



Merlink

for VxWorks

Link Aggregation Framework

Merlink, a member of TeamF1's Switchcraft family of layer 2 protocol products, is a robust, lean, high performance implementation of managed link aggregation for VxWorks® 5.x and AE. It enables multiple links to be aggregated together to form a link aggregation group that a MAC client can treat as a single link. This aggregated link acts as a higher-speed link providing increased bandwidth and fault tolerance. In the event of a link failure, **Merlink** redistributes traffic from the failed link(s) to the other links in the group. **Merlink** also balances the load across different active links in an aggregate group, with support for customizable load-balancing algorithms which can handle links of varying bandwidth. **Merlink's** tiny footprint and configurable modules have been specifically designed for use in an

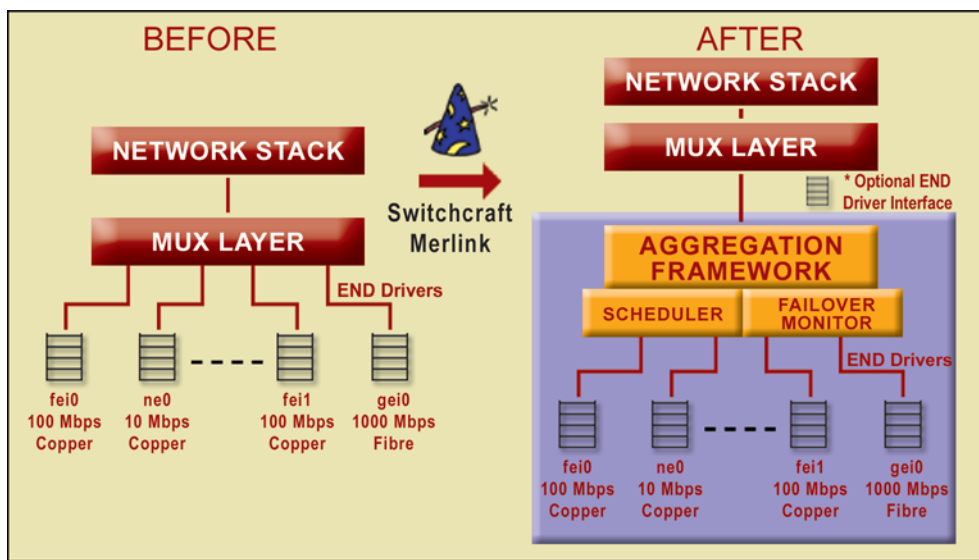
embedded environment. **Merlink** includes built-in support for the IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker protocol, and scales up from basic manual trunking to advanced automated aggregate management. With validated protocol compliance and interoperability with popular switch implementations, **Merlink** is the optimal solution for high-bandwidth and fault-tolerant applications.

Link Aggregation

An increase in data traffic requirements for a link usually involves a difficult choice between reducing services available on the link and dealing with incompatibility between various new physical and data link layer technologies that may provide the higher

Features

- ❖ Link Aggregation for increased bandwidth.
- ❖ Customizable load balancing between different links in an aggregate.
- ❖ Fault-tolerance via fail-over capability.
- ❖ Support for automated aggregation via 802.3ad LACP protocol.
- ❖ Specific ports can be LACP-enabled or disabled.
- ❖ Support for passive and active mode LACP.
- ❖ Handles heterogeneous links of different physical types and bandwidths.
- ❖ Includes Marker Protocol implementation.
- ❖ Supports custom event hooks detection.
- ❖ Support for CPU types of either endian-ness including PowerPC, MIPS, X86, ARM/XScale.



VxWorks Edition Features

- ❖ Designed exclusively for VxWorks and VxWorks-based Platforms
- ❖ Installs between the MUX and END layers without modifying them.
- ❖ Works with standard END drivers or switch chip drivers
- ❖ Can present an END interface for aggregates to the MUX
- ❖ Requires no special VxWorks source modifications
- ❖ Integrated with VxWorks build & install methodology
- ❖ Project Facility componentization
- ❖ Choice of net or system pools for PDUs
- ❖ Native support for VxWorks 5.4.x, 5.5.x, and AE

bandwidth. Combining existing parallel, physical layer links into a single logical link becomes increasingly attractive in such situations. This enables linearly scalable bandwidth per network connection while simultaneously providing high availability and reliability through multiple link redundancy. Further, it leverages economies of scale for existing or legacy physical layer technologies without having to adopt new ones. **Merlink's** framework contains APIs to manually or automatically add and remove existing links from aggregates, and has the capability to make an aggregate appear as a virtual END driver to the higher network layers providing for maximum application portability. On the lower end, the framework builds upon standard END drivers requiring no special OS or networking source code and supports links of different physical types and bandwidths transparently.

Fail Over

Besides the immediate cost and scalability benefits of utilizing multiple lower-bandwidth links instead of one higher bandwidth link, aggregates also enable increased system resiliency and redundancy. Failure of a single component link does not disrupt communications between connected devices, and the loss of a link within an aggregate just reduces available capacity while maintaining the connection and data flow. **Merlink** adds a unique twist to this, additionally providing for sets (sub-aggregates) of ports that can be marked for use only during failure conditions. When the regular ports in an aggregate are active, these ports may not participate in traffic handling. However, when a failure condition occurs, as determined by a configurable function that depends on the state of various links, the fail-over ports take over the data flow. When the regular ports are active again, the framework enables a seamless transition back to normal conditions, setting the pre-designated fail-over ports back into passive mode. **Merlink** thus helps achieve high-availability and fault-tolerance in critical embedded systems.

Load Balancing

To improve the traffic distribution characteristics for an aggregate, **Merlink** provides standard load balancing across links so that traffic is distributed across all links without overwhelming any single link. Normally, link aggregation assumes that a uniform distribution of traffic will occur over the links in an aggregate (except for ones connected to fail-over ports). However, this assumption of a trivial load-balancing may not work for specialized applications, as well as for links of different bandwidths and physical types. For this, **Merlink** provides a plug-in mechanism to introduce a custom load-sharing algorithm between various links, which defaults to a round-robin scheme if no such algorithm is present. The default scheme provides for configurability such as the relative priority of each link and quantum of each time slice across which balancing is scheduled and is a

802.3ad compliant distribution function that can be used as a reference implementation for custom load-balancing algorithms. Addition of custom load balancing schemes requires no knowledge of **Merlink**'s underlying implementation, yet allows for full flexibility via the provided APIs.

IEEE 802.3ad LACP

While **Merlink** supports manual deletion and addition of links to aggregates, its automated configuration facility also provides for a convenient way to manage and detect capabilities and connectivity of individual links within aggregation groups. **Merlink** includes full support for the LACP protocol based on IEEE 802.3ad specification. 802.3ad specifies automatic and dynamic configuration of appropriate sets of links within an aggregation group in the absence of manual overrides. Rapid configuration and reconfiguration are also possible in event of changes in physical connectivity of the links. **Merlink** allows selective configuration of LACP only on certain ports, if required, and supports both active and passive mode LACP based on whether the **Merlink**-enabled system wants to act as a source of control packets (PDUs) or wants to only respond to incoming packets.

Built for VxWorks

Merlink, like other members of the Switchcraft family, is a drop-in component for *VxWorks 5.x* and *AE*. It has been extensively validated on a variety of CPU architectures, which minimizes development and integration efforts. **Merlink** is especially optimized for *VxWorks* with support for standard END drivers and the ability to have the aggregation framework appear as an END device. It supports a variety of switch fabric chipsets and provides a standard reference implementation to glue in a custom switch chip. It enables highly sophisticated management of link aggregation groups in networked embedded systems with the fewest changes to the application and driver code. It is well integrated with the *VxWorks* build & install methodology and includes Project Facility (IDE) integration for its various modules.

Also Available

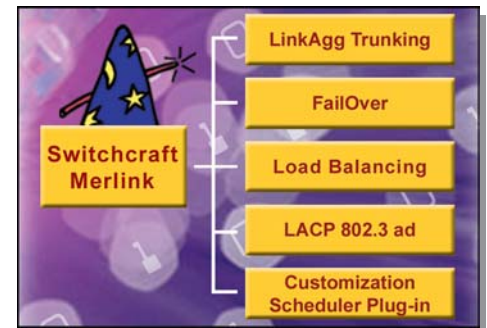
- ❖ **Switchcraft Spantasmic**
Spanning Trees (STP, RSTP, MST)
- ❖ **ClassHopper**
Alternate Queuing Disciplines
- ❖ **SShield**
Secure Shell
- ❖ **NetF1**
High Performance Network Stack

Custom Solutions

TeamF1's professional services can provide the resources and expertise to build customized implementations of **Merlink** solutions for different switching hardware, as well as support for custom load balancing schemes required by specialized applications.

Customization Flexibility

Merlink is available in full-source format and is highly customizable. Hooks are provided for application specific event-handlers to be called when links are added or removed from aggregation groups, or when an aggregation converges after a network connectivity change. With the APIs provided, any custom switch fabric chipset can be glued in for LACP control of hardware based aggregation. This, in addition to its custom load balancing algorithm capability, initialization time and run-time choice of fail-over ports, and complete control over the LACP activities of individual ports and the complete system, makes **Merlink** an ideal choice for critical network applications requiring bandwidth and high availability.



Email: sales@TeamF1.com
Web: www.TeamF1.com
Ph: 510-505-9931 ext. 5
Fax: (510) 505-9941



© 2002-2004 TeamF1, Inc