Network General Corporation Managing and Implementing Switched Networks

Network General Corporation
4200 Bohannon Drive
Menlo Park, CA 94025
Tel (415) 473-2000

Network General

Table of Contents

The information in this guide has been provided by Network General Corporation. To the best knowledge of Network General Corporation, these companies offer the types of products described. These companies are solely responsible for their software, distribution, and support services. Network General Corporation disclaims any and all liabilities for and makes not warranties, expressed or implied, with respect to these products, including, without limitation, the implied warranties of merchantiability and fitness for a particular purpose. Distribution of these products, or information concerning these products, does not constitute Network General Corporation's endorsement of the products, the companies, or support services. Product Information is subject to change without notice.

Introduction

A study by Infonetics suggests that the number of internetwork devices on the average corporate network will grow by 140% between 1992 and 1996. During that same period, the number of network nodes will grow by 150%. This, combined with the current popularity of bandwidth critical applications such as multimedia and graphical applications, rapidly increases the requirement for large amounts of bandwidth and high rates of throughput.

Support for the heightened necessity for bandwidth requires rethinking the way we design our basic network infrastructure. Dividing and segmenting networks is successful to a point, but the constant introduction of additional internetworking devices between workstations and their resources can impact performance.

One industry response to this crisis is the introduction of switching technologies into the network infrastructure. Switching technology promises to provide dedicated bandwidth to the individual resources and devices that need it. Embedding this technology into concentrator and hub products, and offering it at an attractive price-per-port could potentially help ease the migration of today's networks to support tomorrow's applications.

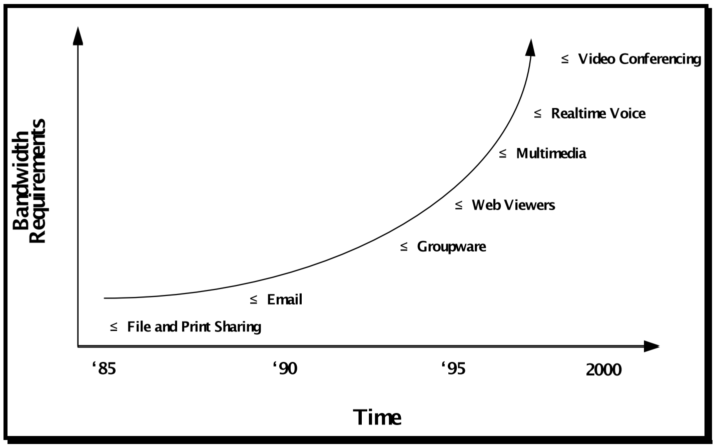

Figure 1: Applications are demanding greater bandwidth

Purpose

This document is targeted at organizations that are interested in switch-based networks. It will provide basic information meant to assist the network manager who is either implementing a switched network, or considering the benefits of switched networks.

This document will describe the basic principles of switched networks, and define some of the technologies associated with switched networks. It will compare some implementation strategies in order to determine the best way to include switched network products into a network environment.

Finally, it will describe how to effectively include network monitoring, management, and analysis tools into the network fabric, once switched products have been introduced.

Switch Design and Implementation Issues

Introduction to Switching Architecture

Traditional Local Area Network (LAN) technologies commonly used shared media to allow internetwork devices to communicate. In such a scenario, each device in a domain shares the same basic resources to communicate with each other. As illustrated in *Figure 2*, when a device communicates on shared media, the shared media hub propagates its signal to each other connected device. Each device in a shared environment is expected to monitor the network to determine when it is legal for it to send its message. The available bandwidth is then divided by all active stations in the domain.

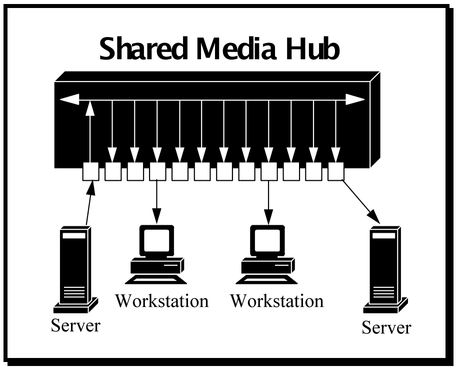

Figure 2: Shared Media hubs propagate traffic to all ports

This technique is analogous to CB radio. Each user must listen to the CB network to determine if any other user is communicating and talk only when other users have completed their message.

The switched networking environment attempts to provide equal and unrestricted network bandwidth access to each device on the network. When a device transmits data across a switched network, the switched hub determines precisely which devices require this information and where to forward it. Each port on a switched hub replicates packets destined only for the device connected to this port, or broadcast and multicast packets destined for all devices. This technique comes close to providing dedicated bandwidth for each device.

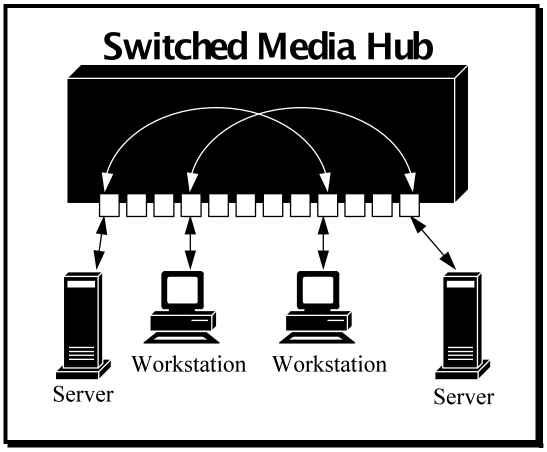

Figure 3: Switched Media Hubs pass traffic directly to the designated port

Unlike the CB radio analogy, switched networks are more closely analogous to placing a telephone call. Devices need not be concerned with which other devices are communicating on the telephone network.

Another analogy would be that shared media hubs are similar in function to multi-port repeaters. They regenerate incoming signals on all ports, without determining which ports actually require the data. A switched media hub is similar to a multi-port bridge. As seen in *Figure 3*, the switch determines where to send incoming data according to the destination hardware address of the packets.

The overall benefit of this technology is that multiple conversations can occur simultaneously on a single switched hub, providing the user with almost dedicated bandwidth.

Switched Networking

Switched Networking (such as switched Ethernet or switched token ring) provides a simple solution to existing networks suffering from traffic congestion. Strategic placement of Switched hubs or locating shared network resources (such as file servers, printers, and routers) on Switched Media hubs can dramatically reduce bandwidth utilization on the network segment. Switched networking provides the simplest upgrade path for an existing shared media network.

There are several differentiating features between Ethernet switch models and manufacturers. Each feature can impact performance, functionality and implementation strategy. The following is a list of some of the more common features:

Forwarding Technique: Does the LAN Switch use Cut-through or Store-and-Forward technology?
Latency: How much delay does the Switch introduce?
Management: How much control does the Switch provide to the user?
Single- vs. Multi-MAC: Does the LAN Switch associate a single address or multiple hardware addresses with each port?
Virtual LAN (VLAN) Support: Does the Switch support logical grouping of devices into separate collision domains?
Spanning Tree: Does the switch support the spanning tree algorithm, or another technique that detects and eliminates topology loops?
Full Duplex: Does the switch allow ports to send and receive data simultaneously?

Forwarding Techniques: An Ethernet Switch manufacturer usually employs one of two techniques to allow the switch to make forwarding decisions about incoming data. Each technique has unique benefits and disadvantages.

Cut-through: This technique requires the switch to hold the packet until the switch receives the destination address (approximately six bytes into the frame). At that point, the switch has enough information to make a forwarding decision. The benefit of this feature is that the switch can forward packets at a higher rate, reducing latency, and increasing overall throughput. The drawback to this technology is that the switch will start to forward a packet before it has seen enough of the packet to determine if there is an error in the frame. The switch could potentially propagate error frames, or even congest the network with “garbage” that the switch may misinterpret as a broadcast.

Store-and-Forward: This technique requires the switch to accept the entire frame before making the forwarding decision. The benefit of this feature is the switch can also determine whether there is an error in the packet before making a forwarding decision. Some switch manufacturers are also introducing collision-avoidance techniques to minimize the propagated network errors. The drawback *to Store-and-Forward* is that the switch must wait much longer before it is allowed to make a forwarding decision. This can increase latency, delay, and impact overall network throughput.

Latency: Switch latency is the measurement of time between the point at which the switch started receiving data, and the point at which it propagates that data to the destination port. There are many factors that can affect latency, such as the forwarding technique that the switch uses.

Switched Ethernet hubs that employ the *Cut-through* method often have a fixed latency delay. A cut-through switch will always make a forwarding decision after seeing the destination address of the packet, regardless of the overall length of the packet. The latency of a *store-and-forward* switch is dependent on the packet size, as the switch must receive the entire packet before it begins to forward it.

Management. Switch management is how much control the user has over the switch. This can range from simple SNMP MIB I/II statistics, to viewing active ports on a hub, to altering the address tables and forwarding information.

Many switch vendors offer a "Monitoring Port" to allow a network analysis device to be connected to the switch, although their implementation techniques vary from vendor to vendor. There are three common implementation techniques that ease the introduction of network analysis devices:

Port Tap: Sending all traffic to or from one port on a switched hub to a designated monitoring port on the hub.

Circuit Tap: Sending all traffic exchanged between two ports on the switch to a designated monitoring port on the hub.

Switch Tap: Sending all traffic that occurs on any port on the switch to a designated monitoring port on the switch.

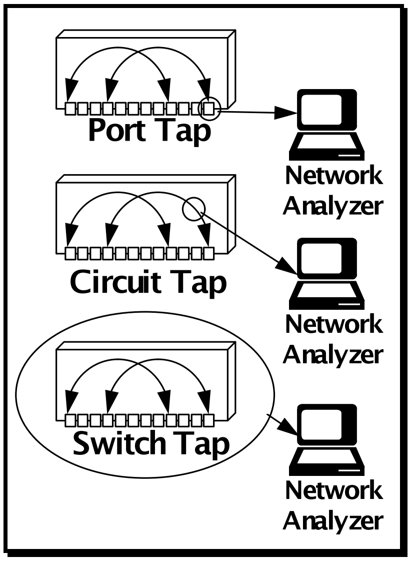

Figure 4: Integrated switch monitoring ports

Each of these techniques simplifies network analysis, and allows monitoring of a hub or connected device without the necessity of additional hardware. Many switched hubs implement one or more of these techniques to assist network analysis.

The port tap and circuit tap techniques limit the amount of traffic that is delivered to the network analyzer. This feature could filter important information from the analysis device if not used correctly. For example, an administrator may use the circuit tap feature to focus on a problem connection between a user and a resource. The source of the problem may actually be another user accessing the shared resource. Focusing on the connection between the first user and the resource may not provide enough detail to identify the problem.

At the other extreme, the switch tap feature does not filter any information. This feature may work well in circumstances where overall use of the switch is low. However, if several of the switch users are demanding high amounts of bandwidth individually, their combined traffic may be greater than the switch can effectively process through a single monitoring port.

Single- vs. Multi-MAC. Single-MAC switches associate a single hardware address with a port on the switch. Multi-MAC switches can associate multiple hardware addresses on each port. Understanding the address limitations of the selected switch is crucial to useful implementation on the network.

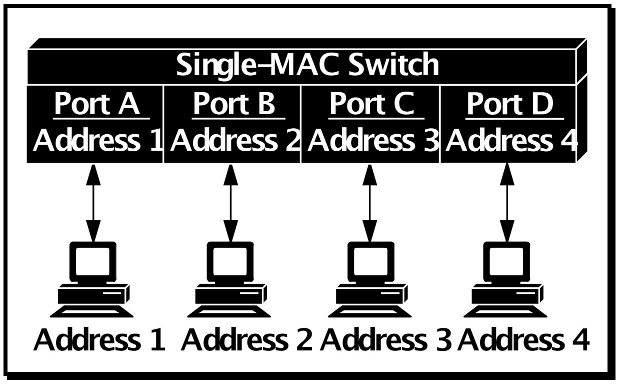

Figure 5: Single-MAC Switches support one device-per-port

Single-MAC switched hubs are primarily designed to connect directly to either an end-user, a shared resource (such as a server or a group of servers), or to interconnect multiple routers that do not implement bridging. They were not designed to inter-connect shared media hubs, or network segments containing multiple devices. *Figure 5* illustrates the single device to single port relationship of the Single-MAC switch.

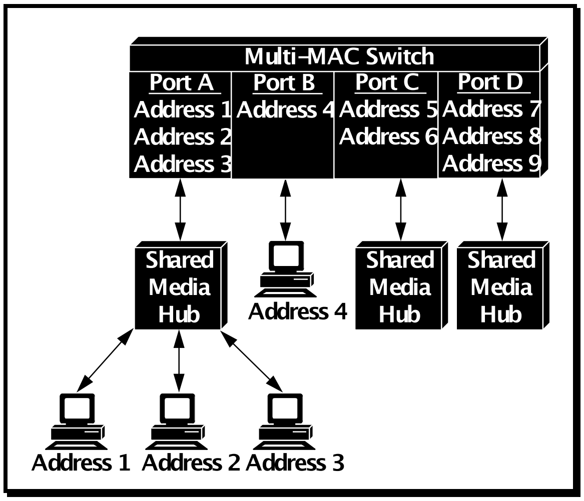

Figure 6: Multi-MAC Switches support multiple devices-per-port

Multi-MAC Ethernet Switches have enough memory to associate multiple hardware addresses with a single port. These Ethernet Switches can be implemented as a logical "hub of hubs", or in backbone architectures. The total number of hardware addresses that the switch can buffer differs from vendor to vendor. It is important to be aware of the limitation of the switch to ensure that the switch is not exposed to more hardware addresses than it is capable of buffering. Once the number of learned addresses exceeds the available address space, the switch may propagate packets with new addresses to all ports on the switch or even drop the frames altogether. *Figure 6* demonstrates how Multi-MAC switches can associate several addresses with each port.

Virtual LAN (VLAN) Support: Virtual LANs (VLANs) allow the administrator to create a flexible, dynamic logical network within a fixed physical network infrastructure. This allows these devices to share common resources, such as bandwidth. *Figure 7* shows several users connected to separate hubs logically grouped into the same virtual LAN.

This also allows the user to create logical subnetworks of different traffic types, isolating certain traffic types from others. For example, the administrator can place all management traffic (such as SNMP) on a separate VLAN so that the amount of management traffic will not interfere with normal operations.

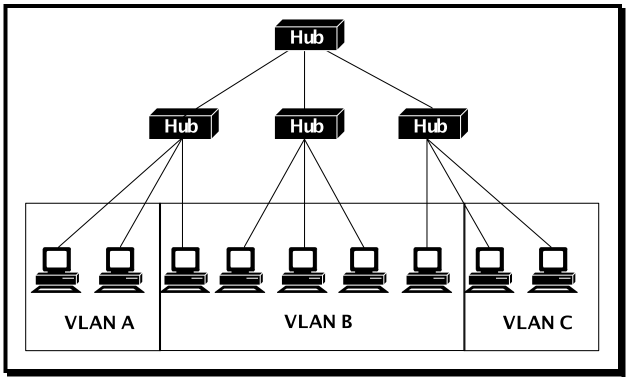

Figure 7: Virtual LANs (VLANs) allow users to divide devices into separate logical workgroups

Spanning Tree: Just as shared media hubs operate similarly to multiport repeaters, Switched Ethernet hubs operate similarly to multiport transparent bridges. They are also susceptible to the same types of problems as bridged networks. This includes topology loops.

Topology loops occur when traffic is transmitted on one segment, bridged to another segment, and returned to the originating segment via a different route. The Spanning Tree algorithm and protocol eliminate this occurrence by allowing the bridged devices on a network to be aware of each other. As seen in *Figure 8,* the switch automatically disables one of the ports that would complete this loop until the topology loop is no longer present. Some switch manufacturers do not implement the Spanning Tree algorithm, but do implement a proprietary solution to prevent network topology loops.

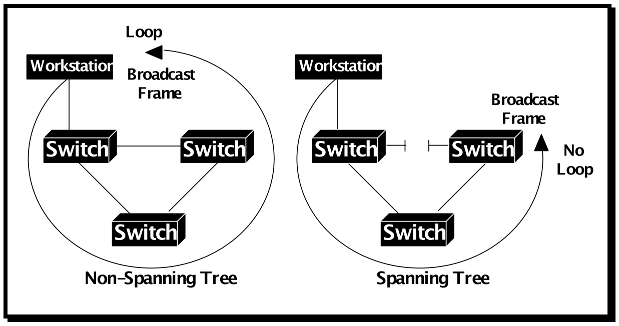

Figure 8: The Spanning Tree Algorithm prevents loops in network topology

Topology loops may be introduced into a network design by accident or intentionally to provide a redundant backup data path. It is often useful to take advantage of redundant links in mission critical applications. If the switched hub does not participate in the Spanning Tree process, proper placement of the switch should be considered carefully.

Full Duplex: One technology that is complementary to switching is full duplex communication. A switch that maintains full duplex ports has the capability to simultaneously transmit and receive data across these ports. To take advantage of this feature, both the switch's port and the connected device must support the full duplex technology.

There are several benefits of this technology:

- *Throughput:* If the switched media hub and the connected device can send data to each other simultaneously, theoretically throughput can double.
- *Collision avoidance:* If the switched media hub and the device can communicate simultaneously without restriction, collisions do not occur.
- *Improved Distance Limitations:* In some circumstances, distance limitations are imposed on some media types because of the potential for collisions (for example: Ethernet over Fiber Optic cable). If collisions can not occur, distance is only limited by media restrictions.

High-Speed Integration - 10/100 Switches

In a small workgroup or domain, connecting all users and servers to Ethernet Switches can definitely improve performance. Many networks cannot take advantage of this technology on their own due to their geographic layout. If a network primarily had workstations distributed on various hubs and located all server resources on a backbone segment, replacing the users normal shared media hub with the same speed switched hub would not provide any benefit. The single connection between the switched hub and the backbone would prohibit more than one user from connecting to the backbone at one time. This eliminates any benefit that the switch might provide.

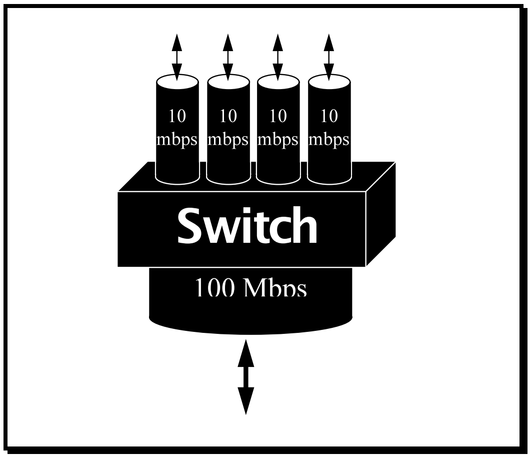

Figure 9: 10/100 switches

One solution to this potential problem is to provide a greater bandwidth "pipe" between the switched hub and the backbone segment. Providing a dedicated 10Mbps connection from each user to the switched media hub, and a high-speed connection from the switched media hub to the backbone improves or eliminates the previous bottleneck between the user and the backbone network (as shown in *Figure 9*).

Common Ethernet switch implementations often provide one or more high speed ports to provide connectivity between the switched media hub and the backbone or other resources. The most popular high speed media include FDDI, Fast Ethernet, and ATM. Each media type has its own benefits.

- FDDI has been available for the longest period of time and is widely supported.
- Fast Ethernet is extremely similar in format to standard Ethernet, and requires little translation by the switch between the two topologies. This can reduce latency between the normal switched ports and the high speed port. Fast Ethernet is subject to the same problems in utilization scaleability as regular Ethernet.
- ATM is, by definition, a switched topology. ATM also offers the potential for greater speeds than either FDDI or Fast Ethernet.

Optimizing Switched Network Design

When considering network design, it is often helpful to draw parallels between networks and highways. Bandwidth is analogous to the size of the highway. Bandwidth does not define the speed with which data is transferred between two devices, but rather the amount of data that can be transferred between two devices. In the same way, the size of the highway does not define the speed at which cars travel, but rather how many cars can travel between two points. In both circumstances, high volumes of traffic can congest the paths, and slow throughput.

To expand on this analogy, most devices that make intelligent data-forwarding decisions on a network (bridges, routers, switches, etc.) are similar to traffic lights. Do traffic lights improve the rate in which cars move through road systems? It depends on the specific circumstances. On interstate roads, traffic lights may decrease the throughput of the road. In downtown city areas, traffic lights are the only device standing between order and chaos.

In the same way, data-forwarding devices such as switches can dramatically improve or detract from the performance of your network, depending on how and where they are implemented.

The optimal, however unrealistic, design for a data network would provide each device with a dedicated high speed connection directly to each service with which the device needed to communicate. The further away from this model the design gets, the worse the network performance will become.

As seen in *Figure 10*, Network B has no performance advantage over Network A. In this simple example, all workstations communicate back to the same, single server. The “bottle-neck” in this example is the single connection from the server to the switch. In Network A, all users were forced to share the network to gain access to the server. In Network B, all users were forced to share the server’s connection to the switch. Only two things have changed between Network A and Network B: the point on the network being shared, and the fact that the switch is now inserting a degree of latency in each transmitted packet.

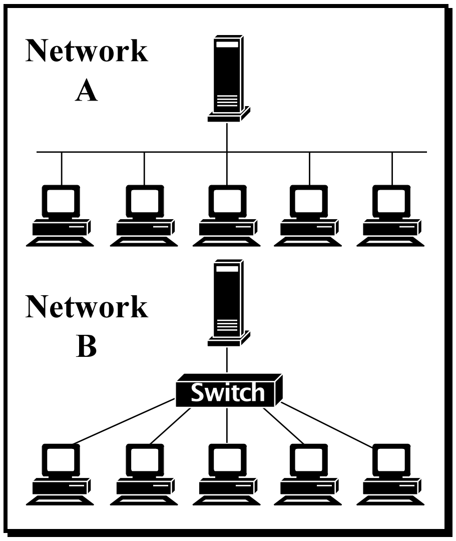

Figure 10: Despite the switch, Network B has no performance advantage over Network A.

Most networks experience a majority of their traffic between workstation and shared resources (such as, a file server or router). Workstations rarely communicate directly to one another, with the possible exception of peer-to-peer networks. If a user sends a mail message to another user on a network, the message is sent from the workstation to the file server. The message is then stored on the file server until it is retrieved by the recipient. The only benefit that Network B may provide is to improve performance from workstation to workstation. Depending on the type of network, this type of traffic may never occur.

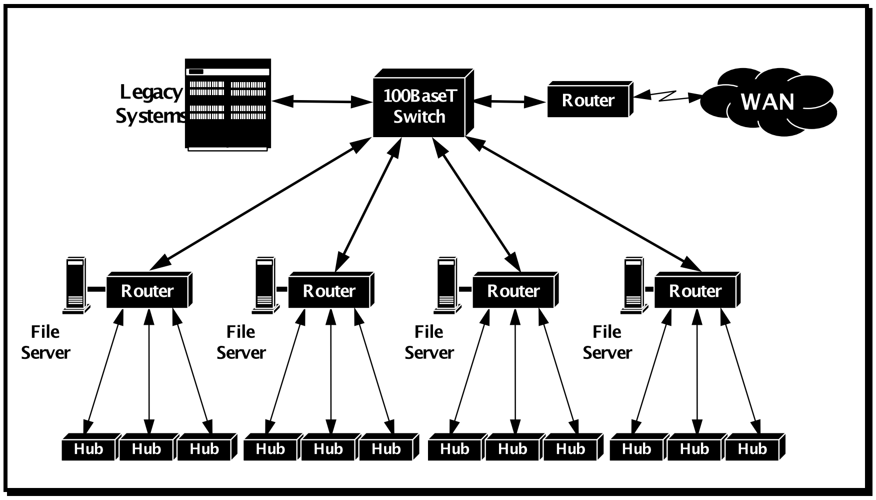

Figure 11: Using a high speed switch as a collapsed backbone to interconnect routers

There are several changes that would make Network B perform better. If the server were connected to the switch using a higher speed connection, such as Fast Ethernet, FDDI or ATM, the bottle-neck could be considerably reduced or eliminated. If the server maintained multiple connections to the switched hub, this would allow several simultaneous messages to be received by the server.

The end result is that wholesale replacement of existing hardware with new Switched technology may not yield a performance improvement. Taking advantage of this technology relies on careful consideration of implementation.

Switch manufacturers design switches with different strengths. Some are better suited for specific tasks than other switches. Some switch designs lend themselves better to a high-speed "backbone-in-a-box" type of use, whereas others may be designed with the intent of micro-segmentation.

Figure 11 illustrates the use of a high speed switch as a collapsed backbone. In this scenario, each router is used to isolate segment traffic from other segments. Connecting all of the routers to a *shared* medium, even if it is a high speed medium, could cause congestion as bandwidth demands of the backbone increase. Directly connecting the routers to each other would be effective as well, but potentially more expensive and not as scaleable.

Provided that the routers are only routing traffic and not bridging, the backbone switch in this diagram may be a "single-MAC" device. If the routers are bridging traffic, the backbone switch must be "Multi-MAC" and have an address table large enough to buffer the addresses of all devices that may propagate bridged traffic. If the switch is a "single-MAC" device, it should only be connected to either the workgroup routers, or resources that are shared by the entire enterprise network, such as external routers or servers.

Figure 12 illustrates the use of switches for the purpose of workgroup segmentation. This example uses switches as a "hub-of-hubs", micro-segmenting each hub into its own collision domain. Additionally, it provides high speed access to shared resources. The design assumes that all traffic on this network is sent either to or from the file servers on the backbone.

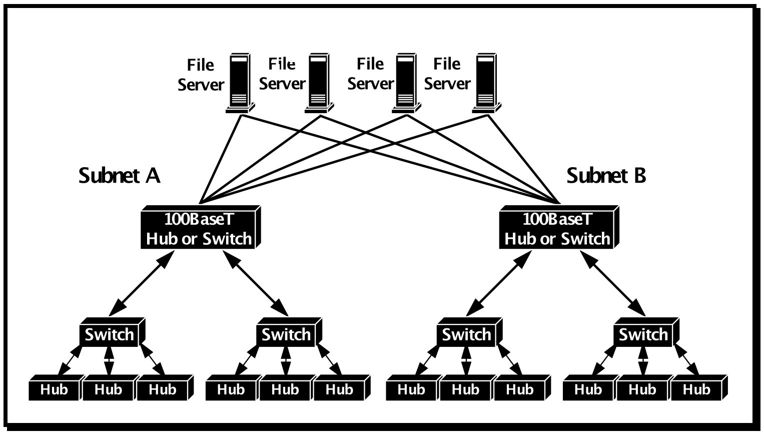

Figure 12: Using a switch to segment workgroups and provide high bandwidth pipes to shared resources

The depicted network is actually two separate, flat networks. The network has no dedicated routers, although the servers could potentially act as routers.

The switches in this diagram must be "multi-MAC", with the ability to associate multiple hardware addresses with individual ports. It is important for the switch to have an address table large enough to track each hardware address of every device attached to the hubs which are connected to it. This includes the workstations connected to the shared media hubs, as well as the servers connected to the high speed switches. The benefits of this design include redundancy (two separately operating backbones) and throughput.

Figure 13 illustrates the use of switching in small, high capacity workgroups. This design is optimal in environments where peer-to-peer networks are present, and high bandwidth capacity is required between devices.

This design provides dedicated bandwidth between each individual device, as well as a high capacity path to shared resources. The peer switches in this design can be "single-MAC" switches, as each device on the network has a dedicated connection to the switch. The backbone switch must have "multi-MAC" capability.

Each of these designs have strategic benefits that are dependent on the network traffic and layout. The first step in determining how to migrate a network to such new technologies is to fully understand the unique characteristics of the network. The key to a successful implementation plan is understanding which devices are using which servers or other shared devices, and understanding the volume of traffic between them. Effective monitoring and analysis of the network *before making design changes is essential*.

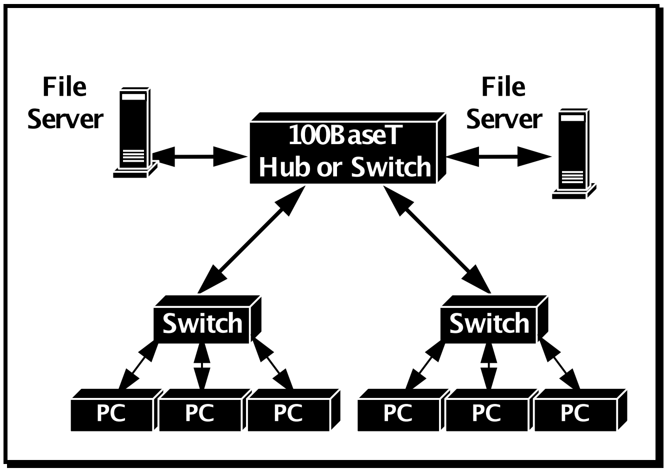

Figure 13: Example of micro-segmenting with workgroup switches

Network traffic patterns can be chaotic systems, much like weather patterns. Long-term accurate prediction of changes can be difficult. Even after a switched network has been successfully implemented, monitoring and analyzing the network must be an ongoing function to ensure that the solutions that were initially introduced to solve congestion problems continue to be effective.

Managing Switched Environments with the Distributed Sniffer System and RMON

Connecting a Network Analyzer or RMON probe (such as a Sniffer Network Analyzer, Distributed Sniffer Server (DSS), or Foundation Probe) to a *switched* hub requires more planning than connecting it in a standard *shared* media environment. As shown in *Figure 14*, if a Sniffer Network Analyzer were connected to a switched hub in the same way a user would be connected to a shared media hub, the user would not see most of the traffic. Switched hubs, by design, only forward packets to the specific port that it needs to. This would limit the Sniffer Network Analyzer to capturing broadcast packets, multicast packets or packets in which the switch has not been able to resolve a hardware address.

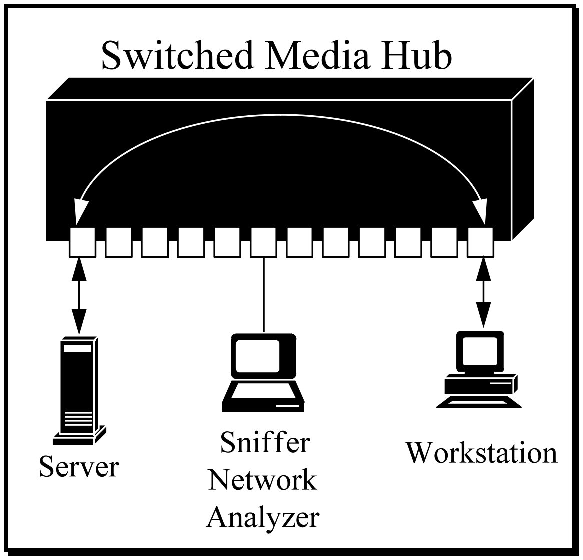

Figure 14: Switched hubs filter traffic from the directly connected Sniffer Network Analyzer

There are several techniques in effectively connecting a Sniffer Network Analyzer to a switched hub. Some of these techniques are specific to a particular hub vendor, using techniques such as *port tapping*, *circuit tapping*, or *switch tapping*. Vendor-specific techniques will be described in detail in the following section.

There are techniques for monitoring a switched hub even if the vendor has not implemented any monitoring technology. Analyzing a switched hub requires the use of an additional shared media hub with a crossover cable, or a hub with a crossover port. There are many small, portable hubs that are ideal for this operation, such as the Cabletron MR9T or the 4-port Pocket Hub by Transition Engineering.

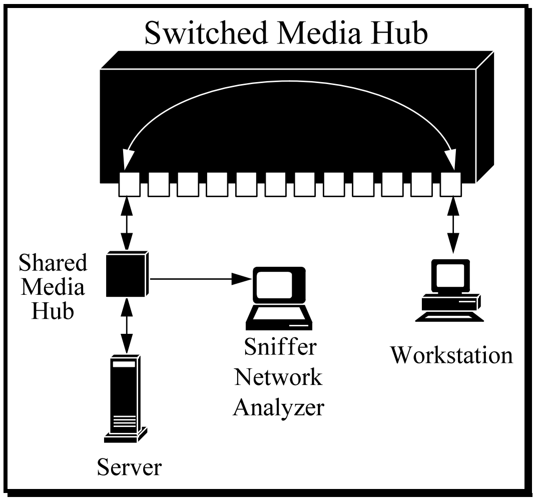

Figure 15: Inserting a shared media hub between the switch and server provides the Sniffer with an access point

The first step in monitoring a switched network involves strategically locating the analysis device. In many network environments, most traffic flows either to or from a shared resource, such as a file server or a router. Placing a network analyzer strategically between the switched hub and the shared resource will allow you to trace all traffic to and from that shared resource. As illustrated in *Figure 15*, placing a small shared media hub between the switched hub and the shared resource provides the Sniffer Network Analyzer with an access point in which to monitor the network without impairing the performance benefits of the switched hub.

Focusing on strategic areas will aid in the network analysis procedure. For example, an administrator wishes to diagnose a problem on a connection between a user workstation and a file server. Connecting the Sniffer Network Analyzer to the shared media hub between the switch and the file server would allow him to determine if the user's data was reaching its destination, and how the server responded.

A distributed monitoring solution would work similarly. Both the Distributed Sniffer System (DSS) Servers and Foundation Probes (NGC's RMON agents) have the ability to maintain separate monitoring and communications ports. This means that the analysis device does not need to communicate with the management console using the same connection to the network that is used for monitoring. The benefit of this architecture is that the distributed analysis tool can be attached transparently to the connection between the shared resource and the hub without creating traffic on the link, or interfering in any other way with the communications of the shared resource. This also allows DSS servers or Foundation Probes to be used in environments that use "single-MAC" switched hubs. Single-MAC switched hubs would not operate correctly if a monitoring device attempted to communicate on the same segment that another device was connected to.

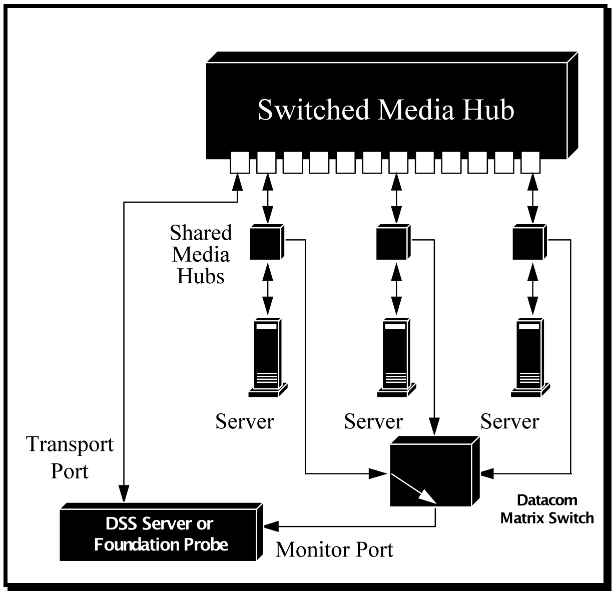

Figure 16: Using a port switch can leverage your DSS server investment across multiple connections

Dedicating a DSS server to each shared resource may not be possible due to budgetary or other availability constraints. Maintaining the ability to capture data on all connections is still needed. As depicted in *Figure 16*, a single Sniffer Server can be connected with as many as eight separate shared media hubs, using a matrix switch (not to be mistaken for a switched hub or packet switch) made by Datacom. DSS software allows the server to remotely control the switch in order to connect the monitoring port of the server to the desired shared media unit to troubleshoot the desired connection. This also allows the user to leverage the DSS server investment across eight separate segments.

The use of additional equipment, such as a shared media hub or matrix switch may not be needed if the switched hub supports some method of port monitoring (I.E. *Port Tap*, *Circuit Tap*, or *Switch Tap*). Port monitoring allows the remote monitoring or analysis device to be connected directly to the switched hub. These switched hubs can often be configured to direct traffic to the monitoring or analysis device using either their proprietary management software, an SNMP umbrella management console (such as SunNet Manager or HP OpenView), TELNET, or even an "out-of-band" serial connection.

In *Figure 17*, the Sniffer Server or Foundation Probe is connected directly to the switched hub. The switched hub in this diagram supports *port tapping*. The switch has been configured to use port #4 as the tap port. The tap port is designated to receive any traffic going to or from port #1. The file server in this diagram is connected to port #1. This allows the Sniffer Server or the Foundation Probe to receive any traffic going to or from the file server.

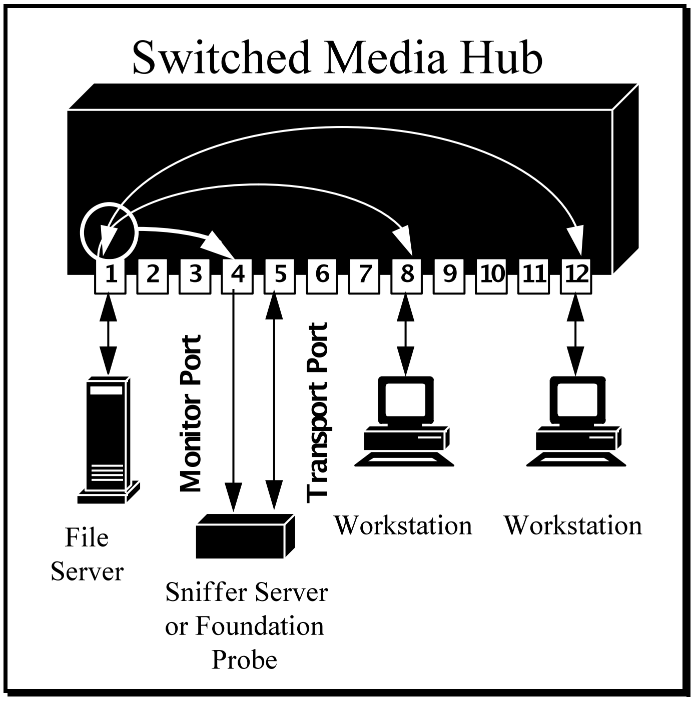

Figure 17: Sniffer Server or Foundation Probe monitoring a specific connection between a workstation and server

Note that the Sniffer Server is also connected to port #5. This allows the Sniffer Server's transport connection to communicate to the central console without interfering with the file servers communication. If the monitoring port (port #4) is ever reconfigured to monitor a different port on the switch, the transport connection does not need to be changed from port #5.

Network Management Traffic

As network fabric increases in complexity, and management traffic becomes a greater percentage of total traffic, some network designs include support for a dedicated network management backbone. Using VLAN technology, SNMP and other management traffic can be isolated from production network traffic. This ensures that, no matter what quantity of management traffic is created, it will not impact the production network. No matter what problems may occur on the production network, management traffic can be available.

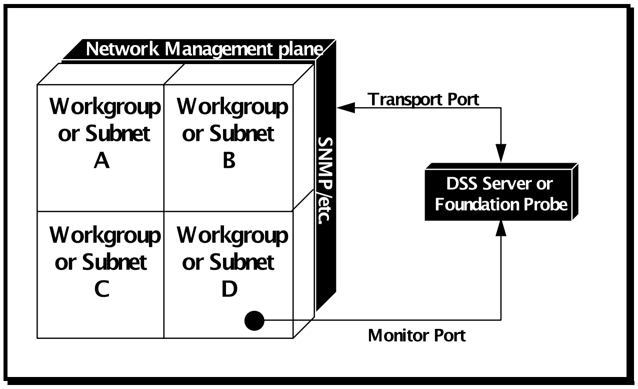

Figure 18: Logical diagram of Virtual LAN for Network Management traffic.

Figure 18 shows a logical diagram of a network that is segmented into four separate VLANs for users, and one VLAN exclusively for network management. The DSS Server or Foundation Probe's monitoring port is logically associated with Workgroup or Subnet D. This allows the monitoring device to track any data transmitted across this workgroup or subnet. If, at some other time, another segment needs to be monitored, the user can associate the monitor port with any other workgroup using the hub's network management software.

Vendor Specific Solutions

The following pages describe how to integrate Sniffer Network Analyzers, Sniffer Servers, or Foundation Probes into selected vendors switched hub products.

Alantec

70 Plumeria Drive
San Jose, CA 95134-2134
(408) 955 9500

Product:
- PowerHub 3000 Series
- PowerHub 5000 Series
- PowerHub 6000 Series
- PowerHub 7000 Series

Monitoring capabilities
- Port Tap
- Multiple Port Tap
- Circuit Tap

Management
- In-Band SNMP, TELNET
- Out-of-Band Serial

Management Strategy

The Powerhub supports Port Tapping, Circuit Tapping, and Switch Tapping (by repeating the process of Port Tapping for each port on the switch). The Switch can be configured through SNMP, TELNET, or a serial connection. Telnet may be the most convenient to use from the SniffMaster Console or Foundation Manager. The following is the Syntax for the commands to administrate the port monitoring functions:

Syntax

Port Tap / Switch Tap Syntax:

```
port-monitor|pm view <options>|all
<monitored port> on <monitoring-port>
[r]
```

`<options>|All`	Specifies which packets are to be monitored
`incoming|i`	Monitors all traffic received by the port
`forwarded|f`	Monitors all traffic forwarded by the Powerhub
`generated|g`	Monitors all traffic generated on the ports by the Powerhub.

Example

port-monitor view all 1 on 4

This sends all traffic from or to port #1 to port #4 (the port that the monitoring device is connected to).

Circuit Tap Syntax

```
port-monitor|pm viewpair <port1>
<port2> on <monitoring-port> [r]
```

port-monitor viewpair 1 8 on 4

This sends all traffic between port #1 and port #8 to port #4 (the port that the monitoring device is connected to).

Port Tap

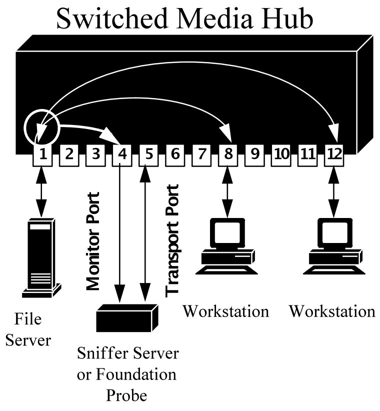

Circuit Tap

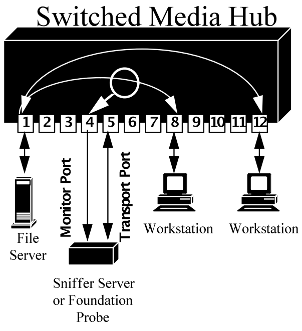

Bay Networks	4401 Great America Parkway Santa Clara, CA 95052-8185 (408) 988-2400
Products Bay (SynOptics) Model 3000 w/3328 host module Bay (SynOptics) Model 5000 Bay (SynOptics) Model 28000 Bay (Centillion) TR Switch	***Monitoring Capabilities*** Centillion TR - *Circuit Tap* Model 3000 w/3328 Module - *Hardware Based* LattiSwitch Model 28000 - ***Management*** In-Band SNMP Out-of-Band Serial
Product	***Management Strategy***
Model 3000 w/3328 Host Module or Model 28000 LattisSwitch hub 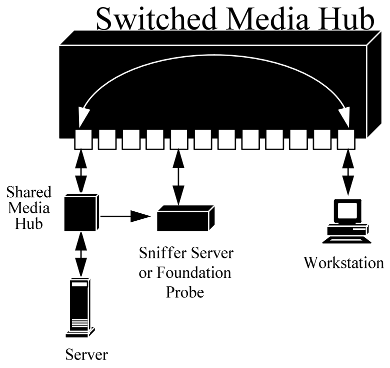 Switched Media Hub Shared Media Hub Sniffer Server or Foundation Probe Workstation Server	The model 3000 hub supports internal switching with the 3328 Host Module. The 28000 is a stand-alone switched hub. The 3328 and the 28000 do not support internal port monitoring, so an external shared media hub is needed to monitor a port. Attaching a small shared media hub on the connection between the shared resource and the switched media hub allows the analysis device to passively monitor the connection. The administrator should consider inserting a shared media hub on each connection that he believes that may need to be analyzed. This preventative measure may eliminate the need to break the connection to the shared resource during periods of normal operation to insert the shared media hub.
Centillion 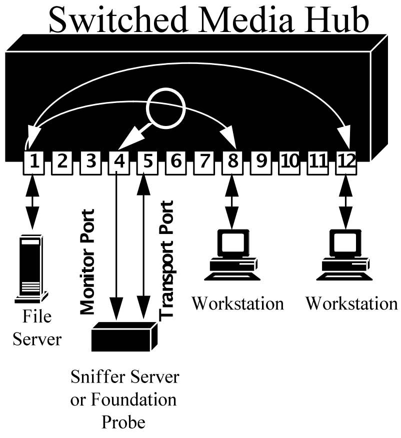 Switched Media Hub 1 2 3 4 5 6 7 8 9 10 11 12 Monitor Port Transport Port File Server Workstation Workstation Sniffer Server or Foundation Probe	The Centillion Token Ring Switch Allows Circuit tapping.

Cisco/Catalyst

170 West Tasman Drive
San Jose, CA 95134-5504
(408) 526-4000

Products

Catalyst 5000

Monitoring Capabilities

Port Tap
Switched Port Analyzer "SPAN"

Management

In Band - SNMP, Telnet, RMON
Out-of-Band Serial

Management Strategy

The Catalyst Switch embedded management provides a feature called Switched Port ANalyzer (SPAN), as well as an implementation of the RMON MIB. The SPAN feature allows the administrator to direct traffic that is going to and/or from a port on the switch to a designated monitor port. This feature is accessible through either SNMP and their own management software, Telnet, or through an out-of-band serial connection.

Using the feature through Telnet or out-of-band serial requires establishing a connection with the switch and issuing the command from the command prompt. the following is a definition of the syntax for the SPAN feature.

SPAN Syntax

```
set span [<source_port> <destination_port>
[rx|tx|both]] [enabled|disabled]
• source_port - The port number to be
  monitored
• desitnation_port - The port number where
  the tapped data is to be redirected
• [rx|tx|both] - Receive, transmit, or
  both.  The default is both.
• enabled|disabled - Enable or Disable
  SPAN.  The default is Enabled
```

Examples:

```
>set span 1 4 both enabled
```

This sends all traffic that is going to or from port #1 to port #4 (the monitoring port).

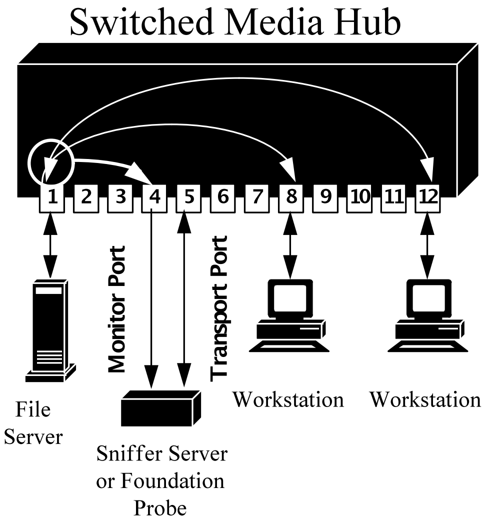

CrossComm	450 Donald Lynch Boulevard Marlborough, MA 01752 (800) 388-1200 (508) 481-4060
Products ClearPath XL 10 ClearPath XL 20 ClearPath XL 80	***Monitoring Capabilities*** Hardware Based ***Management*** In Band - SNMP, RMON
ClearPath XL 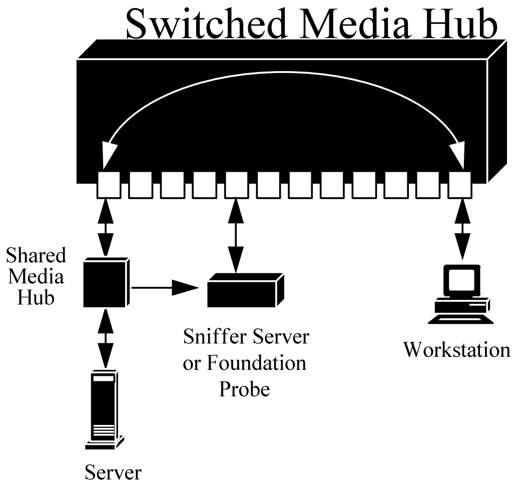 Switched Media Hub Shared Media Hub Sniffer Server or Foundation Probe Workstation Server	***Management Strategy*** The ClearPath XL series is the name of the CrossComm switched hub product line. The ClearPath XL series does not support internal port monitoring, so an external shared media hub is needed to monitor a port. Attaching a small shared media hub on the connection between the shared resource and the switched media hub allows the analysis device to passively monitor the connection. The administrator should consider inserting a shared media hub on each connection that may need to be analyzed. This preventative measure may eliminate the need to break the connection to the shared resource during periods of normal operation to insert the shared media hub.

Grand Junction Networks, Inc.

47281 Bayside Parkway
Fremont, CA 94538
(510) 252-0726

Products	***Monitoring capabilities***
FastSwitch 10/100	Port Tap (Monitoring Port) Multiple Port Tap ***Management*** In-Band SNMP Out-of-Band Serial

Management Strategy

The Grand Junction Fast Switch 10/100 provides the ability of sending traffic from one or more ports on the switch to a designated monitor port. This feature is available through either SNMP or out-of-band serial connection. The simpler way of accessing this feature is through a serial connection.

Once a serial connection has been established, the switch guides the user through several menus that provide access to this and other features. The following is an example of the menu system to access the port monitoring features.

Fastswitch Message or menu	***You type:***
`Enter Password:`	*password*

`FastSwitch 10/100 ES - Main Menu`
 	
`View and Modify Configuration`
`[L] Logon Password`
`[F] Firmware`
`[R] RS-232 Port`
`[S] System`
`[I] Internetwork Connection`
`[N] Network Management (SNMP)`
`[P] Ports`
`[M] Monitoring and Security`
 	
`Status and Performance`
`[A] Address Status Report`
`[D] Detailed Port Statistics`
`[E] Exceptions Report`
`[U] Utilization Report`
 	
`[H] Help`
`[X] eXit Management Console`
`Enter Selection`
	Enter M for Monitoring and Security
	
	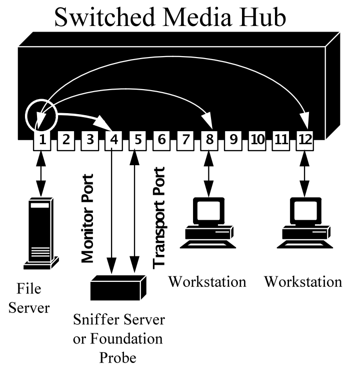
	 Switched Media Hub
	1 2 3 4 5 6 7 8 9 10 11 12
	Monitor Port
	Transport Port
	File Server
	Sniffer Server or Foundation Probe
	Workstation
	Workstation

`FastSwitch 10/100 ES - Monitoring and Security`
 	
`----Frame Capture Support----`
`[C] Capture frames to the Monitor    Disabled`
`[M] Monitor port Assignment     B1-4`
 	
`[A] Add port to capture list`
`[R} Remove port from capture list`
`[E] Exclude all ports from capture list`
`[I] Include all ports in capture list`
`Current Capture list: A, B1-4`
	C (until reads Enabled)
	
	A (to add ports to monitor)
	Select ports to monitor

`Enter Selection:`	

Kalpana, Inc

1154 East Arques Ave
Sunnyvale, CA 94086-4602
(408) 749-1600

Products

EPS-2115M EtherSwitch

Monitoring Capabilities

Port Monitoring (SwitchProbe Port)

Management

In-Band TELNET, SNMP (through SwitchVision Software)
Out-of-Band Serial

```
     SwitchProbe Configuration Menu

Exit to Previous Menu

SwitchProbe Port Number    1

Traffic to Probe           Half Duplex



          Display the Main Menu
Press <CTRL><P> to return to Main Menu
```

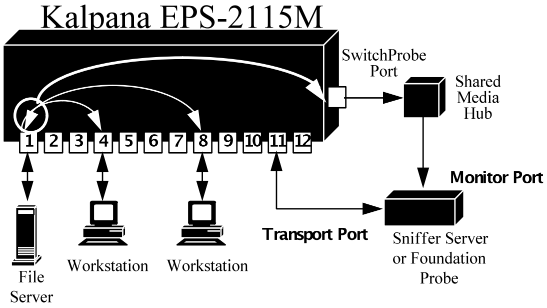

Management Strategy

The EPS-2115M maintains a dedicated “SwitchProbe port” on the back of the switch unit. This port must be connected to a small shared media hub. The Sniffer Server or Foundation Probe must then be connected to the shared media hub.

The SwitchProbe Port can then be enabled through either the SwitchProbe Configuration Menu available through TELNET or a direct serial connection. The SwitchProbe port can also be controlled through SNMP using Kalpana’s SwitchVision management software.

The SwitchProbe Port number identifies which port’s traffic will be replicated to the SwitchProbe Port.

The Traffic to Probe identifies what kind of traffic is replicated to the SwitchProbe port. Half duplex allows all traffic sent or received from the selected port to the SwitchProbe port. Full duplex allows the user to select either sent or received traffic to the SwitchProbe port.

Network Peripherals

1371 McCarthy Boulevard
Milpitas, CA 95035
(408) 321-7300

Products:

EIFO Client/Server Switching Hub

Monitoring Capabilities

Hardware Based

Management

In-Band SNMP
Out-of-Band Serial

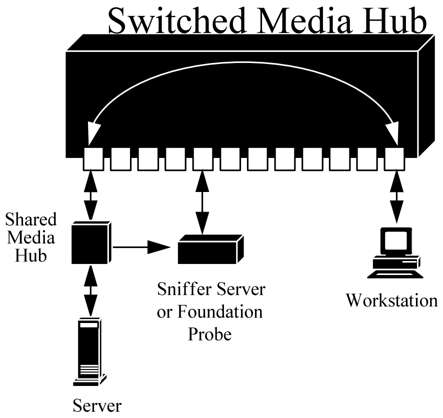

Management Strategy

The EIFO (Ethernet In/Fiber Out) series Switching Hub is offered by NPI. The EIFO series does not support internal port monitoring, so an external shared media hub is needed to monitor a port.

Attaching a small shared media hub on the connection between the shared resource and the switched media hub allows the analysis device to passively monitor the connection.

The administrator should consider inserting a shared media hub on each connection that may need to be analyzed. This preventative measure may eliminate the need to break the connection to the shared resource during periods of normal operation to insert the shared media hub.

PlainTree

59 Iber Road
Stittsville, ON Canada K2S 1E7
(800) 461-0062   (613) 831-8300

Products:
WaveSwitch 100

Monitoring Capabilities
Hardware Based

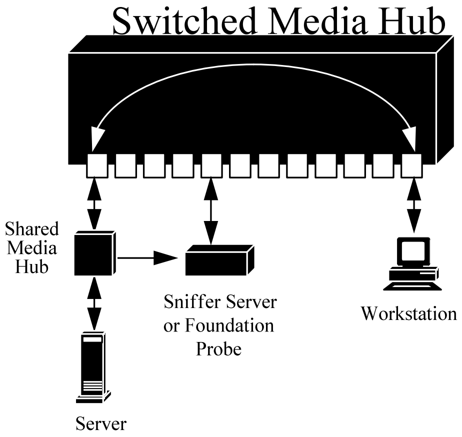

Management Strategy
The WaveSwitch series is the name of the CrossComm switched hub product line. The ClearPath XL series does not support internal port monitoring, so an external shared media hub is needed to monitor a port.

Attaching a small shared media hub on the connection between the shared resource and the switched media hub allows the analysis device to passively monitor the connection.

The administrator should consider inserting a shared media hub on each connection that may need to be analyzed. This preventative measure may eliminate the need to break the connection to the shared resource during periods of normal operation to insert the shared media hub.

SMC Hauppauge, NY
(800) 762-4968

Products:
EliteSwitch ES/1
TigerSwitch XE

Management Capabilities
Hardware based

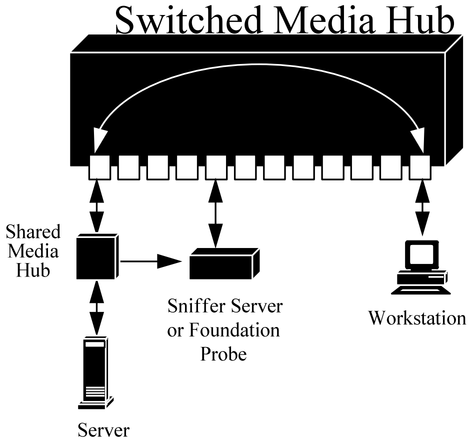

Management Strategy

SMC maintains two series of switched hub products. The TigerSwitch is self contained "workgroup" type hub. The EliteSwitch is a concentrator, allowing the installation of several host modules. Niether the EliteSwitch nor the TigerSwitch series supports internal port monitoring, so an external shared media hub is needed to monitor a port.

Attaching a small shared media hub on the connection between the shared resource and the switched media hub allows the analysis device to passively monitor the connection.

The administrator should consider inserting a shared media hub on each connection that may need to be analyzed. This preventative measure may eliminate the need to break the connection to the shared resource during periods of normal operation to insert the shared media hub.

3Com Corporation

5400 Bayfront Plaza
Santa Clara, CA 95052-8145
(408) 764-5000

Products

LinkSwitch 500, 1000, 1200, 2200, 2700
LANplex 2016, 2500, 6000

Monitoring Capabilities

Port Tap (Called Roving Analysis)

Management

In Band SNMP, Telnet
Out-Of-Band Serial

Management Strategy

The LinkSwitch and LANplex series of switches support a Port Tapping technology that they call Roving Analysis. This feature can be managed through 3Com's Transcend Network Management software or a normal TELNET connection. The following describes syntax for controlling this feature through a Telnet connection. The examples assume that you have an active administrative TELNET connection to the switched hub.

Syntax

```
analyzer add
Select Ethernet Port (1-16):port-number
```

(***port-number*** designates the port number of the connected DSS server or Foundation Probe. The device will respond with the hardware address of the monitoring device.)

```
analyzer start
Select Ethernet Port (1-16):port-number
```

Enter the address the Analyzer is located on - type 'q' to return to previous menu
Address: ***probe-address***

This starts the monitoring process. ***port-number*** refers to the port that you wish not monitor. ***probe-address*** refers to the hardware address of the monitoring device.

Example

```
>analyzer add
Select Ethernet Port (1-16):4
Probe address is 00-00-65-01-02-03
>analyzer start
Select Ethernet Port (1-16): 1
Enter the address the Analyzer is located on - type 'q' to return to previous menu
Address: 00-00-65-01-02-03
```

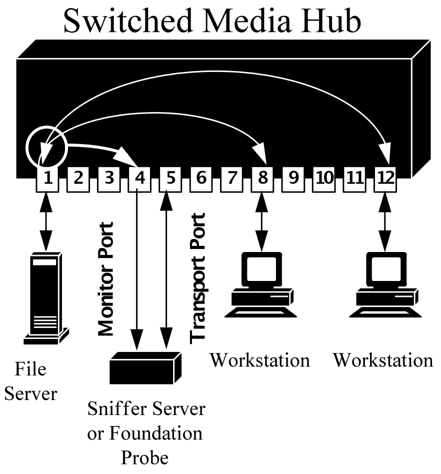
Switched Media Hub
1
2
3
4
5
6
7
8
9
10
11
12
Monitor Port
Transport Port
File Server
Sniffer Server or Foundation Probe
Workstation
Workstation