

Monthly Savings and Outage Prevention with Cisco MATE Design and MATE Collector

This major Latin American service provider used the Cisco MATE portfolio to gain centralized failure analysis, optimization, and visualization of their network to prevent or quickly repair outages, load-balance traffic across transoceanic links, and run links hotter: all leading to huge savings.

EXECUTIVE SUMMARY	
Company: Major Latin American Service Provider	
Industry: Telecommunications	
CHALLENGE(S)	<ul style="list-style-type: none"> Major Latin American service provider could not perform accurate load-balancing of peering links after network failures, causing suboptimal usage of expensive network infrastructure
SOLUTION(S)	<ul style="list-style-type: none"> Cisco MATE™ Design and MATE Collector
PHASES	Design, test, and deploy an MPLS-TE solution over a partial mesh: <ul style="list-style-type: none"> Accurate traffic and topology simulations at very frequent intervals Fast and reliable mechanism to move traffic based on load-share metric Centralized provisioning to deploy recommended configuration changes
RESULTS	<ul style="list-style-type: none"> Reduction in fault MTTR in many cases from days to minutes Optimal traffic load-balance achieved, especially over international links Service downtime practically eliminated, and quality of service greatly improved, resulting in much happier customers

Executive Summary

A typical issue for peering between service providers is to deal with failures on links that may cause severe congestion and suboptimal link utilization. A major Latin American service provider was experiencing this phenomenon with respect to traffic coming from North American content service providers and cable operators. The Cisco MATE portfolio was used as part of an operational solution to emulate failures on the peering links, investigate where the traffic was rerouted in the event of those failures, and make modifications (Multiprotocol Label Switching [MPLS] Traffic Engineering[TE] metrics) to help ensure that the traffic was properly load balanced.

After completing exhaustive failure analysis and optimization, the service provider validated configuration changes derived from MATE Design and then pushed to the production network using a provisioning tool. The service provider re-runs these simulations and corrects problems periodically to help ensure that traffic changes and occasional link failures do not affect service for their customers.

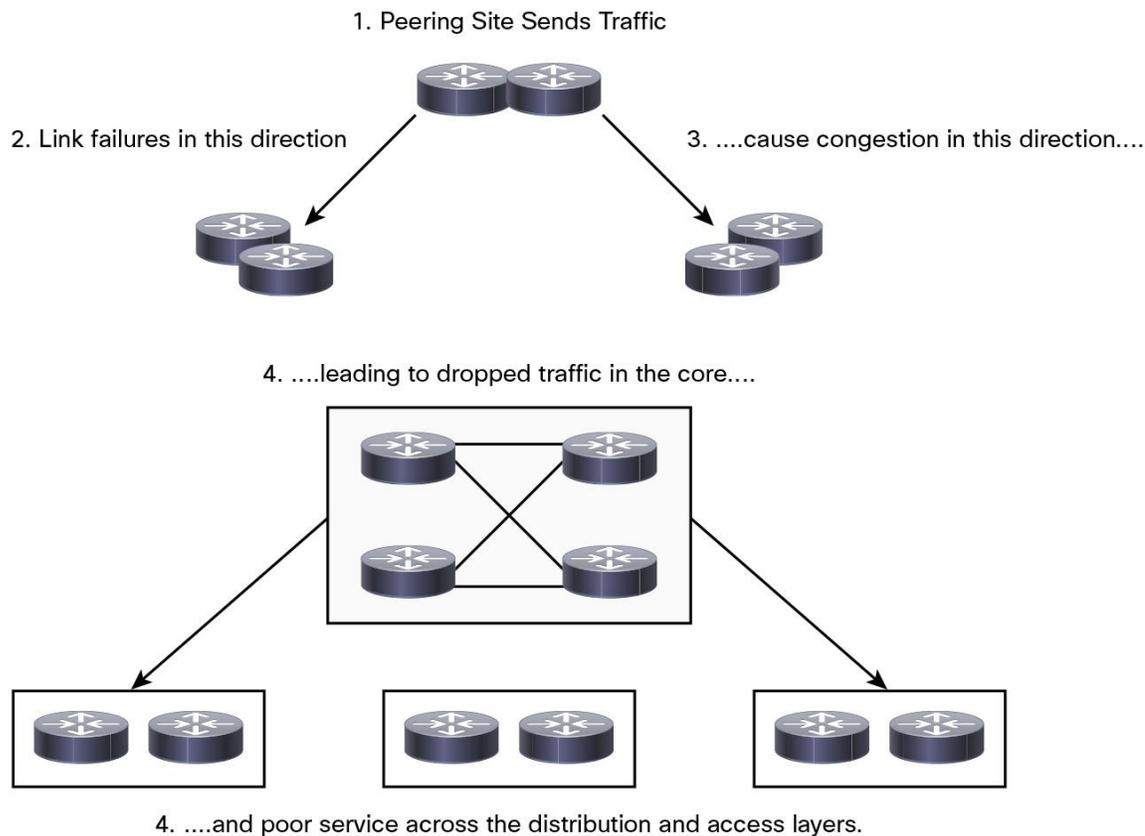
Introduction

A highly successful service provider in Latin America is offering a variety of services to consumer and business markets (as well as other service providers), including IP connectivity, various Layer 2

and Layer 3 VPNs, MPLS, Metro Ethernet, Voice over IP (VoIP), quad-play (video, voice, data, and mobile) and security. Eighty percent of its traffic is best effort, and the remainder requires prioritization.

The IP core backbone network topology has a partial mesh wherein all the national POPs connect to routers that have transit cores dealing with international traffic. There is asymmetrical traffic on these international links. In the event of a failure, this traffic asymmetry is accentuated by the natural behavior of Open Shortest Path First (OSPF) Protocol, which chooses the shortest path as a contingency route to deliver traffic, ignoring the bandwidth on available links. As a consequence, an overload occurs on the links used in the alternate path, which leads to packet drops (see Figure 1).

Figure 1. Regular Link Failures Cause Unbalanced Traffic



To handle the issue described above, the service provider needed to know the exact changes to make in order to properly load-balance the traffic and optimize the network. MPLS TE techniques help to guarantee traffic delivery in the event of link failures as well as optimize link utilization. IGP-based solutions are often complex to deploy, especially after failures.

The service provider sought assistance from Cisco for peering on its MPLS-TE core IP backbone; it needed to cope with unbalanced traffic and high congestion following national and international link failures. With the assistance of Cisco MATE Design and Collector, this service provider was able to optimize and deliver reliable services from peering points in North America. Its two main network control goals were:

- International link optimization
- Peering optimization

Over time, the service provider wanted to be able to fine-tune the network with frequent simulations (hourly if necessary) with specific load-share sets.

Finding the Primary Cause of Congestion

The root cause of congestion is often very difficult to detect, especially when multiple service providers are involved. In particular, this service provider was often not able to analyze the root cause of traffic congestion that resulted in dropped packets. In this case, a cable operator was sending revenue-generating video traffic, which this service provider handed-off to other networks (a carrier of carrier situation) to be viewed by consumers.

Cisco MATE Design and Collector were employed to simulate traffic flows and topology every 15 minutes. This way, increases in congestion could be captured immediately. Using a current model of the network, the service provider performed the following steps with MATE Design:

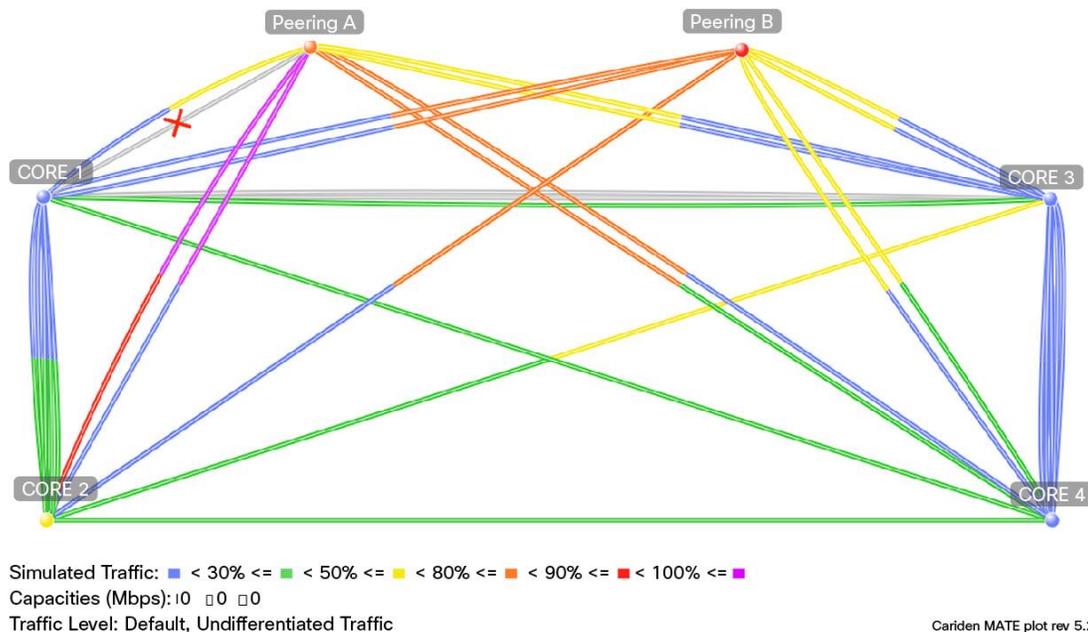
- Isolated the point of congestion
- Applied optimization to reconfigure traffic flows on the network to alleviate congestion
- Established a sequence of network changes

These tasks were (and continue to be) performed repeatedly, with changes applied to the operational network based on MATE Design's recommendations.

The View from MATE Design

Figure 2 shows that a single broken link from Peering "A" leads to congestion on many links. The purple and red links are the most congested.

Figure 2. Single Failure Leads to Congestion on Many Links



As the flow of traffic from peers cannot be controlled, the service provider needed to determine a way to keep its own links balanced, and to maximize the quality of its own services, while avoiding service impact from traffic coming in from other networks.

In short, the main objectives that led the service provider to employ the MATE portfolio were as follows:

- **Load Balancing:** Achieve the best load balancing possible, in order to help ensure minimal congested links.
- **Avoid IGP Metric Changes:** Interior Gateway Protocol (IGP) metric changes are effective, but are potentially too coarse and could cause unintended ripple effects in the network.
- **Avoid Border Gateway Protocol (BGP) Changes:** The IP blocks were not defined to allow specific traffic control, and the traffic behavior for these blocks changes during the day.

- **Survive SRLG Failures:** A shared risk link group (SRLG) identifies an area where a failure at Layer 1 prevents a backup path at Layer 3 from being enabled. The SRLGs for Layer 1 failures affecting Layer 3 must be considered.
- **Meet Service-Level Agreements (SLAs):** SLAs are achievable by reducing packet loss, latency, and jitter.
- **Minimize Costs:** If the links can run hotter overall, the investment on current network infrastructure can be maximized.

To achieve these goals, the Cisco MATE portfolio received network statistics, the IGP link state database, Label Distribution Protocol (LDP) statistics and flows, and BGP information. A plan file was created based on this information and then utilized by both the engineering and operations teams. One team simulates possible network changes, while the other team validates configurations recommended by MATE Design before inputting them into the provisioning system for deployment to the production network.

Simulation Methods to Ensure Network Uptime

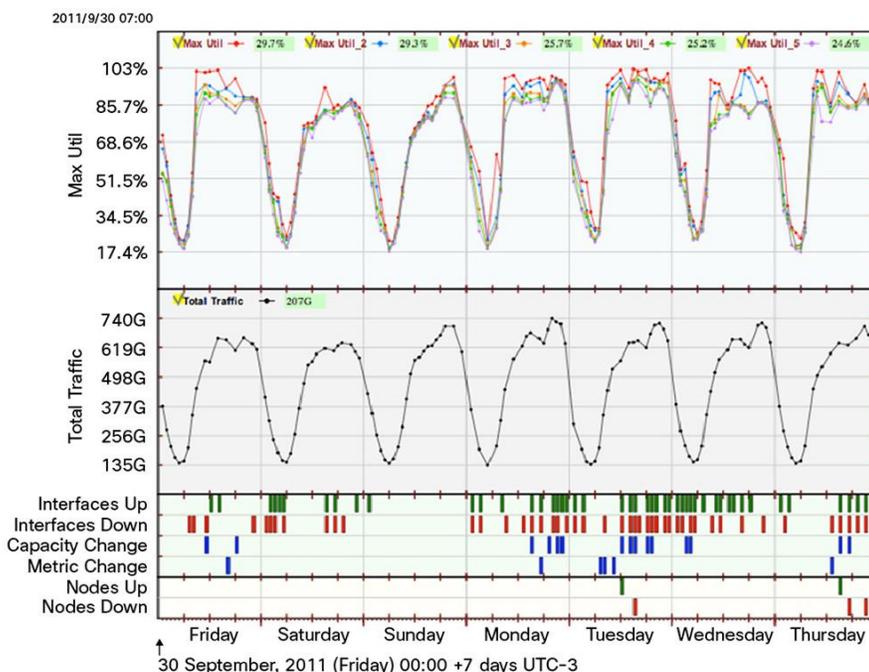
Since there are many MPLS Label Switched Paths (LSPs), with a lot of traffic, they are optimized for worst-case scenarios, to ensure that traffic drops will be minimal under all circumstances.

This service provider used the information generated from MATE Design to accomplish its network engineering objectives by load sharing or by changing LSP affinities. Load sharing can distribute traffic so that no individual router is overburdened, while affinities can either restrict or exclude traffic types from being carried on certain link types. For instance, MPLS TE affinities can be used to explicitly keep certain traffic from “tromboning,” or going out of its way.

Failure cases have to be considered for nodes, links, and SRLGs. Using MATE Design, operators ensure that traffic loss as a result of these failures will be kept to a minimum.

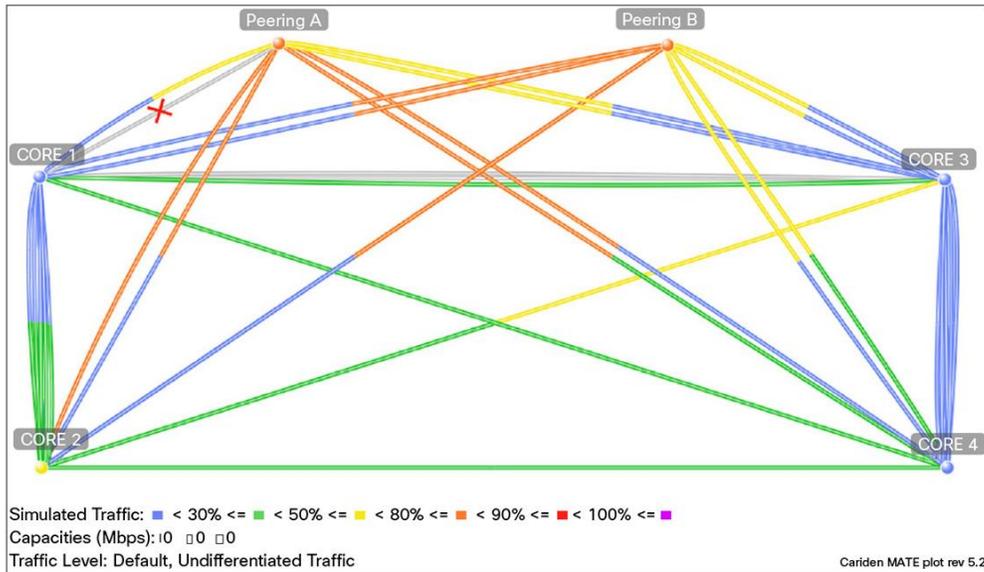
As a result of the analysis with MATE Design, network traffic remains intact even after failures (Figure 3).

Figure 3. Zero Traffic Loss Under Node, Link and SRLG Failures



As shown in Figure 4, international and national links were load balanced, and the service provider now enjoys cost savings by minimizing link upgrades and network failures. The orange and yellow links are not congested, and traffic is clearly more balanced across the two peering sites.

Figure 4. Links Are Uncongested Even Under Failure Conditions



Conclusion

The following results were attained:

- The mean time to repair (MTTR) dropped from days to minutes
- The customer also uses MATE to load-balance across transoceanic links; this allows maximum utilization to increase from 50 to 75 percent, yielding high annual savings



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.

Recycling symbol: Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: 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. (1110R)