Crescendo, Kalpana, and ALANTEC

Third party networking add ons for SPARC servers

Sun Microsystems
Computer Corporation
A Sun Microsystems, Inc. Business

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

© 1992 Sun Microsystems, Inc.—Printed in the United States of America.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A

All rights reserved. This product and related documentation is protected by copyright and distributed under licenses restricting its use, copying, distribution and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® and Berkeley 4.3 BSD systems, licensed from UNIX Systems Laboratories, Inc. and the University of California, respectively. Third party font software in this product is protected by copyright and licensed from Sun's Font Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS
Sun, Sun Microsystems, the Sun logo, are trademarks or registered trademarks of Sun Microsystems, Inc. UNIX and OPEN LOOK are registered trademarks of UNIX System Laboratories, Inc. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCworks, and SPARCompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK® and Sun™ Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

X Window System is a trademark and product of the Massachusetts Institute of Technology.

Please
Recycle

Contents

Abstract

This paper contains information about products from Crescendo, Kalpana, and Alantec. It is meant to provide a snapshot of the products from these three companies, and how these products can add value in a Sun environment. The paper attempts to go to a level of depth between the "glossy level" and the "User's Guide" level. Included are product descriptions, performance information, pricing, contact information, and other relevant information in the 5-8 pages devoted to each company.

The fact that this paper includes information from these three companies does not imply any type of exclusive Sun endorsement. Do not hesitate to contact the vendors with questions.

Crescendo

Overview

Crescendo Communications manufactures wiring concentrators and SBus adapters which provide FDDI connectivity via optical fiber, or via unshielded twisted-pair (UTP) wire, or shielded twisted- pair (STP) wire using a technology called Copper Distributed Data Interface (CDDI). CDDI is simply FDDI with a different physical media interface. Therefore, FDDI and CDDI segments may be used interchangeably within the same ring. This conversion process adds no delay, and different media may be used in any combination without a performance penalty. Typically, wiring concentrators are interconnected with optical fiber. Twisted-pair wire is then used to connect each of these concentrators with workstations and servers. Alternatively, a CDDI-FDDI translator may be used to join twisted-pair links with fiber links.

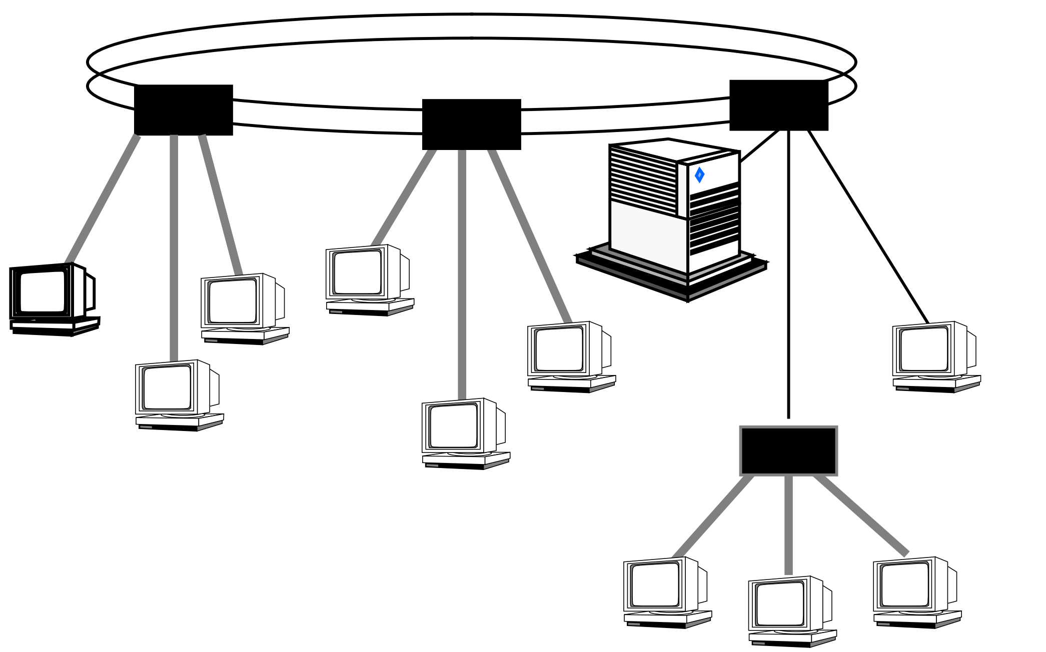

Products

CDDI Sbus Adaptor

The Crescendo CDDI SBus Adapter allows any SBus compatible workstation or server to connect to an FDDI network per the FDDI specification for a single-attach station (SAS). This is done by connecting to a Crescendo CDDI Workgroup Concentrator with twisted-pair wire. Like the concentrator, the adapter is fully compliant with SMT 6.2 and includes an SNMP agent for network management. A number of architectural features ensure maximum performance by providing faster internal data flow and minimizing the overhead imposed on the host processor. For upgrades, firmware may be downloaded, eliminating the need to swap PROM chips. The FDDI SBus Adapter is exactly the same as the CDDI SBus Adapter, except that it uses optical fiber instead of twisted pair. A fiber-optic MIC connector is standard; ST connectors are optionally available.

The CDDI Workgroup Concentrator

The CDDI Workgroup Concentrator has eight CDDI master (M) ports plus provisions for two additional ports. These two ports may either be FDDI or CDDI, depending on the type of optional card that is used. All CDDI ports are designed to accept four-pair UTP cables that are terminated with standard RJ45 connectors. An optional adapter is used to connect to STP cabling systems. The concentrator can connect directly to a backbone as a dual-attach concentrator or it may be cascaded from other concentrators. All valid FDDI topologies are supported, including the fault-tolerant dual-homing configuration. The CDDI Workgroup Concentrator is fully compliant with SMT 6.2 and includes an SNMP agent for network management. The FDDI Workgroup Concentrator is exactly the same as the CDDI version, except that it has eight FDDI ports instead of eight CDDI ports. All FDDI ports are designed to support fiber optic MIC connectors.

Crescendo Manager

The Crescendo Manager is a software application that may be invoked directly, or via SunNet Manager. Either way, it is possible to monitor and control any Crescendo wiring concentrator through a graphical user interface, which displays an image of the concentrator, including the current status of all panel indicators. More detailed examination or control of a concentrator's internal processes is accomplished through the use of simple pull-down menus.

Standardizing FDDI over Twisted Pair

Of the four specifications defined in the FDDI standard, only the PMD is modified to allow use of twisted-pair wire instead of optical fiber; the other three remain unchanged. Therefore, only a small portion of a fiber-based FDDI design needs to be modified to allow operation over twisted-pair wire.

In June of 1990, the American National Standards Institute (ANSI) authorized an effort to develop a specification for implementing FDDI over twisted pair wire. Subsequently, the TP-PMD Working Group was established. Every two months, representatives of over forty leading networking companies assemble to determine how best to accomplish this task. The Working Group is expected to submit a draft standard to the X3T9.5 Committee of ANSI by December of this year.

In 1991, the TP-PMD Working Group decided that a single PMD specification should be developed which supports both STP and high performance data grade UTP wire, effectively ending any possibility of a standard that would only support STP, including the approach recently advocated by IBM, called SDDI.

Recently, the Working Group voted to narrow the field of choices to two options: a proposal by Crescendo Communications, and one proposed by Cabletron and National Semiconductor. Both approaches use the scrambling scheme originally proposed by Crescendo. However, the two approaches differ in the type of encoding scheme that is used and the method used to compensate for the distortion of the signal by the wire. While the Cabletron approach is expected to be slightly simpler to implement, the Crescendo approach is generally regarded to be technically superior. Regardless of the outcome, Crescendo has a policy of upgrading customers to the standard at a cost not to exceed 10% of the product price.

Performance Testing

Independent benchmark testing by Scott Bradner at Harvard University found that Crescendo's SBus adapters can realize an effective throughput of 23 Mbits/sec when using TCP and 30 Mbits/sec when used with UDP. Both tests used the ttcp program to transfer data between two SPARCstation 2's using a buffer size of 16,384 bytes. Network Peripherals' FDDI SBus adapter was also tested under identical conditions, and was found to offer effective throughput of 27.6 Mbits/sec with TCP, and 51.6 Mbits/sec with UDP.[1] In contrast, 8.8 Mbits/sec was possible with Ethernet when used with TCP.

Pricing

CDDI Workgroup Concentrator (Order Number C1000)	$7,995
FDDI Workgroup Concentrator (Order Number C1001)	$12,995
CDDI SBus Adapter Card (Order Number C300)	$1,495
FDDI SBus Adapter Card (Order Number C301M)	$2,395
CDDI-FDDI Translator (Order Number C701)	$995
Crescendo Manager (Order Number C1800)	$895

Features

- In addition to optical fiber, unshielded and shielded twisted pair wire is supported using CDDI(TM) technology.
- An ANSI standard is being developed for FDDI over twisted pair.
- Once a standard exists, Crescendo will upgrade customers at a cost not to exceed 10% of the product price.
- Performance testing at Harvard Univ.: 23 to 30 Mbits/sec.
- Lowest cost per FDDI connection on the market: approx. $2500

1. Updated Information from Mitch Strobin at Network Peripherals as of 6/1/92, from new Bradner study.

Warranty and Support

All of Crescendo's products carry a one year warranty, and technical support is available through a toll-free "800" number. Since the cost per port on the concentrator is about $1000, and the cost of the adapter is under $1500, the total cost per connection is about $2500, a very cost effective FDDI solution. And it's only available on the Sun platform.

Crescendo Communications 710 Lakeway Drive Sunnyvale, CA 94086
(408) 732-4400, FAX 408-732-4604

Contact: Noel Lindsay (e-mail: lindsay@crescendo.com)

Kalpana

Overview

Kalpana Inc. manufactures products that use circuit and packet switching technologies to increase the throughput of Ethernet networks. Aimed at large organizations, Kalpana's products are designed for Ethernet environments where large numbers of users, workstations, file servers and bandwidth intensive applications are creating network overload. By creating up to seven parallel 10 Mbps data paths, the products increase Ethernet throughput 600% to 70 Mbps. With cascaded units, the throughput can exceed FDDI's 100 Mbps at a fraction of FDDI's cost.

Products

The Kalpana EtherSwitch is available in two models:

- EPS-700, with seven ports, and 30 Mbps of throughput
- EPS-1500 with fifteen ports, and 70 Mbps of throughput

Designed much like a telephone system's PBX, the EtherSwitch creates parallel connections between any nodes on any of the Ethernet LANS. The two reasons that the EtherSwitch is a fast throughput device are:

1. Parallel Paths: The EtherSwitch creates multiple parallel 10 Mbps data paths which multiplies the data throughput.

2. Lower Latency: The EtherSwitch uses "on the fly" packet switching which is much faster than a bridge or router's store and forward technology. The bridge or router must store the entire data packet while it processes it. The EtherSwitch looks at

only the first six byes of information to determine which output port to send it to, and the packet is never “stored” inside the box - the front of the packet is exiting the box before the back of the packet is all of the way in.

Table 1

Device	Technology	Latency
EtherSwitch	On the fly switching	40 microseconds
Bridge	Store and forward	800 microseconds
Router	Store and forward	1800 microseconds[1]

1. The ALANTEC router has first packet latency of 1200 microseconds, successive packets are sent out back to back (9.6 microseconds).

Applications

Below are two specific applications for the EtherSwitch in the Sun Environment.

1. File Server Front End.

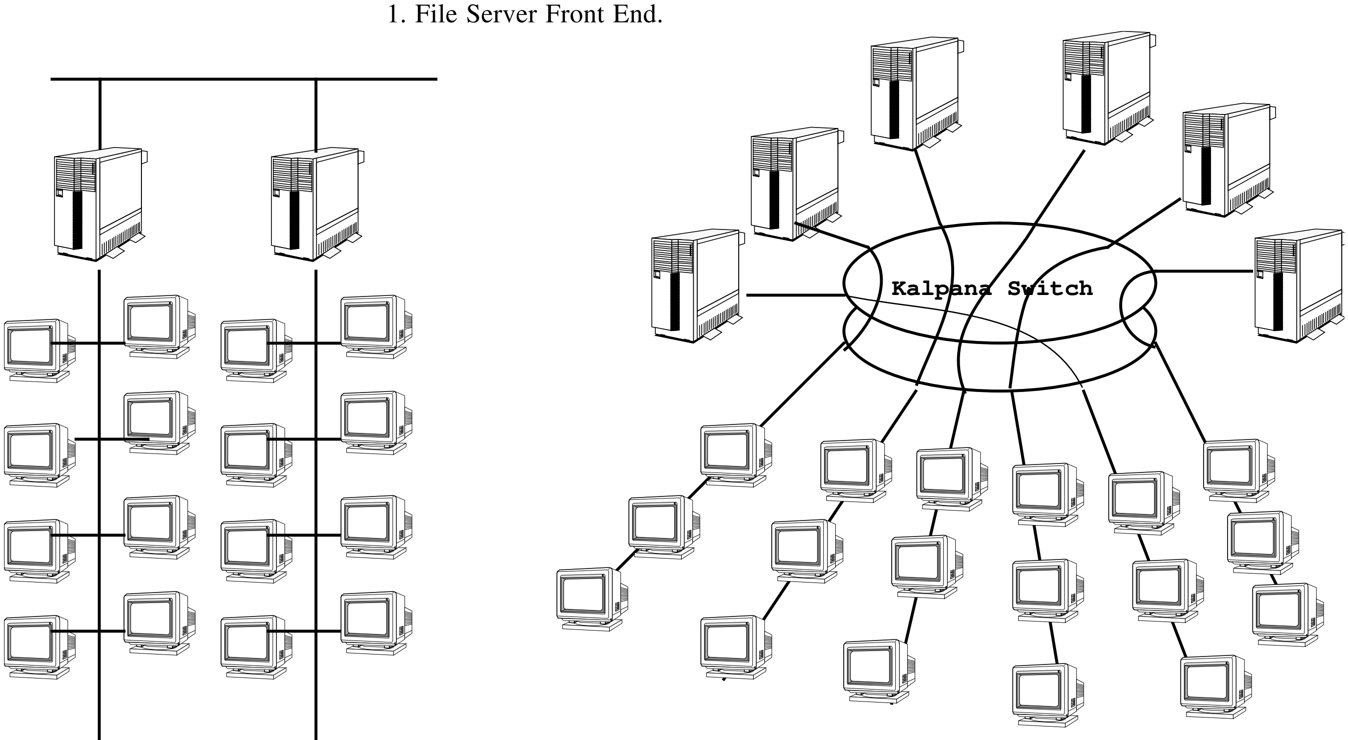

One data conversation at a time (10 Mbps) **Six simultaneous conversations (60Mbps)**

2. Multiple Connections per server:

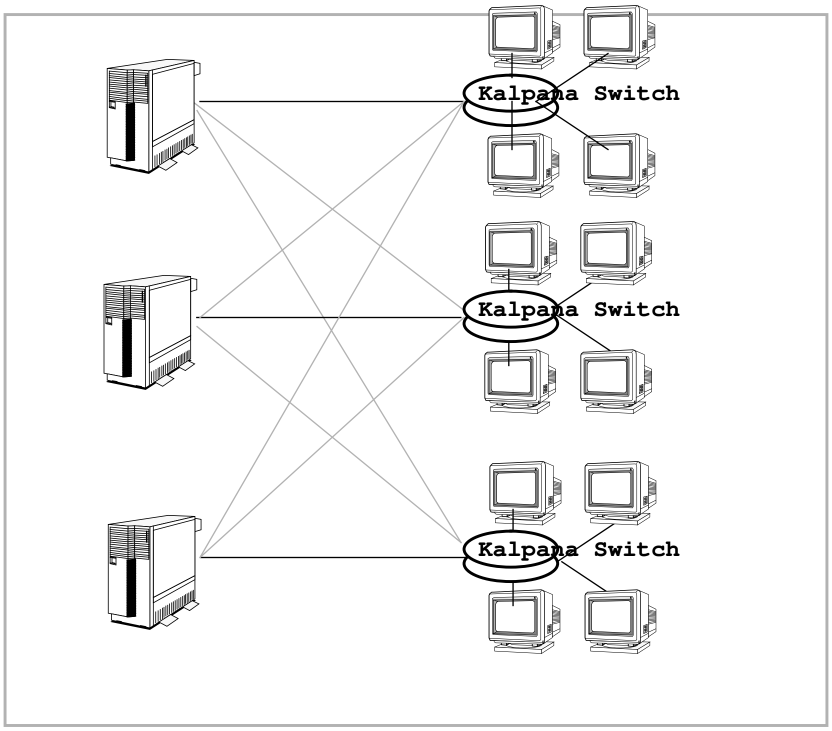

- Each server (in this diagram) can have three simultaneous sessions
- Each of the EtherSwitches has three server connections and four user segments.

Note: Using the EtherSwitch to front end Sun servers can help you beat competitive (NFS) servers that have up to 8 direct Ethernet connections. The 15 port EtherSwitch can provide 11 user segments in addition to four server segments, or alternatively 9 user segments and 6 server segments. The EtherSwitch also off-loads the intersegment traffic management formerly handled by the server. This frees up processing power for running applications. For an example of how Sun servers and the Kalpana EtherSwitch combine to provide a high performance NFS network read "Design and Implementation of a Large Scale High Performance NFS Network".[1]

1. White paper written by Sun SE Michael O'Connor April 24, 1992

Performance

Scott Bradner of Harvard published results of a throughput test that compared the EtherSwitch with products from Alantec, Synernetics, and Cisco. There were two tests which sent maximum traffic through the systems.

1. Four Stream test: This test passed four streams of packets to four distinct output ports so that there were no collisions.

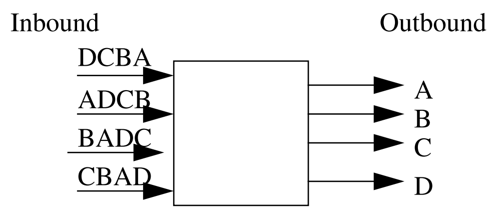

2. Six stream test: This test passed six streams of data to output ports in a way that guaranteed collisions because two of the six incoming streams always targeted a single output port.

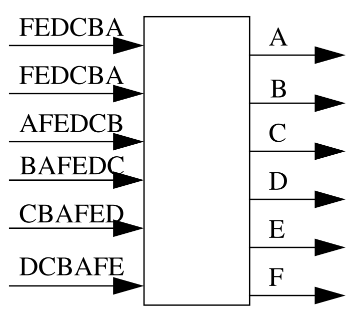

Both tests were conducted multiple times with different sized packets, ranging from 64 byte packets to 1518 byte packets.

The chart below shows the number of lost packets on both the 4 stream and 6 stream tests, using three different packet lengths.

% Lost Packets

Table 2

	Kalpana	Alantec	Synernetics	Cisco	
4 stream	0%	8%	11%	NA	64 byte packets
6 stream	24%	31%	24%	27%	
4 stream	0%	0%	6%	NA	128 byte packets
6 stream	0%	0%	9%	6%	
4 stream	0%	0%	8%	NA	1518 byte packets
6 stream	0%	0%	2%	0%	

(The Cisco product was not included in the 4 stream test.)

Theory of Operations

There is nothing to configure on the EtherSwitch, the only decision is regarding the number of port cards to install (up to 15 maximum). Each port card interfaces to an Ethernet segment which could have as many or as few servers or users as is desired, depending on how much contention of the 10 Mbps Ethernet is tolerable. In the case where one server is assigned to a port, there is a full 10 Mbps data path dedicated to that server. On segments that have users on them, one power user could be assigned a a full 10 Mbps Ethernet, or more likely multiple users would share the 10 Mbps Ethernet.

Each device on the network is assigned an IP address. When the EtherSwitch is plugged in, it builds its own address table so that it knows which port is associated with specific users or servers, or other devices. When the very first packet enters a newly installed EtherSwitch the EtherSwitch enters the sender’s source address in the address table and checks the table to see which port is associated with the destination address in the first six bytes of the packet. After not finding it in the address table, it broadcasts the packet out all EtherSwitch ports. The one port with that device on it responds and the EtherSwitch enters that address in the address table. This address building process takes about 200 microseconds per address, and after the address table is built (in 1 to 3 minutes on a busy network depending on the length of the packets), the latency in processing a packet is 40 microseconds.

When two packets are destined for the same server port on the EtherSwitch, one packet is sent to the server and the EtherSwitch buffers the other packet and then sends it to the server. The latency for the buffered packet ranges form 64 microseconds for a 64 byte packet to 1213 microseconds for a 1514 byte packet. Keep in mind that the majority of the packets on most networks are in the 64-500 byte range. Each EtherSwitch port can buffer up to 256 full sized packets.

Single Point of Failure

Kalpana is aware that there is demand for a redundant EtherSwitch. Until one is available, there are steps a customer can take to address this issue:

1. Where less than the full 15 ports are needed to handle the network configuration, one or more spare port cards can be installed so that if an active port card fails, that card's segment can be unplugged and moved to the spare card.
2. A second EtherSwitch can be completely wired to the same segments (duplicate wiring is limited to the wiring closet) and left inactive. If the active unit fails, it can be turned off and the duplicate unit can be turned on. Within a very short time it will have built its own address tables and be fully on-line.
3. Some customers have purchased an empty chassis to keep as a spare.

Features

- Parallel processing of Ethernet packets
- Boost Ethernet throughput 600% for a 14 port configuration
- Near zero port to port switching delays (40 microsecond latency vs. a bridges 800 microseconds of latency and a router's 1800 microseconds)
- 100% packet filtering and forwarding efficiency
- Automatic learning of network configuration
- Transparent to all high level protocols
- Fully IEEE 802.3 compliant
- Scalable architecture
- SNMP network management

Pricing

Table 3

Empty 15 port Chassis	$1,**500**
Ethernet Packet Processors (one port per card)	$800
Sample Configurations	
EPS-1500 with 4 ports - 40 Mbps capacity	$4,700
EPS-1500 with 8 ports - 80 Mbps capacity	$7,900
EPS-1500 with 15 ports - 150 Mbps capacity	$13,500

Summary

The Kalpana EtherSwitch will help sell Sun workstations and servers because it:

- Gives current users increased access to servers, decreases the wait for downloaded files, and increases network response time,.
- Increases network capacity by partitioning the network segments and increasing throughput. This enables the network manager to add more workstations/servers to what used to be an overloaded ethernet.
- Helps Sun compete against vendors who have up to eight ethernet connections. The EtherSwitch provides up to 11 ethernet segments in addition to four 10 Mbps connections to a Sun Server (or 9 segments and 6 server connections.)

Contacts

Regional Sales Managers:

Pete Wissowaty	Dave Carther	Mike Centrella
Santa Clara, CA	Chicago, IL	New York, NY
(408)988-1900 x116	(708)297-2990	(212)563-7070

International Sales Manager:

Karen Willner 408-988-1900 x109

If you have difficulty contacting the RM in your area, call Marc Cooper, Director Channel Marketing at (408) 988-1900 x113.

Overview

The ALANTEC PowerHub tightly integrates the Ethernet hub and FDDI technology with bridging and routing to create a high performance server front end.

Unlike previous generation Ethernet hubs which shared Ethernet's 10 Mbps bandwidth among multiple workstations, the PowerHub can deliver up to a full 10Mbps of dedicated bandwidth to every workstation on the network. The PowerHub does this while providing the flexibility to bridge, route, or both, between any two workstations or group of workstations attached to the network.

The PowerHub provides:

- Twelve Ethernet segments supporting 10Base5, 10Base2, 10BaseT, or 10BaseF/FOIRL network interfaces.
- Connection to one or two FDDI rings.
- Integrated bridging and routing over each Ethernet segment and FDDI ring.
- Ability to bridge, route, or both on a per port basis.
- Packet forwarding performance at Ethernet's theoretical maximum rate.
- Open management via SNMP, Telnet, and RS-232 connections.

Local Internetworking Hub

Ethernet networks are becoming bottlenecks. While Ethernet's 10 Mbps of bandwidth provides a very fast channel, it is shared by all devices attached to the network and is subject to collisions, significantly reducing effective throughput. Traditional 10BaseT hubs provide an attractive Ethernet wiring solution but do nothing to alleviate network congestion. All devices attached to a traditional 10BaseT hub share a single 10 Mbps channel.

One common way to increase available bandwidth is to use bridges and routers to partition a network into multiple segments, each with fewer users. Newer, faster wiring technologies such as FDDI are also being implemented as a way to enhance networking performance.

Ethernet hubs, bridges, routers, and FDDI all have their roles in modern, high performance network design. The ALANTEC PowerHub tightly integrates these technologies into a single scalable platform, creating a high performance, local internetworking hub.

Unlike traditional 10BaseT hubs, the PowerHub does not share a single 10Mbps channel among multiple users. The PowerHub delivers 120Mbps (12 Ethernets) which can be configured to deliver up to a full 10Mbps of bandwidth to individual users. The PowerHub can also share a single Ethernet between multiple users.

In addition to providing Ethernet bandwidth expansion, the PowerHub contains a high performance internetworking engine. Bridging and routing can be individually configured for each of the PowerHub ports.

For users that need more than (multiple) Ethernet bandwidth, the PowerHub provides connectivity to 100 Mbps FDDI. The PowerHub's two FDDI modules can be used to connect to an FDDI backbone, a local FDDI ring, or to provide single attached connections to the highest performance computing devices.

Application

Two of the most important PowerHub applications are as a high performance internetworking front end for servers, and as a workgroup accelerator for Power Users requiring up to 10Mbps of bandwidth to every workstation on the network.

In many instances today's high performance workgroups and file servers have network connectivity needs that far exceed what can be met in a traditional shared Ethernet environment. For this reason many servers have been equipped with multiple Ethernet adaptors, each of which supports it own IP subnet. While this type of configuration helps to relieve Ethernet congestion, it has a disadvantage in that it unnecessarily burdens the server with routing functionality and thereby decreases the server's overall performance.

An alternative network topology would place the PowerHub between the server and its clients, creating a server front end. The PowerHub could then provide a single dedicated 10Mbps to the server, multiple 10Mbps connections, or even a 100 Mbps FDDI connection.

This type of topology enhances performance in a variety of ways:

- Server performance is increased as CPU cycles previously used to implement routing are made available for NFS and other server functions.
- Peer to peer performance is increased as local routing and/or bridging are handled by the PowerHub.
- Power User performance is increased as they are provided with dedicated 10Mbps channels.
- Workgroup performance is increased as overall available bandwidth is increased from 10Mbps to 120Mbps or even to 320Mbps with the use of FDDI.

Other PowerHub applications include:

- A hub of hubs for connecting together traditional 10BaseT hubs.
- A local network hub providing “backbone in a box” functionality for small high performance networks.
- A migration path to FDDI for workgroup and backbone connectivity.

Performance

In October of 1991, Scott Bradner of Harvard University conducted a series of tests designed to determine the performance of a variety of commercially available bridges and routers. The findings of these tests were made available over the internet and were partially reported in the February 1992 issue of Data Communications magazine.

The testing methodology was to send the theoretical maximum number of packets for a given number of Ethernet streams to a bridge or router and measure the number of packets lost by that device. A stream is defined to be a data transfer between two devices separated by the bridge or router being tested. Thus, a six stream test measures the maximum performance by the device for 12 Ethernet ports.

The results of these tests showed that there exists a class of very high speed bridges and routers that can forward packets at or near “wire speeds”, i.e. at the theoretical maximum Ethernet rate, even when as many as twelve ports are active. In nearly all cases these devices can perform at wire speeds if the packets being forward are 128 bytes or larger, and suffer only a slight performance degradation if the packets are smaller.

Both the ALANTEC PowerHub and the Cisco AGS+ could achieve these results when bridging or routing. Bridge only products with similar performance include the Synernetics LANplex 5000 and Kalpana's EtherSwitch.

The following charts give the maximum packet forwarding rates and the rates achieved by the PowerHub and the EtherSwitch. These tables indicate that both devices forward packets at Ethernet's maximum rate, except in the case of the smallest Ethernet packets. When transferring only 64 byte packets over six streams the EtherSwitch outperforms the PowerHub by 9%. (This testing also showed the PowerHub achieving similar routing performance.)

Bridging Performance Over Four Streams

Table 4

	64	**128**	**1024**	**1518**
Maximum	59,520	33,780	4,788	3,248
PowerHub	54,758	33,780	4,788	3,248
EtherSwitch	59,520	33,780	4,788	3,248

Bridging Performance Over Six Streams

	64	**128**	**1024**	**1518**
Maximum	89,280	50,670	7,182	4,872
PowerHub	61,603	50,670	7,182	4,872
EtherSwitch	67,853	50,670	7,182	4,872

Features

Ethernet

The PowerHub's 12 Ethernets can be configured as dedicated pipes, or they can be shared by multiple users. Features includes:

- On board connectivity to the full range of Ethernet cabling types:
  - 10Base5 (AUI)
  - 10Base2 (BNC)
  - 10BaseT (RJ45)
  - FOIRL/10BaseF (50, 62.5, 100 micron multimode cable with ST connectors)
- Support Ethernet Type 2 format

FDDI

- One or two FDDI interface cards, each providing either a single dual attach or two single attach connections. (self configuring SAS or DAS)
- Support for bridging, routing, or both between FDDI and all twelve Ethernet ports.
- Bi-directional Ethernet/FDDI translation.
- SMT 6.2 with SNMP translation.
- Support for 62.5 micron multimode cable with MIC connectors.

Bridging

Some or all PowerHub ports can be configured to provide protocol independent bridging. Features include:

- Aggregate bridging performance of 62,000 packets per second for 64 byte packets.
- Bridging performance at wire speed with packets 128 bytes or greater.
- Automatic learning of Ethernet addresses with user configurable aging.
- Permanent bridge table entries.
- Support for the IEEE 802.1 Spanning Tree Protocol
- User definable logical filtering on any part of the Ethernet packet. Up to 98 templates and 62 rules using these templates may be specified.
- Support for up to 8,192 MAC addresses.

Routing

Some or all PowerHub ports can be configured to route IP. Features include:

- Aggregate routing performance of 51,000 packets per second for 64 byte packets.
- Routing performance at wire speed with packets 128 bytes or greater.
- Ability to configure multiple ports as a single subnet.
- Ability to configure a single port as multiple subnets.
- Support for Routing Information Protocol (RIP)

Management

- An on board SNMP MIB II agent.
- Support for two Telnet connections and two directly connected RS232 devices for configuration, diagnostic, and management activities.
- Complete port level statistics.
- Support of ICMP PING, incoming and outgoing.

Security

- Password level security
- Optional IP routing security (RFC 1108)
- Configurable network groups, allowing communication only between selected ports.
- Ability to block learning of bridge table entries.

Pricing

Table 5

Model	Price (US $)	Connectors
3100	$16,800	12 10BaseT (RJ45) & 2 10Base5 (AUI). (Use of an AUI connector disables an RJ45 connector.
3200	$19,800	12 (10Base5) AUI & 12 (10 Base2) BNC. A software selection is made to use either the AUI or the BNC connector for each of 12 Ethernet segments.
3300	$23,800	12 FOIRL/10BaseF (ST) & 6 10Base5 (AUI). Individual Ethernet segments are software configured to utilize FOIRL or AUI.
3500	$18,800	36 10BaseT (RJ45) & 2 10Base5 (AUI). Three RJ45 connectors per Ethernet are provided, allowing one, two, three devices to share a 10 Mbps Ethernet. (Use of an AUI connector disables an RJ45 connector set.)

Contact Information

ALANTEC 1-800-alantec

Sun Microsystems
Computer Corporation
A Sun Microsystems, Inc. Business

For U.S. Sales Office locations, call: 800 821-4643
In California: 800 821-4642

Australia: (02) 413 2666
Belgium: +32 2 759 38 11
Canada: 416 477-6745
Finland: +358-0-502 27 00
France: (1) 30 67 50 00
Germany: (0) 89-46 00 8-0

Hong Kong: 852 802 4188
Italy: 039 60551
Japan: (03) 3221-7021
Korea: 822-563-8700
Latin America: 415 688-9464
The Netherlands: 033 501234

New Zealand: (04) 499 2344
Nordic Countries: +46 (0) 8 623 90 00
PRC: 861-831-5568
Singapore: 224 3388
Spain: (91) 5551648
Switzerland: (01) 825 71 11

Taiwan: 2-514-0567
UK: 0276 20444
Elsewhere in the world, call
Corporate Headquarters:
415 960-1300
Intercontinental Sales:
415 688-9000

Printed in USA

4/92