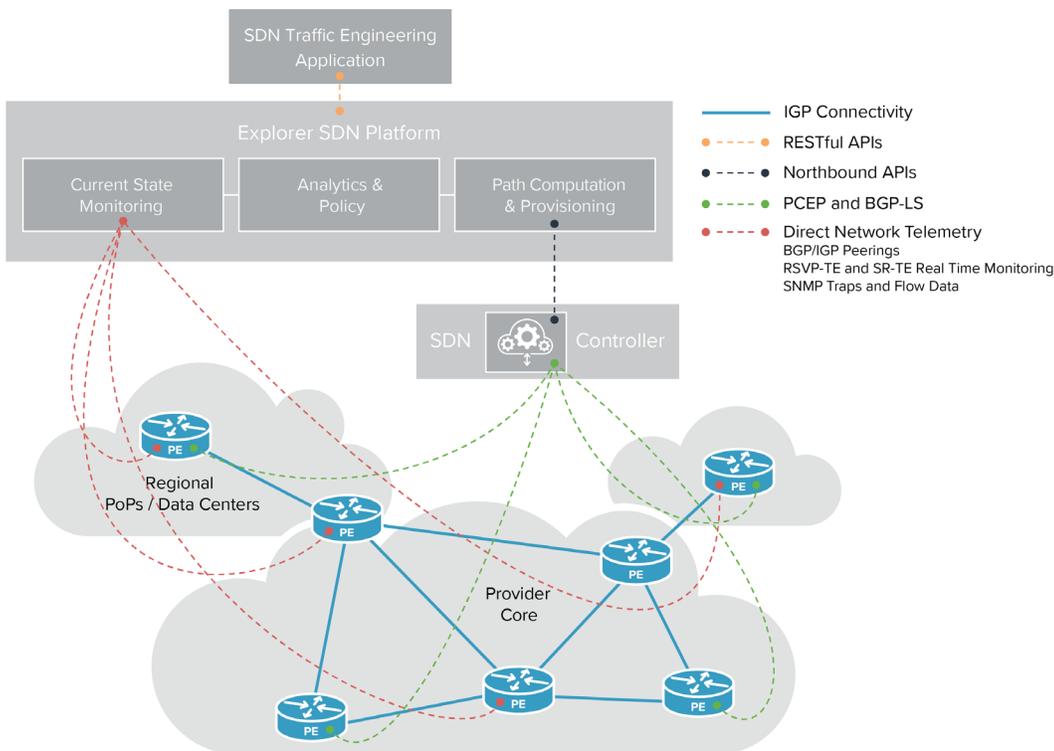


## EXPLORER SDN TRAFFIC ENGINEERING

The Explorer SDN Traffic Engineering (SDN-TE) application is an example of the intelligent automation capabilities that can be achieved with the Explorer SDN Platform. The SDN-TE application leverages the Platform's rich network telemetry and analytics to automate the difficult and time-consuming task of balancing network loads. It does this by creating tunnels to shift loads from heavily-congested links to lightly-used links. This results in better use of network resources and smoother service delivery.

The SDN-TE application taps into several components of the Platform, including real-time and historical routing, traffic and TE Tunnel telemetry, predictive modeling, a path computation engine, optimization algorithms, policy setting, and provisioning via third-party SDN controllers. Like all applications built on the Platform, it accesses these components via open RESTful APIs.



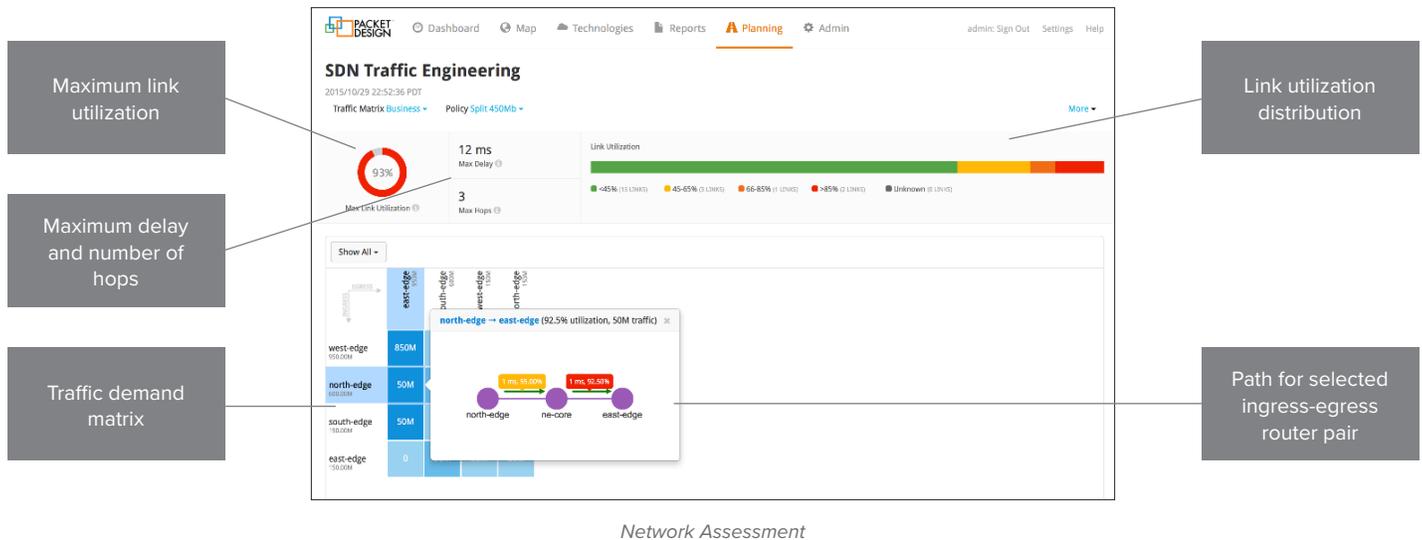
The Explorer SDN Traffic Engineering Application and SDN Analytics and Automation Platform

### Benefits

- Optimize and balance the use of network resources by redistributing traffic loads
- Quickly resolve service slowdowns caused by congested links
- Increase business agility and competitiveness by accelerating TE tunnel-based service activation
- Reduce the operational expense of complex traffic engineering tasks
- Eliminate the risks of service impacts from traffic engineering miscalculations

## Assessing the Opportunity for Optimization

The SDN-TE application assesses the current network state to determine if optimization is needed and possible. Overall network performance is summarized in a single display for the current set of paths and traffic demands, and graphically shows the maximum link utilization, maximum delay and the maximum number of hops. It also charts the number of links that are uncongested, moderately congested and heavily congested. The goal of the SDN-TE application is to reduce the utilization of these overly congested links.



## Traffic Matrices

The SDN-TE application calculates traffic demand matrices that show how much traffic is flowing from each ingress router to each egress router. The cell colors indicate the utilization of the most congested link along the path for each ingress-egress pair. Clicking on one of the cells shows the path for that pair and the congested link along the path.

The SDN-TE application supports full-mesh traffic engineering as well as tactical traffic engineering to manage isolated congestion. Both RSVP-TE and Segment Routing (SR-TE) are supported. The user may select a specific daily, weekly or monthly period for the traffic engineering calculations and use either peak utilization or the 95<sup>th</sup> percentile traffic levels. By specifying a “Maximum utilization per link” value, any link with utilization above this threshold is selected and its ingress and egress routers are automatically inferred.

The application also presents the user with a list of router groups to be included in the matrix. For example, a full-mesh traffic demand matrix will be created if all edge routers are selected. By selecting other router groups, it is possible, for example, to create a demand matrix for just the core routers or for a specific geography.

The SDN-TE application presents the user’s requests to the Explorer SDN Platform and the resulting traffic demand matrix for the selected time period and router groups is displayed. Multiple traffic demand matrices can be created simultaneously and they can either be optimized for the same time or for different times of the day to satisfy the needs of the multiple services running on the network.

Small traffic demands may be excluded from the matrix by specifying a bandwidth value as a filter. This causes only large demands in the matrix to be rerouted. However, even though they are not rerouted, the smaller demands are still factored into the SDN-TE application’s link utilization predictions when making new path recommendations.

Different services require different traffic demand matrices. The definition process allows specification of either a traffic class-of-service, using IP DiffServ or MPLS EXP, or a traffic group definition. The latter allows finer granularity in service specification, including the IP addresses and ports of the service end points.

## Policy

The SDN-TE application user can set policies to define allowable path characteristics. For example, it is possible to specify whether or not traffic may be split across multiple paths or if non-shortest paths should be created. Foremost, a policy specifies an algorithm to use with the default is a fast, global constraint-based optimization algorithm. Others include propagation delay-based and artificial intelligence-based algorithms.

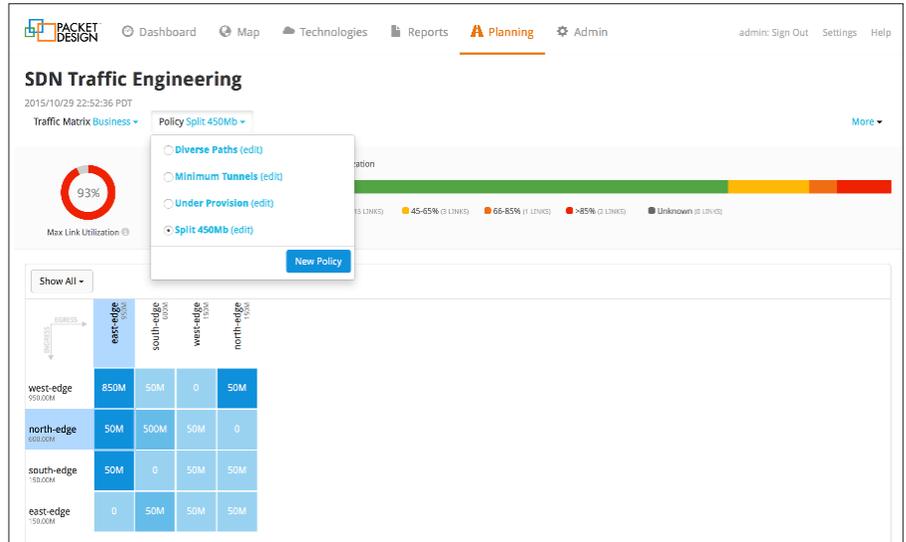
Policy settings also determine whether or not traffic may be split among multiple paths between ingress and egress routers. This is accomplished by specifying a bandwidth value. For example, if this value is set to 450Mb and the traffic volume is 950Mb, three tunnels will be used. If the resulting recommendations yield the same path for two or more of these tunnels, they will be merged to the minimum number of tunnels. Splitting the traffic into multiple paths is useful when no single path has sufficient capacity to satisfy the demand. It is also possible to request that the paths are link-diverse.

Policy is also used to specify over- or under-provisioning of a specific service's demand by a percentage value. For example, the user may elect to over-provision critical business services while under-provisioning less critical services.

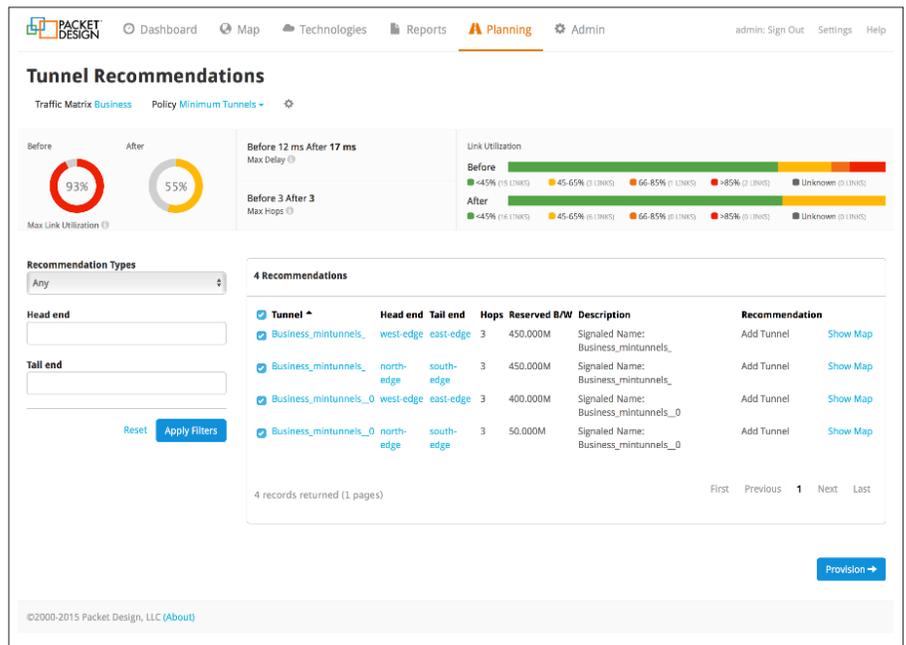
## Intelligent Automation

Once the traffic demand matrices are calculated and the policies established, it can be determined if there are traffic engineering optimization opportunities. With a click of the mouse, the SDN-TE application presents the user with a list of new tunnel recommendations and updates to existing tunnels.

The application displays before and after measures for maximum link utilization, delay and number of hops, and the link utilization distribution. These metrics provide an at-a-glance view of the impact that the calculated new paths will have on network utilization, and if the recommended tunnel changes will eliminate the heavily congested links by redistributing traffic to the under-utilized links.



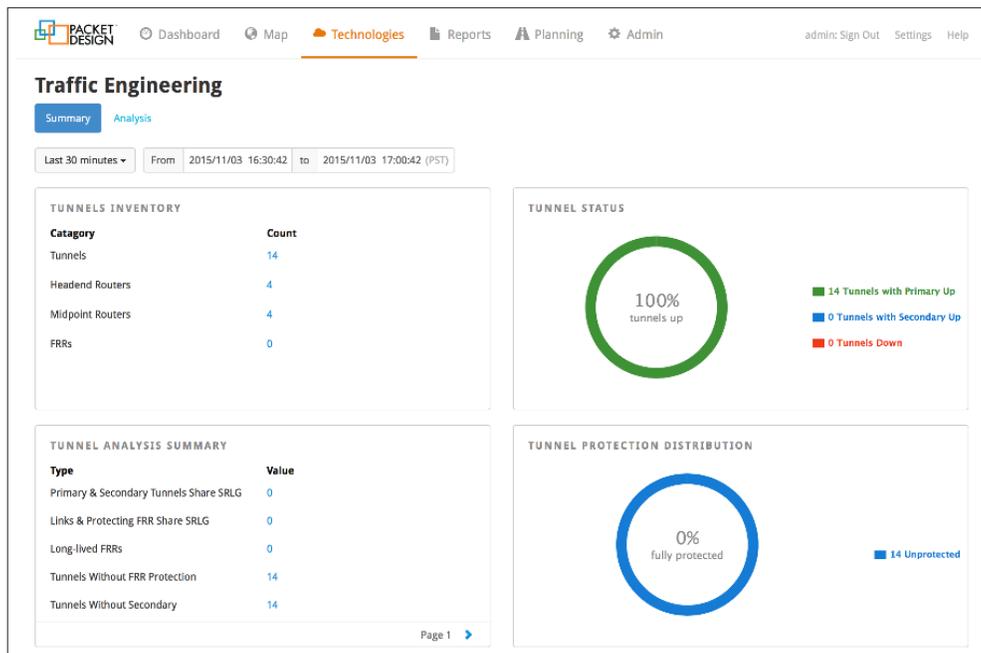
SDN-TE Policy Setting



Tunnel Recommendations and Predicted Impact

Once the user is satisfied that the recommendations will have the desired impact, clicking the SDN-TE application's Provision button signals the Explorer SDN Platform to communicate these paths to the SDN controller via the controller's northbound API. The controller uses a southbound protocol to configure the paths in the network devices and ensures these paths are persistent. The user may update or un-provision them at any time using the SDN-TE application.

The newly-provisioned tunnels are monitored by the Explorer SDN Platform. Any errors, perhaps due to anomalies in the network devices that SDN controllers cannot detect and analyze, will be detected and displayed. The status of the tunnels is shown and whether or not they are protected by fast re-route (FRR) techniques. The Platform's current state monitoring provides real-time telemetry on the provisioned tunnels, not relying on the SDN controller.



Monitoring the Provisioned Traffic Engineered Tunnels

