

Planning and Designing Networks with the Cisco MATE Portfolio

Key Usage Scenarios for Planners, Design Engineers, and Operations

What You Will Learn

The Cisco MATE™ portfolio, which consists of the MATE Design, MATE Live, and MATE Collector products, delivers the network manageability required for simplifying processes and for delivering cost-efficient, reliable services. Each of these tightly integrated products simultaneously supports planning, engineering, and operational tasks. This paper describes the portfolio and outlines key use cases for planners and designers.

Cisco MATE Portfolio

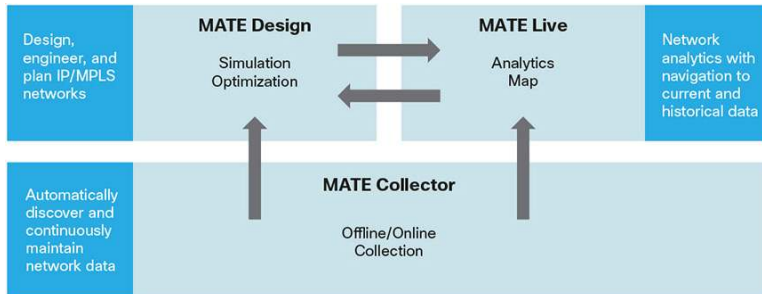
The Cisco® MATE portfolio helps enable service providers to achieve and maintain operational excellence by bringing manageability to their networks. This means operating and growing networks efficiently, having the ability to scale operations, and providing services without disruption. These goals are essential to the success of telecommunications providers, whether competing for customers based on cost or based on differentiated services.

Companies differ in their emphasis. For instance, many Internet service providers (ISPs) rely mainly on increasing capacity to avoid congestion. In contrast, traditional telecommunication providers often emphasize design and analysis to ensure resiliency, while some service providers consciously forego resiliency and work around congestion operationally.

The Cisco MATE portfolio works independently and collectively to meet cross-functional needs ([Figure 1](#)). Cisco MATE Live typically (though not exclusively) focuses on the operational scenarios; conversely, the planning and architecture or engineering scenarios typically feature a combination of Cisco MATE Design and MATE Collector.

- **Cisco MATE Design** is the market-leading integrated system for design, engineering, and planning of IP/MPLS networks
- **Cisco MATE Live** rapidly delivers in-depth network analytics with efficient navigation to both current and historical data for making critical business and technical decisions
- **Cisco MATE Collector** automatically gathers and continuously maintains information on infrastructure elements, topology, operational state, and traffic statistics for network planning and analytics

Figure 1. Key Purpose of Each MATE Product



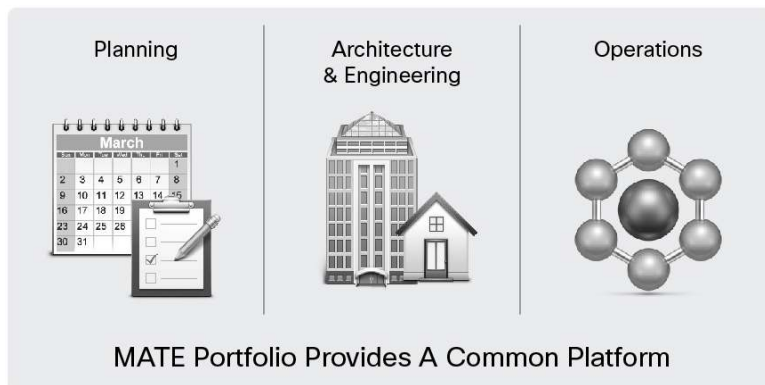
Usage in Cross-Functional Environments

An underlying principle of the MATE portfolio is that many cross-functional activities would be more efficiently performed using a single common platform (Figure 2). The groups that benefit from the MATE portfolio - **the planning, design, and operations groups** - have complementary goals to ensure that networks remain robust and cost-effective to maintain, and that service agreements are maintained along with profitability.

Not everyone neatly divides responsibilities. Rather, there is often an overlap, with groups having common traffic management goals that are achieved through different means.

- **The planning group** ensures there is sufficient capacity to avoid congestion and to handle both current and expected traffic growth. Key questions that a planner asks are, “When and where will my network run out of capacity? What will be the impact of adding a new service or customer to my network?”
- **The engineering group** designs the network and routing to ensure robustness and service delivery, even during failures, in the most economical manner. A typical concern for the architect or engineer is, “Where is my network most vulnerable to failure, and how can I mitigate it?”
- **The operations group** ensures any congestion is immediately mitigated. Operations staff ask questions such as, “What did the network look like before a failure or congestion? At what rate was bandwidth utilization increasing over the past day, week, month, or year?”

Figure 2. The MATE Portfolio Reaches Across Functional Boundaries



Planning Usage Scenarios

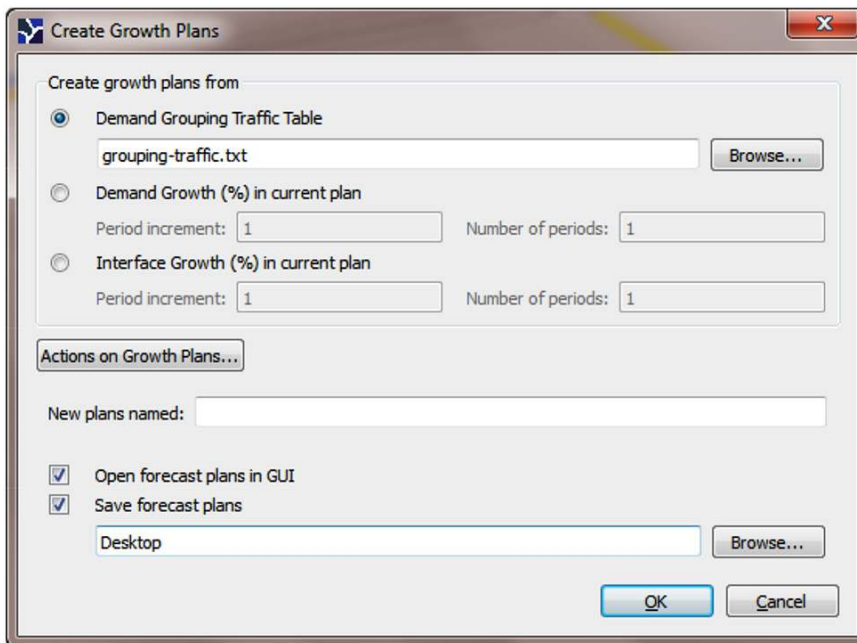
Within planning groups, key areas of concern are growth planning, what-if analysis for topology changes and networks, and testing the impact of new services or customers. Both Cisco MATE Design and MATE Live can be used for this analysis.

Simulating the Impact of Traffic Growth

Network planners must ensure that usage growth is accounted for cost effectively within the risk goals of their company. As such, a planning group often receives or is charged with generating traffic projections based on growth rates so as to determine the right level of build out for accommodating network growth. Predicting growth rates is a difficult, but required facet of success. Scaling network traffic and calculating the network impact is nearly impossible without advanced tools.

Cisco MATE Design helps planners to create growth plans based on measured or simulated traffic. The estimates can be manually entered, for example, based on knowledge of new services being introduced. For simulated traffic (called **demands**), options exist to group these demands to create aggregations (**demand groupings**) that meet explicit needs. For instance, planners could aggregate the traffic sourced from a specific site or the traffic destined for a particular autonomous system AS.

Figure 3. Traffic Scale Based on Demand Groupings



After creating these growth plans, planners can analyze the new plans using the Simulation Analysis tool, and then use the GUI to make further modifications to meet capacity requirements.

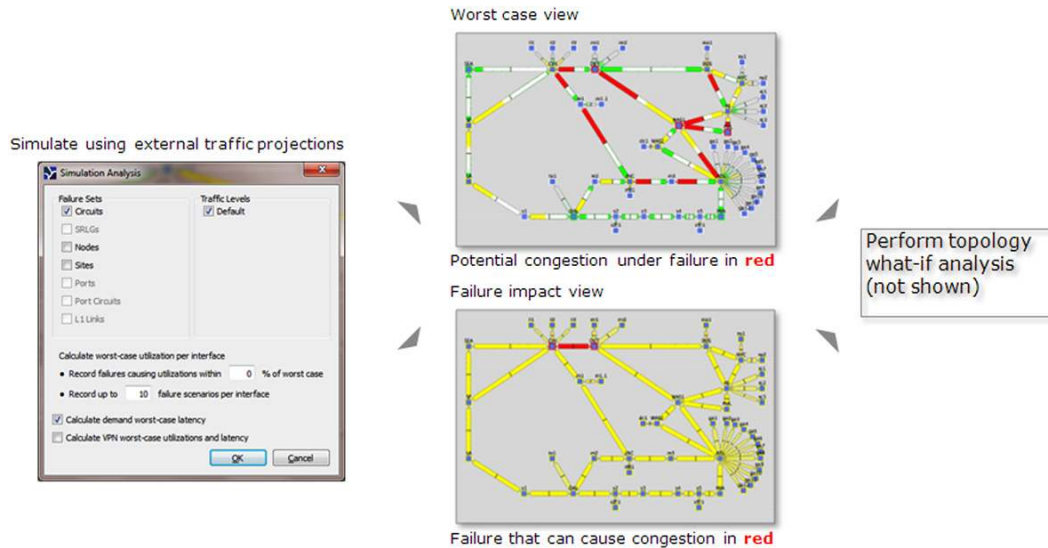
Note: Traffic can also be imported into a Cisco MATE Design growth plan tool from a Cisco MATE Live trending report; this can be very useful for determining future needs such as site-to-site traffic trends per class of service.

In [Figure 4](#), MATE Design uses traffic growth projections and runs analysis on both single circuits and shared-risk link groups (SRLGs) on the network. This allows network planners to determine the effects of worst-case failures,

such as where congestion will occur, as well as determine where there is heightened risk of failures. Planners can also analyze the impact of metric changes or circuit additions.

The return on investment (ROI) comes from the ability to defer upgrade costs until they are absolutely necessary and to maintain peak staff efficiency.

Figure 4. Simulating and Analyzing Worst-Case Failures and Failure Impact



Impact Analysis of Topology Changes or Network Upgrades


A frequent question in the planning process is how to test the value of a minor change that may **obviate** the need for a larger one. For instance, rather than a huge upgrade, a planner may consider whether adding a single link between two sites can mitigate congestion over part of the network.

The complete impact of a network upgrade might not be obvious, even to the best engineers, without rigorous evaluation of changing traffic within the network. Ensuring that an upgrade meets requirements without a negative impact is imperative to ensuring that the upgrades are effective uses of capital.

Being able to emulate and predict the impact of these changes promotes service level agreement (SLA) adherence and staff efficiency. Making a topology change could impact traffic flows, congestion, and latency, so knowing the effect such changes would have is valuable for operators, planners, and designers alike, and is critical for those with penalty clauses in customer SLAs.

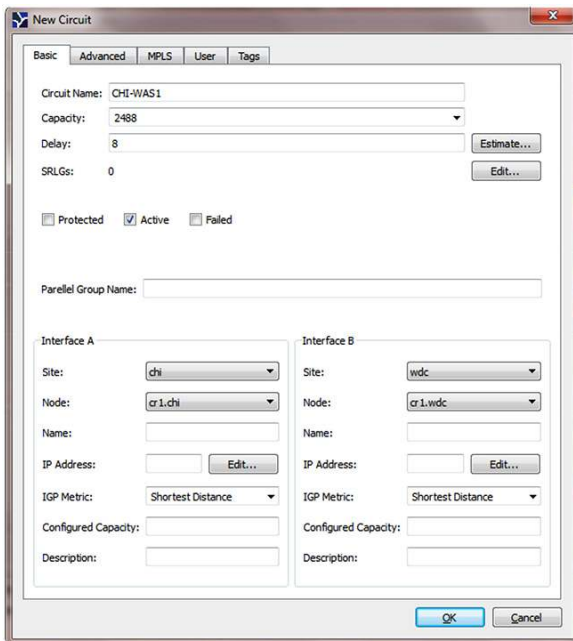
In the following scenario (Figure 5), the link between Chicago (CHI) and Detroit (DET) is congested, as indicated by the red interface.

Figure 5. Congestion Between Chicago and Detroit




Using MATE Design (Figure 6), a planner emulates the addition of the new circuit and specifies parameters (such as capacity and metrics) for it.

Figure 6. Adding a New Circuit and Its Parameters Is Simple



In this case, viewing the resultant traffic (Figure 7) in this part of the network indicates that the congestion is relieved.

Figure 7. Congestion Relieved as a Result of New Circuit



If congestion had not been relieved, a larger circuit or different metric could have been analyzed with a simple change to the circuit properties in the GUI. As such, Cisco MATE Design allows a rigorous evaluation of the proposed changes without having to perform trial-and-error maintenance in a live network.

Another area in which SLAs can be enforced is with respect to quality of service (QoS). For instance, worst-case analysis can consider QoS parameters, such as interface queue policies, weight, and police limit thresholds.

Impact of a New Customer or Service

Another role for the planning department is to determine if more capacity is needed to handle new services or customers. In the example shown in Figure 8, the sales department has inquired whether the network can support a new 4-Gbps customer (or service) originating in San Francisco.

Figure 8. Using MATE Design to Gauge the Impact of a New Customer or Service

Identify traffic sourced from new customer

Add 4Gbps of traffic for that customer

Simulate results

Congested link in red

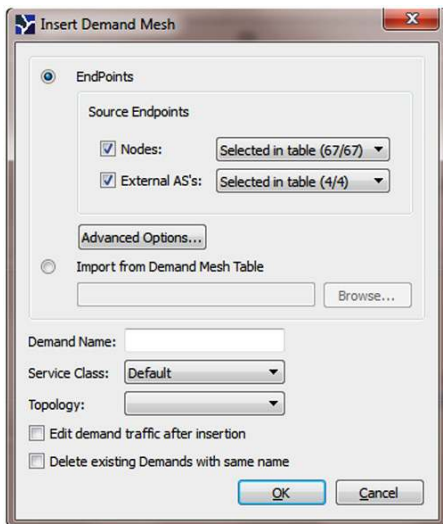
The screenshot displays the MATE Design interface. On the left, a filter dialog is open with 'Name' set to 'contains', 'Source' to 'contains', 'Destination' to 'contains', and 'Service Class' to 'contains'. The 'Current filter' shows 37/296 rows. In the center, the 'Modify traffic for selected demands' dialog is open, showing 'Traffic Level: 2004 stats', 'Number of Selected Demands: 26 / 296', and 'Total Traffic (Mbps): 7157.35'. The 'Add' radio button is selected with a value of 4000 Mbps. On the right, a network diagram shows a congested link highlighted in red.

MATE Design can be used to reduce the evaluation time to accurately investigate the impact from hours to minutes. First, a planner identifies the traffic for the new customer, and modifies the traffic totals for those

demands, and then simulates the results to identify congested links. The planner can also determine whether it is possible to mitigate the congestion by changing metrics or by performing MPLS traffic engineering operations. These exercises help determine if a new circuit is needed.

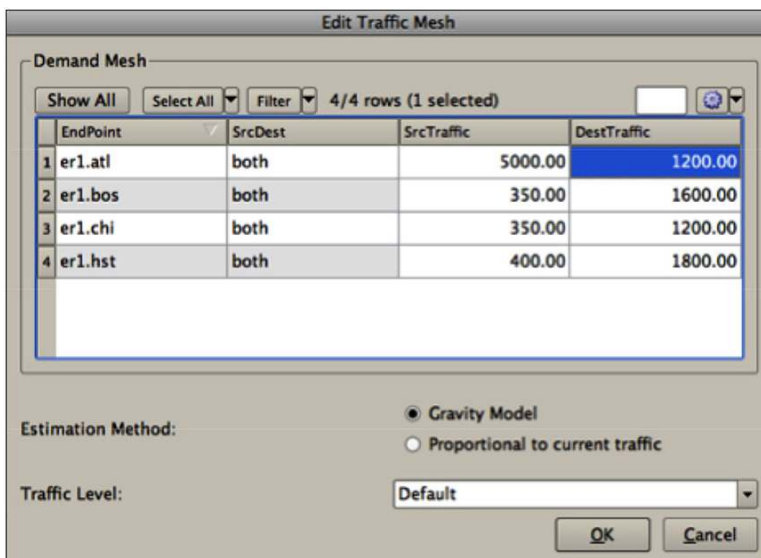
Similarly, a planner can insert a representative traffic matrix for the new customer or service, and then add new traffic to the network model before simulating the results of these changes (Figure 9).

Figure 9. Adding New Demands to the Network Model



New traffic can then be added to these demands for the simulation. Here, the Cisco-patented Demand Deduction¹ can provide extremely accurate simulations if measured traffic is available. Otherwise, Cisco MATE Design offers both manual and automated methods of simulating demand traffic.


Figure 10. Adding Traffic to New Demands



¹ See "Building Accurate Traffic Matrices with Demand Deduction" at <http://www.cisco.com>.

Finally, the simulation is run (Figure 11) to analyze the impact to resiliency and capacity of the network. Since adding new customer demands could impact existing route traffic through congestion, planning methods that use this feature help ensure SLA adherence.

Figure 11. Simulation to Analyze Impact of New Demands

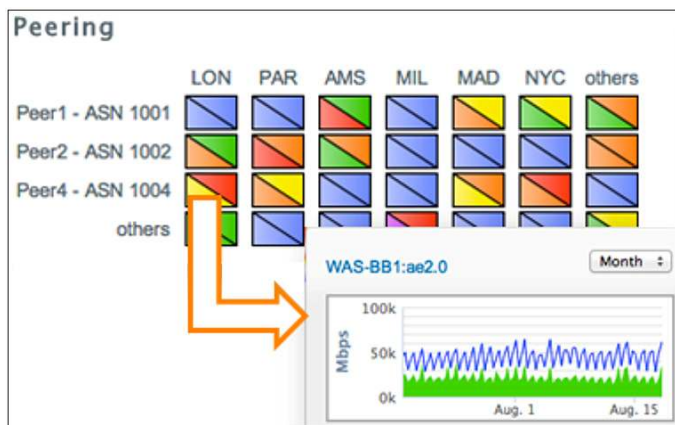


Coordinating Peering: Planning and Operations

Service provider peering² is becoming increasingly complex. Peering arrangements between service providers are a necessity on the Internet. They can reduce costs for transit services and increase resiliency and capacity for all parties involved. However, peering relationships must be monitored closely to ensure that they are fair, and commensurate with the original agreement.

Cisco MATE Live allows operators to closely monitor peering on multiple dimensions. In the example given here, there are both planning and operational elements. As shown in Figure 12 a peering interface displaying a disparity (and congestion due to too much traffic coming into the network) can be investigated to view the traffic in both directions. Thus, operators can monitor the high-level health of peering connections across sites, and can quickly navigate to specific interfaces to analyze issues.

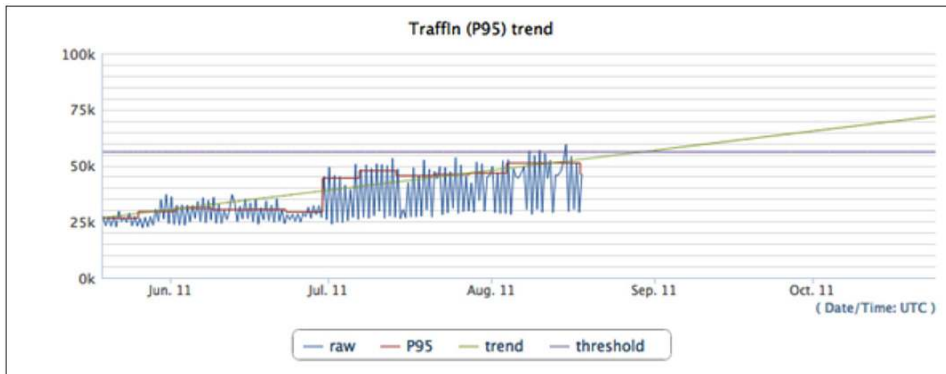
Figure 12. Select a Peering Interface to Immediately Determine Ingress and Egress Traffic



² Peering is the business relationship where service providers provide reachability to each other on predefined terms and to predefined portions of their routing tables. Peering thus provides connectivity to a subset of a provider's customer destinations.

Analytics are immediately available with MATE Live, listing the to/from traffic data. This can enable operators to identify and analyze issues with interfaces and nodes, such as changes to their capacities or routing metrics. For instance (Figure 13), operators can project individual peering interface growth and plan which interfaces will need to be upgraded and when.

Figure 13. Trends for Incoming Traffic on a Peering Interface



This planning is further augmented by the ability to report periodically (Figure 14) on aggregated traffic coming in and out of peering interfaces. These results can be used as the basis for ongoing peering negotiations.

Figure 14. Report on Total Traffic at Peering Locations

| Peer ▼ | Count ▲ | TraffOut (P95) ▲ | Traffin (P95) ▲ | Traffin over TraffOut (P95) ▲ |
|----------|----------|------------------|-----------------|-------------------------------|
| <Filter> | <Filter> | <Filter> | <Filter> | <Filter> |
| Peer 4 | 4 | 44861.09 | 30289.17 | 1.45 |
| Peer 3 | 3 | 53324.62 | 69816.62 | 1.84 |
| Peer 2 | 4 | 74690.76 | 60315.23 | 0.98 |
| Peer 1 | 4 | 82132.84 | 18962.18 | 0.36 |

Architecture and Engineering Usage Scenarios


The engineering group needs to ensure that the network is robust and that service delivery can be assured - even during failures - in the most economical manner possible. There are several techniques used by this group; a key one is to load balance traffic.

Balancing Traffic Using MPLS Traffic Engineering

Load balancing allows for careful management of traffic levels to ensure protection in the event of failure. In some circumstances it can be used to offload traffic from expensive or poorly protected routes onto cheaper or more robust ones. Savings come from being able to delay investment in additional bandwidth, and being able to more efficiently use the bandwidth reserved for maintaining service in the event of network failures. Finally, there are savings from staffing efficiencies in calculating routing and bandwidth requirements.

In Figure 15, there are a total of three intercontinental circuits from North America to Europe: one from Boston to London and two from Washington, DC to Amsterdam.

Figure 15. Evaluating Three Intercontinental Circuits (North America to Europe)



An architect might wonder whether there is value in load balancing the MPLS Labeled Switch Paths (LSPs) over these three expensive links, and whether a different placement of LSPs would be more efficient.

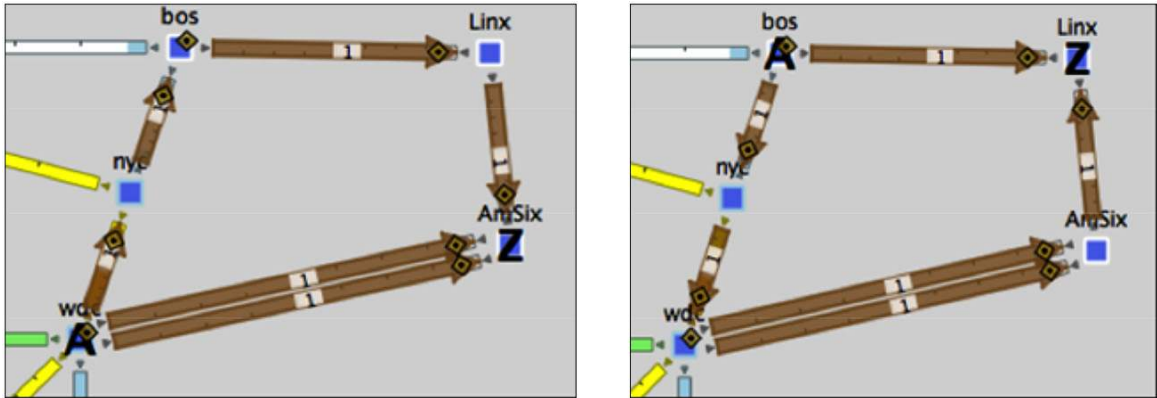
In this example, an architect can determine whether all of the affected links run within acceptable fill ranges. It appears that the local and remote terrestrial links would likely take the brunt of traffic impact, but having these links saturated is more cost-effective and less risky than congestion over transoceanic links.

Effectively balancing network traffic over expensive links is often more complex than anything that can be accomplished with Interior Gateway Protocol (IGP) metrics. The operational goal is to increase the utilization of these circuits to the highest possible level, while also maintaining very accurate traffic balancing, even under failure conditions.

The steps below can be used to understand this approach, but this is not the only option to distribute traffic across a set of network links. The recommended steps assume a network with minimal MPLS, but with the ability to construct explicit LSPs with load sharing.


First, architects should identify the set of nodes on each end of the set of costly links. Next, they would build explicit LSPs from each device on each side of the Atlantic to its adjacent node on the other side using all possible transoceanic links. [Figure 16](#) shows three separate paths from Washington DC to Amsterdam. It then shows that LSPs are built between Bos (Boston) and Linx (London Internet Exchange).

Figure 16. Build Explicit LSPs and Place over All Available Paths



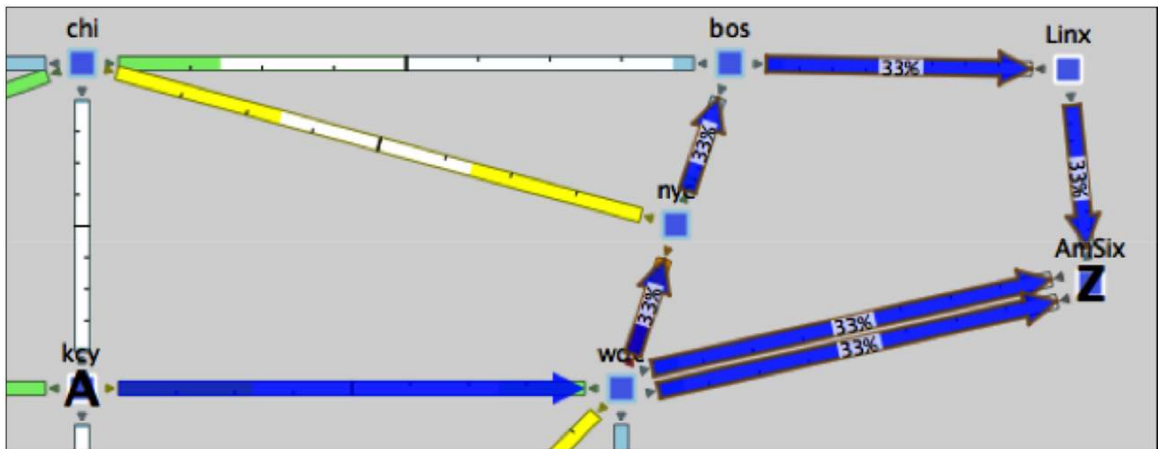
Following this change, an architect can determine whether load sharing is occurring over these LSPs and whether there is an even distribution of traffic over all links. [Figure 17](#) shows load balancing over the northern path.

Figure 17. Load Balancing over the Northern Path



By load balancing over both paths and watching the resulting flows, you can analyze the impact to the terrestrial and the transoceanic circuits to ensure that no new capacity issues have been created. [Figure 18](#) shows load balancing over the southern path.

Figure 18. Load Balancing over the Southern Path



Strategic Traffic Engineering (IGP Metrics)

The engineering group is often faced with the question of whether it is possible to delay a capacity build-out in order to defer capital costs. Often, service level agreements need to be assured with costs kept to a minimum. [Figure 19](#) shows some options you can explore using MATE Design.

Figure 19. Metric Optimization for Strategic Traffic Engineering

The screenshot shows the 'Metric Optimization Settings' dialog box on the left and a report of network metric changes on the right.

Metric Optimization Settings (Normal (Non-Resilient))

- Minimize maximum interface utilization
- Minimize # of interfaces with utilization > 80 %
- Minimize average latency
- Enforce latency bounds

Failure (Resilient)

- Minimize maximum interface utilization
- Minimize # of interfaces with utilization > 80 %
- Enforce no-failure utilization bound: 80 %
- Failures to consider:
 - Circuits
 - SRLGs
 - Nodes
 - L1 Links

Optimization Type: Incremental
 Set Metrics on Interfaces: Core
 Interfaces to Optimize: [List]
 Traffic level: Selected: Default

Network metric changes that need to be provisioned

List of intrasite metrics different from target :

| Node | Remote Node | Interface | Old Metrics | New Metrics |
|----------|-------------|---------------|-------------|-------------|
| cr2.was2 | cr1.was2 | {to_cr1.was2} | 1 | 1064 |
| cr1.ph1 | cr2.ph1 | {to_cr2.ph1} | 1 | 34 |
| cr2.ph1 | cr1.ph1 | {to_cr1.ph1} | 1 | 35 |
| cr1.atl | cr2.atl | {to_cr2.atl} | 1 | 6 |

List of intersite metrics different from target :

| Node | Remote Node | Interface | Old Metrics | New Metrics |
|----------|-------------|---------------|-------------|-------------|
| cr1.chi | cr2.det | {to_cr2.det} | 331 | 4128 |
| cr2.was1 | cr2.nj | {to_cr2.nj} | 192 | 1266 |
| cr2.la | vr.sl | {to_vr.sl} | 517 | 1118 |
| cr1.bos | cr2.det | {to_cr2.det} | 1454 | 936 |
| cr2.det | cr1.bos | {to_cr1.bos} | 1454 | 708 |
| cr1.was1 | cr1.ph1 | {to_cr1.ph1} | 40 | 1112 |
| cr2.chi | vr.m1 | {to_vr.m1} | 251 | 412 |
| cr2.jnc | vr.m3 | {to_vr.m3} | 122 | 985 |
| cr1.atl | cr2.was1 | {to_cr2.was1} | 769 | 1843 |
| cr2.was1 | cr1.atl | {to_cr1.atl} | 769 | 774 |
| vr.m3 | cr1.atl | {to_cr1.atl} | 290 | 446 |
| vr.m1 | cr2.jnc | {to_cr2.jnc} | 477 | 1 |

ROUTING SUMMARY

After (Before)

| Scenario: | Max Utilization over Optimized Links (%) |
|----------------|--|
| Worst-case | 80.6 (131.3) |
| Normal | 71.1 (120.7) |
| if non-res opt | 71.1 |

Improved Max utilization for normal and worst cases

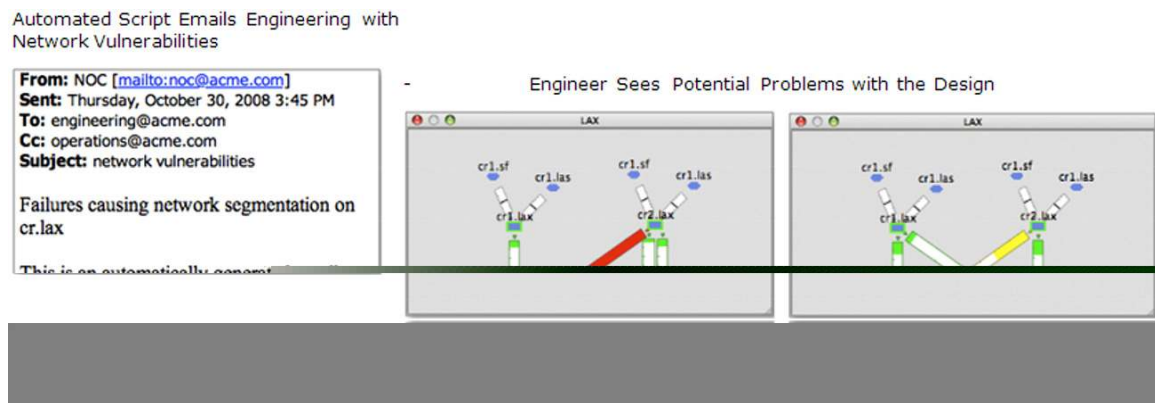
Using the Metric Optimization tool, an engineer selects the option to minimize maximum link utilization while enforcing latency bounds and ensuring availability for circuits and shared risk link groups (SRLGs). MATE Design generates a report of the metric changes that need provisioning to meet these requirements, and shows the maximum utilization that would result for both normal and worst-case failures with these changes. SLAs are assured, and the optimizations result in capital savings.

Design Verification

Network deployments are time-consuming and complicated, and errors can occur at any step. Mishaps can happen from interface port hookup, device configuration, fat-fingering, to bandwidth setup. Cisco MATE Collector can help in ensuring the network is implemented as planned by discovering the deployed network topology and gathering statistics at regular intervals. A visual inspection in Cisco MATE Design can quickly reveal the current state of the network, which can then be verified against the intended design.

Design verification is an area that crosses the functions of engineering and operations. In the example illustrated in [Figure 20](#), an automated script notifies engineering staff that there is a potential problem with a router not being installed as designed in that it was not properly dual homed. An engineer studies the topology and traffic to determine where and how the design is flawed.

Figure 20. Identification of Need to Dual Home Edge Router



Conclusion

Building and running a large-scale network requires a series of activities to be performed; these activities are typically divided into planning (months to years), engineering (weeks to months) and operations (tactical on an ongoing basis). These activities are most efficiently performed using a common infrastructure. The Cisco MATE portfolio is the foundational platform to support planning, engineering, and operations.

For More Information

To learn more about the Cisco MATE portfolio of products visit <http://www.cisco.com>.

For information about other key features of the MATE Portfolio, contact your representative.




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 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)