

White Paper

Performance of the Rapid Spanning Tree Protocol in Ring Network Topology

siemens.com/erstp

Mission critical industrial automation applications require a robust communications network that can recover quickly from cable and equipment failures. The use of Ethernet and IP based fieldbus protocols is growing rapidly and users are gaining confidence. Modbus-TCP, Profinet, Ethenet/IP, DeviceNet, IEC 61850, DNP, and IEC 60870-5-104 are just some of the protocols being used today to interconnect programmable logic controllers (PLCs), intelligent electronic devices (IEDs), and sensors to each other and to central control computers. Applications ranging from motion control, process control, discrete manufacturing, and the electrical utility SmartGrid have one need in common: ensure high availability of the Ethernet network.

There are several approaches to providing high availability; the ring based network topology is the simplest and most pervasive. There are numerous proprietary ring-based protocols available today from several vendors; these methods do not interoperate with each other and lack the scrutiny of an open standard. The Rapid Spanning Tree Protocol as defined by IEEE 802.1D-2004 has equal if not better performance than such ring protocols and provides other benefits such as the ability to support any network topology. Unfortunately, most public information regarding RSTP performance is out of date and misleading as it is based on older version of RSTP. This paper will provide an in depth analysis of RSTP performance along with simple equations for estimating network failover and recovery times so that informed decisions can be made about its efficacy for a given industrial network application. For a ring of twenty switches, worse case failover times on the order of 100 milliseconds is quite realistic which makes RSTP effective for all but the most demanding of applications.

Authors

Michael Pustylnik Mira Zafirovic-Vukotic Roger Moore Siemens Industry Inc.



Contents

Brief history of STP and RSTPBackgroundBrief History of STP and RSTP	3
 Typical architectures in industrial networks Symbols and assumptions Derivation of a single link failover time Derivation of a root bridge failover time Link and root bridge recovery 	5
Detailed examples of the failover and recovery process • Single link failure • Root bridge failure	15
Empirical performance data	16
Comparison of STP, RSTP, and eRSTP with proprietary solutions	17
Conclusions	18
References	19

LXXX III

19.00

RSTP operation – IEEE 802.1D-2004

Background

Ethernet switches operate by storing and forwarding traffic between their ports. The switch examines each Ethernet frame and records the MAC source address and the port on which it resides. Subsequently, when a frame arrives for a given MAC destination address, the switch knows on which outgoing port to send the frame. If a frame arrives and its destination MAC address is unknown or is a multicast address, the switch will flood the frame out all of its ports. If switches in an Ethernet network are connected in a loop a broadcast storm will ensue where a single broadcast frame will circulate endlessly. This condition consumes all available bandwidth on the loop making the network unusable. The Spanning Tree Protocol (STP) allows the physical network to contain loops by forcing some links into a hot standby mode.

Brief history of Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP)

STP [1]

The Spanning Tree Protocol was defined in the IEEE Standard 802.1D editions prior to year 2004. It was designed to solve the fundamental problem of traffic loops and prevent accidental loops in poorly structured and managed wiring closets. The key idea in STP is to force some links into a hot standby mode in order to reduce the network topology to that of a tree. The resulting tree spans (i.e. connects) all switches, but eliminates loops. The steps in order to best accomplish this process are:

- Allow all switches to send messages to each other that convey their identity and link cost.
- 2. Elect a single switch, among all the switches in the network to be a root, or central switch.
- Let all other switches calculate the direction and cost of the shortest path back to the Root using messages received from switches closer to the root. Each switch must have only one best way to forward frames to the Root.
- 4. If two switches servicing the same LAN exchange messages with each other, the one with the lowest cost to the Root will service the LAN. The other switch will discard all frames received from that LAN, thus opening the link and blocking a traffic loop.

STP introduced a few terms which are frequently used below in this paper:

Bridge Protocol Data Unit (BPDU): a specially formatted Layer 2 frame used by STP to exchange information between switches.

Bridge diameter: the maximum number of switches between any two end stations.

Root port: the port that offers the lowest cost path to the root bridge.

Designated port: the port that propagates Root information to the attached network segment.

Alternate port: the port that offers the next best cost path to the root bridge and will become Root Port, if the current Root Port loses connectivity with the root bridge.

Discarding port state: the state in which the port is only sending and receiving STP BPDUs while blocking any regular network traffic.

Forwarding port state: the state in which the port is sending and receiving both STP BPDUs and regular network traffic.

RSTP - IEEE 802.1w [2]

The STP protocol was first published in the IEEE 802.1D-1990 standard and has proven to be a reliable method for providing path redundancy while eliminating loops. However, STP was not originally designed for speed; when a link fails or a failed link returns to service, STP requires at least 30 seconds to restore network connectivity. RSTP is an evolution of STP. It was introduced in the standard extension IEEE 802.1w, and provides for faster spanning tree convergence after a topology change. The 802.1D terminology remained primarily the same, and most parameters have been left unchanged. However, RSTP uses several new concepts:

Taking advantage of physical link failure/recovery detection:

While STP was passively waiting for a timer to expire to react to a change in a link state, RSTP may act immediately upon a link failure/recovery detection.



Proposal-agreement mechanism:

This is a feedback mechanism that takes place between RSTP-compliant bridges. While STP was passively waiting for the network to converge before turning a port into the forwarding state, RSTP is able to actively confirm that a port can safely transition to forwarding without relying on any timer configuration. This leads to a faster convergence.

Edge ports:

All ports that have been configured as edge ports are placed in forwarding state without checking for loops.

The new enhanced mechanisms allow RSTP to reduce failover and recovery times to just a few seconds.

Enhanced RSTP (eRSTP)

Although RSTP offered a significant performance improvement compared to the legacy STP, it still had several weaknesses:

- Even the failover and recovery time of a few seconds was not good enough for mission critical industrial Ethernet applications
- 2. RSTP doesn't support LANs with a bridge diameter greater than 40

Siemens Ruggedcom developed an enhanced version of the RSTP algorithm referred to as eRSTP which is fully compatible with the IEEE 802.1w RSTP protocol while enhancing it in several aspects:

- eRSTP reduces failover and recovery times to just a few milliseconds (5ms per a pair of bridges involved in the topology change)
- eRSTP is able to operate in larger LANs with a bridge diameter greater than 20

Being a proprietary enhancement, the eRSTP algorithm was never published.

RSTP - IEEE 802.1D-2004 [3]

The IEEE Standard 802.1D-2004 edition is a very important step in the STP/RSTP evolution because it:

- Obsoletes the legacy STP.
- Addresses weaknesses of the IEEE 802.1w RSTP and defines a significantly revised and highly optimized version of RSTP. The new RSTP provides for very short failover and recovery times (identical to those of eRSTP).

Although not adopted yet by most networking equipment vendors, the optimized RSTP seems to exceed the performance of different proprietary solutions.

NOTE: Siemens Ruggedcom eRSTP has been enhanced to incorporate the strengths of the optimized IEEE 802.1D-2004 RSTP, while still supporting 4 times longer bridge diameters than those supported by a standard implementation.

Analytical Method for Calculating Ring Failover Times

RSTP is a complicated protocol as it allows for any network topology from a ring to a full mesh. Analytical determination of the failover and recovery performance for an arbitrary network and fault scenario is a non trivial exercise. However, a ring topology is simple enough to perform such analysis which is detailed in the rest of section 3 and the final result is summarized here. The worst case ring network failover time in case of a single link failure can be calculated using the following formulae:

TL + (N - 3)*TPA , if N is even

TL + (N - 2)*TPA , if N is odd

where:

N - Number of switches in the ring

TL - Time required by a switch to detect a link failure

TPA - Time required by a pair of switches to perform RSTP Proposal-Agreement handshaking; equal to the sum of the BPDU processing times in both switches of the pair.

The worst case failover time in case of a root bridge failure can be calculated using the formulae:

TL + (2*N - 5)*TPA , if N is even

TL + (2*N - 4)*TPA , if N is odd

However, the worst root bridge failure case can be easily avoided by adjusting some RSTP management parameters and thus the failover time can be reduced to a value identical to that of the single link failure case.

TL and **TPA** values may differ from vendor to vendor, from product to product, and for different port types. For RUGGEDCOM products, these values are:

TPA = 5 ms

TL = 4-6 ms for 100Base-TX and 100Base-FX links

= 20 ms for 1000Base-X links

= 700 ms for 1000Base-T links (defined by the IEEE Standard 802.3)

Symbols and assumptions

The symbol legend defined in the IEEE 802.1D standard will be used in network diagrams to help illustrate the state transition changes incurred by RSTP during a topology change Port Role Port State Legend Designated Forwarding Discarding Root Forwarding Alternate Discarding Transmitted BPDU Legend Designated Proposal Root Agreement Root with TC

Port Role	Port State	Legend
Designated	Forwarding	•
	Discarding	•
Root	Forwarding	°
Alternate	Discarding	

Port Role	Legend
Designated Proposal	→→
Root Agreement	
Root with TC	TC >

For the matter of our analysis we use the following assumptions:

- The ring contains an even number of switches N = n*2. Derivation of the formulas for an odd number of switches is very similar and, therefore, is omitted.
- 2. Links are not longer than a few kilometers, so BPDU frame propagation time on the wire can be completely ignored.

A time variable **TTC** is used in the analysis. It is the time required by a pair of switches to deliver a Topology Change notification from one switch to another. As BPDU carrying a Topology Change notification is only propagating in one direction and no handshaking is involved, **TTC** is roughly half of **TPA**.

Derivation of a single link failover time

Figure 1 shows a simple ring network with all links intact in a steady state condition; switch S0 is the root bridge and switch Sn has one port in the discarding state. In order to determine the worst case one should perform the analysis for each possible link failure in the ring. Performing the analysis on all possible links would show that the worst case is: the failed link is the one furthest from the switch with an Alternate port; in this case, it is the link between S0 and S1. Similarly, the "best" case would be the link between S(n-1) and Sn. The worst case scenario involves the maximum number of switches required for propagating and negotiating RSTP information around the ring.

After the link failure, the ring is divided into two segments S0-S-(n-1) and S1-Sn because connectivity between them is blocked by the Sn's Alternate port Sn-S-(n-1). (Figure 2)

The following activities are required to restore the ring connectivity – note that the activities are different for the two segments:

- Segment S0-S-(n-1) does not require any topology reconfiguration because the old path to the Root is still valid for all switches in that segment.
- Switch Sn must change its former Alternate port role to Root and its former Root port to Designated; all other switches in segment S1-Sn should swap their Designated and Root port roles.
- Switch Sn must initiate a Topology Change notification which must be propagated from Sn up to S0 so that S0 will flush its MAC address table on the failed link S0-S1.

Table 1 describes how the above activities will proceed at different time points.



Figure 1: Steady state before link failure

Figure 2: Ring segments after link failure

Time	Switch acting	Action Description	Figure
Τι	S ₀	So doesn't do anything because it is the Root. So detects the link failure and immediately ages out its Root info on port S1-S0. S1-S0 is S1's only path to the Root, so S1 will declare itself as a new root bridge and start	Figure 3
Tl + Tpa	S2	Proposal-Agreement handshaking with S ₂ . New Root information comes to S ₂ from its only path to the previous Root, so S ₂ replaces the previous Root info with the new one, puts S ₂ -S ₃ to Discarding state, sends	Figure 4
		Agreement BPDU to S1 and starts Proposal-Agreement handshaking with S3. In a similar fashion, the Proposal-Agreement handshaking	
	{S ₃ ,,S _{n-2} }	will occur on each hop up to S _(n-1) .	-
	Sn-1	{S1,,Sn-1} all agree about recognizing S1 as Root.	Figure 5
	Sn	Successful Proposal-Agreement continues until S_n is reached. As S_n has a better Root (S_0) information, it turn its Alternate Port into Root Port and reply with its own Proposal rather than Agreement. NOTE: When changing the former Alternate port role to Root the switch will put that port to forwarding and send a Topology Change notification to $S_{-(n-1)}$.	Figure 6
	S _{n-1}	S_{n-1} starts Proposal-Agreement handshaking with S_{n-2} . Also, approximately at this time Topology Change notification also reaches $S_{-(n-1)}$.	_
	$\{S_3,,S_{n-2}\}$	The Proposal-Agreement handshaking occurs on each hop back to S_1 (and the Topology Change notification propagates to S_0).	_
	S ₁	All switches agree about the new path to the Root (S_0) and all ports are in Forwarding state.	Figure 7
TL + (n-2+1)*TPA + (n-1)*TTC	So	Topology Change Notification reaches S ₀ and causes it to flush its MAC address table on port S ₀ -S ₁ (this is the purpose of TCN). Network connectivity restored.	Figure 8

 Table 1: Time-line of RSTP actions by different switches in the ring (single link failure)



Figure 3: S1 declared itself a new route

Figure 4: S₂ agrees about S₁ being Root



Figure 5: $\{S_1, ..., S_{(n-1)}\}$ all agree about S_1 being Root

Figure 6: S_n turns its Alternate Port into Root Port and replies with proper Root (S_0) information



Figure 7: All switches agree about new Root Path

Figure 8: Steady state after network recovery

So the overall failover time is

TL + (n-2+1)*TPA + max((n-1)*TTC , (n-2)*TPA)

As we explained above, TPA is significantly longer than TTC, and then the overall failover time is

TL + (n-2+1+n-2)*TPA = TL + (2*n-3)*TPA = TL + (N-3)*TPA

Derivation of a root bridge failover time

Failure of the root bridge requires electing a new root bridge which adds further complexity compared to the single link failure analysis. Figure 9 shows the network in steady state before the root bridge fails. After the root bridge failure, the ring is divided into two segments S-1-S-(n-1) and S1-Sn because connectivity between them is blocked by the Sn's Alternate port Sn-S-(n-1) (Figure 10). Each switch can have a different bridge priority configured. Determining the worst case failover time demands analyzing all cases where the new bridge could become any one of the switches in the ring. Repeating the analysis for all switches would show that the worst case has S1 and S-1 becoming the next best root candidates with S1 taking precedence over S-1. Table 2 describes how RSTP will proceed at different points in time for the worst case.



Figure 9: Steady state before root bridge failure

Figure 10: Ring segments after root bridge failure

Time	Switch acting	Action Description	Figure
Τι	S1 and S-1	Each of S_1 and S_{-1} detects the link failure, ages out its Root info and advertises itself as a new root bridge.	Figure 11
	$\{S_2,,S_{n-2}\}$ and $\{S_{-2},,S_{-(n-2)}\}$	Proposal-Agreement handshaking occurs in parallel on each hop in two segments, up to $S_{(n-1)}$ and $S_{(n-1)}$	_
Т∟+ (n-2)*Тра	Sn-1 and S-(n-1)	${S_1,,S_{n-1}}$ all agree about S_1 being Root, while ${S_{-1},,S_{-(n-1)}}$ agree about S_{-1} being Root	Figure 12
	Sn	S_n receives proposal BPDUs from S_{n-1} and $S_{(n-1)}$ at virtually the same time. For the worst case scenario, we assume that the proposal from S_{n-1} is received and processed first (as you will see below, this case causes certain network "confusion" and thus requires more RSTP actions to resolve it). Since Sn doesn't have any information about the So failure yet, it will change its old Alternate Port role to Root Port and "reject" the S _{n-1} 's proposal by sending its own proposal with better but actually obsolete information about So. This info	Figure 13
		will "confuse" all $\{S_{n-1},,S_1\}$ switches. Right after that, S_n will process $S_{-(n-1)}$'s proposal about S_{-1} as Root. As S_n receives that proposal from its only path to the Root, it immediately ages out the old Root information and starts Proposal-Agreement handshaking with S_{n-1} again – this time proposing S_{-1} as Root. The new Proposal is now fixing the "confusion" just caused by the previous Proposal about S_0 .	Figure 14
TL+(n-2+1+1)*TPA	S _{n-1}	$S_{n\text{-}1}$ starts Proposal-Agreement handshaking with $S_{n\text{-}2}$ about $S_{\text{-}1}$ as Root.	_
	{S _{n-2} ,,S ₃ }	The Proposal-Agreement handshaking occurs on each hop back to S_2 .	_
TL+(n-2+1+1)*Тра + (n-3)*Тра	S2	All switches in the ring except S_1 erroneously agree about S_{-1} being Root. S_2 sends Proposal to S_1 about S_{-1} as Root.	Figure 15
TL+(n-2+1+1)*Тра + (n-3)*Тра+Тра	S ₁	Being a better Root candidate, S1 responds with its own Proposal about itself as Root.	Figure 16
	S ₂	Final round of handshaking starts from S2 towards S-1	

Table 2: Time-line of RSTP actions by different switches in the ring (root bridge failure)

* Note that in the analysis Topology Change notification was not even mentioned because it propagates along with the initial round of handshaking, i.e. much before the final topology is achieved.



Figure 11: S_1 and S_{-1} declare themselves as Root

Figure 12: {S1,...,Sn-1} agree about S1, while {S-1,...,S-(n-1)} agree about S-1



Figure 13: S_n sends obsolete proposal about S_0 being Root



Figure 14: S_n sends new proposal about S_{-1} being Root



Figure 15: All switches except S_1 agree about S_{-1} being Root

Figure 16: S_1 rejects proposal about S_1 and declares itself as Root



Figure 17: Steady state after network recovery

So the overall root bridge failover time is:

TL + (n-2+2)*TPA + (n-3)*TPA + TPA + (2*n-3)*TPA = TL + (4*n - 5)*TPA = TL + (2*N - 5)* TPA

Please note that the described scenario is an uncontrolled theoretical worst case. This extensive failover time can be easily and significantly reduced by taking a controlled case approach. If Sn's Bridge Priority is configured in such a way that it is the best root bridge candidate after the failed S0, the scenario gets much simpler. Events in the network will proceed like in the uncontrolled case until the step in Figure 12, but in the Figure 13 step Sn will advertise itself as a Root to both half-ring segments, thus starting a final round of handshaking. So the failover time gets reduced to:

TL + (n-2+1)* TPA + (n-2)* TPA = TL + (2*n - 3)* TPA = TL + (N - 3)* TPA,

which is identical to the formula derived for the single link failure case.

Link and root bridge recovery

The connectivity recovery mechanism in case of a link or root bridge recovery is different from the link or root bridge failure case in the following aspects:

- Link detection time is not included in the network outage time – RSTP activities only start after the link is detected
- 2. All links are available for the Root information propagation, so multiple "reconfiguration" of certain network segments does not occur

As a result, network outage time in the case of a link or root bridge recovery is always shorter than in the case of a link or Root Ridge failure. Since only the worst case is of importance, link or root bridge recovery case analysis is omitted.

Extrapolation to Meshed Networks

As meshed network topology analysis is more complex than that of a ring, the first impression could be that RSTP performance must be always worse in meshed topology compared to a single ring. However, that is not generally true although it is in some scenarios.

Single link failure

A meshed network can be looked as a ring network with some inter-switch connections which result in shortcuts for BPDU propagation. Figure 18 presents a simplified example of this.

Analyzing this network behavior in the same case of the link S0-S1 failure, shows that RSTP handshaking steps don't have to go all the way from S1 to Sn. Information about the root bridge (S0) location can already be obtained from the switch S(n-2) or S-2 which will make the network recovery faster. An observation can be made that meshed networks provide shorter paths for RSTP information propagation which allows handshaking sequence to complete faster. As a result, a meshed network single link failover time should be same or better than the failover time of the largest outer ring in that network topology (assuming the given topology does have such an outer ring).

Root bridge failure

Unfortunately, the additional complexity of meshed topology causes root bridge failure failover times to increase compared with the simple ring. The problem of the root bridge failure scenario is that, after the failure, every switch holding the obsolete Root information is feeding it back to the network, thus "confusing" all other switches and making them "reject" the right new Root information. The same problem exists in the ring topology as well but, due to the topology simplicity, there are no loops where the obsolete information would circulate and the "confusion" is fixed very quickly. This is not the case in meshed networks.

A meshed network can be looked at as multiple smaller rings interconnected with each other at multiple points. This topology allows the obsolete



Root information to circulate multiple times back and forth inside and between the smaller rings. The network will not be confused forever according to the RSTP standard, the old information will be aged out after it traverses the maximum allowed number of hops (normally 20). However, the network outage time can last as long as seconds. It is caused by the fact that multiple switches in the network are continuously busy with processing and forwarding contradictory BPDUs being continuously received on different ports.

Root bridge failure in a meshed network is very hard to analyze and predict, and a result may be totally different for every specific topology. The common conclusion is though, that the root bridge failover time grows exponentially as more redundant paths are added to the network topology. The root bridge failure in meshed topology is a well recognized problem. Although some proprietary mechanisms are offered to indirectly improve the RSTP performance, no common solution is known for this case. This is the price paid for the high level of link redundancy provided by meshed network topology.

Empirical Performance Data

To corroborate the RSTP analysis, testing of failover performance was done on a test network. The same worst case setup is used on ring networks with 4 to 40 switches as illustrated in Figures 1 and 9. The network consists of RUGGEDCOM RS900 and RUGGEDCOM RSG2100 switches interconnected with 100Base-TX links. The switches were using the Rugged Operating System (ROS[®]) v3.4 that has an IEEE 802.1D-2004 RSTP implementation. A SmartBits Ethernet packet blaster was used to generate high rates of traffic and determine the number of dropped frames during the failover event. 45 Mbps unidirectional traffic was generated. No other application was running on the network. The switches were configured in such way that MAC address tables are not purged by link loss, making the measured failover times due to RSTP only.

Failover performance was measured for the worst case link failure and root bridge failure. Link failure was accomplished by simply disconnecting the cable while root bridge failure was done by powering off the switch. The failover time was calculated by multiplying the number of dropped packets (determined by the SmartBits) by the time between start of transmission of two consecutive messages (which is a constant). The measured and analytical failover times are shown in Figure 19 and Table 3.

The measured failover times in case of a link failure are smaller than the theoretical upper bound that is analytically derived in Section 3.2 which is a result of somewhat conservative choices for the failover equation constants. Measuring the slope of the measured failover times yields TP =2.1 ms to give more accurate results from the failover equation. However, it is prudent to use the more conservative time constants for any network planning activities.



Figure 19: Measured vs theoretical RSTP Failovertime in ring network if single link fails

Time	Switch acting	Figure
40	93.50	381
39	142.53	376
36	244.69	341
35	124.53	336
30	67.59	281
29	100.31	276
26	59.85	241
25	89.71	236

Table 3: Measured vs theoretical RSTP failover

 time in ring network if the root bridge fails

Comparison of STP, RSTP and eRSTP with Proprietary Solutions

For many years, there was no standard redundant LAN solution that would provide short network recovery times sufficient for Industrial Ethernet applications. As a result, some networking equipment vendors offered different proprietary protocols designed to solve the problem for ring network topology. The table on the next page compares some of such proprietary solutions with STP/RSTP and with each other.

Protocol	Vendor	Can be used in multi-vendor	Max Bridge Diameter	Topology	Single ring link failover time (for different number of switches		er time of switches)
		environment	Diameter		10	15	20
STP	IEEE Standard	Yes	40	Any	>30 s		
RSTP (802.1w)	IEEE Standard	Yes	40	Any	Several Seconds		
HiPER Ring [4][5]	Hirshmann	No	Virtually unlimited	Ring	200-500 ms, independent of number of switches		
Turbo Ring [4][6]	Moxa	No	Virtually unlimited	Ring	<200 ms	<250 ms	<300 ms
S-Ring [8]	GarrettCom	No	Data Not Available	Ring	<250 ms		
RS-Ring [8]	GarrettCom	No	Data Not Available	Ring	<100 ms		
RapidRing [7]	Contemporary Controls	No	50	Ring	<300 ms		
RSTP (802.1D-2004)	IEEE Standard	Yes	40	Any	<50 ms	<75 ms	<100 ms
eRSTP	Siemens RUGGEDCOM enhancements to IEEE Standard	No	160	Any	<50 ms	<75 ms	<100 ms

* Note: Proprietary solutions



Conclusion

RSTP performance is both predictable and repeatable or failure and recovery of switches and cabling in a ring topology. Precise equations can be used to determine the network outage time; for a ring of twenty switches, worse case failover times on the order of 100 milliseconds are quite realistic - typical RSTP performance is much better. Unfortunately, most literature to date states RSTP failover performance of several seconds based on results from the older IEEE 802.1w RSTP. The analysis in this paper hopefully will put this misconception to rest and give confidence to designers of such networks; RSTP is more than capable of being deployed in very demanding automation networks.

RSTP has other advantages over the 'ring' protocols such as the ability to support any network topology including mesh, which allows for an even greater degree of redundancy. RSTP tends to have faster not slower network recovery times. RSTP works in a multi-vendor environment and is supported by all the major switch vendors. Finally, RSTP was created by and is supported by an international standards organization, which ensures scrutiny by peers and future harmony with the myriad of other Ethernet standards under development. RSTP is an excellent solution for many mission critical industrial Ethernet applications, which is why it has seen so much success to date and why it will continue to be the dominant Ethernet redundancy protocol in the future.



References

- [1] ANSI/IEEE Std 802.1D, 1998 Edition
- [2] ANSI/IEEE Std 802.1W, 2001 Edition
- [3] ANSI/IEEE Std 802.1D, 2004 Edition
- [4] "High Availability Networks", The Industrial Ethernet Book
- [5] M. Schaub, H. Kell, "Redundancy Process with Hirschmann Switches", ComConsult Kommunikationstechnik GmbH, 2003
- [6] "Stay Connected with Turbo Ring", Moxa Technologies, Inc.
- [7] "RapidRingTM Redundancy from Contemporary Controls", Contemporary Control Systems, Inc., 2004
- [8] "Secure Web Management User Guide for Magnum 6K family of Switches, Release 3.6", GarrettCom, Inc.

Siemens AG Process Industries and Drives Process Automation Postfach 48 48 90026 Nürnberg Germany

Siemens Canada Limited 300 Applewood Crescent Concord, Ontario, L4K 5C7 Canada

siemens.com/ruggedcom

© Siemens AG 2007 Subject to change without prior notice PDF Whitepaper Produced in Canada The information provided in this document contains merely general descriptions or characteristics of performance which in case of actual use do not always apply as described or which may change as a result of further development of the products. An obligation to provide the respective characteristics shall only exist if expressly agreed in the terms of contract.

In order to protect plants, systems, machines and networks against cyber threats, it is necessary to implement – and continuously maintain – a holistic, state-of-the-art industrial security concept. Siemens' products and solutions only form one element of such a concept. For more information about industrial security, please visit www.siemens.com/industrialsecurity