CHAPTER **18**

Metric Optimization

The Metric Optimization tool optimizes metrics in the network core. It replaces the IGP metrics of core interfaces in a network plan with new metrics chosen to maximize the throughput achievable by the resultant routes, under a specified set of failure scenarios.

The Tactical Metric Optimization tool enables you to quickly fix local congestion by doing as few modifications as possible on selected interfaces to push utilizations under a given level.

Note   The core of the network must be connected, and all nodes must be contained in one AS. That is, there must be a path between any two core nodes that passes only through other core nodes in the same AS, and not through edge nodes. If this is not the case, Metric Optimization signals an error

Core Versus Edge

WAE Design plan files allow nodes to be defined as either core nodes or edge nodes. Note that these definitions are not related to any explicit router configurations. A similar classification of circuits is determined implicitly by the nodes they connect. If a circuit is attached to one or more edge nodes, it is defined to be an edge circuit. Otherwise, it is a core circuit. Edge metrics and core metrics mean metrics of edge or core circuits, respectively.

This core and edge distinction helps define certain policies followed by WAE Design tools. In Metric Optimization, these are as follows.

- Only core metrics are changed.
- The objective is to minimize utilization in the core. Utilization in the edge is not taken into account.
- The routes selected by the metrics are restricted by the tool so that any route between a source and a destination will enter the core at most once, and leave the core at most once. That is, *edge leakage* of routes is prevented.

A consequence of these policies is that larger edge metrics in the plan file to be optimized might result in better optimized core routes. With larger given edge metrics, the Metric Optimization tool has more flexibility in how it can set the core metrics without leakage of routes into the edge.

Metric Optimization

There are many trade-offs in the optimization performed by Metric Optimization. The main trade-off is between performance under normal operation versus performance under failure. Metrics can be selected to do well under the former while ignoring the latter, or to attempt to balance performance between the two.

Metric Optimization optimizes undifferentiated traffic, but does not optimize traffic per service class or per interface queue. If there are policies set on the undifferentiated traffic, then WAE Design attempts to minimize both the interface Utilization and the policy violations.

To open the tool, select the Tools->Metric Optimization menu.

Normal (Non-Resilient) Options

The optimization objectives for normal operation are as follows, in order of priority.

- Minimize maximum interface utilization—This default is usually the primary objective, and concentrates on reducing the utilization of the interface with the highest utilization.

- Minimize # of interfaces with utilization > *x* % —Often a small number of interfaces are bottlenecks in the network, and their utilizations cannot be reduced through rerouting. However, it is still desirable to reduce the utilization of other, under-utilized interfaces. This option attempts to reduce all interface utilizations to below a specified utilization level, or, if this is not possible, to minimize the number of interfaces over this level.

- Minimize average latency—If after the optimizations flexibility is available in the choice of metrics, then WAE Design uses this flexibility to minimize the average latency of the routes across the network. This is not user-configurable and is always on.

- Enforce latency bounds—If selected, the choice of paths for each demand is restricted, if possible, to those demands with latency below the latency bound. For information on setting latency bounds, see the Traffic Demand Modeling chapter.

Failure (Resilient) Options

For planning purposes, the most benefit from Metric Optimization is that it is able to optimize routing over a large number of failure scenarios. This capability controls that optimization, as well as controls the trade-off between normal and failure optimizations.
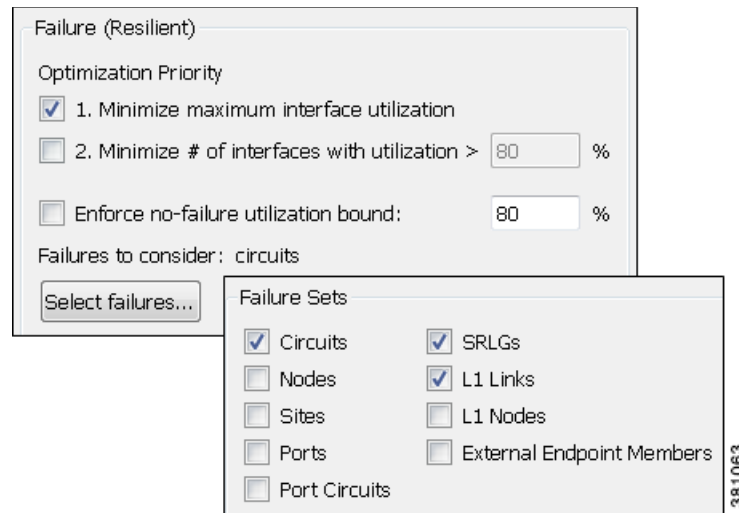
- Minimize maximum interface utilization—If selected, WAE Design selects metrics to minimize the maximum interface utilization over all interfaces and all failure scenarios. If not selected, only normal operation is optimized, and the remainder of the failure section is ignored.

- Minimize # of interfaces with utilization > x % —The number of interfaces with utilization (under any failure scenario) above a specified percentage is minimized.

- Enforce no-failure utilization bound—If not selected, maximum failure utilization is minimized with no consideration for utilization under normal operation. It is often convenient to select this option and place a bound on the utilization under normal operation, which Metric Optimization then satisfies, if possible.

- Failures to consider—This option determines the failure scenarios over which optimization is performed. You can select any combination of failures.

  – All unprotected circuit, node, Layer 1 (L1) node, and L1 link failures

  – All SRLG failures

  This option does not consider port or port circuit failures. For information on protecting objects from being included in this analysis, see the Simulation Analysis chapter.

Note   To control the number of threads that WAE Design processes in parallel when examining failure scenarios, set the "Maximum number of threads" field.

Other Options

You can use these options to fine-tune the optimization.

- Optimization Type—WAE Design can perform two different types of metric optimizations: global or incremental.

- Incremental optimization uses the current routes as a basis from which improvements can be made. The target routes are not changed as radically as when using global optimizations, and therefore the number of metric changes from the metric targets is typically much smaller using the incremental option.

- Global optimization creates optimal routes from scratch, and targets the old metrics afterward. Because WAE Design can choose routes globally, this option tends to perform better under some circumstances, particularly when the target routes are poorly chosen. However, global optimization typically takes longer than incremental optimization.

- Set Metrics on Interfaces—By default, WAE Design sets or changes metrics on any of the core interfaces in the network to achieve its optimization goals. Alternatively, as provided by this option, Metric Optimization restricts metric changes to the following interfaces: all, core, selected in table, or tag name.

- Traffic Level—The traffic level over which the optimization is performed.

- Interfaces to Optimize—By default, WAE Design optimizes all interfaces. You can restrict the optimized interfaces using this option and selecting the desired value in the Optimize utilization on interfaces drop-down list: all, core, selected in table, or tag name.

  For those interfaces that are not optimized, you can set the following.

  - Ignore

  - Minimize number of interfaces with util > ___ %

  - Minimize number of interfaces with util > current + ___ %

- Target to Current Metrics—There are many sets of metrics that will result in equivalent sets of routes. (Doubling all metrics, for example, does not change routes.) WAE Design selects between these sets of metrics by metric targeting. Metrics can be selected to match as closely as possible the current metrics in the network. This simplifies the changeover from one set of metrics to the next.

Targeting optimized metrics to the current metrics in the plan is performed by default. Various targeting refinements are available as CLI options to specify preferred upper and lower bounds on metric selections.

- Enforce Symmetric Metrics—This option sets the two metrics on each interface of each circuit equal to one another. This constraint can reduce optimization performance. However, it ensures that demands from A (source) to Z (destination) and from Z to A (in the core) take the same path, and so the chances of a session between A and Z being interrupted by a failure is minimized.

- Prevent Edge Leakage—This default option restricts how WAE Design can change demand routes. If a demand starts by routing only through the core, or from the edge through the core to the edge, this option prevents the routing from being changed so that it routes through the edge in the middle of its path. This is usually a desired routing policy restriction.

- Simulate after optimization—This default option performs a simulation analysis after optimizing metrics. The simulation analysis allows a detailed view, for example, of the utilizations under failure resulting from the new metric settings.

- New Plan for Result—This default option creates a new plan as a result of running Metric Optimization. Unless a name is specified, WAE Design attaches the suffix *-mopt* to the plan file name. If not selected, WAE Design changes the current plan file with the updated information.

Tactical Metric Optimization

The Tactical Metric Optimization tool is a reduced version of the Metric Optimization tool that minimizes the number of changes and runs faster. The options for this tool are a subset of those for the Metric Optimization tools.

The tool tries to find the most optimal solution it can. In general, the longer it runs, the better the result. The Time Limit option tells the tool to stop looking for better solutions after the specified amount of time, and use the best solution it has found so far. If the option is not set, the tool continues to run until it runs out of solutions to explore, which can take a long time. If you need all available solutions to be explored, use the Metric Optimization tool.

Select the Tools->Tactical Metric Optimization menu.

Report

Metric Optimization writes a Metric-Opt report to the Design History section of the Report window. To access this information later, select the Window->Reports menu.

Metric Optimization Report

- Options—Summarizes the options used to call Metric Optimization. Some of these options (for example, the Output file and Report file) are only relevant when using the CLI tool.

- Number of demands—Total number of demands in the plan.

- Objectives—Summarizes the optimization objectives in order of priority.

Warnings

Warnings detected by Metric Optimization might lead to misleading or undesirable results.

- No Improvement in Optimization

  ```
  No routing improvement found: metrics unchanged
  ```

  Metric Optimization was unable to improve on the metrics in the network. If the incremental option was specified and latency bounds were enforced, this could be because the target metrics violated latency bounds in ways that incremental optimization could not fix.

- Optimization Exit Diagnostics

  ```
  Optimization performance may be limited by low edge metrics
  ```

  During the optimization, some desirable routes could not be selected because the given edge metrics would have caused edge leakage.

- Optimization Constraints Violated

  - `<n> \{intra. inter\site core metrics exceed \crlf \{intra, inter\site-metric \{upper, lower\-bound`

    One or more intrasite or intersite core metric bounds have been violated.

  - Maximum normal utilization exceeds specified bound

    The original plan had normal utilization above the specified bound, and Metric Optimization was unable to reduce the normal utilization below this bound. If this occurs, Metric Optimization tries to set normal utilization as low as possible. No worst-case optimization is performed because reducing utilization under the normal scenario is the highest priority.

  - `<n> demands with non-zero bandwidth exceed latency bounds`

Since enforcing latency bounds is the highest optimization priority, this situation should not happen unless a latency bound is less than the shortest possible latency for a demand.

– <n> demands with zero bandwidth exceed latency bounds

Metric Optimization does not always enforce latency bounds for zero bandwidth demands. These warnings can occasionally occur even if it is possible to find lower latency paths for the signaled demands.

- Unrouted Demands

    – <n> unroutable demands under normal operation

    There are <n> demands for which no route exists between source and destination, even under the No Failure scenario.

    – <n> unroutable demands under <m> circuit failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> circuit failure scenarios. (Different demands might be unroutable under different circuit failures.)

    – <n> unroutable demands under <m> SRLG failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> SRLG failure scenarios. (Different demands might be unroutable under different SRLG failures.)

    – <n> unroutable demands under <m> source/dest node failure scenarios

    There are <n> demands whose source and/or destination nodes fail under <m> failure scenarios. These demands clearly cannot be routed in these circumstances. (Different demands might be unroutable under different node failures.)

    – <n> unroutable demands under <m> non-source/dest node failure scenarios

    There are <n> demands for which no route exists between source and destination under <m> node failure scenarios, where the nodes that fail are intermediate nodes in the path of the demands, and are not source/destination nodes of the demands. (Different demands might be unrouted under different node failures.)

Routing Summary

Summary statistics are displayed of the routes in network, both before (in parentheses) and after the optimization.

- Core/Edge Max Utilization—The maximum percentage utilization over any interface in the core or edge network. The normal utilization is under the normal (no-failure) scenario, and the worst-case utilization is maximized over all failure scenarios in the failure set.

- Latency—Median and average latencies over all routed demands, both in milliseconds and as a percentage of the latency of the smallest latency route possible.

    The latency of each demand's routing is calculated as a percentage of the latency of the shortest possible route, by latency, for that demand. Under the percentage column, the median and average of these percentages are displayed.

- Num of routes away from shortest path—The number of routes (out of all the demands routed) that do not follow the shortest path, by latency. This statistic is an indicator of how much the routes have been affected by utilization minimization as the primary criterion for optimization.

Metrics

The metric targeting choice is shown, together with the number of metrics different from the target metrics. Lists of interfaces for which metrics differ from the target metrics are displayed, together with both the target metric and the optimized metric. The metrics are separated into those on intrasite and intersite interfaces.

Related Topics

- IGP Simulation chapter
- Traffic Demand Modeling chapter
- Simulation Analysis chapter
- Layer 1 Simulation chapter