

Cisco Locator/ID Separation Protocol and Overlay Transport Virtualization Data Center Infrastructure Solutions for Distributed Data Centers

What You Will Learn

Different applications require different types of connectivity. As data centers are distributed geographically, a combination of Layer 2 and Layer 3 connectivity may be required to support different applications. This document discusses the application profiles that may need to be supported across a distributed data center infrastructure. It also provides recommendations about when to use the Cisco[®] Locator/ID Separation Protocol (LISP) and Overlay Transport Virtualization (OTV) technologies, or a combination of both.

The following scenarios are discussed:

- Support distributed clusters with a combination of LISP and OTV
- Streamline disaster recovery with LISP
- Enable workload relocation to third-party (cloud) facilities with LISP

This document is of interest mainly to technical decision makers (TDMs) and network architects and operators. It may also benefit business decision makers (BDMs) collaborating with their technical counterparts.

Introduction

One of the leading trends in current data center deployments is the creation of virtual data centers that consist of a collection of geographically dispersed data center locations. In general, multiple locations are used as a means to increase resiliency and capacity. Increased resiliency and flexible access to additional capacity are possible only if the multiple data centers are modeled as a single large virtual data center in which workloads can move and be distributed seamlessly across locations.

Another leading trend in the data center space is the growing use of very high densities of low-cost commodity servers that, as independent units, provide enough capacity to encourage the wide adoption of virtualization and, as a group, provide the ideal infrastructure to achieve the high-capacity and resiliency benefits of distributed clustered applications. The result of this trend is a notable increase in the amount of server-to-server traffic.

These trends have posed very demanding connectivity requirements on the data center networking infrastructure, which must provide a connectivity continuum across geographically dispersed data centers that supports large amounts of server-to-server traffic, seamless mobility, and relocation of workloads over a massive number of endpoints.

As virtual data centers become a reality, the movement of workloads across different geographic locations will become common in the form of either live moves requiring state preservation or cold moves that allow longer distances and better planning.

The resulting connectivity demands for Ethernet and IP are being addressed by a combination of next-generation technologies such as Cisco LISP and OTV.

Application Profiles

The type of connectivity required between data center locations in the virtual data center depends strictly on the nature of the applications running in these data centers. The most common application parameters, listed here, provide a good idea of the connectivity required to support the various application profiles:

- **Routable applications:** At present, data centers can assume that most applications are IP based.
- **Nonroutable existing applications:** Some existing applications do not use IP, but are purely Ethernet based or, worse, use an older protocol.
- **Clustered applications:** Some applications are clustered, and as clusters they often use non-IP link-local multicast signaling between the application cluster member servers. All other communication, aside from peer discovery and heartbeats on these applications, is normally routable IP traffic.
- **Virtualized applications:** Applications in general can run on virtual machines and are candidates for live mobility.

Cisco LISP Virtual Machine Mobility

The Cisco LISP Virtual Machine Mobility (LISP VM-Mobility) solution allows any host to move anywhere in the network while preserving its IP address. The capability allows members of a subnet to be dispersed across many locations without requiring any changes on the hosts and while maintaining optimal routing and scalability in the network.

Cisco OTV Technology

OTV allows the secure extension of Layer 2 connectivity across multiple locations. The Layer 2 connectivity extension includes the handling of broadcast traffic and link-local multicast, which is critical to clustered applications. By extending Layer 2 connectivity across locations, OTV is in fact extending the associated IP subnet across multiple locations. This extension creates routing challenges that are addressed by using LISP. LISP and OTV are therefore complementary technologies.

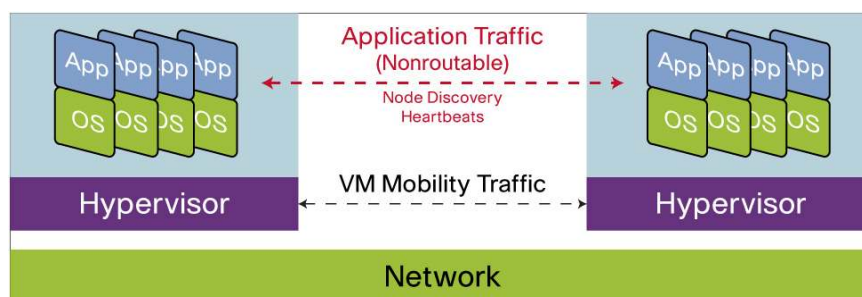
When to Use OTV, LISP, or Both?

Use OTV for server-to-server communication when an application requires non-IP (Layer 2) communication such as link-local multicast or non-IP unicast communication between servers. Today, most applications in this category are clustered applications. For such applications, a LAN must be extended between the various locations in which the application member servers reside. OTV provides the tools to extend a LAN in a secure manner to support this non-IP traffic.

LISP is required to provide optimal routing while supporting mobility and location independence in the virtual data center. This requirement is independent of the need an application may have for extension of Layer 2 connectivity. Alternative route optimization approaches can be used, but LISP is the preferred mechanism.

Using OTV and LISP Together: Live Virtual Machine Mobility and Distributed Clusters


Applications requiring link-local multicast communication among servers are usually clustered applications that use simple link-local multicast mechanisms for peer discovery as well as for the exchange of hello messages and heartbeat signals. Note that the dispersion of the cluster members could be the result of virtual machine movement across sites with virtual machine mobility tools such as VMware vMotion or simply a static distribution of the servers. Regardless of how the cluster members are distributed, the link-local multicast traffic is required by the clustered application independent of any use of virtualization, as shown in Figure 1.

Figure 1. Applications leverage both routable and nonroutable traffic

VMware currently supports live vMotion migration across sites only when an extended LAN across these sites exists¹. The combination of OTV and LISP can be used in a similar way to support the cluster scenario as well as the live VMware vMotion scenario, so both approaches are discussed together here. The combination of these two scenarios is also possible, with a clustered application run on top of a virtualized infrastructure, and live mobility supported for this clustered application.

When members of clusters are dispersed across multiple locations, the LAN needs to be extended across these locations to support the forwarding of the non-IP traffic used for peer-discovery and heartbeats. OTV provides the necessary functions to extend a LAN across multiple sites and to support this non-IP communication between servers in the cluster.

Depending on which layer of the application is being distributed, the cluster members may need to communicate with clients over an IP network. Since the LAN has been extended across multiple locations, the IP subnet associated with this LAN is also extended across these locations. In this case, the subnet no longer retains its location semantics because a subnet traditionally indicates a single location, but this one is extended across many locations. Traditional routing would not know at which of all the locations in the extended subnet a server may be located, resulting in many cases of suboptimal routing to the servers, as shown in Figure 2.

Figure 2. Routing optimization for VM Mobility within extended subnets

When subnets are extended, LISP is instrumental in handling location information at many levels of detail to provide the shortest-path routing to the appropriate locations within the extended subnet and avoid sub-optimal traffic

¹At the time of writing, VMware requires an extended LAN across sites to support live VMware vMotion.

patterns such as the one illustrated in Figure 2. Thus, LISP adds location semantics to an extended subnet that otherwise would not have any location semantics.

The combination of LISP and OTV provide a complementary approach to distribution of applications across multiple locations while preserving optimal routing to every member of the application. Live VMware vMotion is also supported by this combination of LISP and OTV, which not only supports the application requirements, but helps ensure optimal reachability of the roaming virtual machine wherever it moves.

Streamlining Disaster Recovery with LISP


Another common scenario in which servers change location is a disaster recovery situation. Normally a disaster recovery facility is located a long distance from the main production data center and is brought online only when a critical situation forces the main data centers to go offline.

The distance between the main production data centers and the disaster recovery facility is usually large enough that synchronous replication of storage data is not viable. Additionally, the nature of disaster recovery procedures is such that bringing the facility online is usually an unexpected event. For these reasons, workloads and applications are relocated to a disaster recovery facility using point-in-time checkpoints that are designed to reduce the amount of data loss caused by the disaster situation. In other words, live moves of workloads are not possible to disaster recovery facilities, but rather a cold move of the applications is performed using the most recent available copy of the application data.

Bringing a disaster recovery facility online can be a complicated process. One aspect of the restoration procedure has traditionally been the change in IP addressing of the applications as they are relocated to the disaster recovery facility. By decoupling the identity of the application nodes from the location in which they operate, LISP allows applications to preserve their IP addresses as they are relocated to the disaster recovery facility. The capability to preserve the IP addressing of the applications streamlines the restoration procedures of the disaster recovery facility and makes the overall operation much faster. This streamlined process ultimately reduces the amount of downtime caused by a disaster situation.

As mentioned, in such situations live mobility is not a goal, nor are clustered applications distributed between the disaster recovery facility and the main data centers. This scenario involves a full relocation of all applications and all application members. Therefore, Layer 2 connectivity is not required, and OTV also is not required for this application. There is, however, a large benefit in being able to preserve the IP addressing of the various applications and network services, and that is where LISP makes a significant difference in disaster recovery (Figure 3).

Figure 3. Routing Optimization for host mobility across subnets



Enabling Cloud Bursting with LISP

The previous disaster recovery scenario discussed cold moves between locations at which a machine is taken down and brought back up again at a different location. The next scenario concerns planned cold moves aimed at relocating workloads without necessarily preserving existing connections.

In a planned cold move, the application is stopped, all necessary state and storage data is replicated at the target location, and the application is eventually restarted at the new location. This simple approach of relocating workloads at different locations is often referred to as the infrastructure-services in a Cloud model. When the relocation is performed to compensate for capacity exhaustion of the current data center, the process is referred to as cloud bursting: an overflow of the resource capacity in the main data center.

Workloads can burst to data center facilities controlled by the enterprise, or they can be relocated to third-party facilities such as those offered by infrastructure-as-a-service (IaaS) cloud providers. In either case, these are cold moves that do not require Layer 2 connectivity between the data center locations, but the capability to use IP addresses in a manner that is independent of the location in which the IP address is placed. LISP is designed to function in such a way and allow any IP address to exist anywhere without requiring the reconvergence or restructuring of the underlying routing.

As storage and virtualization technologies evolve, live moves to cloud facilities will become a realistic goal. The Cisco LISP VM-Mobility solution is ready today to support such live movement across facilities and provide optimal routing to the moving endpoint while allowing the preservation of active sessions as active workloads move across different subnets at different locations.

Conclusion

LISP and OTV are complementary solutions providing the necessary connectivity between the locations in a geographically distributed data center. Depending on the type of applications supported and the role of the data center locations, either LISP or a combination of LISP and OTV may be required to provide the level of location flexibility required from the IP infrastructure.

For More Information

Comprehensive information about Cisco LISP, OTV, and LISP VM-Mobility solutions can be found at:

- <http://www.cisco.com/go/lisp>
- <http://www.cisco.com/go/otv>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)