on investment

Infinera's Transcend Software Suite enables a series of applications that contribute to these goals. Two of the most relevant applications are end-to-end path computation and service provisioning, which speed up and simplify the deployment of new services.

End-to-end optimal path computation in a large multi-domain network (where a domain can be defined by its geography, technologies, responsibilities in administration or operation, equipment vendor, etc.) is a complex task.

Many of its problems can be addressed with offline service planning, but offline network planning tools frequently require highly skilled planning experts, and, in many cases, optimal results can only be achieved by outsourcing the activity to the equipment vendor, which increases the time and operational costs for planning and provisioning new services. Offline planning tools are also often out of sync with the network's evolving status and its detailed, real-life characteristics, which may limit the efficiency of its solutions.

Additionally, offline planning tools are unsuitable for today's dynamic service routing requirements. With an ever-increasing volume of demands, some of which are short-lived or change at a rapid pace, and expectations of immediate service fulfillment, the need for online, rapid routing and automation of the corresponding provisioning steps becomes evident.

Traditional distributed control planes are one alternative, but they are typically segmented per technology, and therefore network visibility and sharing of information across those boundaries is limited, compromising end-to-end operation. Additionally, not all network elements offer control plane capabilities, and some network elements that support a control plane may not have enough computing power or storage resources to keep up with increasing network size or number of demands.

This is where a multi-domain-capable path computation element (PCE) comes in.

BENEFITS OF TRANSCEND PCE AND SERVICE PROVISIONING APPLICATIONS

- **Accelerate** the activation of new services, expediting revenue and improving customer experience
- **Maximize** network infrastructure investment by allocating network resources to services in an efficient manner
- **Decrease** OpEx with an automation foundation, including open optical networking solutions, that simplifies service deployment in any network
- **Integrate** easily into any software automation environment via open APIs

A PCE is an entity, an application, capable of computing a network path between two endpoints based on a network graph that abstracts and simplifies the network and its connectivity. PCE and PCE-based architectures for end-to-end path computation are defined in IETF RFC 4655.1.

An end-to-end multi-domain PCE that is centralized, consolidating multiple domains and network layers into a global network view, and context optimized (i.e., able to consider real-time live network conditions as well as user-defined constraints when finding the best path), offers a level of scalability, efficiency, and flexibility in service routing that is hard to achieve with manual service planning.

A PCE application integrated with both an operational support system and the network production environment and complemented with a service provisioning engine that automates the configuration of the service in the network enables the operator to automate planning and activation of services and execute autonomously and fast.

The PCE and service provisioning applications are also the pillars for complex automation tasks, such as closed-loop automation processes where events or patterns observed in the network trigger actions in the network. One example is when a trend associated with the degradation of the optical transmission in a fiber (for example, due to aging), or with increasing CPU load in a node, is detected, triggering dynamic rerouting and reoptimization of services and resources, resulting in reprovisioning across multiple network layers, potentially with performance validation and establishment of new wavelengths.

TRANSCEND PCE

Transcend PCE is an end-to-end multi-layer-/multi-domain-capable, RFC 4655-compliant PCE application that is often used together with Transcend service provisioning functionality. For simplification, in the scope of this document, the term Transcend PCE will be used to refer not only to the PCE capabilities as defined by IETF, but also to Transcend service provisioning application, which enables provisioning (configuration and activation) of IP-VPNs, tunnels/LSPs, pseudowires, VLANs, and circuits (OTN and wavelengths) in the network upon path computation.

Transcend PCE is deployed in a centralized server and is able to find global optimal routes across a multi-domain network, while benefiting from “unlimited” computing resources and ensuring scalability. By coordinating and sharing otherwise disconnected information across the multiple layers of a packet optical network (providing visibility, for example, on which optical resources underlay an IP service), Transcend PCE achieves efficient network resource utilization and improved resiliency.

Transcend PCE algorithm takes into consideration a wide variety of data from the network when selecting a path – for example, Transcend PCE can be programmed to factor real-time latency or impairments of the underlying optical layer into the validation of a path search, which is not necessarily the case with other routing implementations (IP/MPLS, Ethernet, GMPLS).

The routing and service provisioning functionalities are exposed by Transcend PCE toward higher-layer applications, an orchestrator, or an operational or business support system through a simplified, abstract, and open interface (RESTCONF-based ONF T-API for Layer 0/Layer 1/Layer 2 services, IETF TEAS for Layer 3 services), which hides the complexity of the underlying network technologies and ensures simple operation. Following an intent-based approach, services can be simply defined by endpoints and a set of service level agreement (SLA) parameters (e.g., latency, resiliency, or restoration) – Transcend PCE computes the path accordingly, ensuring cost-efficient resource assignment, and is able, if desired, to send cross-connection commands to each involved node for service fulfillment. Transcend PCE behavior is programmable, supporting user-defined constraints and customizable abstract cost factors when performing optimal path selection.

LOOKING ONE LEVEL DEEPER

A PCE operates on a graph (i.e., topology) map that represents the current network state and its connectivity (i.e., the forwarding potential between resources). The graph summarizes the network’s existing and potential connectivity across multiple network layers, between all links and nodes, and the structure within each node, and is complemented by information tables that summarize the network policies and capabilities of the resources, as well as user-defined constraints.

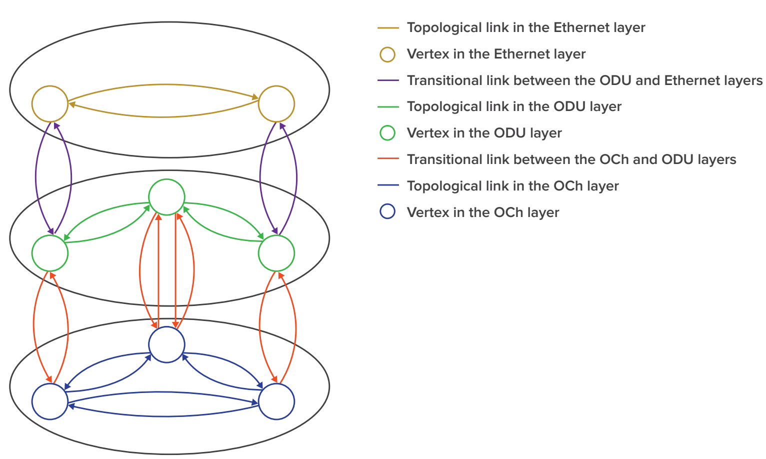

Figure 1: Logical representation of a multi-layer PCE connectivity graph (three-layer example)

The level of abstraction of the graph needs to be such that the path search process is simple and fast, but it still needs to contain enough information to deliver optimal results as the network connections become more complex and scale.

Transcend PCE's path search is based on graph search algorithm, a modified version of the classic Dijkstra algorithm. The search is oriented at minimizing an abstract cost for the path, while guaranteeing that a set of user-defined constraints for that path is fulfilled.

Different parameters such as number of hops, link congestion, latency, underlying optical data rate, link costs, and other user-defined (or machine learning-defined) metrics can be weighted to define the abstract cost to be minimized. Constraints that can be imposed include maximum packet loss, modulation format, wavelength, optical transmission impairment, FEC, latency, shared risk link groups (SRLGs), inclusion/exclusion of links and network elements, and the option to use the same or disjoint paths relative to other services.

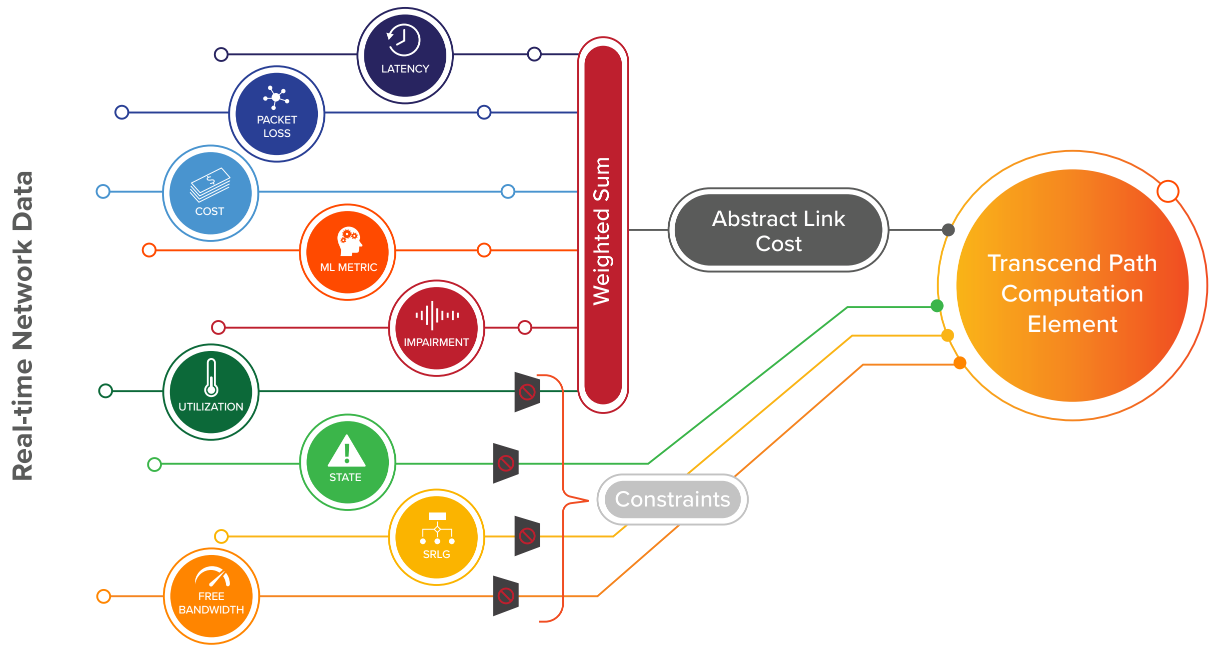

Figure 2: User-defined optimal path

Transcend PCE is able to handle three different types of routing requests:

- "Evaluation request," where the PCE performs a path computation for purely informative purposes, resulting in no impact on the network configuration or model. This allows the comparison of several alternative feasible routes and/or rates before a final choice is selected for reservation or provisioning.
- "Reservation request," where the PCE performs a path computation and reserves the corresponding network resources for future service activation, without provisioning them in the network. The reservation mode can optionally use a timeout mechanism, which frees the resources after a given time period, identified in the service request.
- "Provisioning request," where the PCE performs a path computation and the service provisioning engine immediately configures the corresponding network resources and activates the service.

Transcend PCE is synchronized with the network at all times to ensure that up-to-date topology and resource state information is used whenever processing new requests. The range of objects that may impact the PCE network view includes links between nodes, internal node connections, interface cards, interface ports, and other cards in a network element. When the availability status of a given resource changes (due to a failure, external configuration, etc.), the corresponding network information in the PCE is updated.

Transcend PCE provides the option to work as a stateful PCE, i.e., a PCE that considers the information not only from those paths that have been established in the network, but also from previously reserved paths and associated resources, or as a stateless PCE, i.e., ignoring any previously reserved but not yet provisioned path. Transcend PCE can also act as an active PCE, since it can be configured to actively respond to changes in network resource availability, finding new routes for impacted services and, potentially, rerouting them autonomously.

By default, Transcend PCE will not make use of network resources that are disabled, i.e., objects that are in an unavailable state. However, this behavior can be overridden. Transcend PCE also provides support for routing and provisioning bundled services, i.e., handling through a single request a group of several individual services with common properties.

To meet the required application high availability, Transcend PCE makes use of geographically redundant processors and a resilient management network. Additionally, a microservices-based software architecture and cloud-native technologies allow it to run across containers deployed on a cluster of multiple servers or virtual machines, making use of multiple cores and multiple threads.

MULTI-DOMAIN PCE

Transcend PCE can operate on a single network domain or over multiple domains – the PCE operation is possible as long as the network and connectivity information across all domains and all layers can be abstracted into a single network graph. In a network spanning a variety of domains, each consisting, for example, of different legacy equipment with proprietary interfaces and policies, the ability to mediate all sorts of network elements and summarize their information into one graph may be hard to achieve and/or ineffective.

A common approach to handle this sort of multi-domain path computation is that of a hierarchical path computation element (H-PCE) architecture, introduced in IETF RFC 6805. In this architecture, all paths within a child domain are computed by a child PCE, which has visibility only within the child domain. A parent PCE maintains a domain map containing the child domains and their interconnections. The parent PCE has only abstracted information on network connectivity and resource status within each child domain but is fully aware of the connectivity between the various child domains, as these interconnections are links in its own map. The parent PCE will selectively invoke the necessary child PCEs, possibly in parallel, to assemble an end-to-end path. The inter-domain path computation requires child domain orchestration, i.e., association and coordination of topology, routing, and policy information from multiple children.

Note that this the same concept of orchestrating end-to-end paths across several domains, where each has its own independent entity responsible for intra-domain path computation, may be used for interworking between GMPLS-capable and GMPLS-incapable networks, for example.

Transcend PCE instances can be deployed as a parent PCE in a hierarchical PCE architecture, orchestrating several child PCEs, as well as deployed as a child PCE responsible for a specific child domain under a parent PCE, whether Transcend-based or from a third party. Child PCEs and parent PCEs interface each other through open APIs (RESTCONF-based ONF T-API for Layer 0/Layer 1/Layer 2 services, IETF TEAS for Layer 3 services). Child PCEs may interface their domain devices through proprietary interfaces or through open device APIs, e.g., NETCONF.

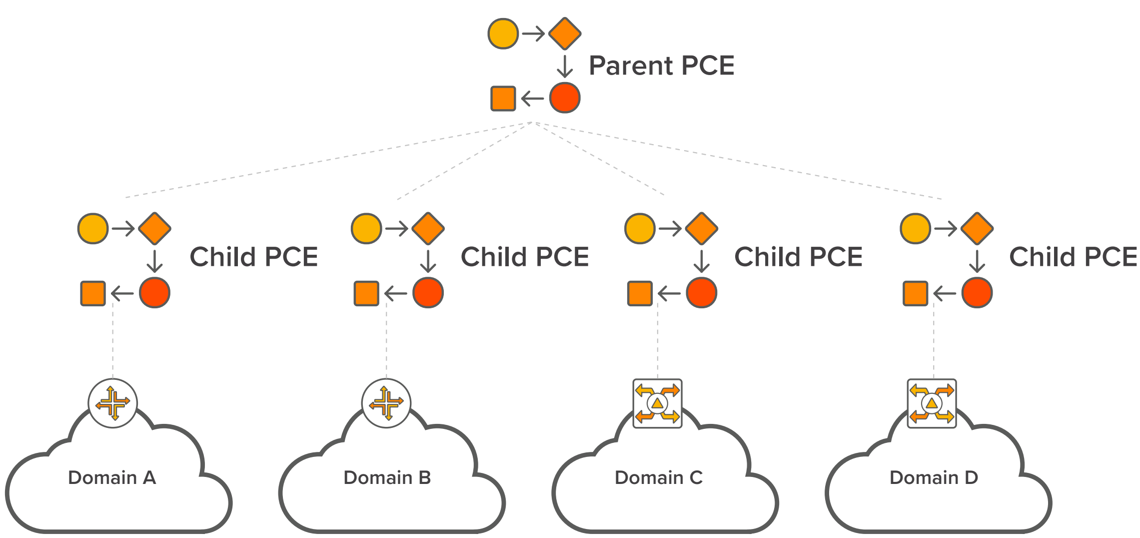

Figure 3: Hierarchical PCE

A NOTE ON OPTICAL IMPAIRMENT AWARENESS IN TRANSCEND PCE

The detailed simulation of optical impairments on a physical fiber and the calculation of fiber-dependent, signal-dependent reach is out of the scope of the PCE application. However, there are alternative ways through which Transcend PCE acquires information that allow it to perform optical feasibility checks.

Transcend PCE can receive as input from a network planning tool, and for a given optical signal type (i.e., considering rate, modulation format, FEC, line interface type), a set of K feasible optical paths between each node pair, including the information on which wavelengths are valid for a given path.

Alternatively, Transcend PCE can interface with an online optical performance engine that validates the optical budget of a path precomputed by the PCE, including the impacts of non-linear transmission effects.

SUMMARY

Transcend PCE supports programmable and context-optimized end-to-end multi-layer, multi-domain path computation and provisioning. It selects optimal paths based on real-time network topology and network status, as well as user-defined constraints and multi-dimensional optimization cost factors.

Transcend PCE application speeds up the activation of new services, accelerates revenue, and improves customer experience. By allocating network resources to services in an efficient manner, Transcend PCE contributes to maximizing the operator's investment in network infrastructure.

Through the support of open APIs, Transcend PCE integrates smoothly into an operator's software automation environment, both overcoming the limitations faced today when planning and validating new services with the use of offline planning tools and manual resource configuration, and providing the foundation for more complex automation tasks, such as closed-loop automation.

Additionally, Transcend PCE provides the means for seamless service deployment in cost-effective open optical networking solutions.

Infinera®