Workgroup: TEAS Working Group
Internet-Draft:
draft-li-teas-hierarchy-ip-controllers-12
Published: 22 October 2023
Intended Status: Informational
Expires: 24 April 2024
Authors: Z. Li
Huawei Technologies
D. Dhody
Huawei
H. Chen
Futurewei Technologies

Hierarchy of IP Controllers (HIC)

Abstract

This document describes the interactions between various IP controllers in a hierarchical fashion to provide various IP services. It describes how the Abstraction and Control of Traffic Engineered Networks (ACTN) framework is applied to the Hierarchy of IP controllers (HIC) as well as document the interactions with other protocols like BGP, Path Computation Element Communication Protocol (PCEP), and other YANG-based protocols to provide end to end dynamic services spanning multiple domains and controllers (e.g. Layer 3 Virtual Private Network (L3VPN), Seamless MPLS etc).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 24 April 2024.

to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction

Software-Defined Networking (SDN) refers to a separation between the control elements and the forwarding components so that software running in a centralized system called a controller, can act to program the devices in the network to behave in specific ways. A required element in an SDN architecture is a component that plans how the network resources will be used and how the devices will be programmed. It is possible to view this component as performing specific computations to place flows within the network given knowledge of the availability of network resources, how other forwarding devices are programmed, and the way that other flows are

routed. The Application-Based Network Operation (ABNO) [RFC7491] describes how various components and technologies fit together.

A domain [RFC4655] is any collection of network elements within a common sphere of address management or path computation responsibility. Specifically, within this document, it means a part of an operator's network that is under common management. Network elements will often be grouped into domains based on technology types, vendor profiles, and geographic proximity and under a domain controller.

Multiple such domains in the network are interconnected, and a path is established through a series of connected domains to form an end-to-end path over which various services are offered. Each domain is under the control of the domain controller (or lower-level controller), and a "super controller" (or high-level controller) takes responsibility for a high-level view of the network before distributing tasks to domain controllers (or lower-level controllers). It is possible for each of the domain to use a different tunnelling mechanism (eg RSVP-TE, Segment Routing (SR) etc).

[RFC8453] describes the framework for Abstraction and Control of Traffic Engineered Networks (ACTN) as well as a set of management and control functions used to operate multiple TE networks. This documents would apply the ACTN principles to the Hierarchy of IP controllers (HIC) and focus on the applicability and interactions with other protocols and technologies (specific to IP packet domains).

Sometimes, service (such as Layer 3 Virtual Private Network (L3VPN), Layer 2 Virtual Private Network (L2VPN), Ethernet VPN (EVPN), Seamless MPLS) require sites attached to different domains (under the control of different domain controller) to be interconnected as part of the VPN service. This requires multi-domain coordination between domain controllers to compute and set-up an E2E path for the VPN service.

This document describes the interactions between various IP controllers in a hierarchical fashion to provide various IP services. It describes how the Abstraction and Control of Traffic Engineered Networks (ACTN) framework is applied to the Hierarchy of IP controllers (HIC) as well as document the interactions with control plane protocols (like BGP, Path Computation Element Communication Protocol (PCEP)) and management plane aspects (YANG models) to provide end to end dynamic services spanning multiple domains and controllers (e.g. L3VPN, Seamless MPLS, etc.).

2. Overview

Figure 1 show examples of multi-domain IP domains under the hierarchy of IP controllers.

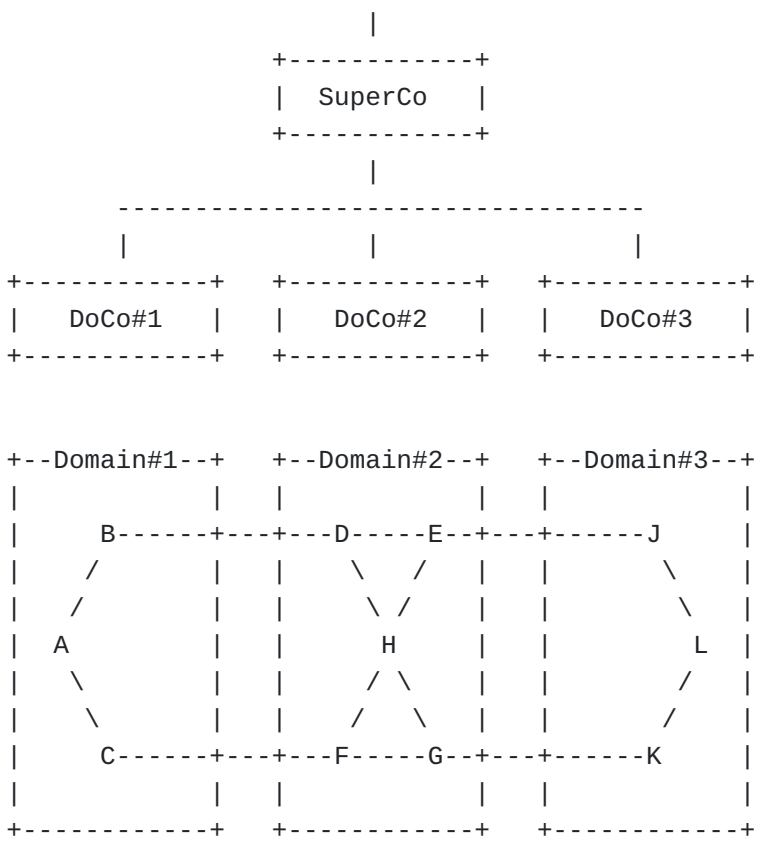

Figure 1: Example: Hierarchy of IP controllers (HIC)

The IP "Super Controller" receives a request from the network/service orchestrator to set-up dynamic services spanning multiple domains. The IP "Super Controller" breaks down and assigns tasks to the domain controllers, responsible for communicating to network devices in the domain. It further coordinates between the controller to provide a unified view of the multi-domain network.

2.1. Mapping to ACTN

As per [RFC8453], ACTN has following the main functions -

- *Multi-domain coordination
- *Virtualization/Abstraction
- *Customer mapping/translation

*Virtual service coordination

These functions are part of Multi-Domain Service Coordinator (MDSC) and/or Provisioning Network Controller (PNC). Further, these functions are part of the controller/orchestrator.

The HIC is an instantiation of the ACTN framework for the IP packet network. The IP domain (lower-level) controllers implement the PNC functionalities for configuring, controlling, and monitoring the IP domain. The "super controller" (high-level controller) implements the MDSC functionalities for coordination between multiple domains as well as maintaining an abstracted view of multiple domains. It also takes care of the service-related functionalities of the customer-mapping/translation and virtual service coordination.

The ACTN functions are part of the IP controllers and responsible for the TE topology and E2E path computation/set-up. There are other functions along with ACTN that are needed to manage multiple IP domain networks.

2.2. Interface between Super Controller and Domain Controller in HIC

The interaction between super controller and the domain controllers in HIC is a combination of Control Plane and Management Plane interface as shown in Figure 2. BGP [RFC4271] and PCEP [RFC5440] are example of the control plane interface; whereas NETCONF [RFC6241] and RESTCONF [RFC8040] are examples of the management plane interface.

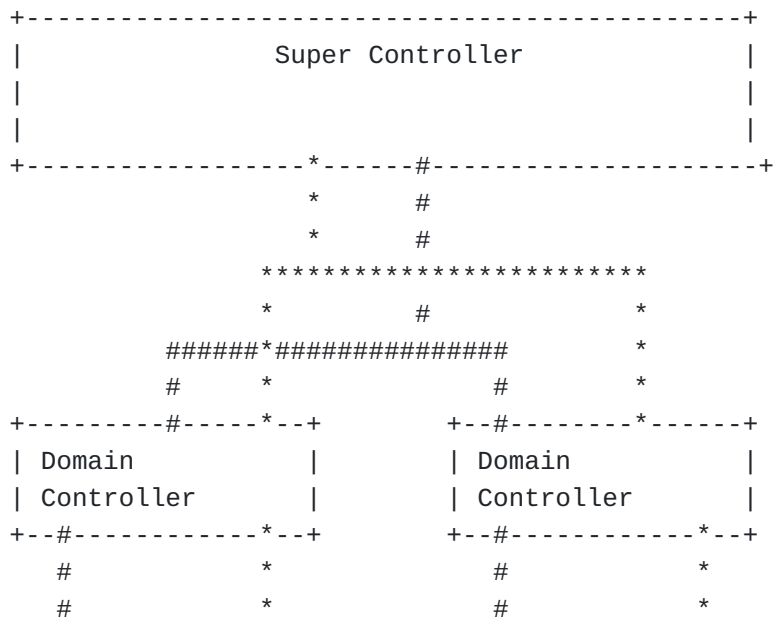

Figure 2: Interface between Super Controller and Domain Controller

Note that ACTN's MDSC-PNC Interface (MPI) could be implemented via management plane interface using YANG models [I-D.ietf-teas-actn-yang] or via PCEP control plane interface [RFC8637].

3. Key Concepts

3.1. Topology

The Domain Controller is expected to be aware of the topology of the network devices in its domain. The domain controller could participate in the IGP ([RFC3630] and [RFC5305]) or use BGP-LS [I-D.ietf-idr-rfc7752bis] by which link-state and TE information is collected and shared with the domain controller using the BGP routing protocol.

An alternate approach would be to rely on the management plane interface which uses the YANG model for network/TE Topology as per [RFC8345] and [RFC8795].

The domain controller is expected to share the domain topology to the Super Controller, as per ACTN (where PNC abstract the topology towards MDSC). A level of abstraction is usually applied while

presenting the topology to a higher-level controller. Topology abstraction is described in [RFC7926] as well as [RFC8453]. BGP-LS, PCEP-LS [I-D.dhodylee-pce-pcep-ls] or management plane interface based on the abstracted network/TE Topology could be used to carry the abstract topology to the super-controller. At minimum, the border nodes and inter-domain links are exposed to the super-controller.

Further [RFC8453] defines three types of topology abstraction - (1) Native/White Topology; (2) Black Topology; and (3) Grey Topology. Based on the local policy, the domain controller would share the domain topology to the Super Controller based on the abstraction type. Note that any of the control plane or management plane mechanism could be used to carry abstracted domain topology. The Super Controller's MDSC function is expected to manage a E2E topology by coordinating the abstracted domain topology received from the domain controllers.

3.2. Path Computation/Path instantiation

The Domain Controller is responsible for computing and setup of path when the source and destination are in the same domain, otherwise the Super Controller coordinates the multi-domain path computation and setup with the help of the domain controller. This is aligned to ACTN.

PCEP [RFC5440] provides mechanisms for Path Computation Elements (PCEs) [RFC4655] to perform path computations in response to Path Computation Clients (PCCs) requests. Since then, the role and function of the PCE has grown to allow delegated control [RFC8231] and PCE-initiated use of network resources [RFC8281].

Further, [RFC6805] and [RFC8751] describes a hierarchy of PCE with Parent PCE coordinating multi-domain path computation function between Child PCE(s). This fits well with HIC as described in this document.

Note that a management plane interface which uses the YANG model for path computation/setup ([I-D.ietf-teas-yang-path-computation] and [I-D.ietf-teas-yang-te]) could be used in place of PCEP.

In case there is a need to stitch per domain tunnels into an E2E tunnel, mechanism are described in [I-D.ietf-pce-stateful-interdomain].

3.3. BGP considerations

[RFC4456] describes the concept of route-reflection where a "route reflector" (RR) reflects the routes to avoid full mesh connection between Internal BGP (IBGP) peers. The IP domain controller can play

the role of RR in its domain. The super controller can further act as RR to towards the domain controller.

BGP can provide routing policies for traffic management, like route preference, AS-path filter policy, IP-prefix filter policy and route aggregation. The controller can distribute these BGP policies into the routers in a single IP domain. For the scenario of multiple domains, the super controller can distribute per BGP Policy into each IP domain controller. Then the IP domain controller trickles down the BGP Policy to the network devices.

[RFC8955] describes the concept of BGP Flowspec that can be used to distribute traffic flow specifications. A flow specification is an n-tuple consisting of several matching criteria that can be applied to IP traffic. The controller can originate the flow specifications and disseminate it to the routers. The flow action includes the redirection to a specific TE tunnel. Also, the IP domain controller could be responsible for collecting the flow sample in its domain and the super controller can act as the Flow Analysis Server.

[RFC7854] describes the BGP Monitoring Protocol (BMP) to monitor BGP sessions. BMP is used to obtain route views with a flexible way. In the fashion of hierarchical architecture, the IP domain controller can be used as the domain Monitoring Station. Meanwhile, the super controller is responsible for a high-level view of the global network state.

4. VPN Service

4.1. Seamless MPLS

Seamless MPLS [I-D.ietf-mpls-seamless-mpls] describes an architecture which can be used to extend MPLS networks to integrate access and core/aggregation networks into a single MPLS domain.In the seamless MPLS for mobile backhaul, since there are multiple domains including the core network and multiple mobile backhaul networks, for each domain there is a domain controller. In order to implement the end-to-end network service provision, there should be coordination among multiple domain controllers.

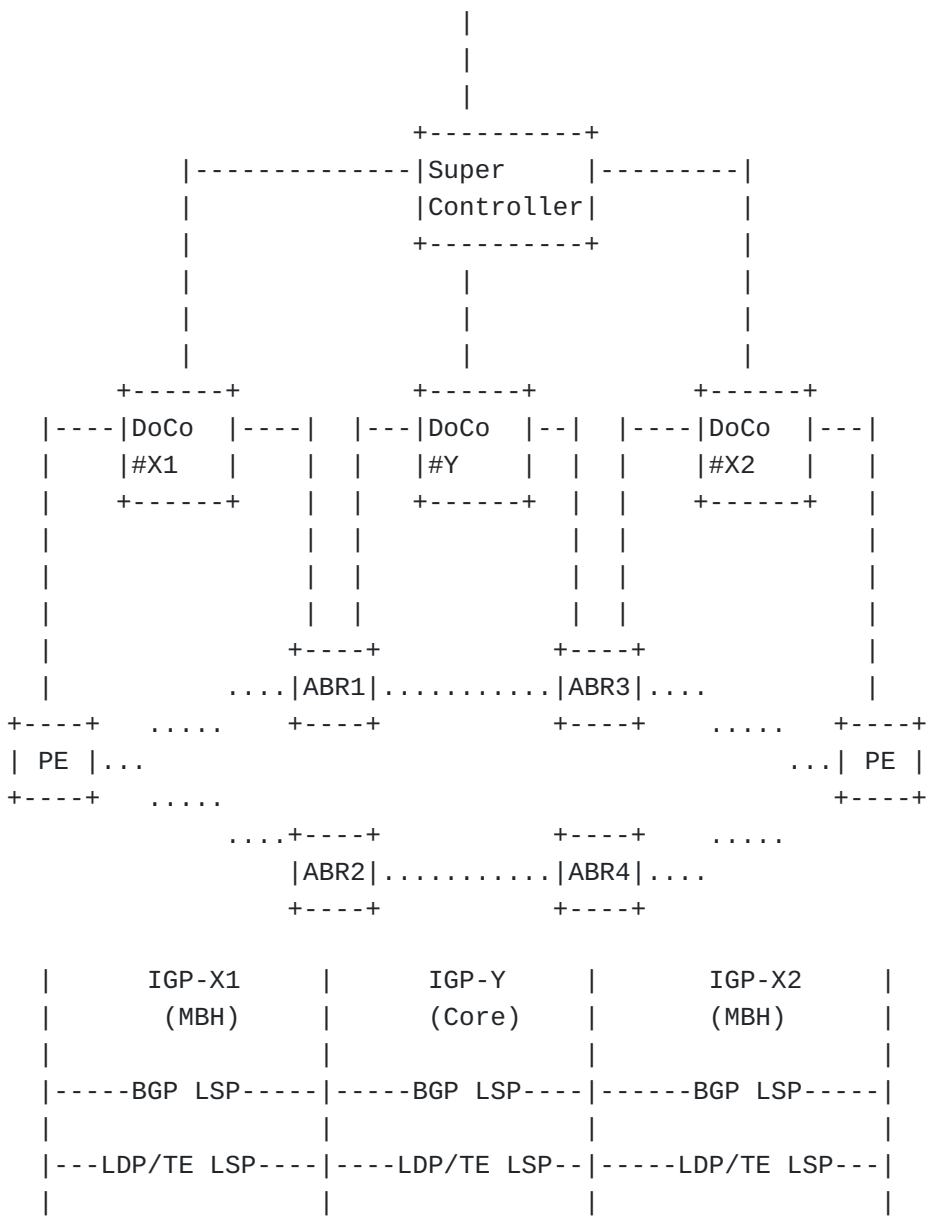

Figure 3: Seamless MPLS

Super Controller is responsible for setting the seamless MPLS service. It should break down the service model to network configuration model [RFC8309] and the domain controller further break it to the device configuration model to the PE/ASBR to make the E2E seamless MPLS service. The selection of appropriate ASBRs and handling of intra-domain tunnels is coordinated by the Super Controller in a similar fashion as shown in Section 4.2.

By enabling BGP sessions between Domain Controller and Super Controller, BGP labeled routes can also be learned at Super

Controller. As Super Controller is aware of the (abstract) topology, it could make intelligent decisions regarding E2E BGP LSP to optimize based on the overall traffic information.

4.2. L3VPN

A Layer 3 IP VPN service is a collection of sites that are authorized to exchange traffic between each other over a shared IP infrastructure. [RFC4110] provides a framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs). [RFC8299] provides a L3VPN service delivery YANG model for PE-based VPNs. The Super controller is expected to implement the L3SM model and translate it to network models towards the domain controller, which in turn translate it to the device model. See [RFC8309] for more details.

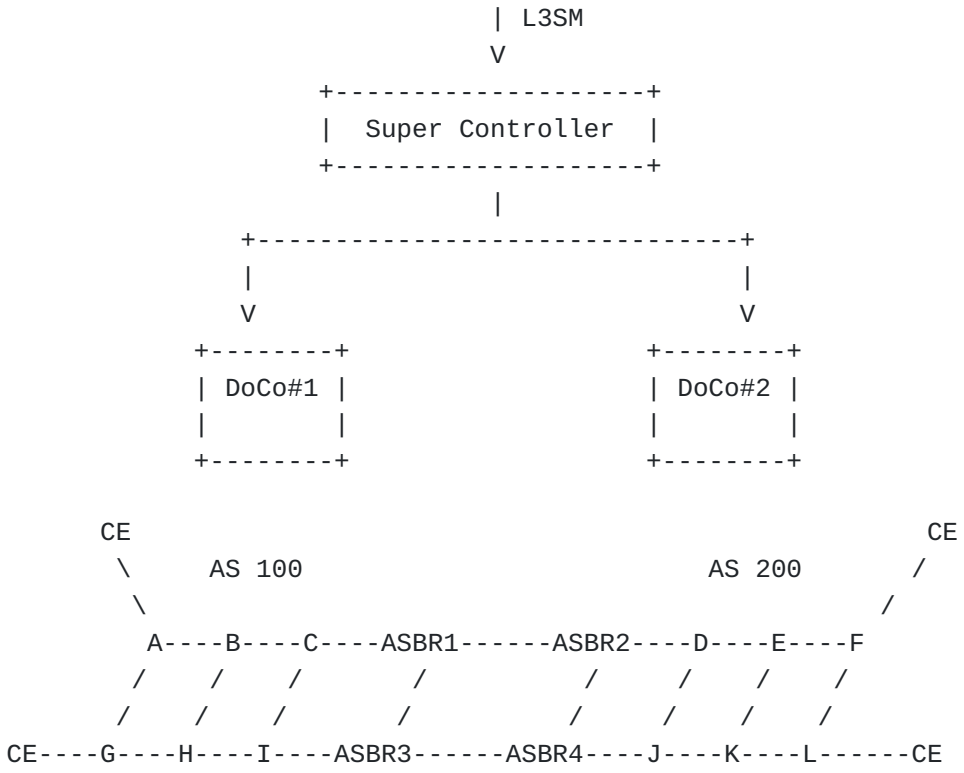

Figure 4: L3VPN

Based on the user data in the L3SM model, the network configurations need to be trickle down to the network device to set up the L3VPN.

[RFC9182] describes the need for a YANG model for use between the entity that interacts directly with the customer (service orchestrator) and the entity in charge of network orchestration and control which, according to [RFC8309], can be referred to as Service Delivery Model. The resulting model is called the L3VPN Network Model (L3NM).

Based on the QoS or Policy requirement for the L3VPN service, the Super Controller may -

- *Set the tunnel selection policy at the PE/ASBR routers so that they could select the existing tunnels
- *Select an existing tunnel at the controller level and bind it to the VPN service
- *Initiate the process of creating a new tunnel based on the QoS requirement and bind it to the VPN service
- *Initiate the process of creating a new tunnel based on the policy

Refer [I-D.ietf-teas-te-service-mapping-yang] for more details from ACTN perspective.

Apart from the Management plane interface based on respective YANG models, the control plane interface PCEP could be used for path computation and setup.

4.3. L2VPN and EVPN service

There are two fundamentally different kinds of Layer 2 VPN service that a service provider could offer to a customer: Virtual Private Wire Service (VPWS) and Virtual Private LAN Service (VPLS) [RFC4664]. A VPWS is a VPN service that supplies an L2 point-to-point service. A VPLS is an L2 service that emulates LAN service across a Wide Area Network (WAN). A BGP MPLS-based Ethernet VPN (EVPN) [RFC7432] addresses some of the limitations when it comes to multihoming and redundancy, multicast optimization, provisioning simplicity, flow-based load balancing, and multipathing, etc.

The handling of L2VPN/EVPN service is done in a similar fashion as shown in Section 4.2.

5. Possible Features/Extensions

This sections list some of the possible features or protocol extensions that could be worked on to deploy HIC in a multi-domain packet network.

1. Simplify the initial configurations needed to set-up the relationship between the super controller and the domain controllers. Note that this could be done via exchanges during initial session establishment, discovery via other protocols, service discovery (such as DNS etc.).
2. The (higher-level controller, lower-level controller) relationship or the role of the controller.

3. The learning and handling of various capabilities of the Super Controller and Domain Controller.

4. Handling of multiple instances of the controller at each level for high availability.

[Editor's Note - This list is expected to be updated in the next version with more details]

6. Other Considerations

6.1. Control Plane

6.1.1. PCE / PCEP

The Path Computation Element communication Protocol (PCEP) [RFC5440] provides mechanisms for Path Computation Elements (PCEs) [RFC4655] to perform path computations in response to Path Computation Clients (PCCs) requests.

The ability to compute shortest constrained TE LSPs in Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks across multiple domains have been identified as a key motivation for PCE development.

A stateful PCE [RFC8231] is capable of considering, for the purposes of path computation, not only the network state in terms of links and nodes (referred to as the Traffic Engineering Database or TED) but also the status of active services (previously computed paths, and currently reserved resources, stored in the Label Switched Paths Database (LSPDB).

[RFC8051] describes general considerations for a stateful PCE deployment and examines its applicability and benefits, as well as its challenges and limitations through a number of use cases.

[RFC8231] describes a set of extensions to PCEP to provide stateful control. A stateful PCE has access to not only the information carried by the network's Interior Gateway Protocol (IGP), but also the set of active paths and their reserved resources for its computations. The additional state allows the PCE to compute constrained paths while considering individual LSPs and their interactions. [RFC8281] describes the setup, maintenance and teardown of PCE-initiated LSPs under the stateful PCE model.

[RFC8231] also describes the active stateful PCE. The active PCE functionality allows a PCE to reroute an existing LSP or make changes to the attributes of an existing LSP, or a PCC to delegate control of specific LSPs to a new PCE.

Computing paths across large multi-domain environments require special computational components and cooperation between entities in different domains capable of complex path computation. The PCE provides an architecture and a set of functional components to address this problem space. A PCE may be used to compute end-to-end paths across multi-domain environments using a per-domain path computation technique [RFC5152]. The Backward recursive PCE based path computation (BRPC) mechanism [RFC5441] defines a PCE-based path computation procedure to compute inter-domain constrained MPLS and GMPLS TE networks. However, both per-domain and BRPC techniques assume that the sequence of domains to be crossed from source to destination is known, either fixed by the network operator or obtained by other means.

[RFC6805] describes a Hierarchical PCE (H-PCE) architecture which can be used for computing end-to-end paths for inter-domain MPLS Traffic Engineering (TE) and GMPLS Label Switched Paths (LSPs) when the domain sequence is not known. Within the Hierarchical PCE (H-PCE) architecture, the Parent PCE (P-PCE) is used to compute a multi-domain path based on the domain connectivity information. A Child PCE (C-PCE) may be responsible for a single domain or multiple domains, it is used to compute the intra-domain path based on its domain topology information.

[RFC8751] state the considerations for stateful PCE(s) in hierarchical PCE architecture. In particular, the behaviour changes and additions to the existing stateful PCE mechanisms (including PCE-initiated LSP set-up and active PCE usage) in the context of networks using the H-PCE architecture.

[RFC8637] examines the applicability of PCE/PCEP to the ACTN framework in detail.

[RFC8283] introduces the architecture for PCE as a central controller as an extension of the architecture described in [RFC4655] and assumes the continued use of PCEP as the protocol used between PCE and PCC. Some related extension to PCEP [RFC9168] and [RFC9050] are also applicable in HIC.

6.1.2. BGP

[I-D.ietf-idr-rfc7752bis] describes a mechanism by which link-state and TE information can be collected from networks and shared with external components using the BGP routing protocol. This is achieved using a new BGP Network Layer Reachability Information (NLRI) encoding format and a new BGP path attribute (BGP-LS attribute) that carries link, node, and prefix parameters and attributes.

BGP-LS is another approach to collect network topology information. It is an extension to BGP for distribution of the network's link-state (LS) topology to external entities, such as the SDN controller. Network's link-state topology consists of nodes and links and a set of attributes. The link-state topology is distributed among the IGP domain. The specific protocol used in an IGP domain could be OSPF [RFC2238] or IS-IS [ISO10589]. Note that, the detailed link-state models of these two protocols are not identical. Therefore, BGP-LS can provide a more abstract topology model that can map the IGP models.

The domain controller acts as a consumer to collect the domain's link-state and TE information via BGP-LS. The domain controller would usually abstract the domain information towards the super-controller and further send it via BGP-LS.

BGP-Flowspec is a solution devised for preventing distributed Denial-of-service (DDoS) attack. BGP-Flowspec distributes specification rules into neighbours. [RFC8955] defines a new BGP NLRI encoding format that can be used to distribute traffic flow specifications. Additionally, it defines two applications of that encoding format: one that can be used to automate inter-domain coordination of traffic filtering, such as what is required in order to mitigate DDoS attacks; and a second application to provide traffic filtering in the context of BGP/MPLS VPN service.

The IP domain controller can act as the traffic sampling node. The super controller can act as the traffic analysis server. When the super controller finds the attack happened, the super controller should distribute the flow rules to associated IP domain controllers. And each IP domain controller should distribute the flow rules into the ingress routers. Additionally one of the actions taken could be "redirect" where flow could be redirected to the TE tunnels created by the controller.

[I-D.luo-grow-bgp-controller-based-ts] describes the traffic steering based on BGP controller. The traditional method for traffic steering depends on the static configuration which is time-consuming and inefficient. With the hierarchical IP controller, the IP domain controller can have the domain network topology view and routing information while the super controller can have the global network topology view and routing information. The super controller can compute the end-to-end paths to satisfy the differentiated service requirement. The IP domain controller may be used to distribute the routing policy into the routers. BGP policy varies in many aspects. Its goal is to meet the customer application and connectivity requirement, and specific service transport needs. So the super BGP controller is responsible for the coordination of multiple domain BGP Policy. And then distribute Policy to the related IP domain

controller. The IP domain controller is responsible for distributing the policy to its network nodes.

[I-D.ietf-idr-rtc-hierarchical-rr] describes the route target (RT) constrain mechanism in the hierarchical route reflection (RR) scenario. [RFC4684] describes the route-target constrain mechanism to build a route distribution graph in order to restrict the propagation of Virtual Private Network (VPN) routes.
[I-D.ietf-idr-rtc-hierarchical-rr] proposes a solution to address the RT constrain issue in the hierarchical RR scenarios. The super controller corresponding to higher level RR can receive the RT-constrain routes from the lower level RR, which is acted by the IP domain controller. The higher level RR will select one of the received routes as the best route. then it should advertise the best route to all the lower level RR to build the route distribution graph. This fits well with the HIC as described in this document.

6.2. Management Plane

6.2.1. YANG Models

This is a non-exhaustive list of possible YANG models developed or in-development that could be used for HIC.

Topology: [RFC8345] defines a generic YANG data model for network topology. [RFC8795] defines a YANG data model for representing, retrieving and manipulating Traffic Engineering (TE) Topologies.

Tunnel: [I-D.ietf-teas-yang-te] defines a YANG data model for the configuration and management of Traffic Engineering (TE) interfaces, tunnels and Label Switched Paths (LSPs).

L3VPN: The Layer 3 service model (L3SM) is defined in [RFC8299], which is a YANG data model that can be used for communication between customers and network operators and to deliver a Layer 3 provider-provisioned VPN service. [I-D.ietf-bess-l3vpn-yang] defines a YANG data model that can be used to configure and manage BGP Layer 3 VPNs at the device. Note that a network configuration model at the Domain Controller level needs to be developed.

L2VPN/EVPN: [RFC8466] defines a YANG data model that can be used to configure a Layer 2 Provider-Provisioned VPN service. This model is intended to be instantiated at the management system to deliver the overall service. [I-D.ietf-bess-l2vpn-yang] and [I-D.ietf-bess-evpn-yang] defines a YANG data model to configure and manage L2VPN and EVPN service respectively. Note that a network configuration model at the Domain Controller level needs to be developed.

OAM: TBD

BGP Policy: [I-D.ietf-idr-bgp-model] defines a YANG data model that can be used to configure BGP Policy based on data centre, carrier and content provider operational requirements. The model is intended to be vendor-neutral, in order to allow operators to manage BGP configuration in heterogeneous environments with routers supplied by multiple vendors. Note that a network configuration model at the Domain Controller level needs to be developed.

BGP Flowspec: [I-D.wu-idr-flowspec-yang-cfg] defines a YANG data model for Flow Specification implementations. The configuration data is described as flow specification rules that can be distributed as BGP NLRI to a network element. The rules can be used to filter Distributed Denial of Service attacks (DDoS) besides other use cases. Note that a network configuration model at the Domain Controller level needs to be developed.

[RFC8969] provides a framework that describes and discusses an architecture for service and network management automation that takes advantage of YANG modeling technologies. This is quite apt for HIC and includes interactions between multiple YANG models as described in [RFC8969].

[Editor's Note - the above list should be extended.]

6.2.2. Protocol Considerations

The Network Configuration Protocol (NETCONF) [RFC6241] provides mechanisms to install, manipulate, and delete the configuration of network devices. The RESTCONF [RFC8040] describes an HTTP-based protocol that provides a programmatic interface for accessing data defined in YANG, using the datastore concepts defined in NETCONF.

Some other mechanism like gRPC/gNMI could also be used between controllers using the same YANG data models.

7. IANA Considerations

There are no IANA concerns in this document.

8. Security Considerations

There are no new security concerns in this document.

9. Acknowledgments

10. Contributing Authors

Dailongfei (Larry)
Huawei Technologies,
Beijing, China

Email: larry.dai@huawei.com

11. References

11.1. Normative References

[RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <https://www.rfc-editor.org/info/rfc8453>.

11.2. Informative References

[I-D.dhodylee-pce-pcep-ls] Dhody, D., Peng, S., Lee, Y., Ceccarelli, D., and A. Wang, "PCEP extensions for Distribution of Link-State and TE Information", Work in Progress, Internet-Draft, draft-dhodylee-pce-pcep-ls-26, 27 August 2023, <https://datatracker.ietf.org/doc/html/draft-dhodylee-pce-pcep-ls-26>.

[I-D.ietf-bess-evpn-yang] Brissette, P., Shah, H. C., Hussain, I., Tiruveedhula, K., and J. Rabadan, "Yang Data Model for EVPN", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-yang-07, 11 March 2019, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-yang-07>.

[I-D.ietf-bess-l2vpn-yang] Shah, H. C., Brissette, P., Chen, H., Hussain, I., Wen, B., and K. Tiruveedhula, "YANG Data Model for MPLS-based L2VPN", Work in Progress, Internet-Draft, draft-ietf-bess-l2vpn-yang-10, 2 July 2019, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-l2vpn-yang-10>.

[I-D.ietf-bess-l3vpn-yang]
Jain, D., Patel, K., Brissette, P., Li, Z., Zhuang, S., Liu, X., Haas, J., Esale, S., and B. Wen, "Yang Data Model for BGP/MPLS L3 VPNs", Work in Progress, Internet-Draft, draft-ietf-bess-l3vpn-yang-05, 13 April 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-l3vpn-yang-05>.

[I-D.ietf-idr-bgp-model]
Jethanandani, M., Patel, K., Hares, S., and J. Haas, "YANG Model for Border Gateway Protocol (BGP-4)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-model-17, 5 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-model-17>.

[I-D.ietf-idr-rfc7752bis]
Talaulikar, K., "Distribution of Link-State and Traffic Engineering Information Using BGP", Work in Progress, Internet-Draft, draft-ietf-idr-rfc7752bis-17, 25 August 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-rfc7752bis-17>.

[I-D.ietf-idr-rtc-hierarchical-rr] Dong, J., Chen, M., and R. Raszuk, "Extensions to RT-Constrain in Hierarchical Route Reflection Scenarios", Work in Progress, Internet-Draft, draft-ietf-idr-rtc-hierarchical-rr-03, 3 July 2017, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-rtc-hierarchical-rr-03>.

[I-D.ietf-mpls-seamless-mpls]
Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", Work in Progress, Internet-Draft, draft-ietf-mpls-seamless-mpls-07, 28 June 2014, <https://datatracker.ietf.org/doc/html/draft-ietf-mpls-seamless-mpls-07>.

[I-D.ietf-pce-stateful-interdomain] Dugeon, O., Meuric, J., Lee, Y., and D. Ceccarelli, "PCEP Extension for Stateful Inter-Domain Tunnels", Work in Progress, Internet-Draft, draft-ietf-pce-stateful-interdomain-03, 4 March 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-stateful-interdomain-03>.

[I-D.ietf-teas-actn-yang] Lee, Y., Zheng, H., Ceccarelli, D., Yoon, B. Y., and S. Belotti, "Applicability of YANG models for Abstraction and Control of Traffic Engineered Networks", Work in Progress, Internet-Draft, draft-ietf-teas-actn-yang-11, 7 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-yang-11>.

[I-D.ietf-teas-te-service-mapping-yang]
Lee, Y., Dhody, D., Fioccola, G., Wu, Q., Ceccarelli, D., and J. Tantsura, "Traffic Engineering (TE) and Service Mapping YANG Data Model", Work in Progress, Internet-Draft, draft-ietf-teas-te-service-mapping-yang-14, 12 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-te-service-mapping-yang-14>.

[I-D.ietf-teas-yang-path-computation] Busi, I., Belotti, S., de Dios, O. G., Sharma, A., and Y. Shi, "A YANG Data Model for requesting path computation", Work in Progress, Internet-Draft, draft-ietf-teas-yang-path-computation-21, 7 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-path-computation-21>.

[I-D.ietf-teas-yang-te] Saad, T., Gandhi, R., Liu, X., Beeram, V. P., and I. Bryskin, "A YANG Data Model for Traffic Engineering Tunnels, Label Switched Paths and Interfaces", Work in Progress, Internet-Draft, draft-ietf-teas-yang-te-34, 1 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-yang-te-34>.

[I-D.luo-grow-bgp-controller-based-ts] Luo, Y., Ou, L., Huang, X., Zhuang, S., and Z. Li, "Traffic Steering Based on BGP Controller", Work in Progress, Internet-Draft, draft-luo-grow-bgp-controller-based-ts-00, 5 March 2018, <https://datatracker.ietf.org/doc/html/draft-luo-grow-bgp-controller-based-ts-00>.

[I-D.wu-idr-flowspec-yang-cfg] Wu, N., Zhuang, S., and A. Choudhary, "A YANG Data Model for Flow Specification", Work in Progress, Internet-Draft, draft-wu-idr-flowspec-yang-cfg-02, 1 October 2015, <https://datatracker.ietf.org/doc/html/draft-wu-idr-flowspec-yang-cfg-02>.

[ISO10589] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, 1992.

[RFC2238] Clouston, B., Ed. and B. Moore, Ed., "Definitions of Managed Objects for HPR using SMIv2", RFC 2238, DOI 10.17487/RFC2238, November 1997, <https://www.rfc-editor.org/info/rfc2238>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <https://www.rfc-editor.org/info/rfc3630>.

[RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4110, DOI 10.17487/RFC4110, July 2005, <https://www.rfc-editor.org/info/rfc4110>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI

10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <https://www.rfc-editor.org/info/rfc4456>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC4664] Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <https://www.rfc-editor.org/info/rfc4664>.

[RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", RFC 4684, DOI 10.17487/RFC4684, November 2006, <https://www.rfc-editor.org/info/rfc4684>.

[RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, DOI 10.17487/RFC5152, February 2008, <https://www.rfc-editor.org/info/rfc5152>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC5441] Vasseur, JP., Ed., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", RFC 5441, DOI 10.17487/RFC5441, April 2009, <https://www.rfc-editor.org/info/rfc5441>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol

(NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <https://www.rfc-editor.org/info/rfc6241>.

[RFC6805] King, D., Ed. and A. Farrel, Ed., "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", RFC 6805, DOI 10.17487/RFC6805, November 2012, <https://www.rfc-editor.org/info/rfc6805>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

[RFC7491] King, D. and A. Farrel, "A PCE-Based Architecture for Application-Based Network Operations", RFC 7491, DOI 10.17487/RFC7491, March 2015, <https://www.rfc-editor.org/info/rfc7491>.

[RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP Monitoring Protocol (BMP)", RFC 7854, DOI 10.17487/RFC7854, June 2016, <https://www.rfc-editor.org/info/rfc7854>.

[RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <https://www.rfc-editor.org/info/rfc7926>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <https://www.rfc-editor.org/info/rfc8040>.

[RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", RFC 8051, DOI 10.17487/RFC8051, January 2017, <https://www.rfc-editor.org/info/rfc8051>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

[RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE

Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/info/rfc8281>.

[RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", RFC 8283, DOI 10.17487/RFC8283, December 2017, <https://www.rfc-editor.org/info/rfc8283>.

[RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8299, DOI 10.17487/RFC8299, January 2018, <https://www.rfc-editor.org/info/rfc8299>.

[RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <https://www.rfc-editor.org/info/rfc8309>.

[RFC8345] Clemm, A., Medved, J., Varga, R., Bahadur, N., Ananthakrishnan, H., and X. Liu, "A YANG Data Model for Network Topologies", RFC 8345, DOI 10.17487/RFC8345, March 2018, <https://www.rfc-editor.org/info/rfc8345>.

[RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <https://www.rfc-editor.org/info/rfc8466>.

[RFC8637] Dhody, D., Lee, Y., and D. Ceccarelli, "Applicability of the Path Computation Element (PCE) to the Abstraction and Control of TE Networks (ACTN)", RFC 8637, DOI 10.17487/RFC8637, July 2019, <https://www.rfc-editor.org/info/rfc8637>.

[RFC8751] Dhody, D., Lee, Y., Ceccarelli, D., Shin, J., and D. King, "Hierarchical Stateful Path Computation Element (PCE)", RFC 8751, DOI 10.17487/RFC8751, March 2020, <https://www.rfc-editor.org/info/rfc8751>.

[RFC8795] Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and O. Gonzalez de Dios, "YANG Data Model for Traffic Engineering (TE) Topologies", RFC 8795, DOI 10.17487/RFC8795, August 2020, <https://www.rfc-editor.org/info/rfc8795>.

[RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <https://www.rfc-editor.org/info/rfc8955>.

[RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <https://www.rfc-editor.org/info/rfc8969>.

[RFC9050] Li, Z., Peng, S., Negi, M., Zhao, Q., and C. Zhou, "Path Computation Element Communication Protocol (PCEP) Procedures and Extensions for Using the PCE as a Central Controller (PCECC) of LSPs", RFC 9050, DOI 10.17487/RFC9050, July 2021, <https://www.rfc-editor.org/info/rfc9050>.

[RFC9168] Dhody, D., Farrel, A., and Z. Li, "Path Computation Element Communication Protocol (PCEP) Extension for Flow Specification", RFC 9168, DOI 10.17487/RFC9168, January 2022, <https://www.rfc-editor.org/info/rfc9168>.

[RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <https://www.rfc-editor.org/info/rfc9182>.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China

Email: lizhenbin@huawei.com

Dhruv Dhody
Huawei
India

Email: dhruv.ietf@gmail.com

Huaimo Chen
Futurewei Technologies
Boston, MA
United States of America

Email: huaimo.chen@futurewei.com