***White Paper***

# Ethernet Protection Switched Rings

## Creating the Survivable Ethernet Network

*9 July 2004, Rev. B*

---

*Delivering services is the focus of modern network communications: services such as broadcast video, POTS, voice over IP, video on demand, and Internet access. These services demand a high-level of quality—a requirement difficult for Ethernet to meet. In order to achieve survivability, Ethernet network equipment must have a technology that performs better than STP or RSTP. The solution is Ethernet Protection Switched Rings.*

---

# Ethernet Protection Switched Rings

## Creating the Survivable Ethernet Network

*9 July 2004, Rev. B*

---

*Delivering services is the focus of modern network communications: services such as broadcast video, POTS, voice over IP, video on demand, and Internet access. These services demand a high-level of quality—a requirement difficult for Ethernet to meet. In order to achieve survivability, Ethernet network equipment must have a technology that performs better than STP or RSTP. The solution is Ethernet Protected Switch Rings.*

---

### Ethernet Protected Switch Rings
### Building the Self-Healing Ethernet Network

*Modern network communications is centered on delivering broadcast video, voice over IP, video on demand, and Internet access.*

The focus of modern network communications is centered on delivering services: services such as broadcast video, POTS, voice over IP, video on demand, and Internet access. Although none of these services is entirely new, the idea of deploying them together over Ethernet-based networks is. And, as it happens, each of these services has a quality that is difficult for Ethernet to provide: survivability. What is survivability? In the services world, survivability means keeping the network in operation; it means failures resolve on their own; it means when failures do occur, they are not noticed by the subscriber. Unfortunately, by these definitions, today's Ethernet is not yet survivable.

*The Spanning Tree Protocol and Rapid Spanning Tree Protocol are effective for preventing loops and assuring backup paths are available. However, that alone does not make Ethernet survivable.*

How can they be made survivable? To a degree both the Spanning Tree Protocol (STP) and Rapid Spanning Tree Protocol (RSTP) have been proven effective for preventing loops and assuring backup paths are available. However, that alone does not make Ethernet survivable. Both protocols are slow to respond to network failures—slow on the order of 30 seconds or more. In order to achieve survivability Ethernet network equipment must have a technology that performs better than either STP or RSTP. And the solution is Ethernet Protection Switched Rings.

Ethernet Protection Switched Rings (EPSR), much like STP, provides a polling mechanism to detect ring-based faults and failover accordingly. But unlike STP, EPSR uses a fault detection scheme that alerts the ring that a break has occurred and indicates that it must take action instead of making a calculation. When a fault is detected, the ring automatically heals itself by sending traffic over a protected reverse path. Because of this ring-fault

Detection scheme, EPSR can converge in less than 50 milliseconds. And because EPSR is a protocol that runs over standard Ethernet interfaces, there is no special interface requirement as in RPR (802.17).

This document provides insight on how EPSR makes Ethernet survivable; how EPSR actually functions; on what platforms EPSR runs; and how EPSR might be implemented in network architecture.

## Building the Rings
## Understanding How EPSR Works

*EPSRs must be extremely flexible and interoperate with other standard Ethernet functions.*

EPSR is a ring resiliency protocol that operates over standard Ethernet ports. Because of this, EPSR must be extremely flexible and interoperate with other standard Ethernet functions.

Before continuing with a discussion of EPSR, there is a bit of new terminology that should be reviewed:

- **EPSR Domain**—The EPSR Domain defines a protection scheme for a collection of data virtual local area networks (VLANs), the control VLAN, and the associated switch ports.
- **Master Node**—The controlling node for an EPSR Domain. The Master node is responsible for status polling, collecting error messages, and controlling the flow of traffic in an EPSR Domain.
- **Transit Node**—All other nodes in an EPSR Domain are transit nodes. Transit nodes generate failure notices and receive control messages from the Master.
- **Primary Port**—The primary port of the Master Node determines the direction of traffic flow. This port is always operational.
- **Secondary Port**—The secondary port of the Master Node remains active, but blocks all protected VLANs from operating until ring failover.
- **Control VLAN**—The VLAN over which all control messages are sent and received. This is the only VLAN that is never blocked.
- **Pre-forwarding**—The term for the state of an Ethernet interface that was previously down, but that is now blocked and awaiting a control message to unblock.

## Examining the Physical Structure

*The least-used node is designated as the master node, assuring optimum reuse of the bandwidth.*

EPSR operates over standard Ethernet interfaces that must be physically connected in a ring or collapsed ring topology. Because traffic flow is based on the master node, its placement is critical. Most scenarios dictate that the least-used node should be designated as the Master node, thereby assuring optimum spatial reuse of the bandwidth. This, however, is not a requirement.

### Converged Transport, Aggregation and Access

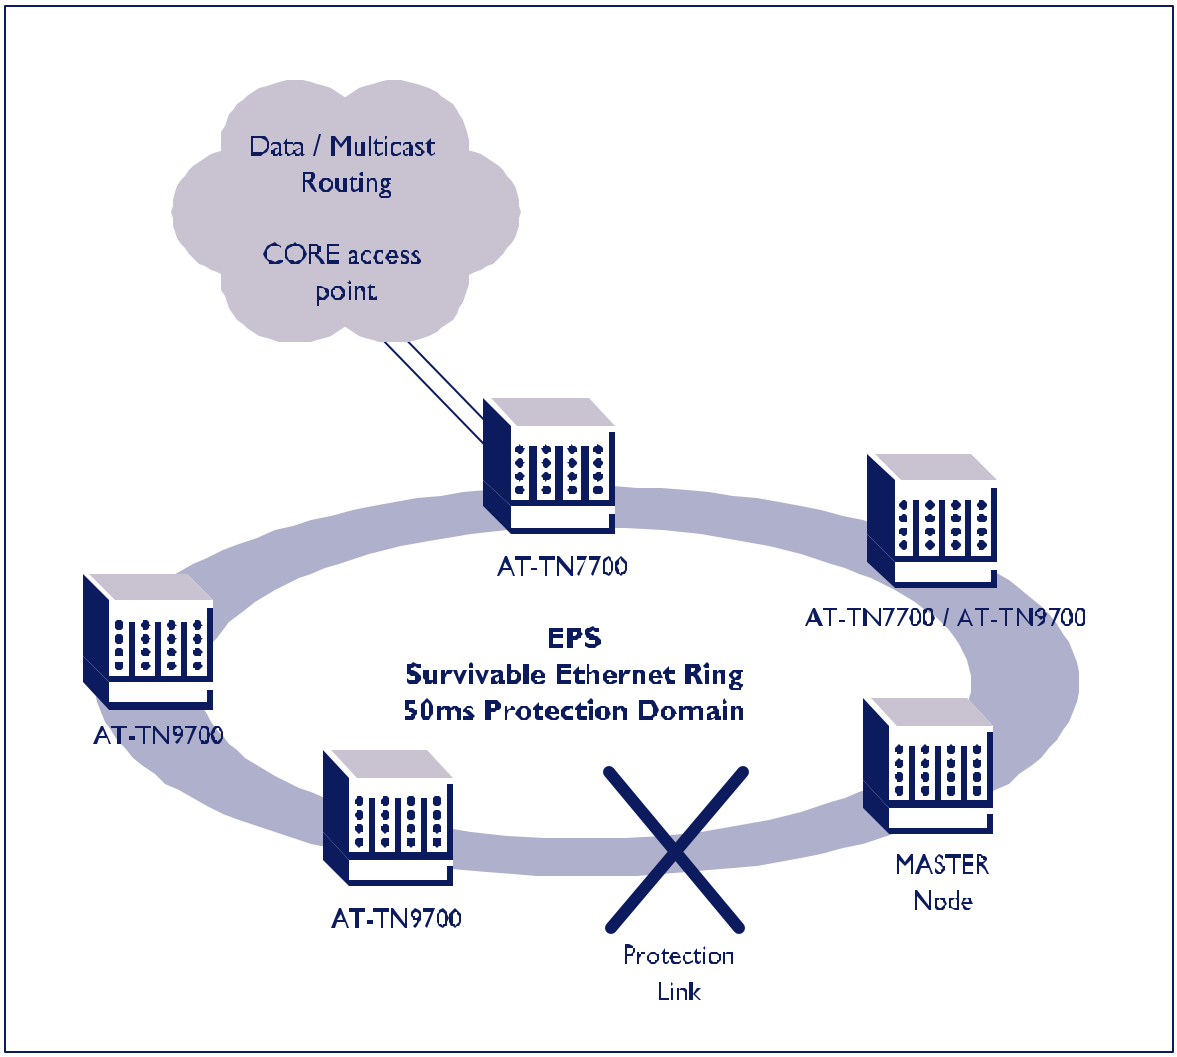

## Looking at Subtended Rings

*The EPSR node interconnects multiple subtended nodes, resulting in a core node.*

EPSR supports multiple physical rings for subtending configurations. This allows the EPSR node to interconnect multiple subtended nodes , resulting in a core node.

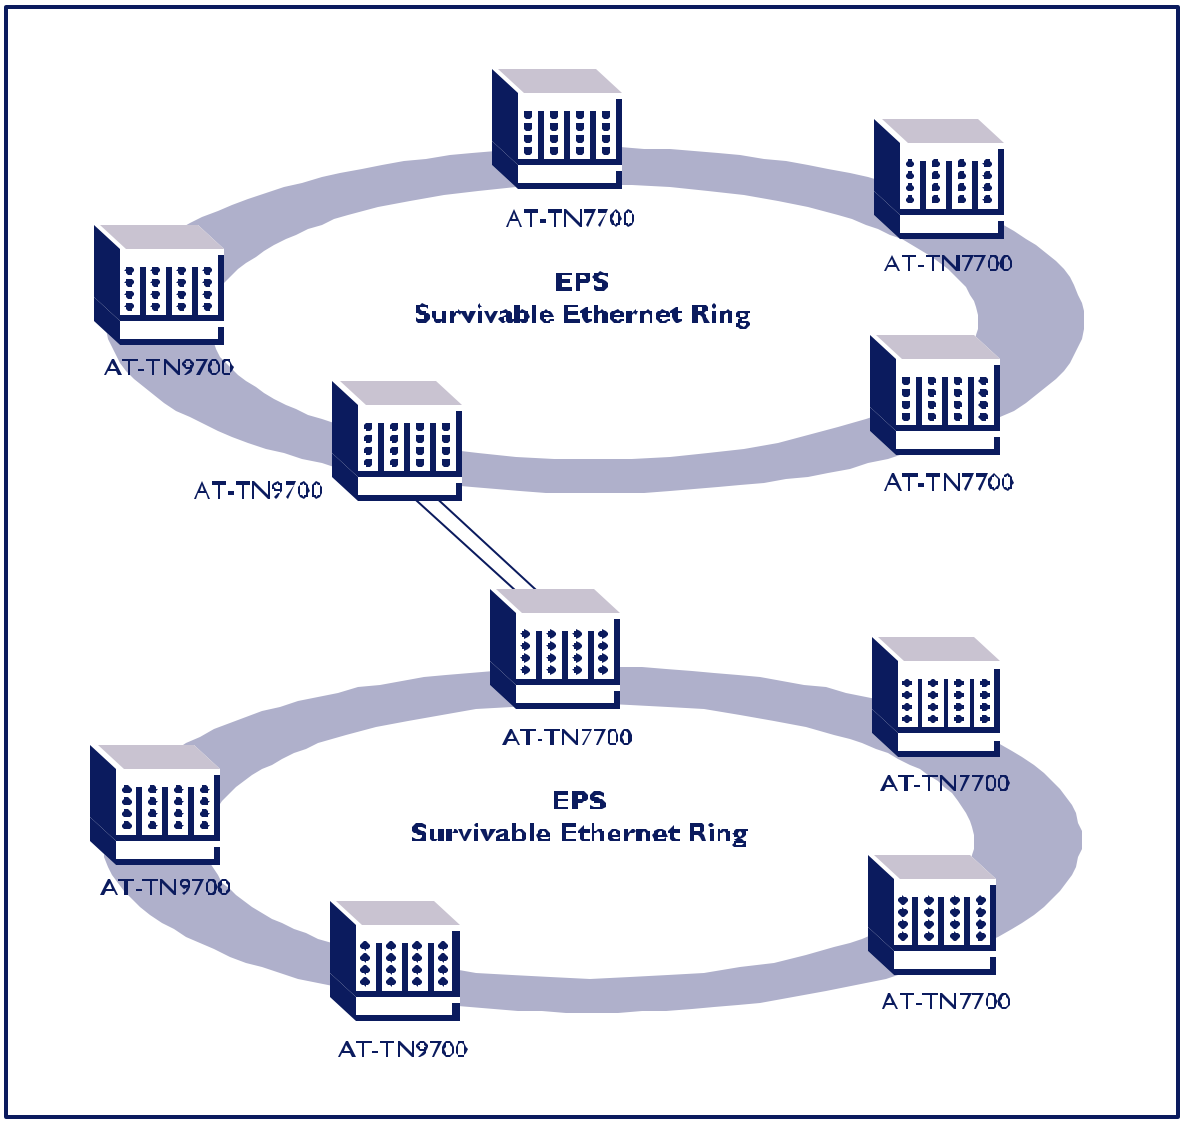

## Creating Multiple Overlapping EPSR Domains

*One EPSR domain can operate through the westbound interface, while a second operates through the eastbound interface.*

Multiple EPSR domains can operate over the same physical ring. For example, while one EPSR domain operates with normal traffic flow through the westbound interface, a second EPSR domain could operate through the eastbound interface. This feature further enhances the spatial reuse capability of EPSR that allows the operator to accurately control the flow of traffic to maximize the available bandwidth of high speed links in both directions.

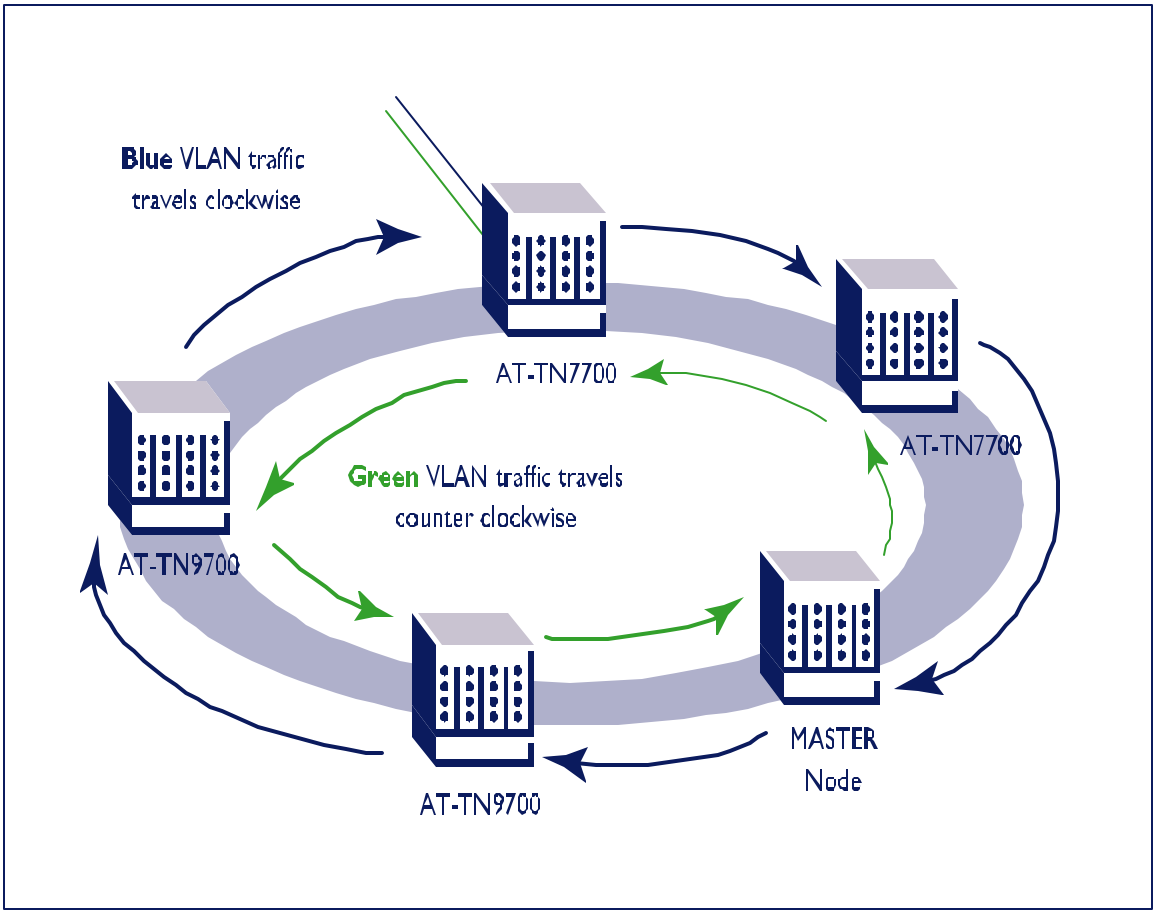

## Feature Interaction
# Blending EPSR with Other Elements

*UFO ensures that traffic is not switched between users but is only directed to designated egress interfaces.*

**Upstream forwarding only**

EPSR is fully capable of supporting Allied Telesyn's Upstream Flood Only (UFO) implementation. UFO ensures that traffic is not switched between users but is only directed to designated egress interfaces. By using the Master node as the egress interface for UFO-protected VLANs, EPSR networks are capable of supporting this enhanced security feature.

**Spanning Tree**

EPSR can be used with either the Spanning Tree Protocol (STP) or the Rapid Spanning Tree Protocol (RSTP). The only restriction is that STP/RSTP cannot be configured on the same interface as EPSR.

## EPSR with STP/RSTP Rings

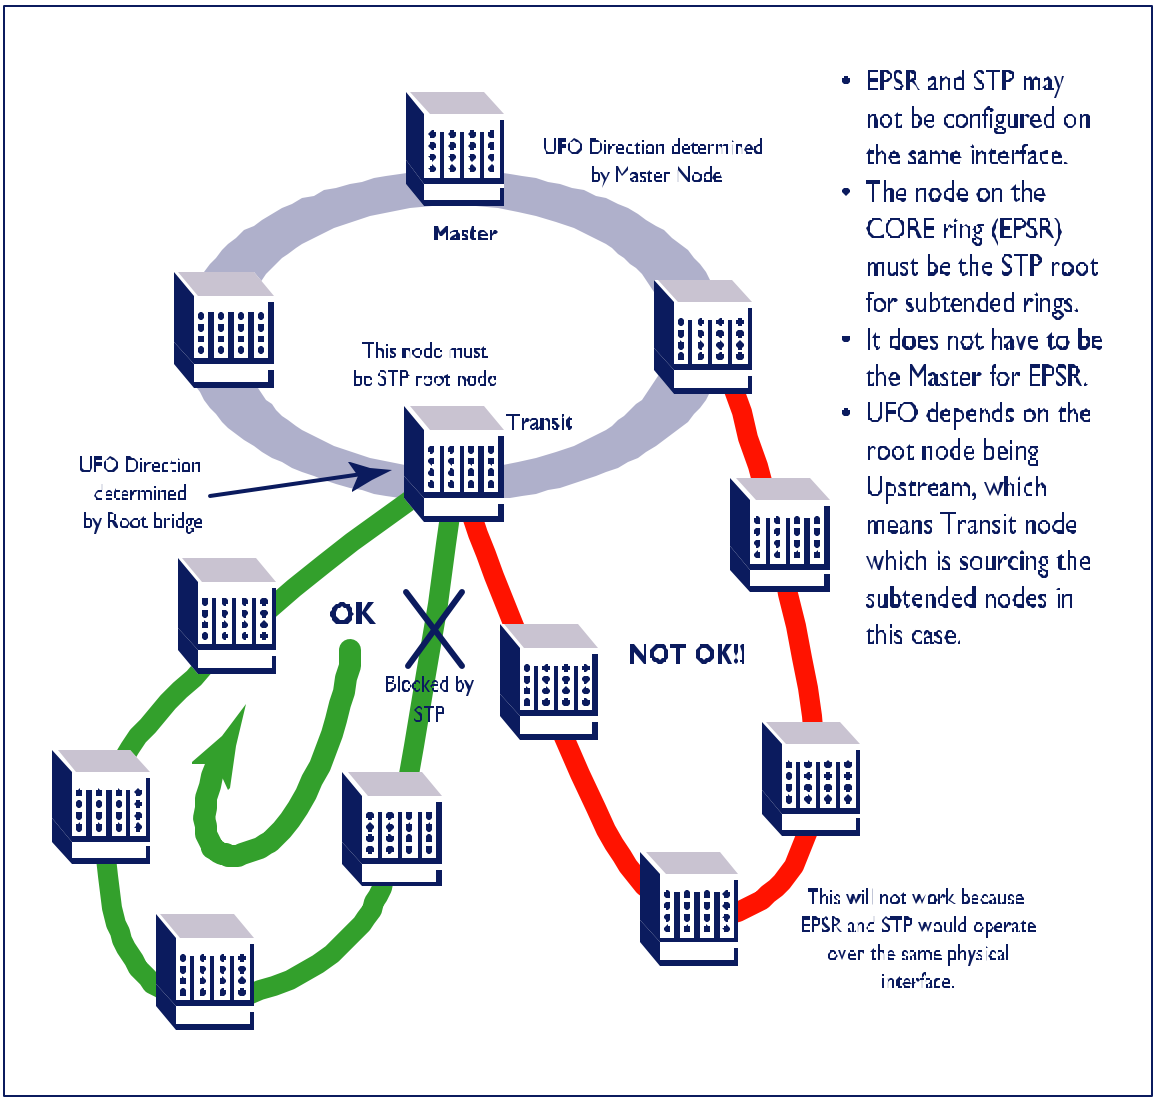

## Protocol Overview
## Looking at Traffic-Flow

*There is no limit to the number of nodes that can exist on any given EPS ring.*

EPSR operates over standard Ethernet ports and requires the use of a specified control VLAN. The control VLAN cannot carry any data traffic, only control messages; therefore only EPSR control packets are sent over the control VLAN. All other VLANs on an EPSR domain are protected VLANs carrying traffic such as Voice, multicast video and data.

EPSR is a ring resiliency protocol. There is no limit to the number of nodes that can exist on any given EPS ring. There is, however, a limit of one Master node per EPSR ring; all other nodes on the ring are classified as transit nodes. The Master node has configurable primary and secondary ring ports.

*The master node periodically sends out health-check messages through its primary port at times specified by the user.*

**Normal Operation**

Under Normal operation the Master node's secondary port is blocked for all protected VLANs; only the control VLAN remains unblocked on the secondary port. The Master node periodically sends out health-check messages through its primary port at times specified by the user. These health-check messages are then received on the Master node's secondary port.

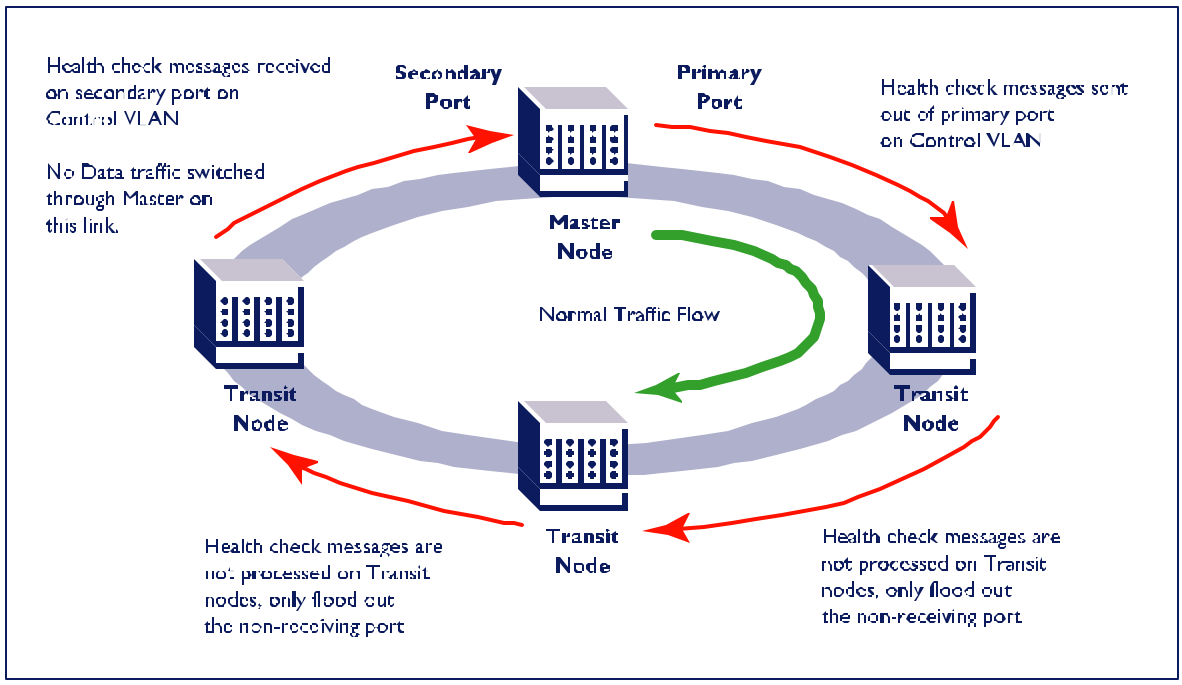


**Ring Failure Operation**

The ring can enter into ring-fault operation in two ways: The Master node fails to receive two consecutive health-check messages on its secondary port; or a transit node sends an EPSR LINK-DOWN control message to the Master node. Note that the Master node continues to send health-check messages through its primary port when in a ring-fault state. When the ring enters into such a state, the Master node unblocks its secondary port and

flushes its forwarding database (FDB). It also sends a RING-DOWN-FLUSH-FDB control message to all transit nodes instructing them to flush their FDBs as well. Then, by normal bridge learning, all paths and communication are restored.

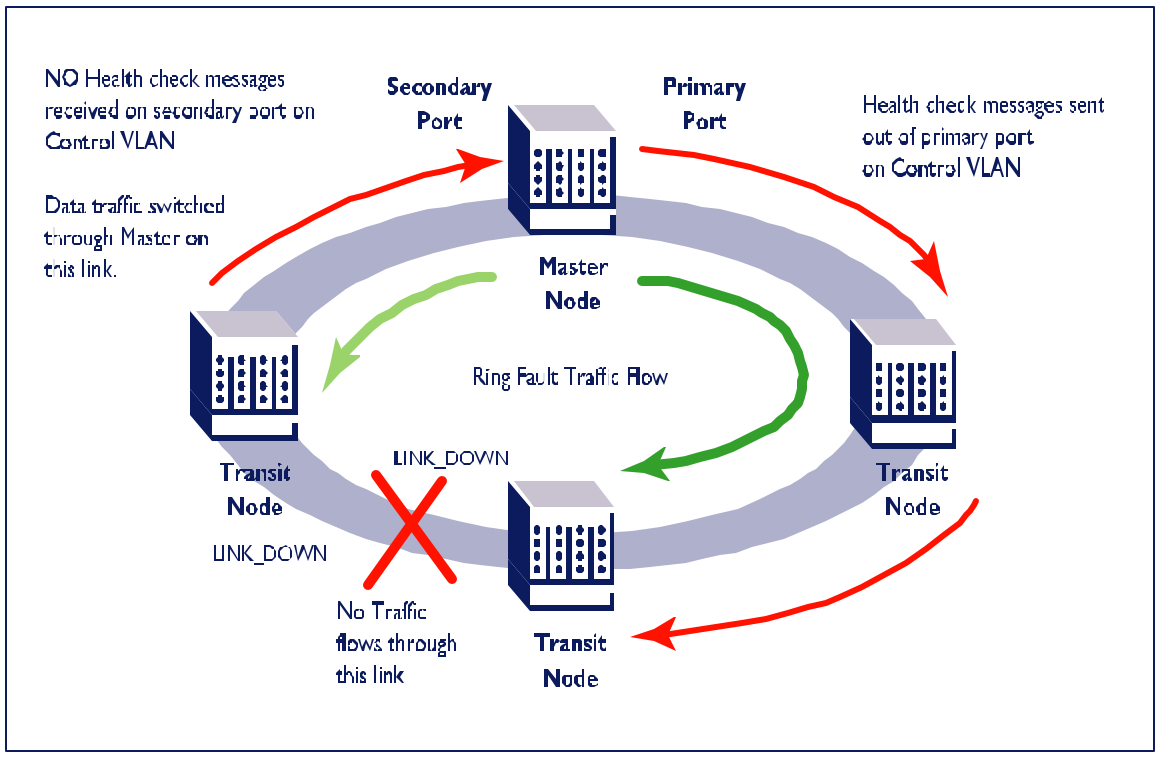


*When the ring enters into a ring-fault state, the master node unblocks its secondary port and flushes its forwarding database.*

**Ring Recovery**

When the Master node either starts to receive its health check messages on its secondary port, or when the failed transit node sends a LINK_UP message, the Master node will then restore the ring to its original topology. The master node accomplishes this by re-blocking its secondary port, flushing its FDB, and sending a control message to all transit nodes instructing them to flush their FDBs.

The transit node whose port has just come back up is responsible for preventing a network loop while the Master's secondary port is still in operation. It accomplishes this by blocking the protected VLANs on the restored port and placing them in a pre-forwarding state. When the ring restoration control message arrives from the Master node, the transit node then flushes its FDB and unblocks the pre-forwarding port.

*EPSR uses fiber topology to provide better-than mesh protection while conserving valuable fiber resources.*

**Network Topology**

Most highly available Ethernet networks require a complicated mesh topology. In carrier access networks this would require a tremendous amount of fiber from each outlying remote location – cabinet enclosures for example – to the CO as well as between each of the outlying remotes. In order to take advantage of fiber deployments utilizing ring or collapsed ring topologies , a survivable ring technology is required. EPSR takes full advantage of the either fiber topology to provide better-than mesh protection while conserving valuable fiber resources.

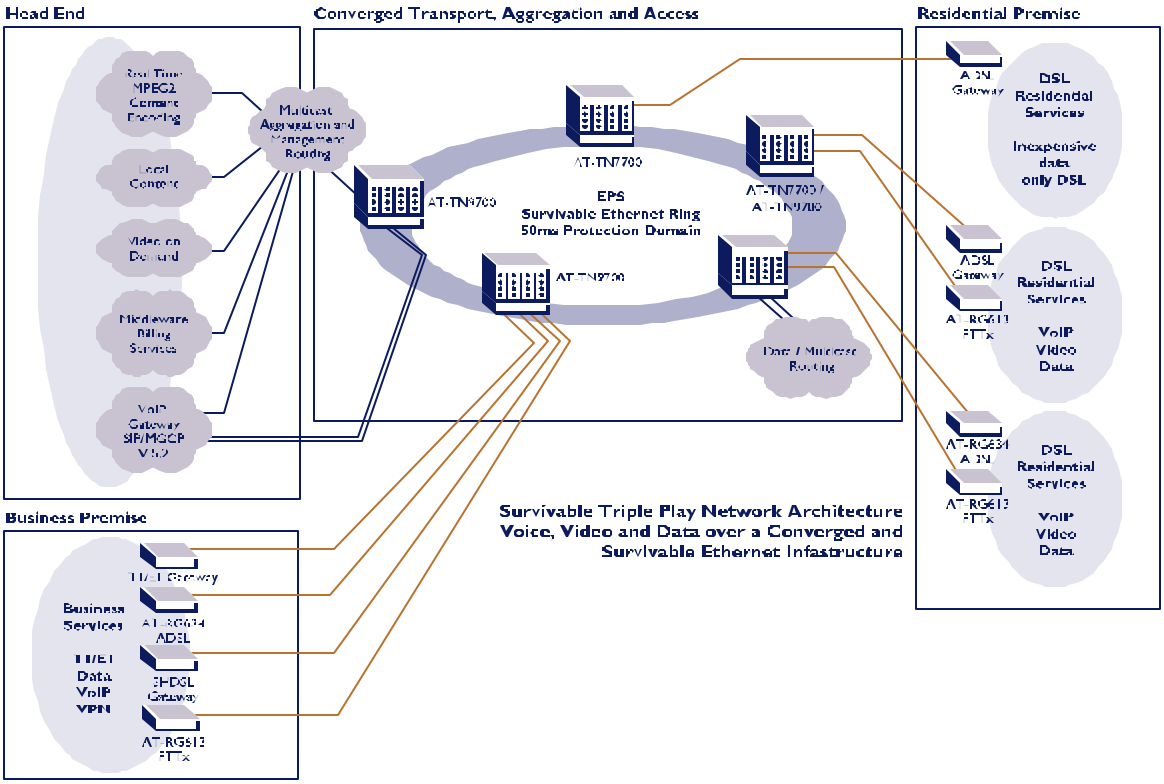


*Allied Telesyn equipment supporting EPSR will integrate seamlessly into an Extreme Networks environment.*

**Network Scalability**

With the advent of 10 gigabit Ethernet, scalable Ethernet networks became possible. When combined with 10 gigabit Ethernet, EPSR allows network providers to merge two traditional network elements—the transport node and the access node—into a single element. By combining these elements, the provider can reduce not only the number of managed nodes in the network, but also the amount of infrastructure and the cost required to manage those nodes.

Additionally, EPSR allows users to expand their networks without incurring downtime. This is possible because an existing ring can be broken, a node inserted, and the ring restored with minimal impact to service.

**Supported Services**

EPSR operates over standard Ethernet interfaces; as such all Ethernet-based services and controls are still supported. This includes:

- QoS
    - Marking and remarking
    - 802.1p priorities
    - Rate shaping / Metering
    - Bandwidth limiting
- IGMP Snooping
- Double Tagging (Tag stacking)
- All ranges of filtering

**Interoperability**

EPSR is based on RFC3619; as such it is fully compatible with other suppliers that adhere to that standard. This means that Allied Telesyn equipment supporting EPSR will integrate seamlessly into an interoperable network environment.

## Company Overview
# Allied Telesyn: It's Our Network, Too.

*A global company with nearly two decades of continuous profitability.*

Founded in 1987 with the goal of producing feature-rich, reliable, standards-based networking products, Allied Telesyn has a proven track record in bridging the gap left by other Ethernet networking manufacturers, whose solutions are often limited in scope or cost-prohibitive.

By taking cues directly from our customers and leveraging our global manufacturing competencies, we've evolved a market-focused approach to system development that is geared entirely to applications, rather than individual components. And by concentrating on battle-tested, end-to-end solutions for vertical market applications we avoid the scattershot, company-focused approach common in the industry. Our tagline: "It's Our Network, Too" is a testament to our high-level of accountability and to our investment in our customers' bottom line success.

*Allied Telesyn focuses entirely on end-to-end, purpose-built Ethernet applications.*

Allied Telesyn focuses entirely on end-to-end, purpose-built Ethernet and IP applications; with a complete line of networking products that includes Layer 2 switches, Layer 3 switches, carrier class fiber/copper Multiservice Access Platforms, wireless access points, wireless adapter cards and residential gateways. No other networking vendor can match Allied Telesyn's breadth and depth of Ethernet products—we are the leading manufacturer of media converters, unmanaged Fast Ethernet switches and hubs, fiber optic network adapters and other feature-rich interconnectivity products, worldwide. Additionally, Allied Telesyn has developed a world class systems engineering and support organization that ensures networks are designed and implemented to handle the stress of providing voice, video and data services.

*A world-class engineering and support organization spanning five continents and more than 30 countries.*

With engineering, manufacturing, sales, and distribution divisions strategically located throughout the Americas, Europe, Asia and Japan, Allied Telesyn is able to deploy solutions anywhere in the world, quickly and efficiently. And by rigorously testing products in design and support centers and leveraging our design and manufacturing competencies, Allied Telesyn is able to offer solutions for the access edge that are both customized and plug-and-play. This ideal combination helps our customers keep costs low, speed network deployment and maximize network uptime.

*The ideal choice for cost-conscious IT professionals who are looking for high-quality, feature-rich network solutions.*

Our customer-driven approach—combined with a pragmatic, value-based pricing scheme and a superlative service organization—has made Allied Telesyn a global networking leader, with more than 17 years of continuous profitability and products deployed in more than 50,000 companies in 30 countries and five continents. Allied Telesyn: the ideal choice for cost-conscious IT professionals who are looking for high-quality, feature-rich network solutions at a lower price.

## www.alliedtelesyn.com