Black Box Corporation
The Source for Connectivity®

FEBRUARY 1996
LB9306A
LB9313A

SuperSwitch-6
SuperSwitch-13

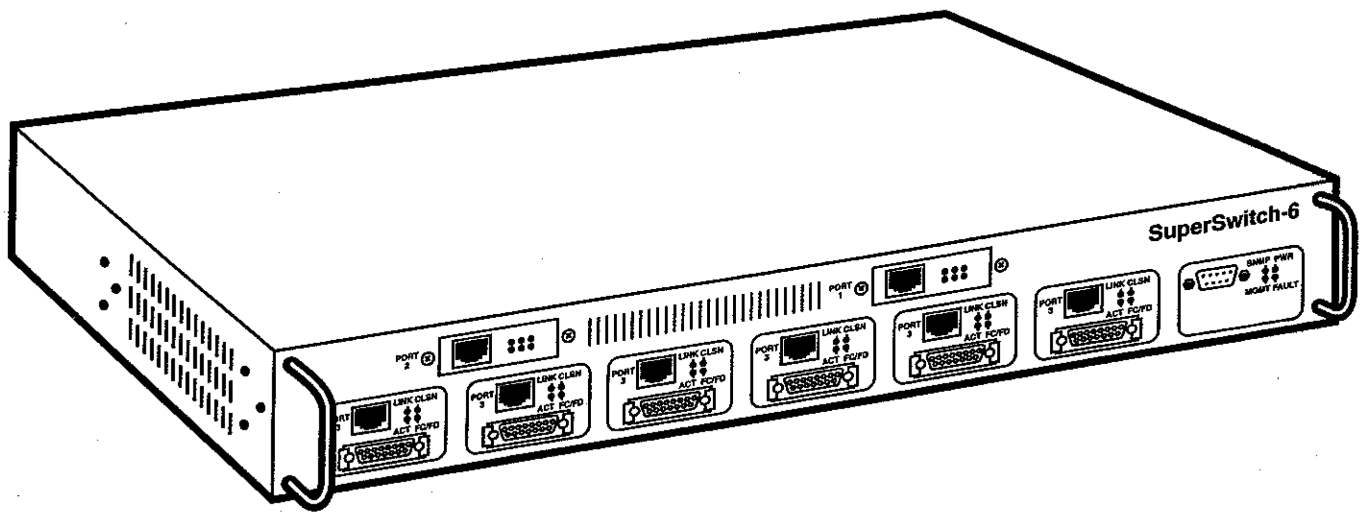

CUSTOMER SUPPORT INFORMATION

Call our Technical Support Specialists to discuss your application.
For 24-hour technical support: Call **(412) 746-5500** Or Fax: **1-800-321-0746**
To order: Call **(412) 746-5500** 8:00 A.M. to 8:00 P.M. EST
Mail order: **Black Box Corporation**, P.O. Box 12800, Pittsburgh, PA 15241

Contents

1. Specifications

Buffers	28 KB per port
Addresses	1024
Addresses Table	Transparent, automatic self-learning at full wire speed. Cache aging time manageable. Custom filtering by hardware address and port.
Store-and-Forward Switching	Provides complete runt and error filtering on all packets. Flow control prevents packet loss.
Network Management	In-band and out-of-band SNMP w/MIB II, private MIB and out-of-band serial console support
Filter/Forward Rate	~500,000/500,000 pps
Learning Rate	~250,000 pps
Status Indicators (per port)	(4) LEDs: LINK, ACT (Activity), CLSN (Collision), FC/FD (Flow Control/Full Duplex)
General Indicators	(4) LEDs: MGMT ACT (Management activity), SNMP (SNMP presence), POWER (power-supply status), FAULT (device-failure detection)
Boot and Configuration	NVRAM configuration loaded on power up and fully downloadable. Firmware local/remote downloadable.
Interfaces	AUI DB15, UTP RJ-45, RS-232 DB9
Standards Supported	IEEE 802.3 10BASE5, IEEE802.3 10BASE-T, 100BASE-TX Fast Ethernet, 100BASE-FX Fast Ethernet, FDSE (Full Duplex Switched Ethernet), IEEE 802.1d, SNMP (RFC 1157, etc.), MIB II (RFC 1213, etc.), Bridge MIB (RFC 1493)
Mounting	Standard 19-inch rack, with mounting brackets. 1" minimum clearance, sides and back
Emissions and Safety	FCC Part 15, Sub. J, Class A, ETL (UL® 1950), TUV/VDE, CSA, BABT, CE
Cooling	(2) DC fans
Operating Temperature	32 to 122°F (0 to 50°C)
Storage Temperature	14 to 149°F (-10 to 65°C)
Relative Humidity	<95%, noncondensing
Power	90 to 260 VAC, 50/60 Hz
Size	LB9306A: 2.5"H x 17.3"W x 14"D (6.4 x 44 x 36 cm) LB9313A: 3.37"H x 17.3"W x 14"D (8.6 x 44 x 36 cm)
Weight	LB9306A: 15 lb. (6.8 kg); LB9313A: 18 lb. (8.16 kg)

2. Introduction

2.1 Description

Ethernet segment switches divide a large, unwieldy local network into smaller segments, insulating each from the other's local traffic and increasing aggregate bandwidth, while still retaining full connectivity. The SuperSwitch is designed to be self-configuring and extremely fast, making installation simple and providing an increase in available bandwidth.

The Switch learns on which port stations are located by remembering the source address of every packet received on every port. When packets are received for stations that the Switch has learned, they can be forwarded to only the port that the station is on. This is different from a hub, which sends every packet to every port.

If the Switch receives a packet addressed to a station that the Switch has not learned, it sends the packet to every port (except the port on which it was received), just as a hub would. This ensures that full connectivity is maintained. Broadcasting packets for stations that are not learned generates unnecessary traffic, but the extra traffic generated by unlearned addresses is insignificant since every protocol requires responses for packets that have been received, causing the switch to learn the address of all active stations.

2.2 Features

- (6) Auto-select AUI or UTP Ethernet ports
- (7) UTP-only Ethernet ports (LB9313A)
- Standard 19-inch rackmount chassis (rackmount kit included)
- Auto-ranging power supply (automatically adjusts to any voltage between 90 and 264 VAC at 50/60 Hz)
- Reverse polarity detection and correction (UTP interface)
- Full duplex selectable on all ports for either AUI or UTP
- Flow control selectable on all ports
- Two interchangeable fast interface ports (optional)
- Up to 1024 address cache entries
- Extensive custom-filtering table
- Serial console port with password protection
- Downloadable management system software (serial or TFTP)
- SNMP support
- Spanning Tree

2.3 The Switch and Network Management Systems

The SuperSwitch can be monitored and controlled through the ONMP Network Management System (NMS), or through a generic SNMP NMS.

The SuperSwitch does not require ONMP or any other NMS; however network-management functions greatly assist monitoring and controlling your network. For example, through ONMP or SNMP you can monitor port statistics, configure individual ports, and view bridging information.

For information on connecting your Switch to an ONMP station or performing network-management tasks using ONMP, refer to your ONMP User Manual.

For information on configuring your Switch for use with an SNMP NMS, consult your Administrative Interface User's manual.

3. Typical Configurations

A Switch boosts network performance by sending network traffic only to the ports that need to see it. In a normally functioning network, this means that most packets will appear on only the segment with the sender and the segment with the receiver. One way to make sure that maximum performance is realized is to configure the network so that for most of the traffic the sender and receiver are on the same segment. Using this simple rule to configure your network can significantly increase the usable bandwidth.

3.1 Single-Server Environment

In a network with a single server, a good way to improve bandwidth is to attach a dedicated 100-Mbps Fast-Ethernet link to the server. This can yield up to a 1000% improvement in throughput. In order to establish a Fast-Ethernet connection to the server, install a 100BASE-TX Fast Ethernet NIC in the server machine and connect it to a port on the Switch which has been configured for Fast-Ethernet operation. The client machines may then be distributed among the other six ports. Note that in this configuration, during peak usage, the Flow-Control indicators may light. This is normal for this configuration. Make sure that Flow Control is enabled on all of the client ports (see Figure 3-1). As clients are upgraded to 100-Mbps operation, they can be added to a 100-Mbps repeater local to the server. If you install a second fast interface into the SuperSwitch, a second 100-Mbps switched segment can be connected.

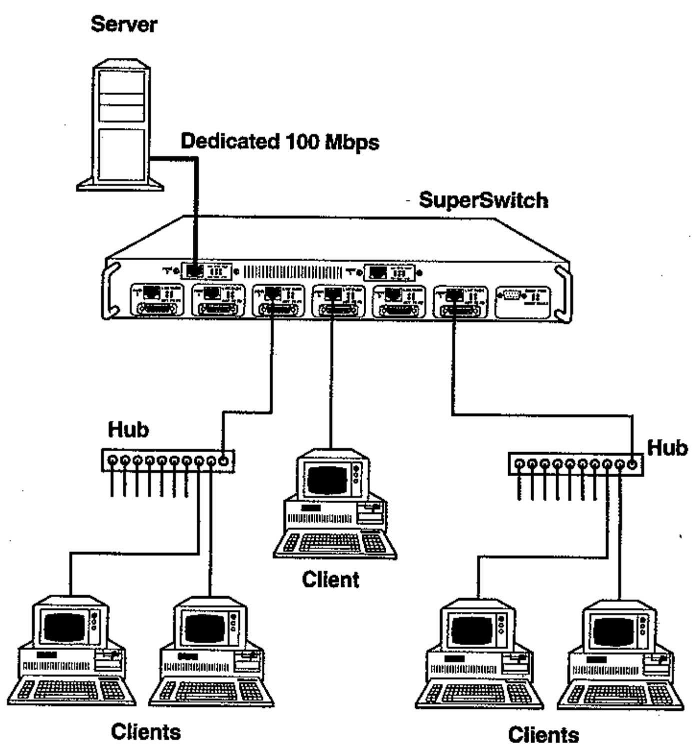

Figure 3-1. Single-Server Environment.

3.2 Multiple-Server Environment

With multiple servers, how you install the Switch depends on your network's traffic patterns.

If your network is made up of distinct groups of users, with each group using mostly that group's own server, then connect each group and its server on the same port of the switch.

For instance, suppose that there are several departments, and each department has its own server. The departments are interconnected primarily for email and backup purposes. If you connect each department, with its server, to its own port on the Switch, only the email and backup traffic has to go through the Switch.

If, however, your users typically use more than one server extensively, then configure your network so that machines that serve more than one group are given their own port on the switch. For instance, suppose that there are several departments, each of which has its own file server, and there is a single firewall router that is connected to a backbone. Each department should have its own port on the switch. The router should also have its own port on the switch, since it used by members of all departments. (See Figure 3-3.)

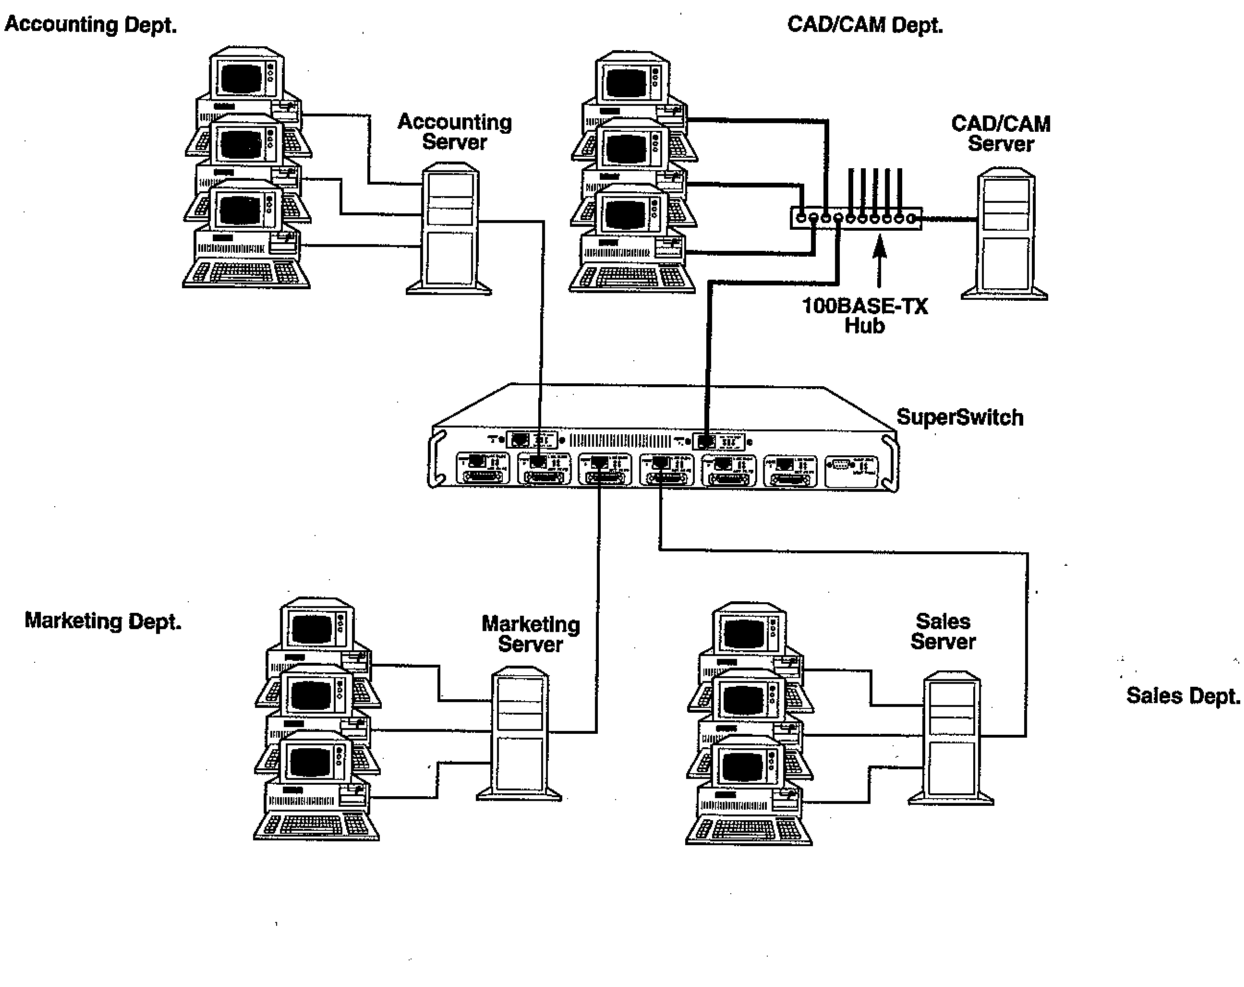

Figure 3-2. Multiple-Server Environment.

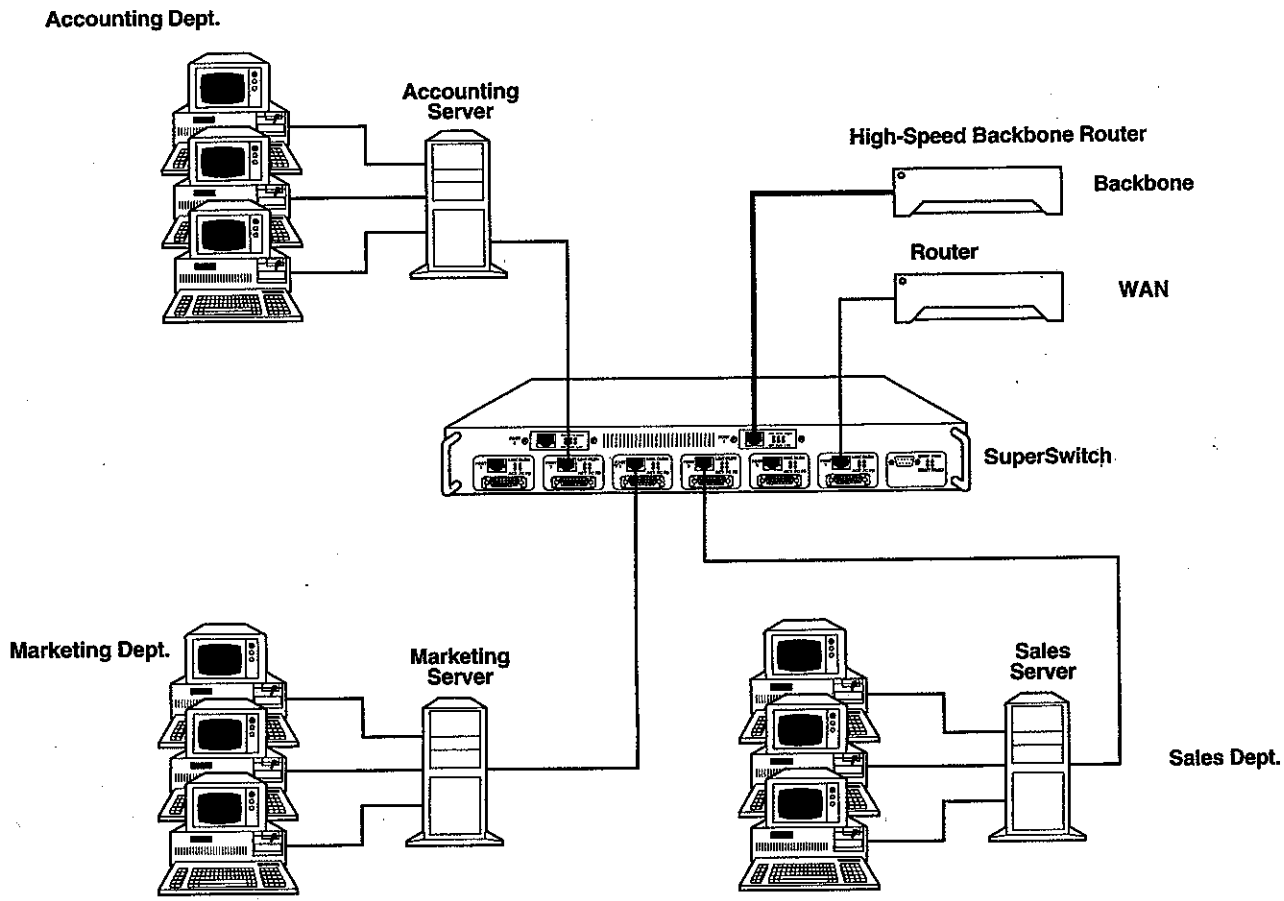

Figure 3-3. Multiple Server/WAN Environment.

3.3 Peer-to-Peer Environment

Peer-to-peer environments can be difficult to configure. If there is equal traffic between all machines, then any division of machines will work. However, if the traffic is not completely uniform, then careful configuration is required. For instance, if a few machines are used primarily as servers, with most other machines as clients, even though they can all communicate with each other, the machines are best arranged as if dedicated servers and clients are being used. A good example of this is a typical UNIX® environment. Frequently, in UNIX workgroup environments, application binaries are shared throughout the cluster. However, since different platforms require different binaries, disks are typically shared primarily among machines of the same type.

A natural configuration would be to place systems from the same vendor on each segment of the Switch. On the other hand, many workgroups consist of machines from several vendors, and since data files are more likely to be shared within the workgroup, it may be best to ignore the machine's vendor and assign ports based upon the workgroup. Ultimately, the way the network is used dictates the optimum configuration; as much traffic as possible should be filtered by the Switch. (See Figure 3-4.)

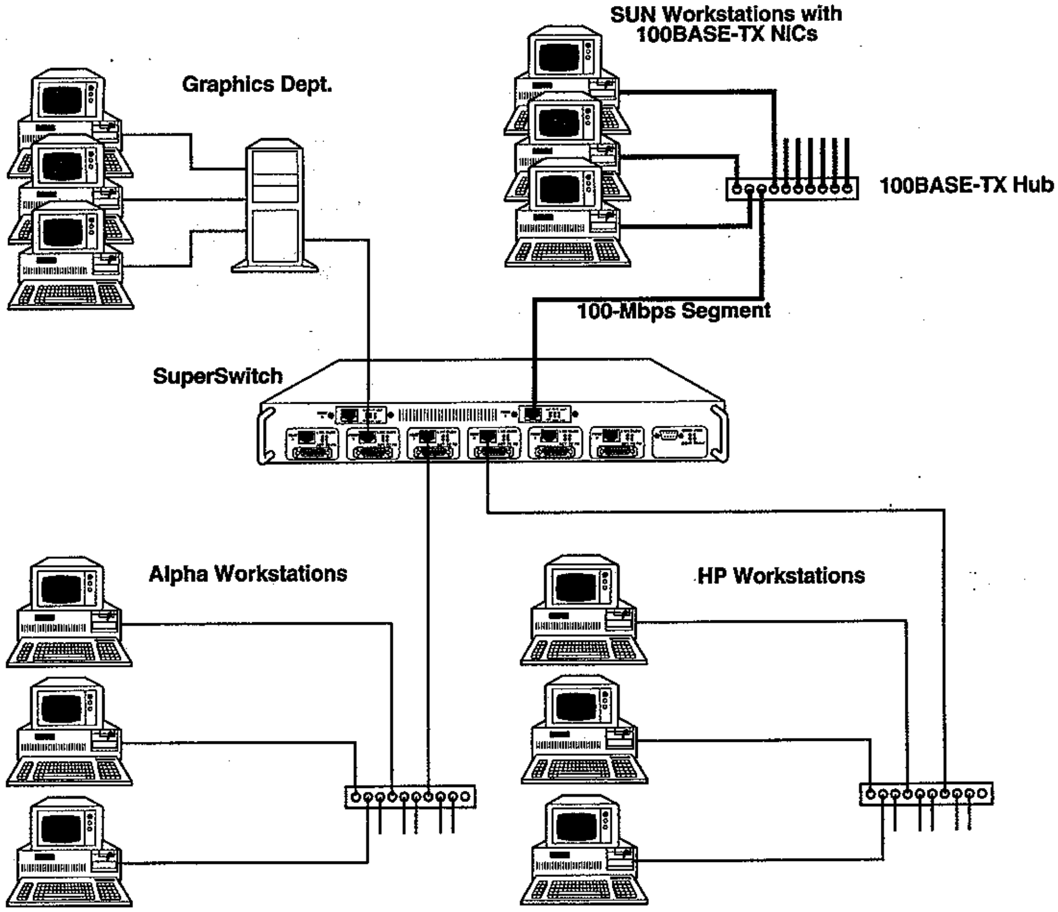

Figure 3-4. Peer-to-Peer Environment.

3.4 Heterogeneous Environment

The simplest way to configure a heterogeneous network is by machine type. Since machines that use different protocols rarely need to communicate with each other, it's best to give each protocol a different port on the Switch. For instance, if your network has AppleTalk®, NetWare®, UNIX, and DECnet™ machines, each should have its own segment. If this isn't possible because of cabling or logistic constraints, try to configure your network as closely as possible to this model; a few machines connected on a less-than-ideal port will not have a very large an effect on the network's performance. (See Figure 3-5.)

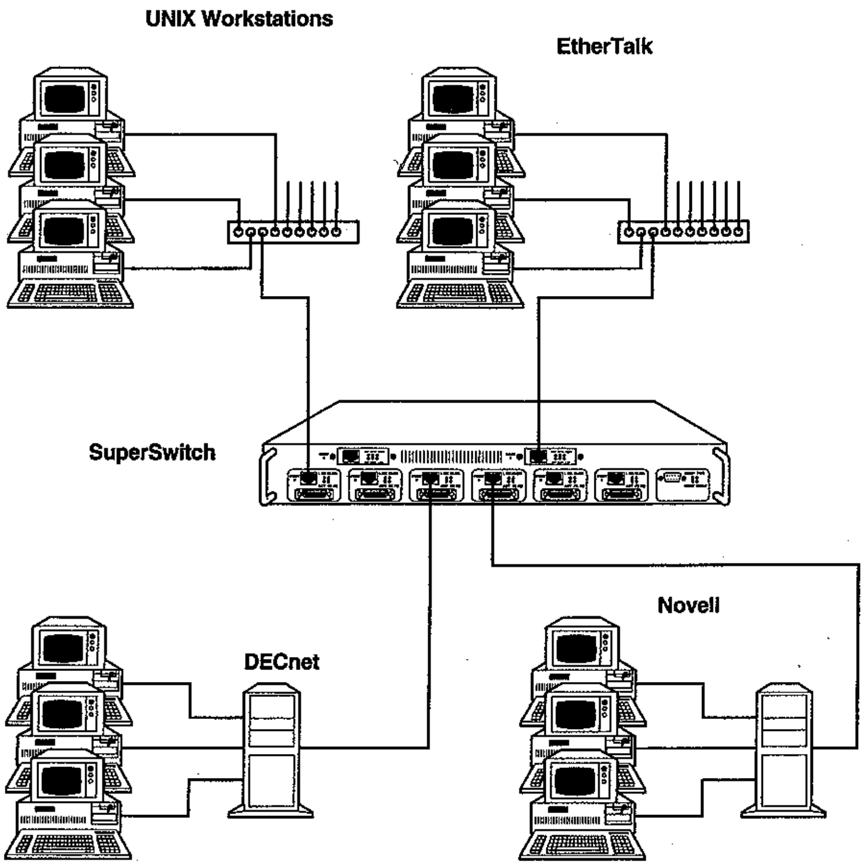

Figure 3-5. Heterogeneous Environment.

3.5 Excessive Flow Control

During times of peak network use, you may occasionally see the Flow-Control indicator flash. This is normal. However, if it stays lit for more than a few seconds at a time (or there is an excessive number of flow controls reported by the NMS), it could indicate a problem with your network configuration.

A port's Flow-Control indicator flashes whenever a packet is received that needs to be forwarded to a port that already has too many packets queued for it. This indicates a temporary over-bandwidth situation on one port (the total traffic to the port was in excess of 10 Mbps). This typically occurs when there are several fast machines on different ports trying to access a machine across the Switch. If this occurrence was due to an unusual event, then no further action is necessary. If this is part of the normal use pattern for the network, then the stations causing the flow control to activate should be identified and moved to the same segment as the machines they are communicating with.

When a situation arises where Ethernet bandwidth is insufficient for the traffic, there are only three possible actions: drop packets, use flow control, or replace the network.

Buffering packets only works for a very short while; an extended over-bandwidth situation will eventually overflow buffers and cause dropped packets.

Flow control is the most reasonable solution, since it relies on the Ethernet collision-detection mechanism to relieve temporary overbandwidth situations.

Replacing the network is theoretically the best solution, but is too big an investment in time, as well as money, to solve a problem which happens only occasionally.

3.6 Repeater Count Limits

The Switch does not count as a repeater. Each of the segments connected to the Switch can support a full Ethernet network. There can be up to four repeaters/hubs between the Switch and any station.

If you have specific questions about your network configuration, or have a particularly difficult network, call your supplier.

4. Installation and Setup

The SuperSwitch can operate as a stand-alone unit or in conjunction with many other Ethernet products. The SuperSwitch can be managed with ONMP software or with any SNMP-compatible NMS, or through an out-of-band console interface.

4.1 Installing the Switch

Step 1: Determine the Best Location for the Switch
Mount the Switch in a 19-inch rack using the included rackmount ears, or place it on a secure flat surface after installing the included rubber feet. Make sure that the unit is within reach of the necessary connections (such as a power outlet, Ethernet connections, and, if the Switch will be monitored via the serial port, a PC, UNIX workstation, or modem).

Step 2: Plug in the Switch
Simply plug the power cord into the Switch and an easily accessible outlet. Turn the power switch to the ON position. A built-in power supply automatically adjusts to any outlet providing between 90 and 264 VAC at 50/60 Hz.

For a 115-volt configuration, the power cord to be used is minimum type SJT (SVT) 18/3, rated 250 volts AC, 10 amps with a maximum length of 15 feet (4.5 m). One end is terminated in a IEC 320 attachment plug, the other in a NEMA 5-15P plug.

The power cord to be used with a 230-volt configuration is minimum type SJT (SVT) 18/3, rated 250 volts AC, 10 amps with a maximum length of 15 feet (4.5 m). One end is terminated in an IEC 320 attachment plug. The other end is terminated as required by the country where it will be installed.

Step 3: Connect the Ethernet Devices
For best performance, the Ethernet segments connected to the Switch must be carefully configured. Generally, the segments should be configured so that machines on a given port primarily communicate among themselves (most traffic does not need to cross the switch). However, there are situations for which this is not the best configuration. Refer to **Chapter 3** for a more detailed discussion of network configuration.

NOTE

The default configuration of all ports is twisted-pair mode only. In order to change this default, the Administrative Interface or an SNMP NMS must be used.

To connect an Ethernet device to a 10BASE-T or 100BASE-TX port:
The 10BASE-T and 100BASE-TX ports on the Switch are designed to be connected directly to a hub using a standard straight-through patch cable. In order to cascade Switches or connect a workstation to the Switch, there must either be a hub in between them, or a crossover cable must be used. Refer to **Section 4.5** for a more detailed description of cabling.

To connect an Ethernet device to an AUI port:
The AUI port is designed to be connected directly to an AUI drop cable. To connect any other type of media (for example, ThinNet or fiber), an external transceiver is required. When a transceiver is used, make sure that the SQE Test is disabled on the transceiver. If the SQE Test is not disabled, the Switch will function normally; however, the collision indicator and the collision-event counters may not be accurate.

To use Full Duplex:
Full Duplex may be used in either UTP or AUI mode. Full Duplex may be selected via management. If Full-Duplex AUI is selected, a full-duplex transceiver must be attached. A half-duplex (normal) transceiver attached to a full-duplex port will cause the Switch to behave abnormally, but will not cause permanent damage.

Step 4. What to Do Next
If you will be using the SuperSwitch as a standalone device (not under ONMP or other NMS control), you have completed the installation and setup of the Switch. Refer to **Chapter 3** to learn how to optimize your network for operation with the Switch.

If the SuperSwitch will be controlled by ONMP, read your ONMP user manual.

4.2 Using a Generic Network Management System

Refer to the Administrative Interface User's manual for instructions on connecting and configuring your Switch for use with an SNMP NMS, and for instructions on using the console Administrative Interface.

4.3 Operating the SuperSwitch

Operation of the Switch requires minimal user intervention. The unit automatically learns the addresses of new stations as they appear, and will relearn addresses dynamically if the network is reconfigured. In order to optimize network performance, the network must be carefully configured. Refer to **Chapter 3** for a discussion of optimizing your network for switching.

4.4 Understanding the LEDs

Each of the Ethernet ports on the Switch has four status LEDs. Each of the Fast Ethernet ports has six status LEDs. There is a separate bank of four LEDs for the management status. Refer to Figure 4-1 for their placement.

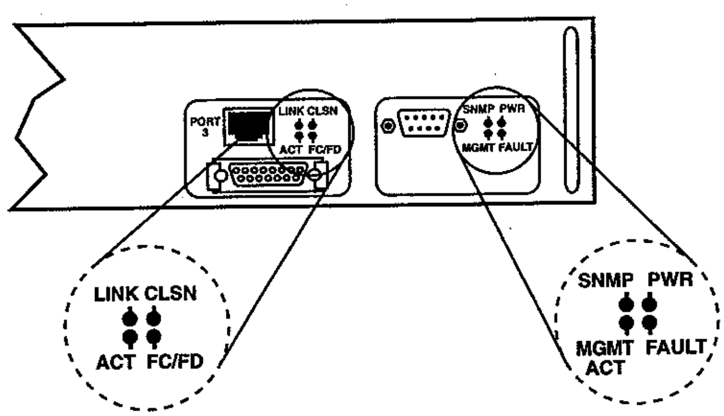

Figure 4-1. A Closeup of the LEDs.

4.4.1 Ethernet Port Status LEDs

LINK indicates whether or not the UTP port is attached to a device. This will only light when the UTP connection on the Switch is active and a connection has been established to another device. The LINK light will never come on if the AUI interface is selected.

ACT indicates when there is activity on the attached segment. This activity can be due to either the Switch transmitting, traffic which will be forwarded, or traffic which will be filtered.

CLSN lights when a collision occurs while the Switch is transmitting.

FC/FD is a multi-purpose indicator. Depending upon the operation mode of the port, this LED has different meanings. If the port is configured for:

- Standard Ethernet with Flow Control Enabled, the LED is normally off, and blinks on whenever a path buffer has become full and Flow Control is used.
- Standard Ethernet with Flow Control Disabled, the LED is normally off, and blinks on whenever a path buffer has become full and a packet could not be queued for forwarding.
- Full Duplex, the LED is normally ON, and blinks off whenever a path buffer has become full and a packet could not be queued for forwarding.

4.4.2 Management Status LEDs

The SNMP LED is normally ON when the SNMP daughterboard is present.

The MGMT ACT LED is normally OFF. It blinks ON at regular intervals when the SNMP board is installed.

The POWER LED indicates the status of the power supply. It may remain off for a few seconds during the power-on self-test. It is normally ON.

The FAULT LED indicates that the system has detected a problem. It may remain on for a few seconds during the power-on self-test. If this indicator blinks or remains lit, contact your supplier.

4.5 Cabling Requirements

The Switch's Ethernet interfaces are of the same type presented on a computer. The AUI interfaces are DTE, and the UTP interfaces are MDI. The pinouts of the connectors are as follows:

Female DB15 (DTE)

1	COL Shield
2	COL+
3	TX+
4	RX Shield
5	RX+
6	V-
7	Unused
8	Unused
9	COL-
10	TX-
11	TX Shield
12	RX-
13	V+
14	Power Shield
15	Unused

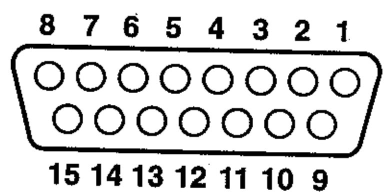

Female RJ-45 (MDI)

1	TD+ (Orange/White)
2	TD- (Orange)
3	RD+ (Green/White)
4	Unused (Blue)
5	Unused (Blue/White)
6.	RD- (Green)
7	Unused (Brown/White)
8	Unused (Brown)

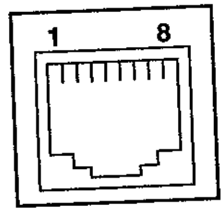

In order to connect the Switch to a hub, a standard "straight across" patch cable (UTP) is required (Category 3 or better for 10 Mbps, Category 5 for 100 Mbps). Remember that UTP cable runs can be a maximum of 100 meters. In order to connect the Switch directly to a station using UTP (in Full or Half Duplex), a "crossover" cable is required (see below). There is no AUI crossover equivalent.

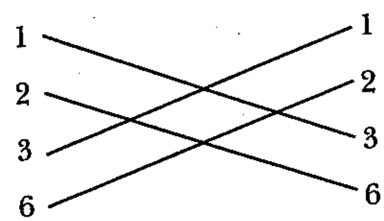

5. Troubleshooting

5.1 Troubleshooting

If your Switch is not working correctly, follow these steps (in order):

1. Make sure that the unit is plugged in to a grounded, functioning AC outlet providing between 90 and 264 VAC at 50/60 Hz. Check the power fuse and replace it if it's blown.

CAUTION

For continued protection against fire, replace the fuse with a fuse of the same type and rating.

2. Review all link LEDs to make sure that those ports you believe should be functioning are properly attached to a cable.

3. Verify that your cables are wired correctly; use a UTP crossover cable to directly connect another switch or any other DTE type-device (such as a workstation) directly to a port. Refer to **Section 4.5** for the proper connector pinouts.

4. Review all link LEDs to make sure that those ports you believe should be functioning are properly configured, and not disabled or partitioned. If the suspect ports are disabled or do not seem configured properly:

   Reconfigure the port through ONMP or the Administrative Interface. If you have problems using Automatic Port Selection, use ONMP or the Administrative Interface to manually set the port to either UTP or AUI.

5. Review all Full-Duplex LEDs to make sure that those ports you believe should be functioning are in the correct mode of operation.

6. If the Flow-Control LED shows excessive activity, refer to **Chapter 3** for discussion of how to best configure your network for operation with a Switch.

7. Make sure the equipment attached to the Switch is properly configured.

5.2 Calling Your Supplier

If you determine that your SuperSwitch is malfunctioning, *do not attempt to alter or repair the unit*. It contains no user-serviceable parts. Contact your supplier.

Before you do, make a record of the history of the problem. Your supplier will be able to provide more efficient and accurate assistance if you have a complete description, including:

- the nature and duration of the problem.
- when the problem occurs.
- the components involved in the problem.
- any particular application that, when used, appears to create the problem or make it worse.

5.3 Shipping and Packaging

If you need to transport or ship your SuperSwitch:

- Package it carefully. We recommend that you use the original container.
- If you are shipping the SuperSwitch for repair, make sure you include everything that came in the original package. Before you ship, contact your supplier to get a Return Materials Authorization (RMA) number.

Appendix A. Fast Ethernet Interfaces

The Fast Ethernet ports on the SuperSwitch are in the MDI configuration for direct connection to a Fast Ethernet hub (MDI-X) using standard Category 5 patch cables. The maximum cable length from the Fast Ethernet port to a workstation or hub is 100 meters. For direct connection to another MDI interface, a crossed patch cable is required. The cable pinout is identical to a 10BASE-T cross cable and is shown in **Section 4.5**.

A.1 LEDs

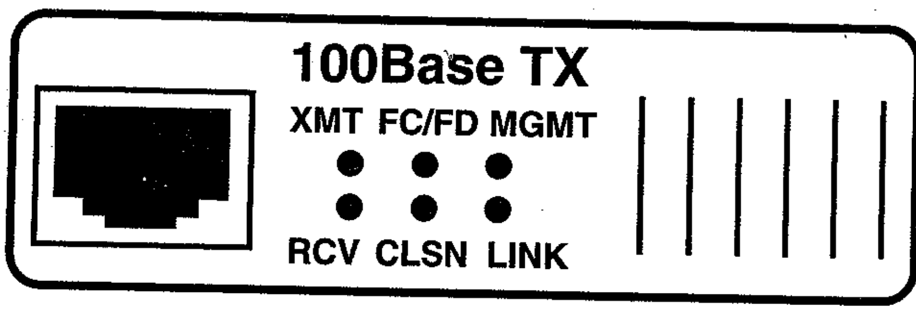

Figure A-1. LEDs.

XMT lights when the switch is transmitting packets on this port.

RCV lights when packets are received on this port, even if they are not forwarded.

FC/FD is a multipurpose indicator. If Flow Control is enabled on this port, the LED will blink whenever Flow Control is activated. If Flow Control is disabled, the LED will blink whenever a packet is lost. If Full Duplex is active, the LED will normally be on, and will blink off whenever a packet is lost.

CLSN lights whenever there was a collision while the switch was transmitting.

MGMT lights whenever the port management is accessed. It will normally blink every few seconds.

LINK lights whenever a link signal is received. This indicates when a connection has been established.

A.2 Configuring Your Network

There are two situations where plain Ethernet switching isn't the right solution. The first occurs when even a dedicated 10-Mbps link isn't fast enough. The second is when planning a network-wide upgrade path.

A bandwidth bottleneck can occur anywhere in a network where there are devices capable of continuously saturating an Ethernet. Two places that immediately come to mind are links to a file server and links between switches.

Since everyone on the network uses the file server, a large fraction of the LAN traffic must pass over the link to it. Modern file servers are typically not limited by their processor power; it is the client/server traffic on the LAN that limits the overall performance. Even a dedicated 10-Mbps link may not be fast enough. Fast Ethernet can provide a simple solution to this problem by allowing the file server to communicate at 100 Mbps to a 10/100 switch, which then distributes the traffic to several 10-Mbps ports.

Links that runs at more than 10 Mbps are essential between switches. If the link between two switches is at 10 Mbps, then only one full-speed Ethernet can be handled at a time. With the 100-Mbps connection of Fast Ethernet, 10 full-speed Ethernet streams can be handled simultaneously.

An Ethernet to Fast Ethernet switch presents the network administrator with a clear, sensible upgrade path. At first, the switch can be used simply to segment the existing Ethernet, immediately boosting performance. Next, with the addition of a single NIC, a file server can be migrated to Fast Ethernet, increasing its availability. As needed, additional file servers or individual users can be moved to Fast Ethernet, while leaving all other parts of the network running as usual. If some of the shared Ethernet segments are still congested but don't warrant a full 100 Mbps, additional switches may be added to further divide the shared segments, creating a collapsed 100-Mbps backbone and small switched 10-Mbps segments.

Fast Ethernet is by far the closest relative to Ethernet. Fast Ethernet is more than just a name: It works just like regular Ethernet, except 10 times faster. Software, except for low-level drivers (which are always card-specific), works with Fast Ethernet without modification. Frames from a device using Fast Ethernet require no translation except for the physical-layer conversion. Other networking technologies use different frame sizes and protocols than Ethernet, requiring frames passing through the converter to be fragmented or translated in at least one direction. This leads to incompatibilities between machines on different networks and between converters made by different vendors, not to mention the software incompatibilities caused by different frame formats. Translating frames also imposes an additional delay penalty, increasing the overall latency and reducing the available throughput.

The SuperSwitch combines switched 10-Mbps Ethernet and Fast Ethernet technology in a seamless integration of regular Ethernet and the IEEE 802.3u 100BASE-TX or -FX standard. In conjunction with a 100BASE-TX hub, it is an ideal base to begin migration to Fast Ethernet.

Illustrated below is a typical application of Ethernet to Fast Ethernet switching technology. In the example, a 100BASE-TX hub is used to join a Fast Ethernet workgroup to a legacy Ethernet segment, a dedicated 10-Mbps station, a remote station, and a Fast Ethernet backbone. Clients that need the most bandwidth to the server are connected directly to the same 100BASE-TX hub as the file server. Slower stations are connected via either a dedicated or shared Ethernet link. One station is connected by a long-distance full-duplex Ethernet link, although an entire segment could be supported there by another switch. The use of the SuperSwitch allows the legacy investment in cabling, hardware, and software to be used fully, while allowing for performance improvements through segmentation and the upgrade of the server only. New stations can be added to either the Fast Ethernet or regular Ethernet segments, depending on the projected need.

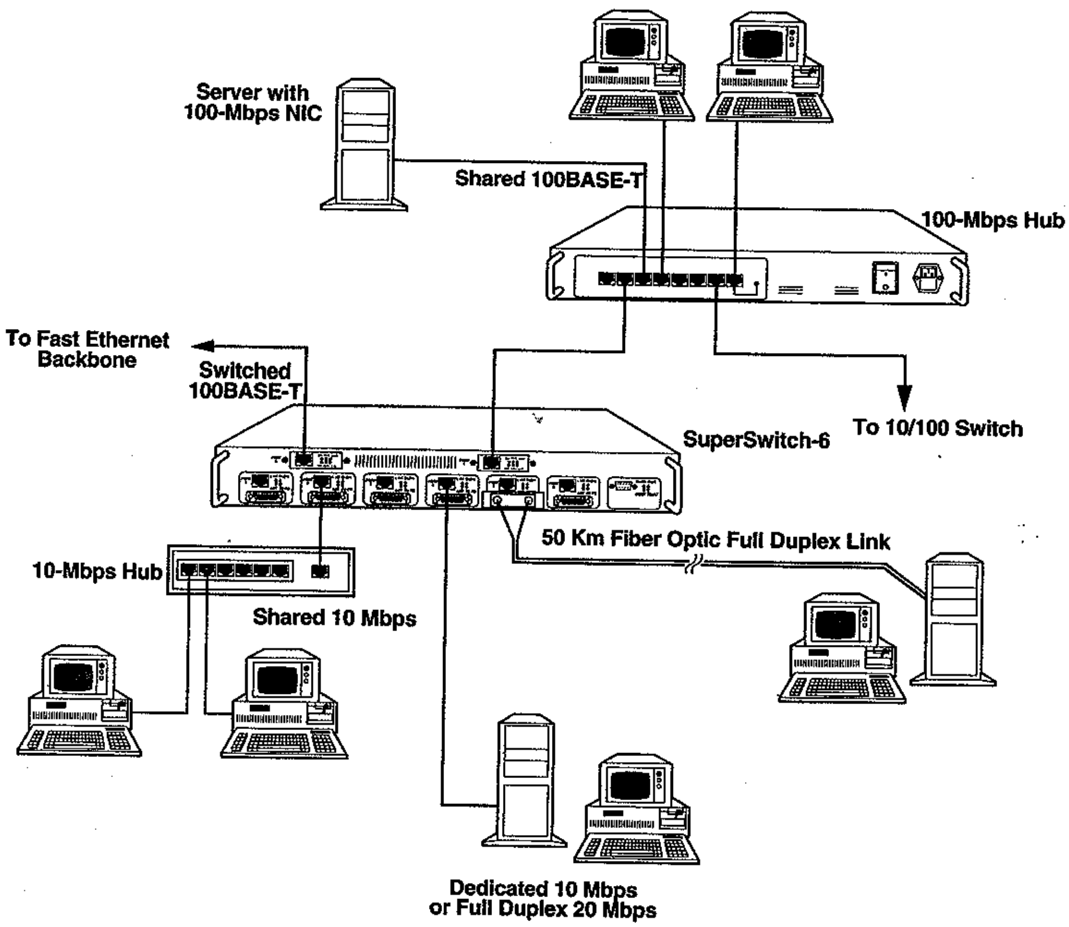

Figure A-2. Ethernet to Fast Ethernet Application.

Fiberoptic cabling is rapidly becoming a cost-effective alternative to Category 5 UTP. Its falling cost, along with greater longevity, virtual immunity to interference, greater transmission distances, and higher bandwidth capability make fiber a strong alternative for new installations. The combination of fiber optics, Fast Ethernet, and full duplex yields a compelling solution for LAN congestion. With single-mode fiber optics, a Fast Ethernet signal may be driven up to 50 km (31 miles) with no repeater. However, this is not possible in half duplex, since the maximum diameter of a Fast Ethernet network is 500 meters (1640 ft.). Full duplex overcomes this limitation by allowing simultaneous transmission and reception, eliminating collisions. With this combination, a 200-Mbps link can be established between two switches separated by over 30 miles. This should be more than sufficient for most LAN applications.