Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 17, 2014

M. Scharf, Ed.
V. Gurbani, Ed.
G. Soprovich
V. Hilt
Alcatel-Lucent

February 13, 2014

The Virtual Private Network (VPN) Service in ALTO: Use Cases, Requirements and Extensions
draft-scharf-alto-vpn-service-02

Abstract

The Application-Layer Traffic Optimization (ALTO) protocol is designed to allow entities with knowledge about the network infrastructure to export such information to applications that need to choose one or more resources from a candidate set. This document provides motivation for using ALTO in a Virtual Private Network (VPN) environment. We discuss use cases, requirements, and possible extensions to the base ALTO protocol that will be needed to support VPN services.

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 http://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 August 17, 2014.

Table of Contents

1. Overview 2
2. Terminology 4
3. Encompassing example 4
3.1. A VPN scenario 4
3.2. Exemplary use of ALTO 6
4. Use cases 9
4.1. Use case 1: Application guidance in an L3VPN 10
4.2. Use case 2: Application guidance in an L2VPN 11
4.3. Use case 3: VPN guidance without addresses 12
4.4. Use case 4: Extending the VPN 13
4.5. Use case 5: Shrinking the VPN 14
4.6. Use case 6: VPN selection 14
4.7. Use case 7: Other use cases 14
5. Requirements and potential solutions 15
5.1. Requirements 15
5.2. Potential Solutions 16
6. Security considerations 17
7. IANA considerations 17
8. References 17
8.1. Normative References 18
8.2. Informative References 18
Appendix A. Acknowledgements 18
Authors' Addresses 19

1. Overview

Virtual Private Network (VPN) technology is widely used in public and private networks to create groups of users that are separated from other users of the network and allows these users to communicate among them as if they were on a private network. According to [RFC4364], the generic term "Virtual Private Network" is used to refer to a specific set of sites as either an intranet or an extranet that have been configured to allow communication. A site is a member of at least one VPN and may be a member of many.

Service providers offer different types of VPNs. [RFC4026] distinguishes between Layer 2 VPN (L2VPN) and Layer 3 VPN (L3VPN) using different sub-types. Virtual Private LAN Service (VPLS) is an L2VPN provider service that emulates the full functionality of a

traditional Local Area Network (LAN) [RFC4762]. A VPLS makes it possible to interconnect several LAN segments over a packet switched network.

Another solution is an L3VPN, which interconnects sets of hosts and routers based on Layer 3 addresses. In this context, a virtual private network is defined in [RFC4364] as follows:

> Consider a set of "sites" that are attached to a common network that we call "the backbone". Now apply some policy to create a number of subsets of that set, and impose the following rule: two sites may have IP interconnectivity over that backbone only if at least one of these subsets contains them both.
>
> These subsets are Virtual Private Networks (VPNs). Two sites have IP connectivity over the common backbone only if there is some VPN that contains them both. Two sites that have no VPN in common have no connectivity over that backbone.

VPNs can also include "pseudo L1/L2" connectivity, such as pseudowire emulation (PWE) carrying legacy TDM or ATM circuits for point to point connectivity. Further examples are integrated optical solutions delivering light paths or integrated optical and Ethernet transport. It is instructive to note that point-to-point VPN services of this type rarely carry VPN edge addresses within the network; e.g., packets are encapsulated and transported without any kind of address facing the customer drop side of the network.

A VPN may also include mechanisms to enhance the level of separation (e.g., by end-to-end encryption), but the discussion of such mechanisms is outside the scope of this document. In the following, the term "VPN" is used to refer to provider supplied virtual private networking.

The ALTO protocol [I-D.ietf-alto-protocol] is designed to provide network information (e.g., basic network location structure, preferences of network paths) with the goal of modifying network resource consumption patterns while maintaining or improving application performance. The most important use case is providing application guidance in the global Internet, so that applications do not have to perform excessive measurements on their own. For the very same reason, topology exposure is also very useful in VPNs. But the constraints for using ALTO in L3VPNs or L2VPNs differ from the public Internet. This document presents these use cases and discusses requirements and extensions to the base ALTO protocol that will be needed to realize the VPN Service in ALTO.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Encompassing example

3.1. A VPN scenario

Below, we present an example for a VPN scenario that describes an environment for an ALTO VPN Service. This scenario is subsequently used to analyze specific use cases.

We consider the following: there are two distinct entities, one, the network service provider (NSP) who owns the network and offers a VPN to the second entity, the customer, who has premises in four different locations that shall be interconnected by that VPN. The sites could be office branches, data centers, etc. Throughout this document, we assume the following four sites:

- Site 1

  Location name: SITE-CHICAGO

  Geography Degree: 41.85 N, 87.65 W

- Site 2

  Location name: SITE-OTTAWA

  Geography Degree: 45.24 N, 75.43 W

- Site 3

  Location name: SITE-SANFRANCISCO

  Geography Degree: 37.75 N, 122.28 W

- Site 4

  Location name: SITE-PARIS

  Geography Degree: 48.86 N, 2.35 E

It is assumed that these sites are interconnected by a VPN that may be identified by the hypothetical name "vpn42". This document

specifically considers two different VPN types for the interconnection:

- L3VPN: The local area networks at each site will have a certain IP subnet ranges, for instance 10.0.1.0/24 at site 1, 10.0.2.0/24 at site 2, etc.

- L2VPN: All sites form part for a flat sub-IP network, e.g. a logical Ethernet segment. Different to a local network, the network potentially interconnects geographically remote sites.

The VPN will not necessarily be static. The customer could possibly modify the VPN and add new VPN sites, e. g., to handle peak-load demand or to consolidate VPN sites to account for reduced traffic. The service provider could offer a Web portal or other Operation Support Systems (OSS) solutions that allow the customer to grow or consolidate the VPN. Details on how the customer can configure VPNs are outside the scope of this document.

Furthermore, we assume that the customer is running at least one application that can benefit from application-level traffic optimization, e.g., using application-internal routing mechanisms or placement functions. For instance, typical uses cases for VPN customers could be:

- Enterprise application optimization: Enterprise customers often run distributed applications that exchange large amounts of data, e.g., for synchronization of replicated data bases. Both for placement of replicas as well as for the scheduling of transfers insight into network topology information could be useful.

- Private cloud computing solution: An enterprise customer could run own data centers at the four sites. The cloud management system could want to understand the network costs between different sites for intelligent routing and placement decisions of Virtual Machines (VMs) among the VPN sites.

- Cloud-bursting: One or more VPN endpoints could be located in a public cloud. If an enterprise customer needs additional resources, they could be provided by a public cloud, which is accessed through the VPN. Network topology awareness would help to decide in which data center of the public cloud those resources should be allocated.

These examples focus on enterprise customers of NSPs, which are typical users of provider-supplied VPNs. Such VPN customers typically have no insight into the network topology that transports the VPN. For instance, the actual delay between two VPN sites may

significantly depend on the routing in the NSP MPLS/IP network. If better-than-random decisions are required, applications have to rely on own measurements. An alternative would by guidance by an ALTO server offered by the NSP.

It is important to emphasize that other scenarios and use cases exist and the examples enunciated so far are merely used to illustrate how ALTO can be used in a VPN context. A common characteristic of these use cases is that applications will not necessarily run in the public Internet, and they will typically not be accessed by residential customers. The internal use of ALTO by a specific application is not considered in this document.

3.2. Exemplary use of ALTO

In the example VPN described in the previous section, it would be beneficial if an ALTO server would expose cost maps or provide a ranking service that represents the costs between different sites, e. g., endpoints of the VPN. Similar to existing use cases of ALTO, this enables an application integrating an ALTO client to use this information for application-level traffic optimization. This results in the following scenario:

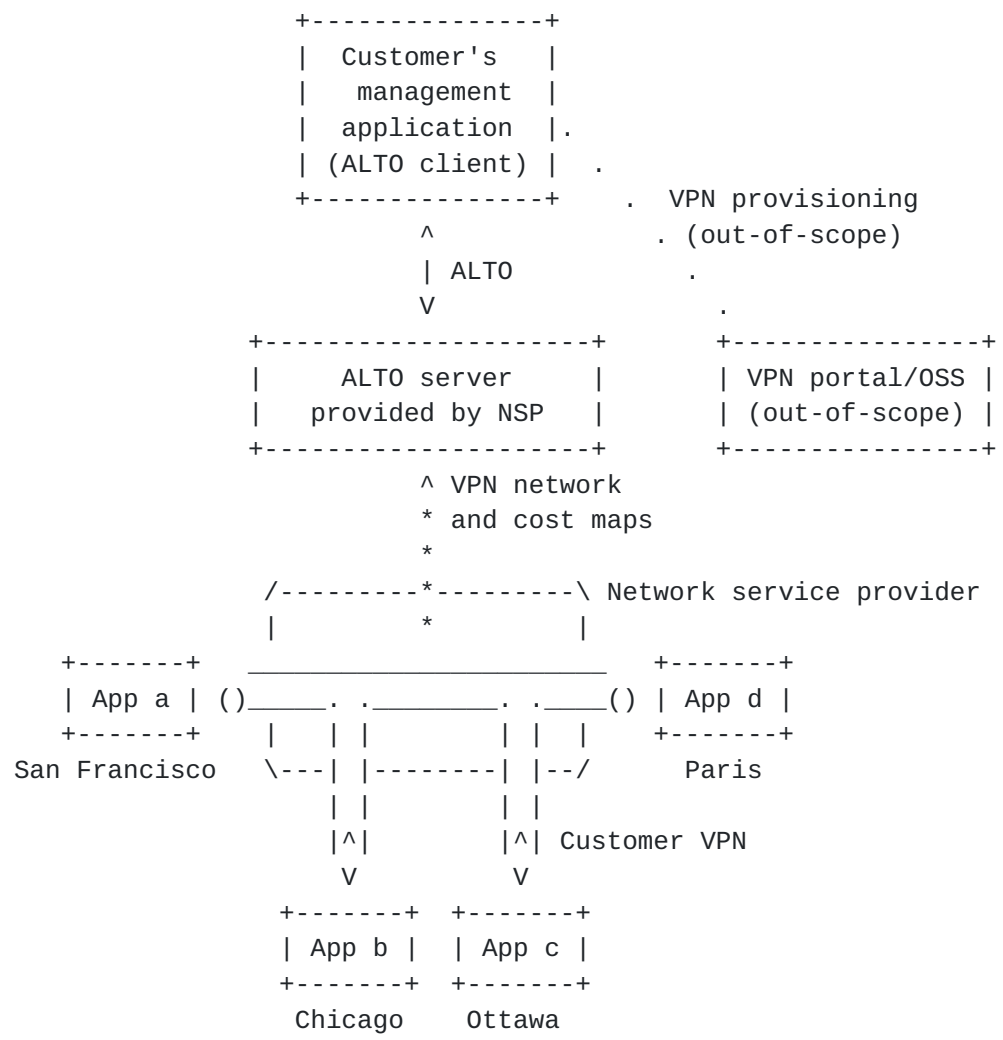

Figure 1: Overview of an ALTO usage scenario

The network service provider could operate an ALTO server. An ALTO client in an application could then retrieve an ALTO cost map by querying a corresponding URI, such as:

uri: http://alto.nsp.org/vpn42/costmap

The NSP can assign PIDs to each of the VPN endpoints; this renders computations at the ALTO server to fit in the current model of using the protocol. A corresponding example would be:

Site 1: PID "pid14"

Site 2: PID "pid21"

Site 3: PID "pid11"

Site 4: PID "pid27"

The example below further expands on the VPN by demonstrating the resulting network topology provided to an ALTO server. The picture corresponds to the VPN of the customer and also includes the Provider Edge (PE) routers and Customer Edge (CE) devices:

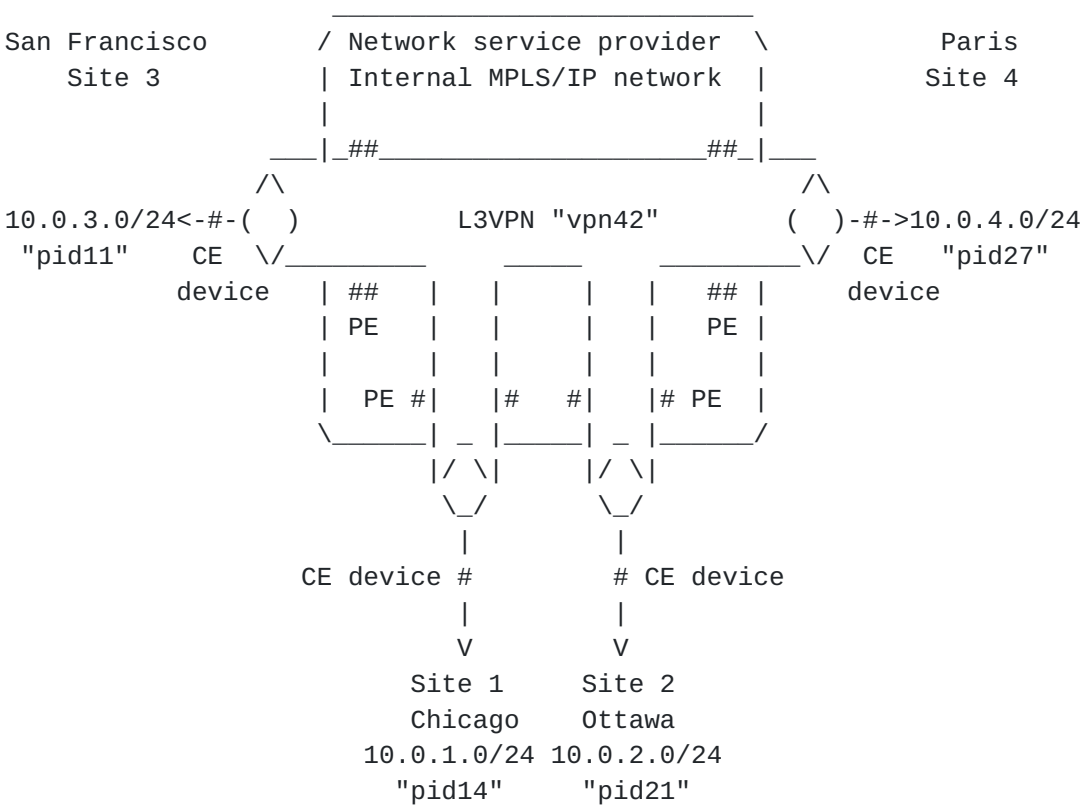

Figure 2: Example for mapping of VPN sites to ALTO PIDs

The costs exposed by the ALTO server can be based on routing costs inside the service provider network or other network topology information, such as delay measurements, traffic engineering (TE) data, etc. As with other use cases of ALTO, the costs can reflect the service provider's preference and policies regarding communication between the involved VPN endpoints.

Generally, two different types of applications can consume the information provided by the ALTO server. The first class can be composed of discrete application instances executing at the various sites that are interconnected by the VPN. ALTO is used to optimize the routing or resource consumption among those application instances. A typical examples is a distributed database, i.e., an enterprise backend system. In Figure 1, these application instances are referred to as "App a", "App b", "App c", and "App d". Generally

speaking, this usage mirrors the canonical use of ALTO in unstructured P2P networks or Content Delivery Networks (CDN) networks whereby a rendezvous is desired between a consumer and a plurality of producers. In this document, we label this class of applications by the term "user applications".

The second class represents management applications that typically work on VPN level. In addition to consulting an ALTO server provided by the NSP, this type of application possibly has its own understanding of what resources are available at different sites, and it could possibly even trigger more complex actions such a building out VPNs, e. g., by contacting a VPN portal of the NSP. In Figure 1, as well as the rest of this document, uses the term "management application" for this use case. An example would be an orchestration solution for cloud computing resources. It could use the topology and cost maps illustrated in Figure 2 to control VPN placement. In principle, management applications have some similarity to centralized resource directories in P2P networks (e. g., trackers), which are an important existing use case for ALTO. Yet, unlike resource directories in the Internet, a VPN typically interconnects mainly sites within one administrative domain.

There may also be an overlap between both types of applications. Furthermore, in particular for the first class of applications, the customer could run an own ALTO server, which could expose topology map and cost maps with further details only visible to the VPN customer (e.g., network segments behind NATs). Since such information is independent of the use of a VPN and typically not known to an NSP, these usage scenarios are not further detailed in this memo.

4. Use cases

Current VPNs provide no clear mechanism to convey information about the network infrastructure to management or user applications using that VPN, e.g. regarding preferences or topological properties of the network service provider network. Applications thus have to rely on other mechanisms such as local measurements to optimize their traffic. The ALTO protocol has been designed to overcome such limitations in the Internet. ALTO, being a well-established, generic, and flexible protocol, can be used in VPNs, too.

We now present various use cases that exhibit the utility of considering a VPN extension of ALTO. Through a series of use cases we demonstrate how a VPN customer and the NSP can use ALTO; we also highlight similarities and differences when using ALTO in the general Internet.

4.1. Use case 1: Application guidance in an L3VPN

The NSP providing the L3VPN service can offer an ALTO server that exposes network and cost information to applications running traffic over that VPN. Since an L3VPN is IP-based, this use case is in principle similar to the use cases already addressed by the ALTO base protocol.

Example 1: Consider the customer in Section 3.1 that has four VPN sites. A user application in one site (say Site 1) would now like to find out which of the other sites (Site 2 to Site 4) are topologically close to Site 1, perhaps to determine where to replicate a certain data set. A corresponding ALTO query would return the costs between those sites. The user application could then select a host in the corresponding subnetwork and connect to that endpoint.

Example 2: In addition to network proximity information, the user application could also be interested in guidance regarding network parameters that cannot be measured directly. For instance, a relevant parameter for a VPN site could be the level of redundancy for that VPN site, e.g., whether there is resilience by network protection schemes in the NSP network.

Example 3: It is quite common for VPN Customer Equipment (CE) to be multi-homed at the Provider Edge (PE). A CE may well home into to several PE routers and thus may have different network cost functions. For instance, assume that in Site 1 the CE will peer to a local PE1 and remote PE2. The cost to reach Site 2 in the VPN could be 1575 for PE1 and 2250 for PE2. The CE will thus choose to steer traffic from Site 1 to Site 2 toward PE1. While the realization of such traffic steering is outside the scope of this document, CE multi-homing places an explicit need to expose more than one set of network costs for a VPN endpoint.

In principle, the existing ALTO services such as network and cost map can provide such guidance. However, it is important to note that a VPN might not run in a public environment. The IP address ranges inside a VPN might not be globally unique or routable. Furthermore, a provider based VPN service normally maintains a strict separation between service provider addressing (such as addresses or Provider Edge routers) and customer addressing. As a result, an ALTO server will not expose the internal IP addressing of the network service provider, making it difficult to identify services using IP addresses in general. In a BGP L3VPN, the VPRN BGP Route Distinguisher could possibly be used as a service identifier, but it is unclear whether an application of a customer or the ALTO client will indeed know such network-internal information of the NSP and whether the NSP would

want to expose it. Also, it would make sense to define an ALTO VPN extension independent of a specific VPN technology.

The network costs in a VPN depends on VPN topology, which needs to be taken into consideration when calculating ALTO information. Given that VPNs are often offered by a single network service provider, ALTO cost information could include information that may be available for a single autonomous system, but difficult to gather in the Internet as a whole. Examples would be the provisioned bandwidth, network-internal latencies, or the path resilience. In a static VPN environment e.g. with a reserved resources in an MPLS/IP wide area network, these costs can be assumed to be rather stable and e. g. reflect the reserved bandwidth between VPN sites. For an application it is simpler and less intrusive to obtain such information about the VPN from the network instead of performing measurements, which would possibly require special probe instances at the different VPN sites (e. g., data centers). But as the encoding of such costs in ALTO is independent of the usage of a VPN, this document does not mandate any specific way how to build ALTO cost maps.

This memo does not argue that ALTO shall be used as a generic data center information exchange protocol. For instance, a general data center resource information model has been suggested in [I-D.lee-alto-ext-dc-resource]. According to that model, the ALTO server also includes data-center information not related at all to the network, such as compute resources, memory, power consumption, etc. While VPNs are an important technology to interconnect data centers, the ALTO VPN service solely focuses on networking cost.

4.2. Use case 2: Application guidance in an L2VPN

The use case outlined in Example 1 also exists for L2VPNs, which are an important technology to transparently interconnect different LAN segments of enterprise users. Again, applications could benefit from getting insight into topological properties of the wide area network providing the L2VPN service, in order to avoid the overhead of own measurements.

Example 4: The user application described in Section 3.1 again wants to find out how well connected (topologically close) Site 1 is to Site 2, 3, or 4. Different from the previous example, all sites are now part of the same Layer 2 subnet. Another example for an application that would benefit from ALTO is a cloud management system. Such a management application could be interested in finding out whether migrating of a Virtual Machine from Site 1 to another site would improve performance, perhaps due to better connectivity or lower latency.

While this use case is in principle similar to the previous one, there is a major difference regarding addressing: Unlike the L3VPN, an L2VPN is not necessarily IP-based; it may use MAC addresses instead of IP addresses. While IP addresses can be aggregated easily and represented succinctly using CIDR notation, MAC addresses do not lend themselves to such aggregation and representation. Furthermore, MAC addresses are not useful to applications themselves. And finally, MAC addresses may not readily be known and available to an ALTO server of the network service provider. And even if they are, an ALTO map using MAC addresses will be very large. In summary, use of MAC addresses is not scalable and nor does it denote any hierarchy that can be used for aggregation. Some other means of identifying services and hosts will be required when using ALTO in L2VPNs.

4.3. Use case 3: VPN guidance without addresses

The VPN interconnects different sites through the network service provider's network. An application might be interested in getting topology information among those sites without knowing actual addresses or identifiers used internally by the VPN. In fact, a VPN site may not even have an address known or visible to applications, e.g., a pseudo-wire VPN.

Example 5: A management application might ask for all VPN sites (i. e., corresponding PIDs) that have a delay less than 40ms or a routing cost less than 55, from VPN Site 1. A specific example for such an application might be cloud management system that uses application-level traffic optimization mechanisms. In the scenario introduced in Section 3.1, such an application may have a-priori information, learned from e.g. a VPN portal, about the VPN type and/or VPN identifiers ("vpn42") as well as some unique site identifier such as "SITE-CHICAGO" but no network addresses. The query could also be more complex or include constraints, e. g., limited to a particular TE class. Note that the ATLO protocol does not necessarily have to support the query constraint itself; if corresponding maps are available, the application can analyze the data itself.

Example 6: In absence of well-known existing network identifiers, a management application might want to query for VPN sites based on yet other attributes, such as geographical distance. For example, an application might want to find all the VPN sites (i. e., corresponding PIDs) within 50 KM of 45.35N, 75.92W. Such geographic queries would be typical of policies bounding delay by geographic distance or administrative and legal requirements.

Such application guidance is obviously similar to existing use of the ALTO cost map or ranking services except that the queries are not based on network addresses.

4.4. Use case 4: Extending the VPN

The customer can possibly grow the VPN to include new sites that are connected at a later time to the VPN. The actual mechanisms for VPN reconfiguration are outside the scope of this document.

Example 7: A management application could be interested in guidance for VPN sites that are currently not part of the VPN, but that would be available e. g. to increase capacity or geographic coverage. Assume that two sites Site 1 and Site 2 are already connected to the VPN. Some time later, scale-out to a third site is required, and the application has to decide whether Site 3 or 4 is better suited for a new application instance. This is an realistic example for a cloud management system that is geographically distributed. Such a system would then have to decide whether Site 3 or Site 4 is topologically closer to the existing VPN endpoints, in order to determine the best location from the network point of view. An ALTO server could provide guidance on the offnet distance of Sites 3 and 4 to the existing VPN sites.

Apparently, the question whether to actually extend the VPN in a specific way may also include decisions outside the scope of ALTO, such as price information or other commercial or legal policies. The actual VPN re-configuration and attachment of a new site to the VPN topology requires back-office interaction and provisioning actions by separate, orthogonal mechanisms such as a Path Computation Element (PCE). Actual path setup by a PCE is independent from the selection of a suitable target site. But it makes sense to use the well-established ALTO methods in order to get at an early stage network proximity information as input information for the selection and configuration process. Applications typically cannot measure the network performance to destinations not already part of the VPN.

For a network service provider, customer guidance for VPN extension by ALTO offers a new possibility to optimize its internal traffic engineering. For instance, an operator could recommend to customers not to connect to a destination operating in protection mode, e.g., after a fiber cut, because in such a case the network may have less sparse resources. Note that a customer is not able to measure such constraints. ALTO is a simple interface to expose such information to applications.

From an ALTO perspective, growing VPN sites possibly results in different types of endpoints, some of which may exist a-priory but not be reachable within the VPN. They could possibly be understood as "shadow" PIDs that become active once the VPN is extended. Once the VPN is modified, new endpoints or PIDs may be created, i. e., the

ALTO network and cost maps may have to be updated accordingly after the VPN is re-configured.

4.5. Use case 5: Shrinking the VPN

Much like a VPN may grow dynamically, it can also shrink when the resources in the VPN are underutilized. Instead of keeping the underutilized resources alive, the VPN operator my decide to consolidate the resources and remove sites from the VPN.

Example 8: Once again, consider the customer in Section 3.1 that has four VPN sites. Based on low resource demand, the management application may wonder whether Site 1 (Chicago) and Site 2 (Ottawa) can be consolidated, e. g., by moving resources between them. One important constraint for such a decision could network proximity information. After such a consolidation, the VPN network and cost maps will be updated to reflect the new topology.

From an ALTO server perspective, this use case is similar to a general application guidance. Yet, there could be a benefit for the service provider to provide special guidance regarding removal of VPN endpoints if there is a benefit for its internal traffic engineering (e. g., consolidation of network resources used by several VPN customers).

4.6. Use case 6: VPN selection

In a more advanced use case, ALTO could also be a selection function to choose VPNs based on network cost criteria.

Example 9: In a multi-homing environment, ALTO could be used to select one VPN out of several candidates to reach a certain destination, taking into account smaller costs, e. g., according to distance or to preferences of the network service provider network.

This use case differs from the previous examples since more than one VPN is involved, i. e., the ALTO guidance is not used to perform application-layer traffic optimization within one VPN, but instead across different VPNs.

4.7. Use case 7: Other use cases

The aforementioned use cases could be complemented by other use of ALTO information. For instance, if applications using the VPN are flexible regarding the timing of data transfers, an ALTO server could provide guidance when and how to schedule such data transfers, possibly with time-shift enhancements. This scenario is further detailed in [I-D.randriamasy-alto-cost-schedule].

5. Requirements and potential solutions

5.1. Requirements

Based on the scenarios listed in Section 4, several requirements can be derived for a VPN Service in ALTO:

REQ 1: The existing ALTO protocol and RESTful interface should be used as far as possible to enable an NSP to expose properties of a VPN.

REQ 2: A VPN Service must not require that network service provider expose internal addressing, such as internal addresses or loopback addresses of the Provider Edge (PE) routers.

REQ 3: A VPN Service must use the PID concept of the base ALTO protocol as far as possible, i. e., the VPNs and network entities in the VPNs can be identified by PIDs. This permits use of the existing ALTO services such as the map service for VPNs, as well as the inherent topology abstraction provided by ALTO.

REQ 4: A VPN Service must be possible for different VPN types, i. e., it must not be limited to L3VPNs only.

REQ 5: The VPN Service must support use cases where IP addresses are not the only form of endpoint identification.

REQ 6: If IP addresses are used, a VPN Service must not assume that IP address are globally routable or unique.

REQ 7: A VPN Service should include certain attributes that are unique to a VPN and that are not represented by the current set of attributes in the base ALTO protocol. Examples include location name of a site, geography coordinates (degree/digital), role, redundancy, default policy, or geography restriction.

REQ 8: The PID must be selectable using standard ALTO filtering. A standard interface query should allow finding resources using, say, the location name attribute or the geography attributes.

REQ 9: The PID should be selectable using a filter that computes matching sites within a certain distance of a particular geographic coordinate based on latitude and longitude, in case that no other address information is known in advance.

REQ 10: Incremental build out (as well as the shrinking) of resources that are part of the VPN must be supported, i. e., the ALTO VPN

service should also be able to expose information about new sites to be attached to the VPN, or provide guidance for removal of sites.

REQ 11: A VPN Service requires that an ALTO server can report the (expected) connectivity between VPN sites regarding different metrics, including for example geographical distance, delay, or provisioned bandwidth.

REQ 12: Information about a VPN must only be exposed to authorized users of that VPN.

5.2. Potential Solutions

In the following we analyze how the requirements of a VPN Service can be addressed by ALTO extensions.

REQ 1: This is an inherent, general requirement for any new use or extension of ALTO.

REQ 2: This requirement can be supported in ALTO today, because it is left to the service provider which information to expose e.g. in ALTO cost maps.

REQ 3: The PID concept itself is generic and thus can fulfill this requirement.

REQ 4: L3VPNs are rather similar to existing use cases of ALTO in the Internet. Insofar as L2VPNs or pseudo-wire VPNs have the notion of some address, ALTO seems to able to handle these through an extension that extends the definition of an address to include other identifiers besides IP addresses.

REQ 5: Use of ALTO with network identifiers that are not IP addresses requires work. There is a need to analyze how to name VPNs and endpoints and how to achieve a mapping to the information stored in the ALTO server. One potential approach would be to characterize an endpoint by any identifier that is unique within a map.

REQ 6: ALTO can be used as of now with IP address ranges that are not globally routable. However, it must be emphasized that private VPN environments without uplink to the global Internet may only have connectivity to a limited number of IP subnets, i. e., the ALTO server will not be able to provide any reasonable guidance for most parts of the IP address space. Also, the ALTO server operator must take into account that IP address ranges in different VPNs may overlap, possibly also with the transport network infrastructure (e. g., PE routers).

REQ 7: Extensions to ALTO will be needed, in particular assignment of properties to PIDs. [I-D.roome-alto-pid-properties] extends the ALTO protocol by defining PID-based properties in much the same way that the original ALTO protocol defines endpoint-based properties. The VPN use case may also benefit from other extensions to ALTO. Representing the network topology as a graph and using TE metrics (e.g., [I-D.wu-alto-te-metrics]) may allow the VPN operator to be better informed about link-level information to grow (or shrink) a new (or existing) VPN.

REQ 8: Assuming extensions in REQ 7, filtering should be fairly easy. If a client has to perform complex queries, it could also download all PID properties and execute complex filter logic itself, if full data retrieval is supported by the ALTO server. There is a tradeoff between the complexity of the server and the client.

REQ 9: Extensions to ALTO will be needed. In complex cases, client-side execution of the filtering is an alternative.

REQ 10: This requirement will possibly require extensions to ALTO, e. g., to distinguish between endpoints that are already attached to the VPN and sites outside the VPN. This can be achieved e.g. by endpoint and/or PID properties. Changes of the VPN topology are likely to change the ALTO maps, i.e., standard ALTO mechanism for incremental updates and push notifications would be of added value.

REQ 11: Registration of corresponding metrics is useful. An subset of metrics is described in [I-D.wu-alto-te-metrics]. Further analysis is needed for additional connectivity aspects, such as resilience parameters, shared risk link group (SRLG), etc.

REQ 12: Existing authentication and access control mechanisms for ALTO could be sufficient to meet this requirement, subject to further analysis.

6. Security considerations

TBD.

7. IANA considerations

TBD.

8. References

8.1. Normative References

[I-D.ietf-alto-protocol]
Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-25 (work in progress), January 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4762] Lasserre, M. and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.

8.2. Informative References

[I-D.lee-alto-ext-dc-resource]
Lee, Y., Bernstein, G., Dhody, D., and T. Choi, "ALTO Extensions for Collecting Data Center Resource Information", draft-lee-alto-ext-dc-resource-03 (work in progress), January 2014.

[I-D.randriamasy-alto-cost-schedule]
Randriamasy, S. and N. Schwan, "ALTO Cost Schedule", draft-randriamasy-alto-cost-schedule-02 (work in progress), October 2012.

[I-D.roome-alto-pid-properties]
Roome, B. and Y. Yang, "PID Property Extension for ALTO Protocol", draft-roome-alto-pid-properties-00 (work in progress), October 2013.

[I-D.wu-alto-te-metrics]
Wu, W., Lee, Y., Dhody, D., and S. Randriamasy, "ALTO Traffic Engineering Cost Metrics", draft-wu-alto-te-metrics-00 (work in progress), October 2013.

Appendix A. Acknowledgements

TBD.

Authors' Addresses

Michael Scharf (editor)
Alcatel-Lucent

Email: Michael.Scharf@alcatel-lucent.com

Vijay K. Gurbani (editor)
Alcatel-Lucent

Email: vkg@bell-labs.com

Greg Soprovich
Alcatel-Lucent

Email: Greg.Soprovich@alcatel-lucent.com

Volker Hilt
Alcatel-Lucent

Email: volker.hilt@bell-labs.com