RingCentral CloudConnect

Background

The RingCentral service offerings have been engineered to offer a full SLA guarantee even when subscribers are connecting across the public Internet. Disruptions may occur, but they are generally isolated to the 'last mile', the physical link between the customer and the Internet Service Provider (ISP).

In return for cost savings, many customers accept the slight risk of disruption. Some customers, however, have business critical communications and cannot afford even the slightest interruption of voice services. They must have the insurance afforded by having one or more links directly connected to Ring Central's Data Centers. This completely bypasses the potential problems associated with routing issues or DDoS attacks that are a threat on the Internet and, additionally, provides end-to-end network Quality of Service (QoS).

RingCentral is present in many major Data Centers around the world. Some locations are fully configured to support any type of customer connection while others may support only a subset of connection types.

There are several connection types which may be used by customers to implement a connection. The three most common are Layer-2 Point-to-Point (P2P) Circuits (known as MetroEthernet E-Line), MPLS circuits, and SD-WAN links. It is critical that RingCentral Custom Engineering work with the customer from the very first sign that a direct connection may be needed. The process is complex and requires carefully scheduled interaction with many different groups.

Note: Please note that although Layer-2 MultiPoint-to-MultiPoint (MP2MP) MetroEthernet E-**LAN** Circuits are supported, they are discouraged by RingCentral. These circuits provide any-to-any connectivity among 3 or more devices. This configuration does not allow for proper QoS Traffic Shaping configuration. Usage of this technology is at the customer's risk.

Also note that several carriers have implemented or are implementing RingCentral Enhanced Internet services. These are directly connected with RingCentral at multiple points and can offer the customer Internet service with guaranteed QoS and bandwidth to RingCentral. This service requires no interaction with RingCentral for provisioning and can be turned up by the carrier just as soon as the carrier can deliver Internet to the customer! This provides almost the same level of service as a CloudConnect without the cost and the numerous implementation headaches. At least one of the carriers offers a highly redundant SD-WAN overlay with automatic failover to a backup commodity Internet link.

Layer-2 (MetroEthernet E-Line) Overview

From a conceptional viewpoint you should consider a Layer-2 MetroEthernet E-Line Circuit to be an extremely long Ethernet cable that runs from Point A to Point Z. Layer-2 Circuits are very simple to implement and can offer great cost savings. The MetroEthernet Forum has a whitepaper (link below) which discusses the details of MetroEthernet.

Link: ***https://www.mef.net/Assets/White_Papers/Metro-Ethernet-Services.pdf***

RingCentral assigns RFC6598 network addresses (a /30 CIDR block) to the link. This space is very large and is reserved for Carriers to use for NAT links to customers. A larger block can be assigned if needed for additional NAT pool space but must be specifically requested.

This connection is set up with the Enterprise customer applying a Source-NAT masquerade (interface overload) policy on all traffic that flows from their side toward the RingCentral side. Thus, all the customer traffic that flows across this link will have a source address of the interface on the customer side of the link.

Care must be taken to automatically disable routing to Ring Central through this interface if it fails for any reason. This is usually done by using dynamic BGP routing across the link, Path Monitoring (Palo Alto specific), or SLA/IP tracking route removal (Cisco routers/ASA). This is documented in the Customer Peering Setup paper, including configuration examples. Do not depend on the link state of the interface as it may not reflect the true status of the RingCentral port, rather it shows the status of the carrier's port even though part of the circuit may be down.

This diagram shows the use of a single Layer-2 E-Line link (green color) connecting the customer's primary site to the RingCentral site in San Jose CA.

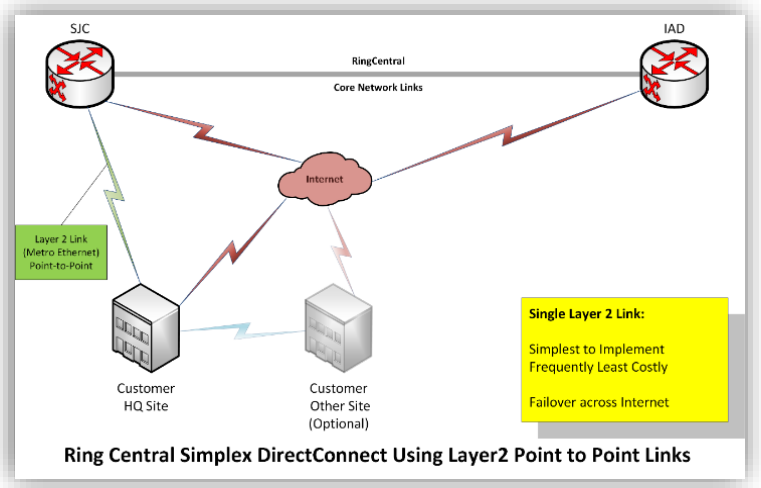

Ring Central Simplex DirectConnect Using Layer2 Point to Point Links

It is possible to have two Layer-2 E-Line links to add redundancy. This is shown in the following diagram which shows connections to RingCentral sites in both San Jose CA and Vienna VA.

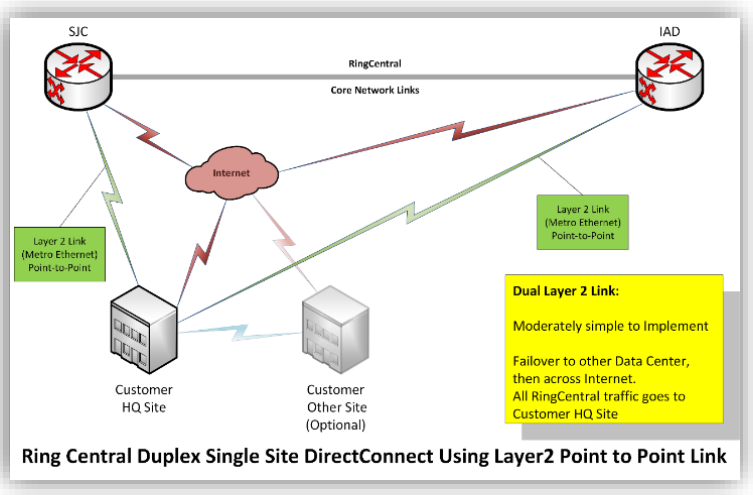

The following diagram shows a more complex arrangement using two Layer-2 Circuits connected to two different Customer sites. It is assumed that the two customer sites are interconnected in some fashion and implement some form of dynamic routing protocol. Setup of this configuration can be complex and should be carefully planned with input from an experienced routing network engineer.

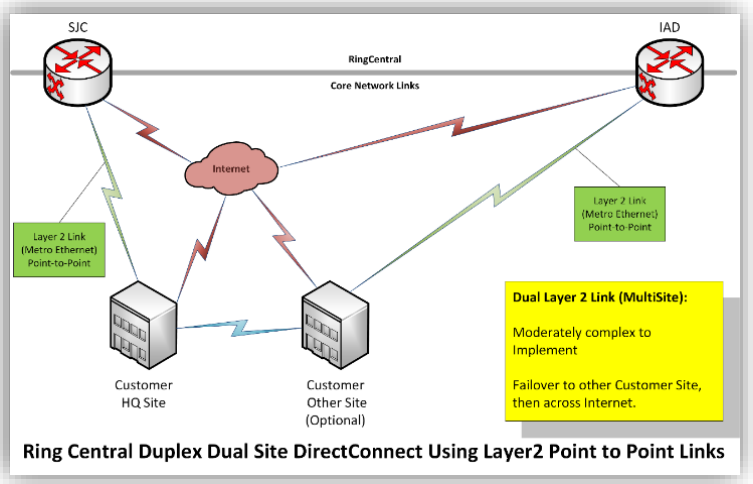

MPLS (Layer-3) Overview

Larger networks with multiple sites often use MPLS links to interconnect their sites. In an MPLS network, each Customer Site is connected back to the core of the carrier network which acts as a large virtual central router for the Customer. This creates a large MultiPoint-to-Multipoint (MP2MP) network. Many large enterprises use their MPLS network as a backbone to carry fully meshed, encrypted tunnels (DMVPN, ADVPN, etc) with their corporate traffic.

The customer may decide to have an MPLS link installed into a RingCentral Data Center location. This effectively extends the customer's network into the RingCentral rack. RingCentral can then engineer a link allowing RingCentral based traffic (only) to flow preferentially across this link. Customer traffic to/from RingCentral will revert back to traversing the Internet upon failure of this link if routing is properly engineered.

Every MPLS link is sized to a specific total traffic rate and is assigned a QoS profile, usually referred to as a COS (Class of Service) profile. Most carriers default to 4 COS classes within a COS profile; some allow the customer to request up to 6 COS classes. Each COS class is assigned a percentage of the total link traffic rate which is guaranteed. Usually COS #1 is reserved for Real-Time (Voice) traffic, COS #2 for video, COS #3 for SIP/Signaling traffic, and COS #4 for all other traffic. Traffic marked as DSCP EF or IP Precedence 5 goes into COS #1. (**Beware** – Ensure that you have enough bandwidth allocated for COS #1 in each link as any traffic over the specified allocation is usually discarded!!!) Traffic in other COS classes that exceed the assigned traffic rate is simply shunted to the Best Effort COS class. Bandwidth allocated for a specified COS class and not used is dynamically reallocated for use by the other COS classes.

A device referred to as a CE (Customer Edge) router is required in order to terminate the carrier links for each location. The CE router talks to a PE (Provider Edge) router inside the carrier's network. In most cases, the carrier supplies the CE router as part of a managed service offering. Some providers support virtual firewall services and Internet Access as part of their MPLS offerings, making this link the only link a Customer needs. (At least in the Carrier's opinion.)

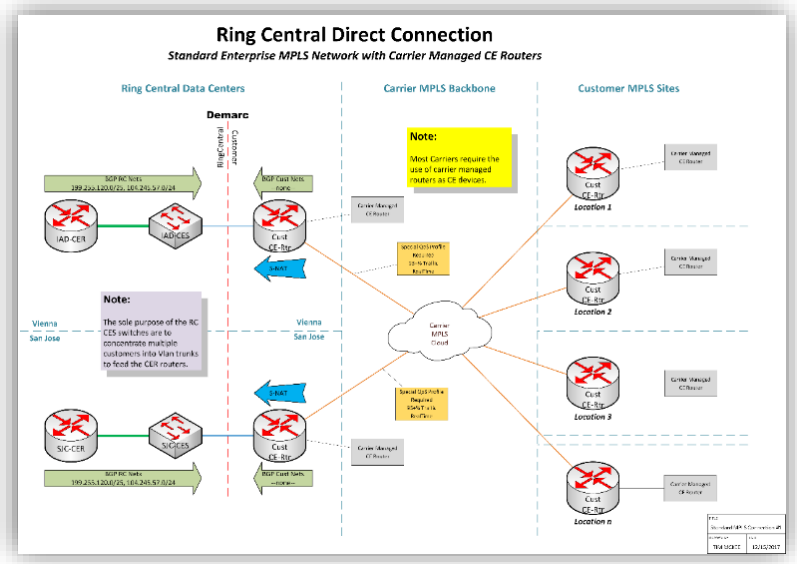

Some carriers will allow the Customer or RingCentral to provide the CE router services. This is ideal for Ring Central as some Data Center locations do not have space available for customer collocated devices.

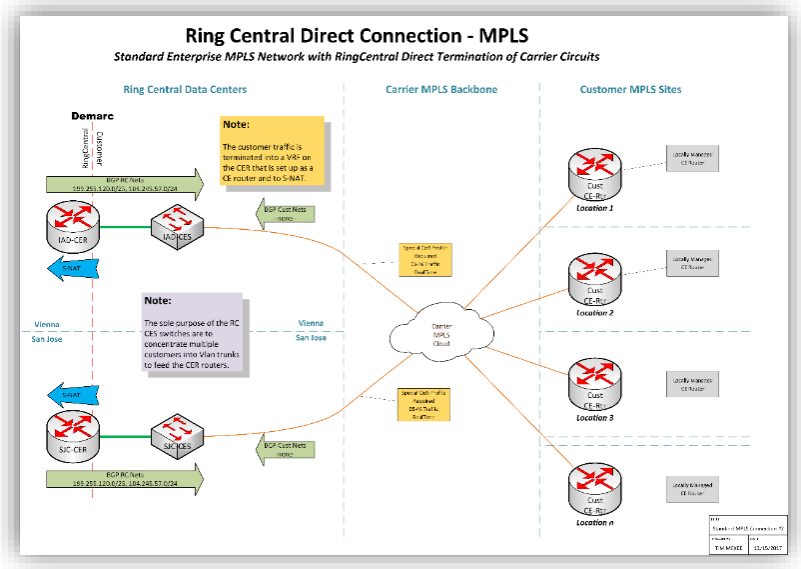

The following diagrams show MPLS networks that are connected at one and two Ring Central Data Centers with Internet failover redundancy:

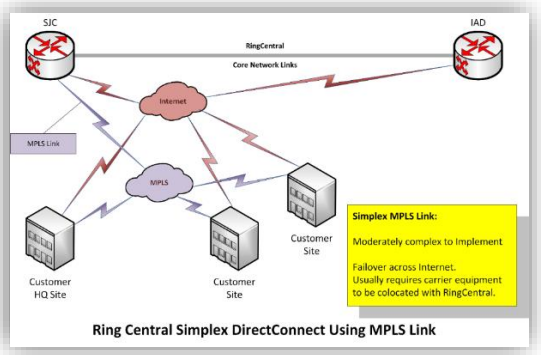

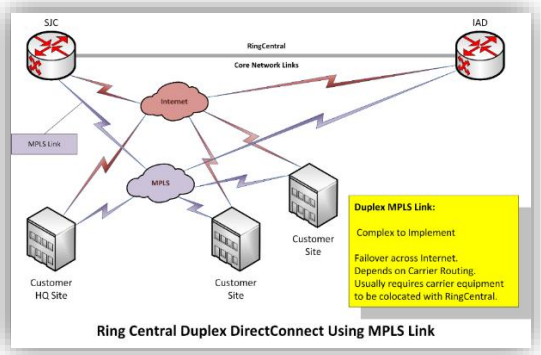

SD-WAN Overview

All of the current marketing hype is pushing SD-WAN as the latest, greatest panacea for IT. It can fix everything!!

In reality, SD-WAN is a very broad term that is used differently by different manufacturers. It ranges from a setup that dynamically routes the most critical traffic across the best of two ISP links to complex equipment that sends duplicated packet streams across multiple links to reduce packet loss and minimize jitter. Some implementations can setup a dynamic full-mesh VPN and interconnect customer sites into a comprehensive LAN, taking the place of an MPLS network.

Some SD-WAN implementations only require a device at the customer site(s) while others require additional 'gateways'. The more complex versions that remediate packet loss always require a gateway device. As an example, one complex device in our lab was able maintain very good voice quality (MOS score of 4.3+) when using a single link with 15% packet loss!

RingCentral will support almost any SD-WAN implementation so long as the SD-WAN vendor is willing to work with us. We are constantly testing implementations in our labs.

Engage RingCentral Custom Engineering anytime SD-WAN is mentioned in concert with our services.