Annex [?] (informative)

1:1 Protection with Load Sharing Use Cases

This material is intended to assist the reader in understanding the use of 1:1 protection with load sharing and in evaluating the benefits of load sharing in particular PBB-TE network configurations. The set of use cases described is not intended to be exhaustive.

To more easily understand the use cases a brief description of load sharing and the associated protection model is helpful. In conventional 1:1 protection schemes traffic is carried between two protection points over a single “working” path. A second “protection” path is reserved to carry the traffic in the event that the working path fails. In the case of PBB-TE protection with load sharing a protection group comprising two or more TE service instances (TESIs) is provisioned to carry traffic between the two protection points, and backbone service instances (BSIs) are distributed (the load is shared) across some or all of these TESIs. In the event that one of the TESIs fails the traffic assigned to the failed TESI is reassigned to the other TESIs until the failed TESI is restored. This reassignment may also be distributed (load shared) across the remaining TESIs. The reassignment is pre-planned to enable rapid protection switching.

Protection with load sharing provides several benefits which are illustrated by the use cases below. Among these are

- The network operator may provision a protection group of TESIs on a set of acceptable paths between protection points, normally between edges of a PBB-TE network domain, and manage the distribution of BSIs to these TESIs to meet changing traffic demands according to traffic engineering policies. The traffic distribution within the protection group can be altered to adjust link loads and optimize network resource use without requiring (or with less frequent) provisioning of new TESIs.
- The bandwidth reserved to protect traffic in the event of failure may be substantially reduced, providing more efficient resource usage.
- If some BSIs are to be protected and some are not, a load sharing protection group can be provisioned to support both protected and unprotected BSIs rather than requiring a protection group for protected BSIs and a separate TESI for unprotected BSIs.
- For cases in which conventional 1:1 protection behavior is desired, it can be accommodated within the load sharing protection model as a subcase in which there are two TESIs in the protection group and all BSIs are distributed to one of the TESIs under normal conditions and redistributed to the other TESI in case of a fault on the first TESI.

The following use cases provide some illustration of these benefits.

[?].1 Use case 1 – 1:1 protection using multiple routes

Figure [?]-1 shows a use case in which four distinct PBB-TE Service Instances (TESIs) have been provisioned along each of four diverse paths between BEB-A and BEB-C.

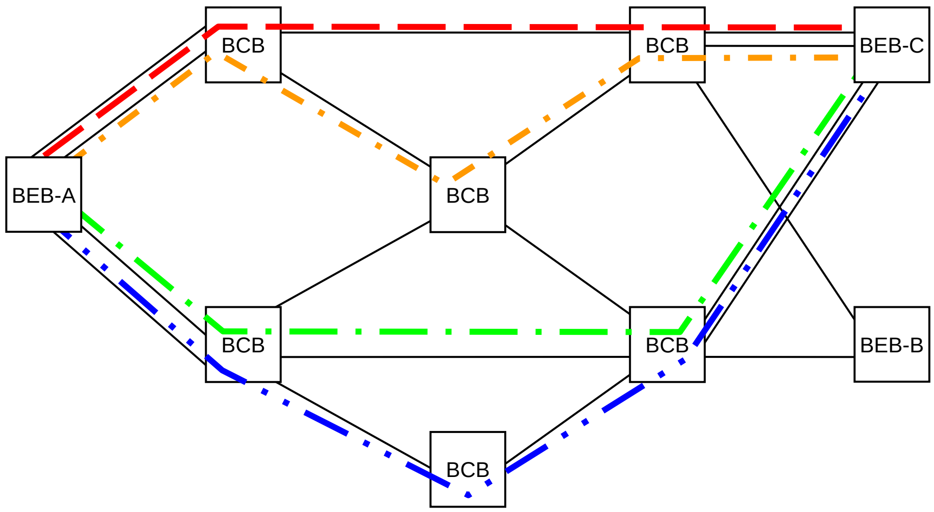

Figure [?]-1 — Multiple Diversely Routed TESIs

Each TESI carries some number of Backbone Service Instances (BSIs) between BEB-A and BEB-C. The assignment of BSIs to TESIs determines distribution of traffic between BEB-A and BEB-C across the core bridges and links in the network. If there is low or moderate traffic load, more traffic may be assigned to the shorter than to the longer paths to consume resources on fewer links. If there is a high traffic load it may be distributed roughly evenly to the various TESIs to avoid overloading links and to reduce the bandwidth reserved for protection.

Using conventional 1:1 protection the bandwidth reserved for protection is 100% of the working bandwidth. Using 1:1 protection with load sharing the bandwidth reserved for protection can be significantly reduced.

For example, if we consider a scenario using the network in Figure [?]-1 in which there are 12 BSIs with equal bandwidth requirements to be transported between BEB-A and BEB-C, these BSIs can be evenly distributed over the four TESIs by assigning three BSIs to each TESI as shown in Figure [?]-2. The assignment of BSIs to TESIs is accomplished by setting the B-VID field in the appropriate CBP BSI table entry to the value corresponding to the selected TESI.

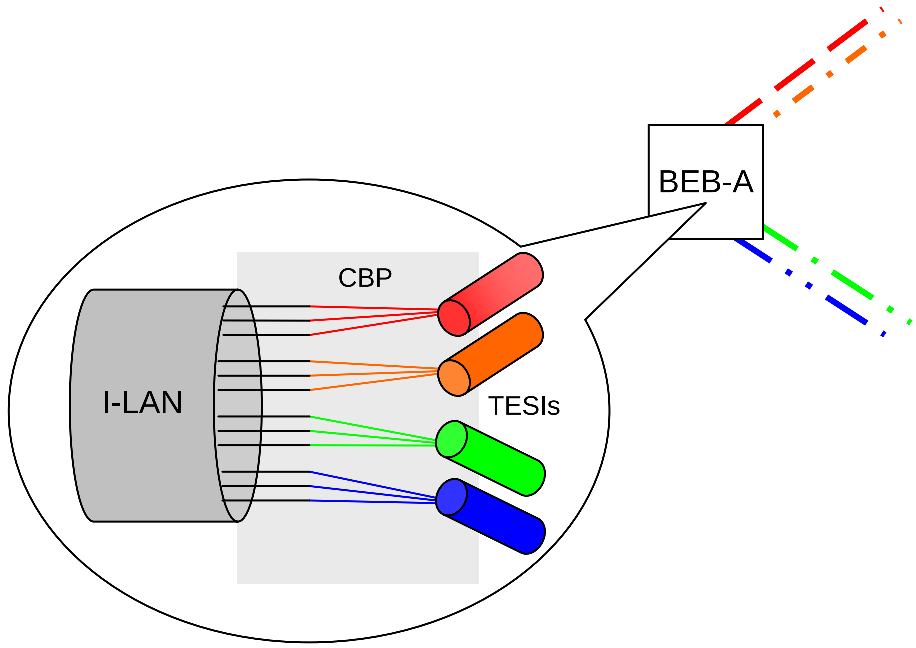

Figure [?]-2 — BSI Distribution over Multiple TESIs

To protect the traffic between BEB-A and BEB-C against the failure of any one of the TESIs, additional bandwidth is reserved for each TESI to provide sufficient capacity to carry some of the BSIs assigned to each of the other TESIs. This protection bandwidth is shared under the assumption that only one TESI in the group will have a failure at any given time[1]. In the current example, each TESI is carrying three BSIs. If each TESI is provisioned with sufficient bandwidth to carry four BSIs (one extra BSI) then there will be sufficient capacity to recover from any single link failure in the network.

To provide 1:1 protection, each of the BSIs is assigned both a working TESI and a protecting TESI to which the BSI will be switched if its working TESI fails. In this example, each of the BSIs assigned to a given working TESI is assigned a different protecting TESI (sharing the protection load for that working TESI). When one TESI fails the three BSIs assigned to it are switched to their protecting TESIs as shown in Figure [?]-3.

[1] In multiple failure cases traffic may be lost, which is also the behavior in conventional 1:1 protection.

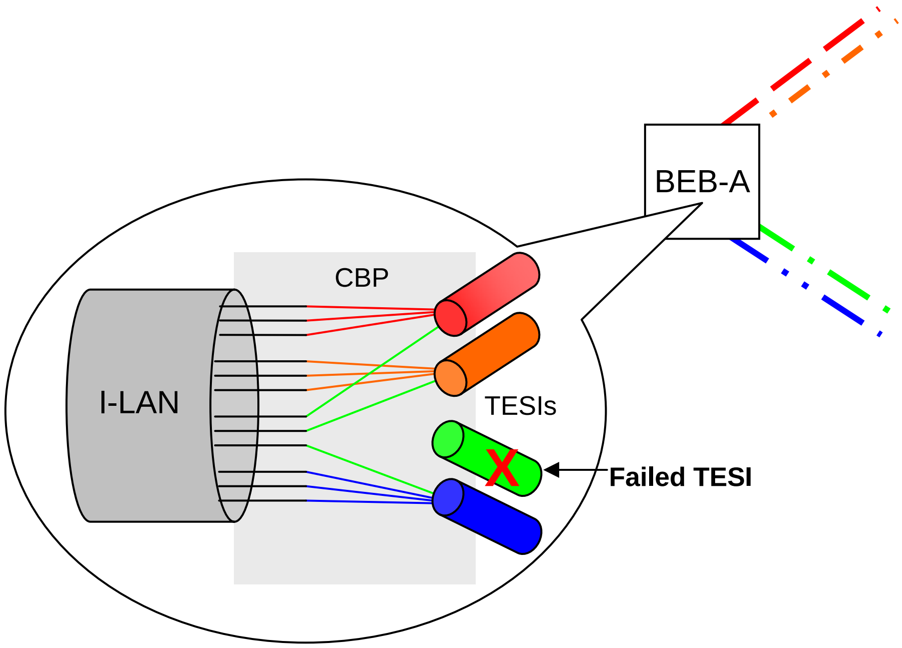

Figure [?]-3 — 1:1 Load Shared Protection Switching

This form of load shared protection switching follows the operational model of 1:1 protection, and the load sharing adds the advantage that traffic can be distributed across the core network to meet overall traffic engineering policies under both normal and failure conditions. Furthermore, the amount of bandwidth reserved for protection is reduced by the load sharing arrangement. In this example conventional 1:1 proteection would reserve 12 BSIs of bandwidth for protection. Using load sharing significantly less bandwidth is reserved for protection, the equivalent of only 4 BSIs. A more general analysis of protection bandwidth savings is provided in section [?].4.

[?].2 Use case 2 – 1:1 protection with TE over parallel links

Figure [?]-4 shows a use case in which there are multiple parallel links along a route between BEB A and BEB C. Again, each TESI carries multiple BSIs. In this use case link aggregation is either not available or does not provide sufficient (deterministic) control over the distribution of traffic across the parallel links to meet the service provider's traffic engineering requirements. Thus the parallel links are not aggregated. Instead, a set of TESIs are configured to provide paths using each of the parallel links. BSIs are distributed across the TESIs to achieve the desired traffic engineered load on each link.

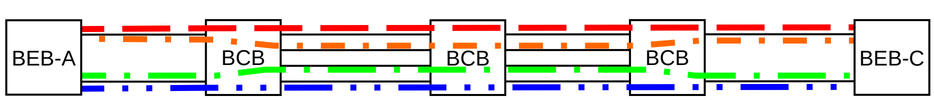

Figure [?]-4 — TESIs over Multiple Parallel Links

As in use case 1 above, each BSI is assigned a working TESIand a protecting TESI. In the event of a failure on one of the parallel links used by the TESIs between BEB-A and BEB-C, the BSIs on the failed TESI are switched to their protecting TESIs by changing the B-VID field in the CBP backbone service instance table.

The potential for roughly even load distribution and sharing of bandwidth reserved for protection in each of the TESIs described in use case 1 applies in this use case as well.

[?].3 Use case 3 – conventional 1:1 protection

In cases where conventional 1:1 protection switching behavior is desired, the load sharing protection mechanism can accommodate this using a protection group containing two TESIs as shown in Figure [?]-5.

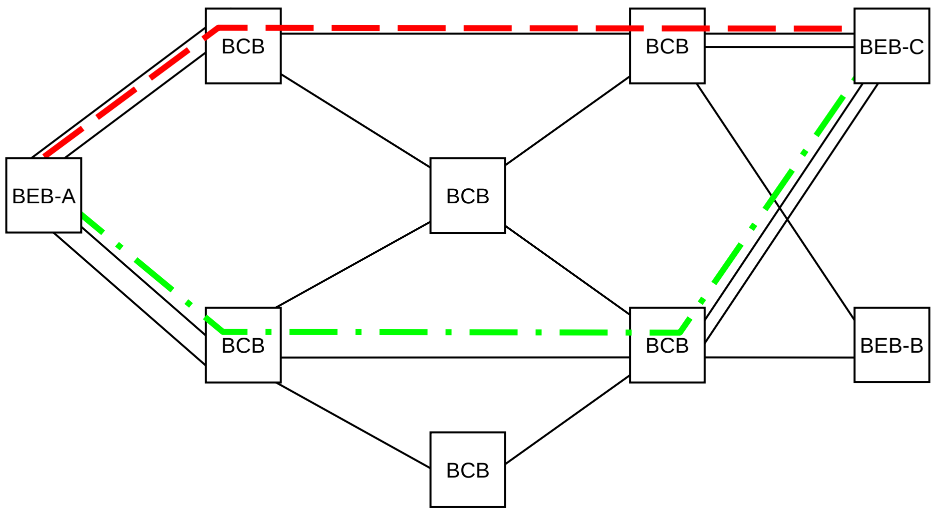

Figure [?]-5 — Two Diversely Routed TESIs

All the BSIs are assigned to one of the TESIs (the “working” TESI) as shown in Figure [?]-6.

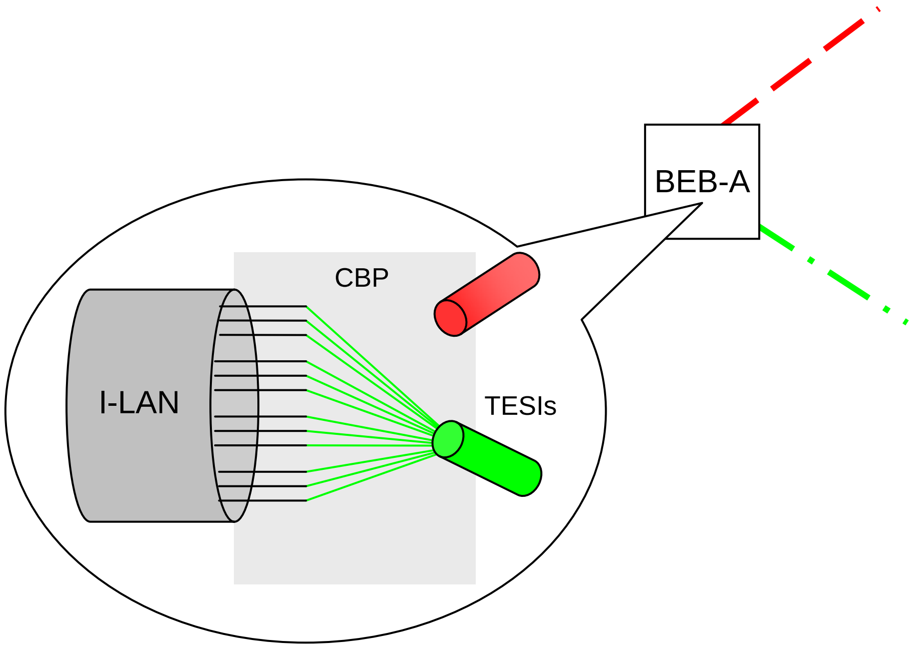

Figure [?]-6 — BSI Distribution to "Working" TESI

In the event of failure of the working TESI all the BSIs are reassigned to the other ("protecting") TESI as shown in Figure [?]-7.

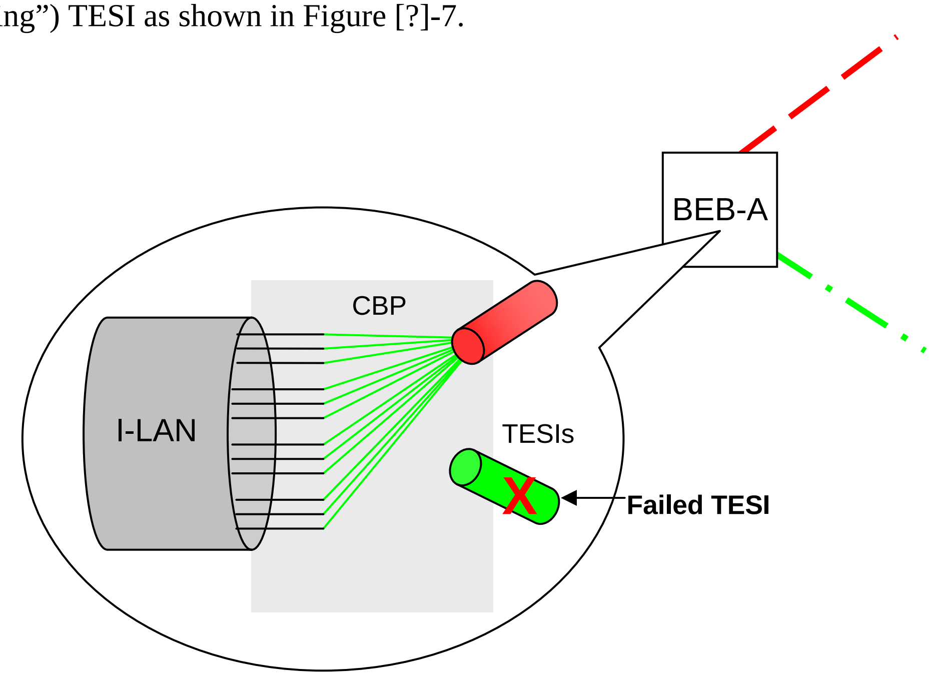

Figure [?]-7 — Conventional 1:1 Protection Switching

Thus PBB-TE 1:1 protection with load sharing can support the conventional 1:1 use case with the same protection model that supports the other use cases.

[?].4 Bandwidth gain using load sharing

In this section traffic engineered networks using 1:1 protection with and without load sharing are compared. It is shown that the use of load sharing can reduce the level of link bandwidth commitment required to provide 1:1 protection of traffic. Thus, for a given link capacity and a given set of parameters described within the analysis, more protected traffic can be carried when load sharing is used than when load sharing is not used.

Variables

S: The number of TESIs providing connectivity between a pair of BEBs[2].
B: The total working bandwidth reserved for traffic between two BEBs.
N: The number of TESIs sharing a link.

Network Description

The PBBN consists of BEBs at the periphery and BCBs in the interior. The number of BCBs and links between BCBs is sufficient to ensure S disjoint TESIs between each pair of BEBs as illustrated by [?]-8. The figure shows a link shared by two TESIs.

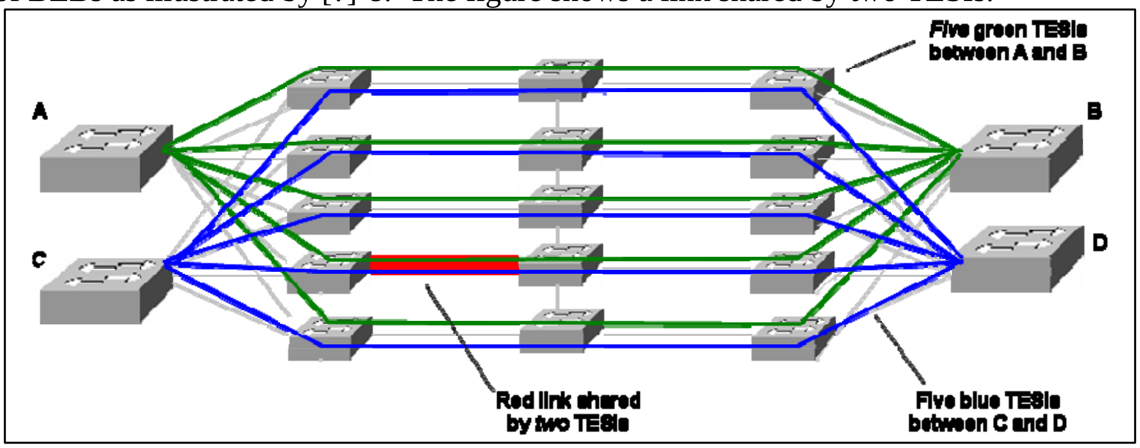

Figure [?]-8: PBBN Example Showing S = 5 and N = 2

Each TESI is assumed to carry a large number of BSIs relative to the value of S as illustrated by Figure [?]-9. That is, we assume that BSI capacity is sufficiently granular that BSIs can be distributed among TESIs such in such a way that bandwidth commitments associated with the TESIs are approximately equal.

[2] For purposes of this analysis, we assume that this number is the same for each pair of BEBs.

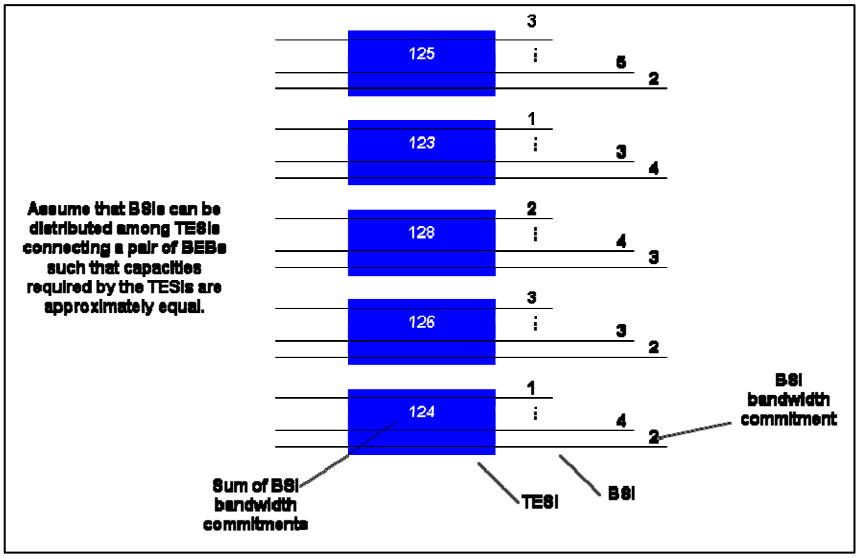

Figure [?]-9: Assume that BSI bandwidth requirements can be distributed nearly equally among TESIs

Bandwidth Analysis

Comparison of Link Utilization with and without load sharing with equal numbers of TESIs.

Load-sharing: Recall that B is the total working bandwidth associated with the set of TESIs interconnecting two BEBs. In the case of load-sharing, each TESI carries B/S working traffic. We assume that the failure of a TESI, e.g. as the result of a link failure, results in the distribution of the component BSIs to the remaining S-1 TESIs. Thus, each TESI must be associated with a reservation of B/S for working traffic and (B/S)/(S-1) for protection traffic, or a total reservation of (B/S)*(1+(1/(S-1))). A link thus carries (NB/S)*(1+(1/(S-1))).

Non-load-sharing: In the non-load-shared case, the working bandwidth associated with a TESI is associated with one link and the protection bandwidth associated with that TESI is associated with a different link. Hence, 2N commitments of bandwidth B can be spread across the set of S links. Assigning bandwidth to links in this fashion results in a worst-case commitment of B*CEILING(2N/S) per link.

It follows that for given values of S and N, we can compute the ratio of link load without load sharing to link load with load sharing as:

B*CEILING(2N/S) / ((NB/S)*(1+(1/(S-1))))

Or

CEILING(2N/S) / ((N/S)*(1+(1/(S-1))))

Or

(S(CEILING(2N/S))) / (N(1+(1/(S-1))))

This ratio represents the bandwidth gain realized by using load sharing. This bandwidth gain is shown as a percentage in Figure [?]-10 below for N in the range of 10..20 and S in the range of 2..10.

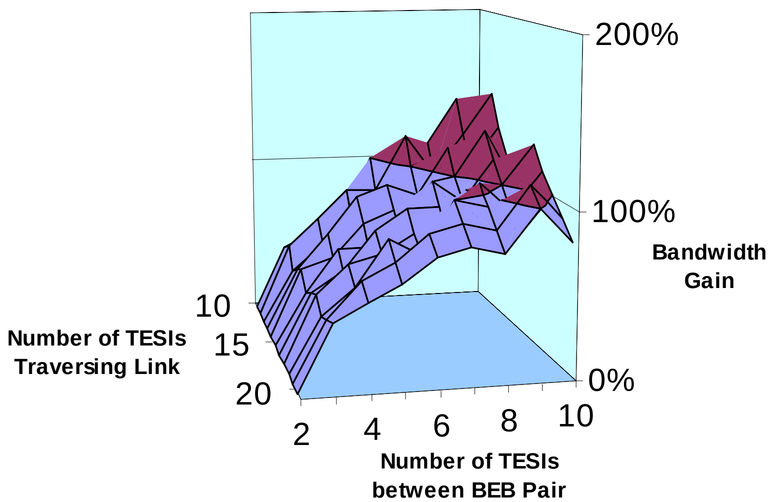

Figure [?]-10: Bandwidth Gain using Load Sharing

Discussion

It can be seen that the benefit of load-sharing increases with the number of disjoint TESIs forming the protection group. As might be expected when only two TESIs are in the protection group there is no advantage. As the number of TESIs in the load sharing protection group increases the bandwidth gain approaches 100 %, with some cases showing better than 100% gain due to the granularity of TESIs. As the number of TESIs per link increases the granularity effect is reduced. That is, load can be somewhat more effectively spread across links in the non load sharing case as the number of TESI's sharing a common set of links increases; however, bandwidth utilization is always significantly improved by load sharing if more than two TESIs are available.