128 TECHNOLOGY

FEATURE BRIEF: SERVICES & TOPOLOGY EXCHANGE PROTOCOL

Introduction

Network operators are under pressure to deliver new services at an ever-increasing pace. Consumers are using more and more services from the network, and the network exists to deliver these services. The network has gone from being a tool for professors to share research to a necessary utility to deliver services globally. Whether it is to connect with friends via social media, learn recipes at the kitchen counter, or to browse libraries for research information, the underlying motive to use the network is to consume services. Network operators need efficient ways to deliver new services and innovative methods for the network to learn and connect these services to consumers. Networks must be service-oriented to support today's (and tomorrow's) networking requirements.

Common Routing Protocols

Many routing protocols have been developed over the course of the last three decades to learn of available routes that exist in the network, build routing tables, and make routing decisions. Some of the most common routing protocols are EIGRP, OSPF, IS-IS, and BGP. All routing protocols have strengths and weaknesses. They convey connectivity information well but have no knowledge of individual services or requirements associated with these services, or path decisions to make based on these requirements for each service. Without this information the network makes decisions once and continues to route packets for all services without being able to provide any specialized treatment to any specific service.

Service Discovery Protocols

There are service discovery protocols that enable automatic detection of services such as DNS-SD, DHCP, SLP, etc. They enable service discovery but are unable to establish path preference based on current network conditions. There are numerous other solutions developed that utilize declarative management models to control and provide declaration of intent regarding behaviors such as OpFlex. Network Intent Composition (NIC), OpenStack Congress, Open Daylight Group Based Policy (ODL-GBP), etc. These projects/solutions are able to convey application intent. However, they lack the ability to convey current network state required to deliver those applications. Routers now require both sets of information to make informed, real-time decisions.

Services and Topology Exchange Protocol (STEP)

The Services and Topology Exchange Protocol (STEP) provides 128T Session Smart™ Routers with the ability to describe connectivity between them to exchange services and reachability to those services. The 128T Session Smart Routers use a service-oriented, session-based paradigm. Network administrators define services to represent capabilities that the network is designed to deliver to consumers. STEP enables the exchange of this service capability to all 128T Session Smart Routers in the network along with reachability and other parameters to connect to these services. This enables network administrators to build an application-oriented, intent-based network that follows business logic and uses real-time information to enable 128T Session Smart Routers to decide how to connect applications.

STEP Overview

The Services and Topology Exchange Protocol (STEP) utilizes a centralized high-speed database that shares incremental updates when needed with local devices. This is known as the 128T Repository. The 128T Session Smart™ Router's available services are published into a 128T Repository via neighborhoods (covered shortly), along with connectivity information between 128T Session Smart Routers in a network.

The 128T Repository is the datastore for the information that gets exchanged between the 128T Session Smart Routers, and is comprised of many discrete repositories — each of which contains a subset of the overall 128T Repository. This also includes two different types of data: topology and service data.

When the consumer connected to a 128T Session Smart Router wants to access a service offered by another 128T Session Smart Router, the provider of the service grants access to the 128T Repository, thereby describing the method for reaching that service. The reachability is computed using a set of one or more neighborhoods. Two routers that both have access to the same named neighborhood are presumed to be mutually reachable via a common L3 network.

STEP Design Goals

Numerous methods of group policy dissemination have led to several theories and implementations. Unfortunately, this has also led to many disparate, one-dimensional systems. The design goals of STEP are based on the following principles:

Automated Exchange: 128T Session Smart™ Routers must be able to publish service, connectivity, and other related information to a single 128T Repository for sharing information. 128T Session Smart Routers must be able to receive notifications from the 128T Repository when there are relevant changes.

The exchanges should only include changes relevant to the localized network rather than sharing unnecessary information everywhere. To minimize the size and frequency of these exchanges, only relevant differences from previous updates are shared.

High-Scale and Performance: The single 128T Repository must be able to support a large scale of 128T Session Smart Routers. The exchange of messages and information should be instantaneous to enable real-time decision making.

Support for Next-Generation Applications: STEP must be flexible and extensible to support next-generation applications. The exchange of information should be via the use of human-readable JSON files to make the messages easy to interpret.

Robust and Inherently Secure: 128T Session Smart Routers should be able to make all decisions independently without requiring constant connectivity to the 128T Repository. The network should be robust and able to reestablish connections and resume updates after losing connections. All exchanges of messages between 128T Session Smart Routers and the 128T Repository must be encrypted and authenticated to ensure complete security.

Extensible to New Definitions: The 128T Repository may be used in the future for other purposes like exchanging multicast group memberships, application identification dictionaries, etc. to ensure that the same infrastructure can be reused for multiple use cases, rather than having different siloed solutions for every vertical.

Neighborhoods

The 128T Session Smart™ Router uses neighborhoods to indicate network connectivity. Each connected interface on a 128T Session Smart Router is assigned a neighborhood identifier, which is exchanged via the 128T Repository to allow participating 128T Session Smart Routers to understand the network topology.

Neighborhoods are simple text labels with the following properties:

- All interfaces belonging to the same neighborhood are intended to be mutually reachable.
- Neighbors are not necessarily connected on the same Ethernet. As with I-BGP, neighbors need not be directly adjacent and there may be a dependence on an IGP (or EGP) in order for a 128T Session Smart Router to reach a peer within a group of routers under a single administrative domain (also known as Authority).

Interfaces which are not explicitly configured with a neighborhood identifier are automatically assigned a neighborhood label, unique to that interface.

Each router within an Authority publishes the neighborhoods associated with all of its interfaces to the Authority's 128T Repository. For each neighborhood, the 128T Session Smart Router includes its waypoint address(es), tenant assignment (if applicable), and other relevant information. This set of information allows consumers of the 128T Repository to create a logical graph of the connectivity within the Authority.

128T Repository

The 128T Repository contains various data about an Authority. Its services that it wishes to expose as well as the connectivity map of its Authority by way of a series of named neighborhoods. The 128T Repository must be reachable via a domain name, or a publicly accessible IP address, so that other 128T Session Smart™ Routers in the same Authority can reach it easily. The 128T Repository can be provisioned directly to influence the dynamically learned information. The 128T Repository serves as a central location to coordinate information exchange among 128T Repositories on the 128T Session Smart Routers.

Topology

The topology information contained in the 128T Repository contains the following information:

- The routers (by name) that are in the Authority.
- For each router, its capabilities such as software version.
- For each interface on a router, the neighborhood assigned to that interface.
- For each interface on a router, a waypoint address used to reach it.
- For each interface/waypoint address, a security policy (public key).
- GPS coordinates for a router.

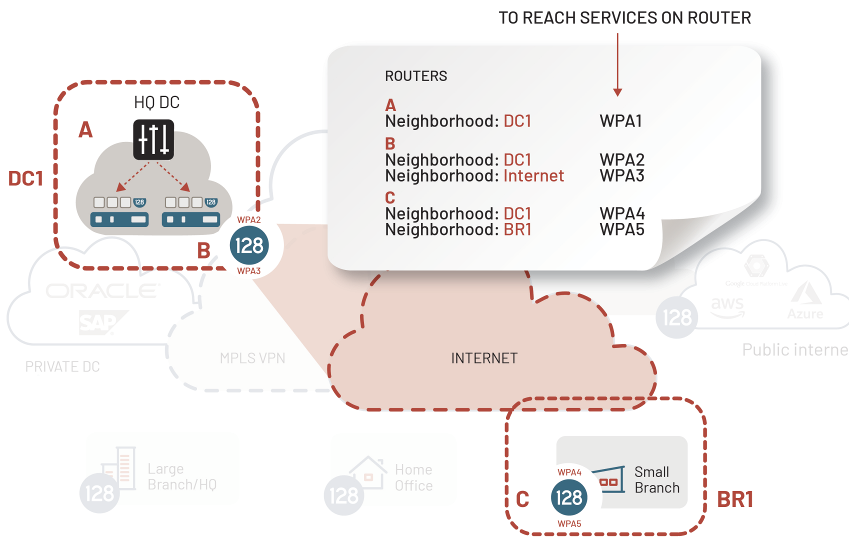

The contents of this 128T Repository describe the entire topology within an Authority, allowing 128T Session Smart™ Routers within that Authority to use the data to build a graph of the connectivity between themselves (i.e., which 128T Session Smart Routers share common neighborhoods and therefore are intended on being topologically connected), as well as describing the waypoints that 128T Session Smart Routers need to use when communicating with their neighboring routers. All routers within an Authority have unfiltered access to the Authority's topology information.

Service

A provider of a service wishing to publish their services, either within their own Authority to other 128T Session Smart™ Routers or to external Authorities, does so by way of the 128T Repository.

The data within the 128T Repository is supplied by the 128T Conductor (the source for all configuration information within an Authority), and includes:

- Each service within the Authority and all of each service's attributes (as defined by the 128T data model).
- Each tenant and its pertinent attributes.
- The security policies referenced in services and tenants (i.e., their public keys).

Each 128T Session Smart Router within the Authority adds service route(s) that are active and usable, along with the associated waypoint(s)/neighborhood(s) into the 128T Repository.

When a 128T Session Smart Router's service-route becomes active, it publishes reachability into the 128T Repository by way of the neighborhood(s) of the interface on which the service-route has become available.

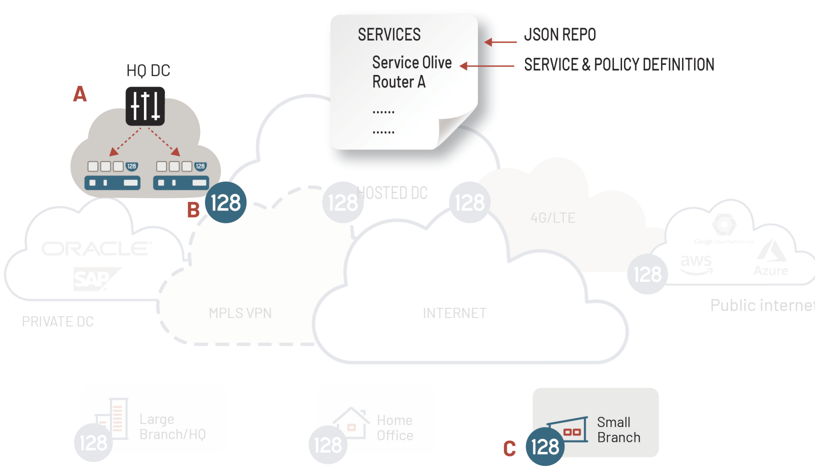

Exchange

The 128T Session Smart™ Routers update their local 128T Repository when there are any changes. These are then synced with the 128T Conductor. The 128T Registry on the 128T Conductor then passes on any deltas with updated information to other 128T Session Smart Routers in the Authority.

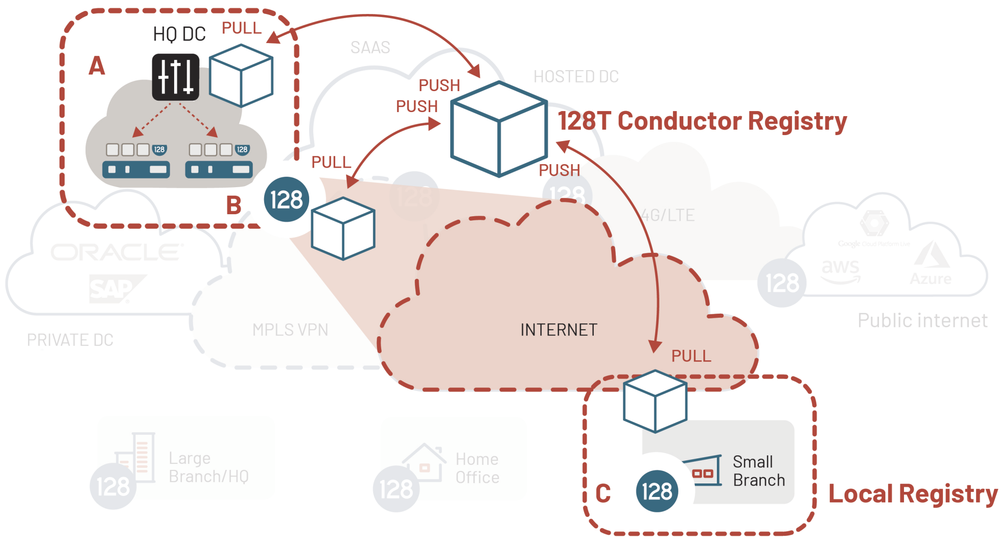

The 128T Session Smart Routers can subscribe to the 128T Repository at the 128T Conductor to get notifications on changes. They can also poll the 128T Repository at the 128T Conductor to get information if required (for example when the a new 128T Session Smart Router is installed and wants to learn connectivity and service information).

The 128T Session Smart Routers build a network connectivity graph based on the information they learn from the 128T Repository. This information is used to reach services and send data along paths efficiently.

Extensions

STEP can be extended to include other relevant information that can assist the 128T Session Smart™ Routers to make intelligent decisions. The 128T Repository can be used to share information about multicast groups and the location of senders and receivers (or which 128T Session Smart Routers are interested in membership of specific multicast groups). This will enable the 128T Session Smart Routers in the Authority to build an efficient multicast distribution tree without reliance on PIM, BIER, LSM, or other complex multicast tree building and distribution protocols.

The 128T Repository can be used to share near real-time information on service loads and usage. This can enable 128T Session Smart Routers to load balance sessions to services efficiently. The 128T Repository can also share information on path conditions, which can be used to ensure that sessions are sent on paths that meet the required SLA. The 128T Session Smart Routers can act as path computations engines (PCE) without relying on an external controller for decision making.

Benefits

The Services and Topology Exchange Protocol (STEP) provides the following benefits to any enterprise building their 128T network:

Service Discovery: STEP provides a simple way to provision and advertise services to routers. Dynamic service discovery enables enterprises to spin up and down new locations for existing services based on loads, add new services, and remove or modify existing services. This reduces time to market and helps grow elastically.

Multi-hop Path Selection: STEP enables routers to learn the current real-time status of the entire network. This allows the 128T Session Smart™ Routers to choose and keep application traffic on paths that meet the demands of that application. This results in superior end-user experience.

Analytics and Heuristics: STEP provides service and network information in the 128T Repository. This can be used to generate heat maps to pinpoint faults and congestion. With network state information over time, big data analytics can predict congestions and outages ahead of time. This reduces time to maintain and debug. Proactive fault prediction enables a self-healing network.

Security and Trust: STEP enables trust through the use of blockchains to share keys and validate ownership. Each Authority in the 128T network is allowed to access services and prefixes that are assigned to it. This reduces human errors and malicious activity.

A dynamic and self-optimizing network enables enterprises to get the network to do what their business needs.

Summary

The Services and Topology Exchange Protocol enables 128T Session Smart™ Routers to dynamically learn connectivity and service information within an Authority. It is simple, efficient, secure, and reliable. It can be extended to include other relevant information that can enable intelligent decision making within the network. Most common protocols developed either rely on complex exchanges to learn connectivity and have no awareness of services. Policy dissemination protocols on the other hand only provide declaration of policies relying on controllers to make decisions. STEP provides a solution that enables true intent-based networking by allowing routers to learn business intent (services) and the ability to connect consumers to business services efficiently.