



Aria Networks

Tech-Brief

Traffic Distribution in Traffic-Engineered Networks



Traffic Distribution



Introduction

The manner in which traffic is distributed across a telecom carrier's network is fundamental to the operational efficiency of the network. Efficient traffic distribution ensures the network is utilised optimally and can lead to a reduction or deferral of capital expenditure.

This technology brief presents the objectives of traffic distribution in traffic-engineered networks, describes the concept of Least Load Routing (LLR) and also briefly describes the LLR capabilities of Aria's iVNT.

Traffic Distribution –The Objectives

The following are the key objectives of telecom carriers when considering distribution of traffic within their network:

- To reduce or defer the need for additional capital expenditure by better leveraging the deployed assets.
- To avoid congestion, within the capacity constraints of the available network resources.
- To minimise the impact to traffic under failure and outage scenarios.
- To maximise available network capacity for future expansion, or in other words, to maximise the potential for adding new services and traffic without requiring additional network resources or the re-optimisation or repositioning of existing services.
- To distribute certain traffic types in order to maximise capacity for 'other' traffic. For example, in MPLS-TE networks, to position MPLS-TE LSPs so that there is spare capacity on all links to carry IGP-routed MPLS and IP traffic.
- In particular in MPLS-TE networks, one objective is to make Fast Reroute (FRR) protection more practicable by ensuring that; no link is operating at maximum capacity which would require that it was protected to its maximum capacity which is unlikely to be possible as most links in a network have a similar maximum capacity; all links have spare capacity to carry detour LSPs.

Least Load Routing – The Concept

Least Load Routing (LLR) is a mechanism that attempts to minimise the utilisation of all the links in a network by distributing the traffic demands across all available network resources.

The principal of LLR, if applied as an absolute requirement in determining traffic paths, potentially creates some very long paths as it seeks to place traffic equally on all links. Consider the network in Figure 1. If two services are requested from A to B and from A to C, it would be wrong to select the two paths [A-B, A-D-E-F-G-H-I-J-K-L-M-N-O-C]. Looking at these paths, one might decide that the load has been evenly distributed over as many of the available links as possible, but this has not achieved the desired result, it has simply increased network utilisation. On the other hand, the solution [AB, ADC] is very reasonable.

So it can be seen that any LLR mechanism must aim to minimise total bandwidth usage while still aiming to distribute the traffic within the network.




Figure 1 – Unnecessary Long Path

When used by telecom carriers LLR is rarely used as an absolute requirement, in practice other path constraints are often factored in. Typically, the bandwidth requirements of the path remain top priority, so that the minimum bandwidth for a path is always delivered and that the attempt to provide LLR is not used as a reason to deliver less than the desired bandwidth.

Cost, delay, jitter and hop-count caps (that is, their maximum values) are additional examples of constraints that are intended to influence the paths selected, and like bandwidth, should optionally constrain the interpretation of LLR. This is, in fact, the most likely usage of LLR by a telecom carrier. For example, telecom carriers may require LLR, but they do not want any path cost to be increased excessively and will typically place a cap on the cost of the paths. Thus LLR should be achieved within the per-path constraints, as these per-path constraints are necessary to protect the integrity of the service being delivered.

The telecom carrier also expects other constraints, such as 'minimise cost', to be traded with the desire for LLR. For example, if the telecom operator requires 'minimum cost' but also requires 'LLR' there is an expectation that imperfect LLR will be achieved, and that some paths might not take the least cost path. Conversely, the expectation is that excessively costly paths will not be chosen just to achieve a slight improvement in LLR.

Consider also the use of point-to-multipoint (P2MP) paths and LLR. P2MP tree shapes and their associated computation techniques are inflexible. This means that if a telecom carrier requires a Steiner P2MP tree, this is what must be computed. However, just like any other constraint, the Steiner constraint could be balanced with other constraints. For example, if a Steiner tree is being computed it is the telecom carrier's expectation that an attempt will be made to minimise the sum cost for the P2MP LSP and that this attempt may be traded with an attempt to achieve LLR.

iVNT – Least Load Routing

iVNT enables LLR to be easily and extensively modelled and provisioned, enabling a telecom carrier to have full control over the ability to load balance traffic in the network. The following briefly describes the LLR capabilities in iVNT.

iVNT allows a telecom carrier to optionally implement LLR. iVNT can achieve placement of traffic paths (e.g. LSPs) that meet the full required bandwidth and other constraint requests for each path, as a priority over the request for LLR. However, this prioritisation is fully tuneable using the configurable constraint-weights of iVNT. A "constraint-weight", is the importance of a constraint, relative to other constraints, in determining path selection. If a telecom carrier decides to weight LLR as more important, relative to other path constraints, iVNT can facilitate modelling using the relative priority of the operator's preference.

iVNT enables the LLR placement of individual traffic paths to be achieved while also achieving the path's minimum bandwidth requirements, and the LLR placement of other paths does not curtail the placement of any other path because its minimum bandwidth cannot be met.

By default within iVNT, cost, delay and hop-count caps have higher priority (if used) and they out-weigh the selection of LLR, but this prioritisation is also fully tuneable using the configurable constraint-weights of iVNT.

iVNT also allows the full trade-off through the configuration of constraint-weights between LLR and all other constraints in the system. Note that both 'minimise hops' and 'minimise cost' are intuitively contradictory to 'LLR', but that this can also be handled by iVNT in the same way that iVNT handles all contradictory constraints that can influence path selection, that is, by the use of constraint-weights.

iVNT allows the relative weightings of 'cost' and 'LLR' to be traded for P2MP LSPs as well as for P2P LSPs.

Summary

LLR and the capabilities provided by iVNT enable telecom carriers to optimally distribute traffic across the network to ensure the operational efficiency of the network, and:

- to reduce capital expenditure
- to avoid congestion
- to minimise the impact to traffic under failure
- to maximise available network capacity for future expansion
- to distribute certain traffic types in order to maximise capacity for 'other' traffic