



Aria Networks

Tech-Brief

**End-to-end Protection in
Traffic-Engineered Networks**



End-to-end Protection



Introduction

Planned or un-planned network outages are unavoidable events in telecom networks. To meet customers' service level agreements telecom carriers often require services to be protected such that traffic can be diverted during outages. There are many schemes that can be applied to achieve end-to-end protection and this technology brief describes the most common.

Aria Networks' iVNT encompasses many capabilities to help telecom carriers assess and avoid the impact of network outage scenarios. The capabilities of iVNT to provide protection mechanisms are described in this technology brief.

End-to-end Protection

End-to-end protection is a circuit-based protection scheme that models methods traditionally used in transport networks. It is appropriate to label switched networks because Label Switched Paths (LSPs) are essentially circuits or virtual connections through the network.

End-to-end protection provides an alternate circuit onto which traffic can be switched if a problem is detected with the preferred circuit. The naming distinction "working path" and "protection path" is usually applied.

It is usual to require diversity between the working and protection paths, i.e. the working and protection paths do not share common network resources, often with the exception of the source and destination nodes (routers/switches) of the network. If the paths are not diverse then they are both vulnerable to single failure points. For example, if both paths transit a single link, then the failure of that link would impact both paths.

There are three resource diversity issues that can apply to the path: link diversity, node diversity, and Share Risk Group (SRG) diversity.

- Link diversity means that the two paths share no common link. Usually this means "share no common link in either direction" – that is, the path A-B-**C**-D-E and the path A-F-**D**-**C**-G-E are not link diverse as they share the link between C and D.
- Node diversity means that the two paths share no common node, with the typical exceptions of the source and destination nodes. Note that node-diverse paths are automatically link-diverse, but that link-diverse paths can share nodes. A node in this context could be, for example, a Provider (P) or Provider-Edge (PE) router.
- SRGs are collections of network resources that have a shared failure risk. Most commonly this relates to links provided by fibres that use the same cable duct and so the links are susceptible to damage to the duct. It is also increasingly used to refer to resources that lie in the same area of terrorist susceptibility. SRGs are also used notably in certain geographies to refer to all links that cross seismic fault lines even if they may be many miles apart. Other examples include; the resources that are contained in a building; the resources that share a power source.

Diversity in any of these three categories may be required or desired. Desiring diversity may be particularly helpful when protectable domains are connected together. An example is two ring topologies that are connected by a single ring-interconnect.

Protection may be provided in a variety of ways:

- **1+1 Protection**
In 1+1 protection, traffic is actively transmitted on both the working and protection circuits. The receiver is responsible for determining which signal to utilise. 1+1 protection is very fast, that is, it typically minimises traffic loss under failure scenarios, but is very expensive because the network resources must be configured and utilised and are therefore not available for use by other traffic.
- **1:1 Protection**
In 1:1 protection, traffic is only transmitted on the working path. When there is a failure, traffic is switched to the protection path. 1:1 protection is slower to switch traffic from the working to the protection path than 1+1 protection, but is almost as expensive. Some saving may be made at the source and destination nodes, but otherwise, the main saving comes from the 'extra traffic' option described below.
- **1:n Protection**
In 1:n protection, traffic is only transmitted on the working path. When there is a failure, traffic is switched to the protection path, but multiple working paths share the same protection path. Therefore a single failure among the set of n working paths causes the one protection path to be used. When the protection path is being used the other working paths are no longer protected. Note that each of the working paths should be disjoint from the others otherwise a single failure could require more than one working path to switch to the protection path. 1:n protection is popular because it requires less network resource per protected path.
- **m:n Protection**
In m:n protection, the behaviour is as for 1:n protection, but there is a larger pool of protection paths available. Note that the m protection paths do not need to be mutually disjoint, and that if m out of the n working paths are not disjoint, protection can still be guaranteed. m:n protection offers more protection than 1:n and handles various diversity problems, but is m times more expensive.
- **Extra Traffic**
Extra Traffic is a protection option available in 1:1, 1:n and m:n protection schemes. In each case, the protection path is allowed to carry other traffic while it is not being used to carry traffic from a failed working path. In the event of a failure, the 'extra traffic' is pre-empted from the protection path so that the protected traffic can be carried. This mode of protection offers a way to generate revenue from protection resources.

- **Restoration**
Restoration is the process of computing a new working path after the failure of an existing path. A restoration path only has to be disjoint from the failed resource (link, node, SRG) and so it is easier to place such a path within the network. Additionally, a restoration path does not need to be pre-provisioned so network resources are not tied up. Note, however, that restoration is slow to restore traffic as new paths must be computed and signalled before the traffic can be restored.
- **Mesh Protection**
Mesh protection is not an end-to-end protection mechanism. In mesh networks it is possible for the protecting paths of two distinct LSPs (for example, A-B-C-D and E-F-G-H) to share common links. In this example, the protection paths might be (A-P-Q-**R-S**-D and E-X-Y-**R-S**-Z-H). If 1:n protection is used, it is possible for the resources on the common protection link (R-S) to be shared between the two protection paths.

iVNT – End-to-end Protection Capabilities

iVNT enables protection to be easily and extensively modelled and provisioned. The following describes the protection capabilities in iVNT.

All service types, including point-to-point, point-to-multipoint and mesh based services, can be optionally protected and each individual service can be optionally configured with one of three settings; "Unprotected", "End-to-end protection desired", or "End-to-end protection required". A request for a protected service results in a working and a protection path being computed.

iVNT supports the 1+1 and 1:n protection models. iVNT also has the ability to compute a path that is SRG diverse, i.e. the working and protection paths do not share any common SRGs.




Figure 1: A Simple Point-to-point Protected Path in iVNT

iVNT also supports link and node diversity preferences for protected paths. Link diversity can be optionally 'required' or 'best effort'. Node diversity can be optionally 'required', 'best effort', or 'don't care'. Note that when working and protection paths are computed that are required to be node diverse, they are automatically link diverse.

iVNT enables modelling and analysis of restoration scenarios for protected paths. Paths that can be restored after a network failure can be modelled by iVNT and those not requiring restoration can be optionally "ignored" using iVNT "pinning" capabilities. The capability to "pin" a service enables a telecom operator to optionally restore a service or retain the service in its failed state.

By default, iVNT ensures the parameters of the working path are inherited by the protection path. This ensures, for example, the protection path is allocated the same bandwidth and other characteristics as the working path. However, the protection path characteristics can be modified by the operator, for example it may be desirable to offer various tariffs for protected services where the characteristics of a service change under failure conditions, such as lower bandwidth under failure conditions.

iVNT enables the operator to have complete control of all the parameters and constraints that influence the path selection. This is achieved using the configurable constraint-weights of iVNT. A "constraint-weight", is the importance of a constraint, relative to other constraints, in determining path selection. One such constraint that can be weight-adjusted is "End-to-end Protection". This means the requirement for end-to-end protection can be optionally ignored allowing a service to be established without the protection path, and the "desire" for end-to-end protection can be scaled against other constraints for the protection path, and the requirements of other paths.

If working and protected paths are computed and they do not have identical results due to the characteristics of the network topology (e.g. different costs, delays, hop counts) iVNT selects as the working path the one with the lower values for the constraints of most concern, i.e. those with the higher constraint-weights.

iVNT is able to compute protected services and show the operator: (a) if the working and protection path requirements have been achieved for the paths, (b) if the working or protection path requirements have not been fully met, (c) if the working path requirements have been fully met, but the protection path requirements have not been fully met (despite the protection being desirable or required).

When the network is unable to meet the requirements or constraints of the protected service, iVNT provides the operator with feedback of the "failed constraints". This gives the reason(s) why the protected service could not be established and allows the network operator to change the network design or relax specific service requirements or constraints to enable the paths to be successfully deployed.

iVNT enables the operator to distinguish the two paths of a protected service by indicating that one is the working path and one is the protection path. iVNT displays full details, properties and routes of the computed working and protection paths in both tabular and graphical form so the operator can fully visualise the paths. iVNT also provides output, in a variety of forms, for example via XML, CSV or via vendor specific formats, so the paths can be provisioned in the network.

iVNT allows protected services to be modified and deleted. iVNT also allows the operator to change a protected service into an unprotected service.

Summary

The capabilities provided by iVNT enable end-to-end protection to be easily and extensively modelled, analysed and provisioned allowing telecom carriers to assess and avoid the impact of network outage scenarios in label switched networks and to meet the contracted customer service level agreements.