

Implementing Robust Ring Networks using RSTP and eRSTP

1 Introduction

This document outlines the recommended parameters of a layer 2 network comprising two levels of interconnecting rings. As shown in Figure 1, a central, "main" ring is built using RuggedCom RuggedSwitch® Ethernet bridges, and "sub" rings may be built using any RSTP-compliant Ethernet bridges. The two-level topology pictured in Figure 1 allows for the implementation of familiar ring networks while at the same time gaining some of the advantages of redundantly interconnected RSTP networks.




Figure 1: Two level ring network

Implementing Robust Ring Networks using RSTP and eRSTP

2 Number of bridges supported in a sub-ring

The maximum number of bridges in the sub-rings is given by the formula:

$$N_{sub} = (MaxAge - 1) - \left\lfloor \frac{(N_{main} - 1)}{4} \right\rfloor \quad [1]$$

where:

- N_{sub} is the maximum number of bridges supported in a sub-ring.
- $MaxAge$ is the "Max Age Time" RSTP configuration parameter on RuggedSwitch® bridges. Note that every bridge in the network must correctly assume the $MaxAge$ parameter propagated downwards from the root bridge, according to the RSTP standard. $MaxAge$ defaults to a value of 20, and may be set to a maximum of 40.
- N_{main} is the number of bridges in the main ring.

The formula in [1] describes the maximum number of bridges supported per sub-ring for a given $MaxAge$ setting and number of bridges in the main ring. The formula takes into account the worst-case scenario in which the network can still heal itself: that of two simultaneous link failures, one in the sub-ring, and one in the main ring, both in worst-case positions (see the example, below).

For the formula to apply, the bridges in the main ring must be running eRSTP, i.e. they must be RuggedSwitch® bridges. The other assumption is that the bridges in the main ring will be configured to have lower IDs than those in the sub-rings, in order to guarantee that the root bridge will be a member of the main ring.

Implementing Robust Ring Networks using RSTP and eRSTP

3 Example

The worst case for a single link failure in the example topology is the failure of the link between a member of a sub-ring adjacent to the main ring (A) and its corresponding neighbour on the main ring (B), as shown in Figure 2. The failure of this link means that once the network has healed, the effective distance from (A) to the root bridge (C) will be $N_{sub} + 1/4$ of the number of hops around the main ring to the root.

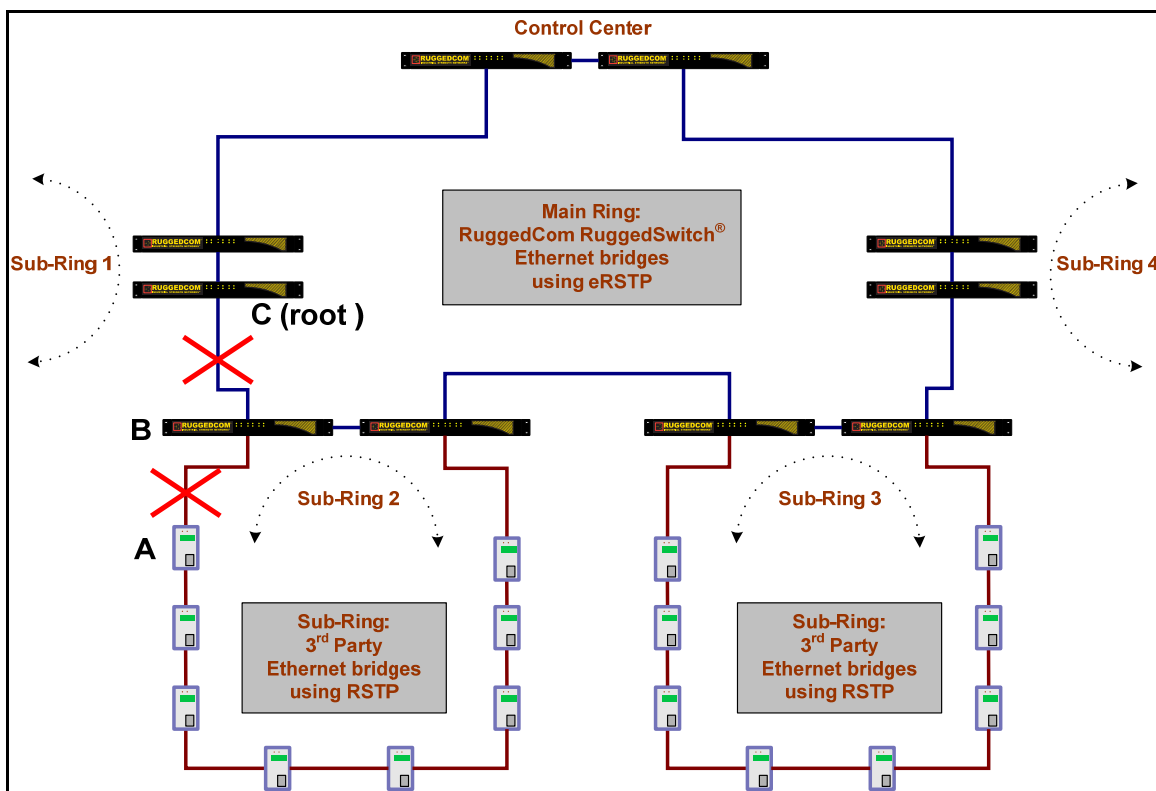


Figure 2: Two level ring network with link failures

If in addition, there is a link failure in the main ring, between (B) and the root bridge (C), the number of hops required to reach (A) from (C) becomes the maximum possible for the topology. This is the case taken into account by formula [1], above. As long as N_{sub} and N_{main} are constrained as prescribed, the ring network depicted can converge on a single spanning tree. If N_{sub} were to be exceeded in any of the sub-rings, the maximum age time would be exceeded by RSTP management packets before they could span the full path to the far end of the sub-ring. The result would be a disconnected tree for each sub-ring in which N_{sub} was exceeded.

For detailed information regarding the configuration of RuggedCom Ethernet bridges, please refer to the RuggedSwitch® Product Family User Guide.

Implementing Robust Ring Networks using RSTP and eRSTP

4 Failures not handled in this topology

If the guidelines above are adhered to, examples of more severe failures that the network would not be able to heal would be:

- More than one link failure in the main ring or any of the sub-rings
- A combination of bridge and link failures so as to effectively create two link breakages, as above

These faults would result in two networks completely isolated from one another. This could be addressed by increasing the level of interconnection in the network, at the expense of extra cabling and/or bridge nodes.

5 Extrapolation to other topologies

The key to any RSTP-based network topology is the *bridge diameter*, which in short, represents the maximum “distance” between any bridge in the network and its RSTP root. This distance is a function of the number of hops and the maximum age setting for the network.

A layer 2 network which at the outset exceeds the bridge diameter will split itself into more than one network. The key point to note is that even a network which does not normally exceed the bridge diameter may, in case of a link failure (and depending on the position of the failed link in the topology), cause certain parts of the network to exceed the bridge diameter. A network which is designed to take this into account may nevertheless break in the presence of multiple failures, and so on.

In general, it is recommended to undertake an analysis of network topology in the face of combinations of link failures, similar to that presented in this paper. This analysis, along with consideration of the risk that unhandled failures could present versus the cost to enhance the network to avoid them, will result in a network with more predictable behaviour in failure scenarios.