# Network Health Check And Baseline

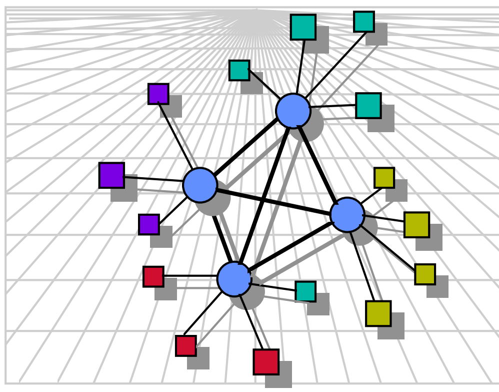

***Submitted To:***

**Hallmark**
**2713 Farnam Street**
**Omaha, Nebraska 68102**

***Submitted By:***

**Network Partners, Inc.**

**(402-434-5395)**

**October 18, 2002**

# Table of Contents

# Executive Overview

Network Partners Inc. (NPI) was engaged by Hallmark to perform a Network Assessment and Baseline of the enterprise data network during the week of August 28, 2002. Network Assessment and Baseline documents are generally viewed as high-level summaries of network resource consumption as of a particular date, and as an indicator of any inefficiency that may exist due to problems or constraints within the network. Initial documents tend to provide a discovery and feedback mechanism that often results in various forms of network change. These early changes tend to significantly impact network traffic and reduce the longer-term value of the data.

The purpose of a network assessment is to identify those areas needing improvement over time. The recommendations contained in this document should not be construed as criticism of current operational practices, but rather to provide a structure to improve operational efficiencies.

Operationally, the Hallmark enterprise network is currently functioning well with a few exceptions. When comparing this network to others analyzed by Network Partners (NPI) during the past two years, the network can be given a grade of “B”. Although the network is operational, this report will discuss and document those areas where improvements in stability and effectiveness can be achieved.

The network infrastructure components (e.g., Cisco switches and routers) serving the firm’s thirteen locations are based on high quality Ethernet components. Hallmark’s implementation of these components is currently relying on many of Cisco’s default configuration parameters. The default values were observed causing:

- Erratic response times due to improper auto-negotiated Ethernet settings
- Unstable Spanning Tree negotiation of Root Bridges
- Switch-to-switch trunk problems in Atlanta

The Frame Relay network was observed to be functioning as engineered, however several changes are recommended to improve the quality of service as seen by Hallmark users. The network has been over-subscribed by a factor of 2.5 causing congestion, discarded packets, retransmissions and erratic response times.

In addition, the management of the Frame Relay network was observed to be less than satisfactory primarily due to bandwidth management responsibilities that are not being addressed by anyone, and Hallmark's inability to access the Cisco router configurations.

As new technologies are deployed from a business perspective, Hallmark (or their agent) must have the ability to manage available bandwidth and the quality of service provided to Hallmark users. As servers and workstations become faster, bandwidth constraints will become more acute and more difficult to manage.

Hallmark may also wish to consider other network management vendors, including those from Qwest; however, the services to be provided should be well understood from a technical viewpoint to ensure the level of service expected can truly be provided.

The changes recommended to the Ethernet local networks and to the Frame Relay network will raise the effectiveness of Hallmark's end-to-end network to world-class levels.

# Local Area Network Findings

## Omaha Configuration Overview

The following high-level diagram represents the Local Area Network configuration for the Omaha office.

### Figure 1

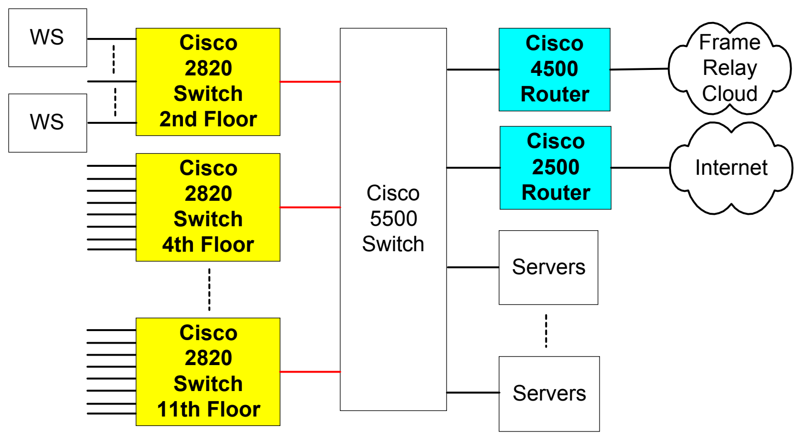


Each remote location is configured in a similar manner with slight variations in specific switch model numbers and quantities.

The local networks in Omaha and the remote sites were observed to be functioning as engineered, however several changes are recommended to improve the quality of service and stability as seen by Hallmark users. The following sections will discuss the findings and recommendations.

## Ethernet Usage Summary

The utilization of each of the Ethernet segments and inter-switch trunk links were observed to be operating well within industry accepted levels. As shown in the following table, the inter-switch trunk links were generally operating at less than 2% utilization with an occasional short-duration peak above 2%.

### Inter-Switch Trunk Links - Omaha

| Switch IP | Port | Ave Util. (%) | Peak Util. (%) | Errors | To Do Actions |
|---|---|---|---|---|---|
| 10.15.172.20 | 4/1 | < 1 | 1 | 0 | No Actions Required |
| 10.15.172.21 | 4/2 | 1 | 2 | 0 | No Actions Required |
| 10.15.172.40 | 4/3 | < 1 | 2 | 0 | No Actions Required |
| 10.15.172.41 | 4/4 | < 1 | 1 | 0 | No Actions Required |
| 10.15.172.50 | 4/5 | < 1 | 1 | 0 | No Actions Required |
| 10.15.172.51 | 4/6 | < 1 | 2 | 0 | No Actions Required |
| 10.15.172.60 | 4/7 | < 1 | 1 | 0 | No Actions Required |
| 10.15.172.61 | 4/8 | < 1 | 1 | 0 | No Actions Required |
| 10.15.172.70 | 4/9 | 1.5 | 2 | 0 | No Actions Required |
| 10.15.172.71 | 4/10 | 1.5 | 2.5 | 0 | No Actions Required |
| 10.15.172.80 | 4/11 | < 1 | < 1 | 0 | No Actions Required |
| 10.15.172.81 | 4/12 | < 1 | < 1 | 0 | No Actions Required |
| 10.15.172.90 | 1/2 | < 1 | 4 | 0 | No Actions Required |
| 10.15.172.91 | 3/1 | < 1 | < 1 | 0 | No Actions Required |
| 10.15.172.110 | 1/1 | < 1 | 1 | 0 | No Actions Required |

The following table reflects seven of the higher usage Omaha switch ports associated with servers or other common network resources. The ports were all operating well within industry-accepted levels.

## Typical Server Ports - Omaha

| Port | Ave Util. | Peak Util. | Errors | To Do Actions |
|---|---|---|---|---|
| 7/2 | 1 % | 2 % | 0 | No Actions Required |
| 7/3 | 1.75 % | 7.5 % | 0 | No Actions Required |
| 7/5 | < 1% | 1.5 % | 0 | No Actions Required |
| 7/9 | 1 % | 6.4 % | 0 | No Actions Required |
| 8/8* | 5% | 15% | 0 | No Actions Required |
| 8/10 | < 1 % | 2 % | 0 | No Actions Required |
| 9/13 | 1 % | 3 % | 0 | No Actions Required |

*Port 8/8 is running at 10 Mb. All other ports listed in table are running at 100Mb.

As noted in the above table, port 8/8 was operating correctly at a 10 megabit per second half-duplex speed. However, the operational statistics (e.g., collisions, runts, giants) indicated that over 2,900,000 collisions have occurred over the past several weeks. Collisions on half-duplex ports such as this are normal and acceptable, however user interactions with the attached device would function at a higher performance level if these devices were upgraded to 100 megabit full-duplex operation where possible.

The following chart reflects the utilization of port 8/8 which is a representative sample of the majority of server and other shared devices in Hallmark's local network.

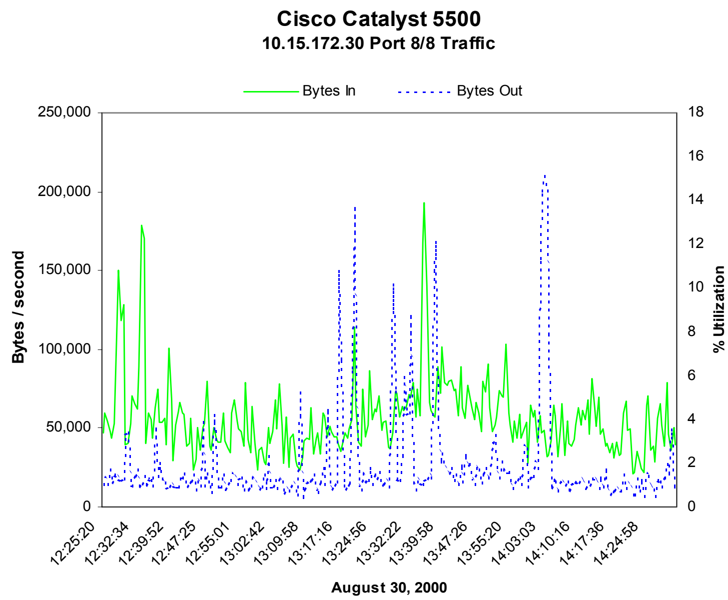


## Managing a Switched Network

To properly manage a switched Ethernet network, a detailed understanding of how a typical switch operates is necessary.

**Figure 2**

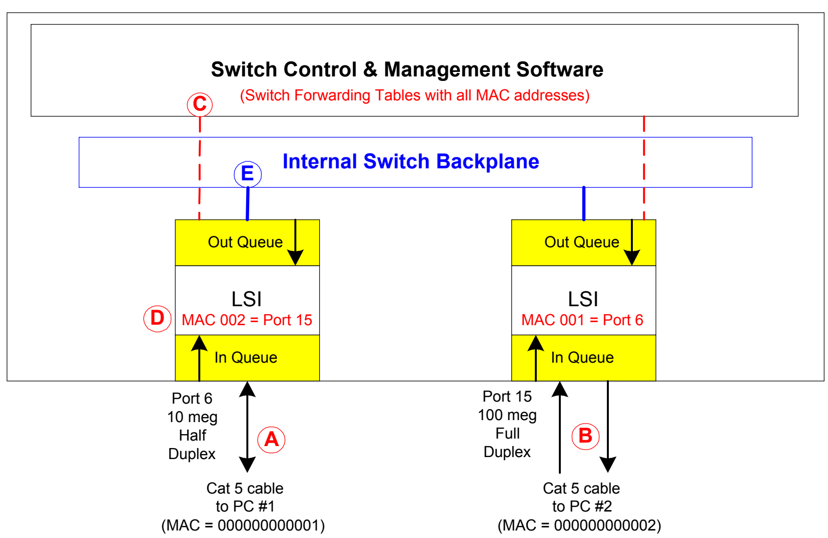


In this example, two PC's are connected to this switch: PC #1 on port 6 (half-duplex) and PC #2 on port 15 (full-duplex).

The session startup between PC #1 and PC #2 occurs as follows:

1. PC#1 sends a broadcast looking for PC#2 across path **A**

2. The Large Scale Integrated circuit (LSI) learns PC #1's MAC address from the Ethernet header (*source address*)
3. The LSI communicates the *source address* and *port observed* to the switch control software via path **C**
4. The switch software populates its Switch Forwarding Table with the source address and port number
5. The LSI requests the port number from the switch software associated with the MAC address contained in the Ethernet *destination address* field
6. The LSI stores the destination MAC and port number in its internal memory, item **D**
7. All communications from PC#1 to PC #2 occurs across the switches internal back plane, path **E**
8. Traffic volumes, error conditions and other operational statistics are kept by the LSI and passed to the switch control & management software on an occasional basis
9. The switch control software waits for another session startup

The two most important points to noted above are:

- The port's *auto-negotiating* functions (e.g., 10 vs 100, half vs full duplex) are usually built into the Large Scale Integrated circuit
- All packet *buffering*, *queuing* and *switching* functions are implemented on Large Scale Integrated circuits

The Large Scale Integrated (LSI) circuits used by the various switch manufacturers (including Cisco) are generally manufactured by some other company that specializes in integrated circuit manufacturing. The LSI chips (like Pentium processors) have been designed to perform very specialized functions and have a very limited amount of memory within the chip. The majority of the LSI circuits are fixed-function and cannot be reprogrammed to add external memory or new Ethernet functions (Quality of Service, QoS). Each LSI chip can support from one to eight physical Ethernet ports depending upon the exact manufacturer design.

### *Ethernet Auto-Negotiating*

The auto-negotiation function within Ethernet switches and a PC's Network Interface Card (NIC) are not based on standards. It is left up to each vendor to determine how they will negotiate 10 megabit verses 100 megabit, half verses full duplex, and whether the port supports any form of Ethernet flow control. Likewise, there are no published standards relative to whether port speed should be negotiated before or after the duplex discovery function, etc.

As a result, auto-negotiation of these important settings are marginal at best, and frequently result in *different settings* each time a physical Ethernet connection is established. This includes each time a server or workstation is rebooted, and, each time the Ethernet cable is de-inserted / inserted into a switch. Incorrect port settings will generate substantial quantities of errors and significantly slow network response time (and therefore, application response time).

The majority of Hallmark's switches, servers and workstations at all locations have been configured for auto-negotiation with few exceptions.

The settings of a Cisco switch can be easily viewed from a standard telnet session:

```
OMA_5500 (enable) show port
Port  Name               Status    Vlan   Level  Duplex Speed Type
----- -------------- --------- ------ ------ ------ ----- ------------
4/11 8th Floor .80    connected 1      normal   full   100 100BaseFX MM
4/12 8th Floor .81    connected 1      normal   full   100 100BaseFX MM
7/1  OAMFS01_A        connected 1      normal a-full a-100 10/100BaseTX
7/2  OMAFS02_A        connected 1      normal a-full a-100 10/100BaseTX
7/3  OMASM01_A        connected 1      normal a-full a-100 10/100BaseTX
```

The above sample display indicates the ports 4/11 and 4/12 providing connectivity to switches on 8th floor have been statically defined, however ports 7/1, 7/2 and 7/3 to various servers are configured to auto-negotiate these important settings. The "a-" in the above display indicates auto-negotiation.

Incorrectly negotiated port settings will result in significant errors as shown in this sample display:

```
Port  Align-Err  FCS-Err    Xmit-Err   Rcv-Err    UnderSize
----- ---------- ---------- ---------- ---------- ---------
 9/12          0          0          0          0         0
 9/13          0          0          0          0         0
 9/14          0          0          0          0         0
 9/15      57507      51228          0          0         0
 9/16     587761     521443          0          0         0
 9/17       5375       5430          0          0         0
 9/18       5378       5926          0          0         0
 9/19     565785     558570          0          0      1804
```

A review of the primary backbone Cisco 5500 switch in the Omaha data center indicated that more than 50 ports had experienced significant errors at one time or another that are indicative of incorrect port settings. The remote locations were similar.
The table below indicates some of the switch ports that were observed with significant *increasing* errors during this network assessment.

| Location | Ports |
|---|---|
| Omaha | 3/1,7/6-9, 7/11, 9/1 9/6, 9/10, 9/13, 9/14, 9/16 |
| Scottsdale (PHX) | 2/3, 2/7, 2/11, 2/12 |
| Atlanta | 2/4, 2/5, 2/6, 2/11 |
| Denver | 3/7, 3/10, 3/12 |
| Washington DC | 2/1-12 with emphasis on 2/1, 2/4, 2/6 |
| Richmond | 2/2, 3/1, 3/2, 3/3 |
| Newport Beach | 1/1, 2/2, 2/6, 2/7, 2/8 |
| Little Rock | 1/1, 2/5 |
| Kansas City | 1/1, 2/1, 2/2 |

Hallmark personnel should statically define each physical port associated with servers, switch trunks, routers and all other common shared devices. Locking the switch port to static settings will resolve a significant number of incorrect auto-negotiating problems, however each attached device (e.g., server) should also be statically defined to ensure reliable operation. Statically defining NIC card settings (e.g., servers) generally requires a reboot even though many NT displays indicate otherwise.

### *Ethernet Switch Queuing*

As indicated above, the LSI circuits used in all major manufacturer's switches contain only limited memory to buffer and queue packets for transmission to an end device, and to queue packets onto the internal switch back plane. The buffering is typically limited to two-to-five packets, depending on the LSI manufacturer's design.

Referring to Figure 2, above, packets arriving from PC #1 over the Ethernet Cat 5 cable are received into an Inbound Queue. If the LSI circuit already knows where to forward this inbound packet, it will transmit the packet on the internal back plane to the LSI chip directly associated with the outbound port.

The back plane speed is significantly faster than each of the physical Ethernet ports, and will seldom become any type of operational issue with current technology. For those interested, the speed is typically measured in millions of packets per second or greater.

Significant Outbound Queues will develop, however, in the LSI circuits when a packet is waiting to be transmitted on the slower Cat 5 cable. This interface to the Cat 5 cable is typically 100 to 1000 times slower than the internal back plane. If two or more devices attempt to send packets to the same port, the outbound port queue will be quickly filled and packets discarded (dropped on the floor). If the LSI circuit discards a packet, it will notify the switch control & management software that it is doing so for reporting and statistical purposes. Discarded packet statistics can be retrieved via an SNMP management system such as CiscoWorks.

Two issues significantly impact whether the LSI circuit can transmit a packet, or, whether it must queue the packet for later transmission.

- Ethernet 10 / 100 / 1,000 megabit speed
- Half verses Full Duplex

If a server is configured to operate in half-duplex mode (which includes the auto-negotiated half-duplex mode), the LSI circuit will not be able to transmit data on a port whenever data is arriving on the same port. As a half-duplex port becomes busier and busier, the probability of the small outbound queue overflowing and discarding packets increases *exponentially*. Therefore, Hallmark should never rely on half-duplex Ethernet configurations for any high-volume or response-time-sensitive devices, including servers and routers.

## Figure 3

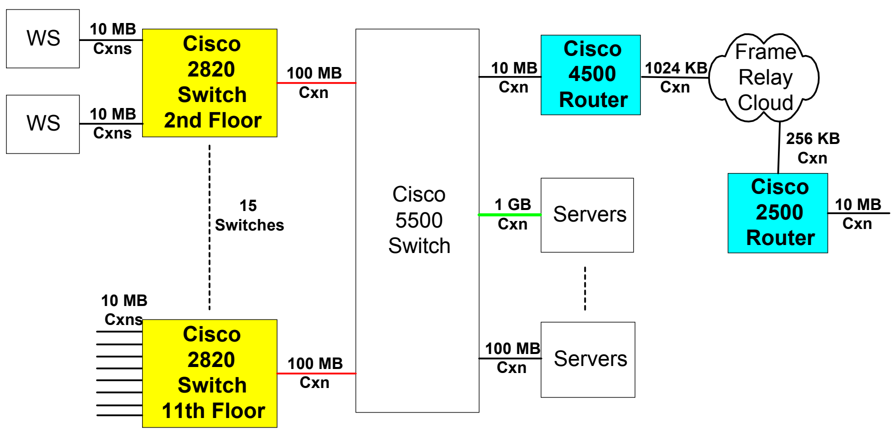


Likewise, the speed of the Ethernet connection determines how quickly the LSI outbound queue can be emptied. In Figure 3, the LSI circuit sending data on the gigabit connection to the middle server can empty the queue quicker and handle more simultaneous user sessions than can the 100-megabit server or the 10-megabit half-duplex router connection.

Hallmark personnel should develop a plan to migrate all high-usage servers to gigabit speeds, and all others to 100-megabit full-duplex statically defined configurations.

### *Spanning Tree Implementation*

All switches in a network environment have the capability of using the Spanning Tree Algorithm to ensure the existence of a loop-free topology. For most major switch manufacturers, the Spanning Tree Protocol (STP) is operational by default, including the Cisco Catalyst series.

The Spanning Tree Algorithm will send Bridge Protocol Data Units (BPDU) packets on each port of each switch. Other switches in the same network listen for these BPDU packets to elect a *Root* bridge (or switch), and each evaluates the *Path Cost* from the root bridge

## Figure 4

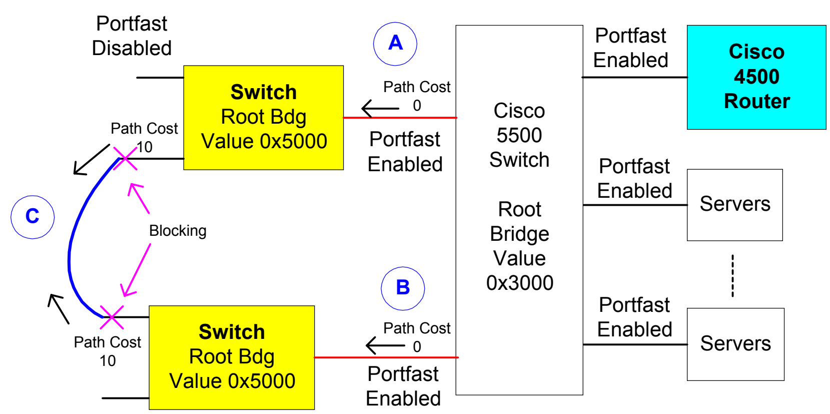


to determine which path to disable when a loop is created. Without Spanning Tree, should an accidental loop occur (e.g., Path A to Path C to Path B in Figure 4), packets would be continuously circulated around the path consuming all available network resources (meltdown).

All switches shipped from the various manufacturers include *default* Root Bridge Values (used to elect the Root) and default Path Costs (used to select operational interfaces). Root Bridge values of 32,768 (or 0x8000) and Path Costs of 10 on each interface are standard. The manufacturers expect the network engineer to change these default values to those supportive of each company's network environment.
All switches used in the various Hallmark local network environments were still configured with default values. As a result, switches other than the intended primary Hallmark switch were elected as the root. Significant spanning tree topology changes were observed in both the Omaha and Denver locations.

When properly configured, should a loop develop (such as Path A, B and C in Figure 4), the intended interfaces (Point C in Figure 4) would be placed in a *non-forwarding (or blocking)* condition automatically eliminating the accidental loop. Without specifying reasonable root bridge values and path costs, the link at Point B (as an example) may be blocked causing all users connected to one switch to pass through one or more other switches significantly impacting performance and response time.

Hallmark personnel should review the Root Bridge values for each switch at all locations, and specify a *lower* value for each *primary* switch. The lowest value will be elected the root bridge at each location.

The Cisco Catalyst series of switches have also implemented a second level of loop protection. This second level essentially blocks data from being forwarded *into* the switch for the first 30 seconds following the activation of a physical port (link up). This feature is known to cause significant (and erratic) problems in the Microsoft NT environment as the station inserting into a switch will attempt to contact the Primary Domain Controller (PDC) before the switch forwards data packets. The 30-second delay causes the server or workstation to believe there are no domain controllers in the local network. (Note: any slight interruption to a Cat 5 cable will also cause the physical port to stop forwarding data for the 30-second interval as well. Should someone accidentally disconnect a server cable for one or two seconds, this feature will cause the server to disappear for 30-seconds.)

This second *feature* can be disabled by specifying *port-fast = enable* for each physical port where this feature is unwanted.

Hallmark personnel should consider enabling port-fast on all switch ports that are not likely to be involved with cabling loops (e.g., servers, routers, trunks to other switches). Some accounts have elected to enable this parameter on all switch ports as this feature also impacts the new high-speed workstations.

### *Catalyst Software Version Levels*

The Cisco Catalyst 5000 switches implemented at all locations are operating with four different versions of switch software as shown in the following table.

| Location | Current Version | Latest Version |
|---|---|---|
| Omaha | 4.4(1) | 5.5(3) |

| | | |
|---|---|---|
| Little Rock | 4.4(1) | |
| Washington, DC | 4.5(2) | |
| Denver | 4.5(2) | |
| Kansas City | 4.4(1) | |
| Newport Beach | 4.5(2) | |
| Atlanta | 2.3(1) | |
| Richmond | 5.1(1) | |
| Scottsdale | 5.1(1) | |

Two major software problems were identified during this network assessment. First, should support personnel use the trace-route diagnostic facility within those switches using the v4.4(1) software, the switch will automatically reboot when support personnel exit their telnet session.

Second, dual trunk links have been implemented between the Catalyst 5000 and a 4006 switch in Atlanta to provide greater bandwidth to the users. The dual trunks are not functioning, as the two software versions are not compatible with each other.

Hallmark personnel should upgrade the operating software in each of the Catalyst 5000 switches to current levels, and review the requirements for each directly attached remote switch (e.g., 4000, 2820) to ensure compatibility with the primary switches.

## Switched Local Network Recommendations

The following items recap recommendations noted earlier in this section of the document.

Ethernet switches, regardless of vendor, have limited capabilities to buffer packets. The limited buffering can cause packets to be discarded and end user response time to become erratic and slow. To minimize this negative impact, Hallmark personnel should:

1. Statically define each physical port associated with servers, switch trunks, routers and all other common shared devices.
2. Never rely on half-duplex Ethernet configurations for any high-volume or response-time-sensitive devices, including servers and routers.
3. Develop a plan to migrate all high-usage servers to gigabit speeds, and all others to 100-megabit full-duplex statically defined configurations.

Hallmark personnel should review the spanning tree Root Bridge values for each switch at all locations, and specify a *lower* value for each *primary* switch. The lowest value will be elected the root bridge at each location.

In addition, Hallmark should consider enabling Cisco's port-fast option on all switch ports that are not likely to be involved with cabling loops (e.g., servers, routers, trunks to other switches). Some accounts have elected to enable this parameter on all switch ports as this feature also impacts the new high-speed workstations.

Hallmark should upgrade the operating software in each of the Catalyst 5000 switches to current levels, and review the software requirements for each directly attached remote switch (e.g., 4000, 2820) to ensure compatibility with the primary 5000 switch.

# Frame Relay Network Findings

## Configuration Overview

The following diagram represents a high-level view of Hallmark's current Wide Area Network.

**Figure 5**

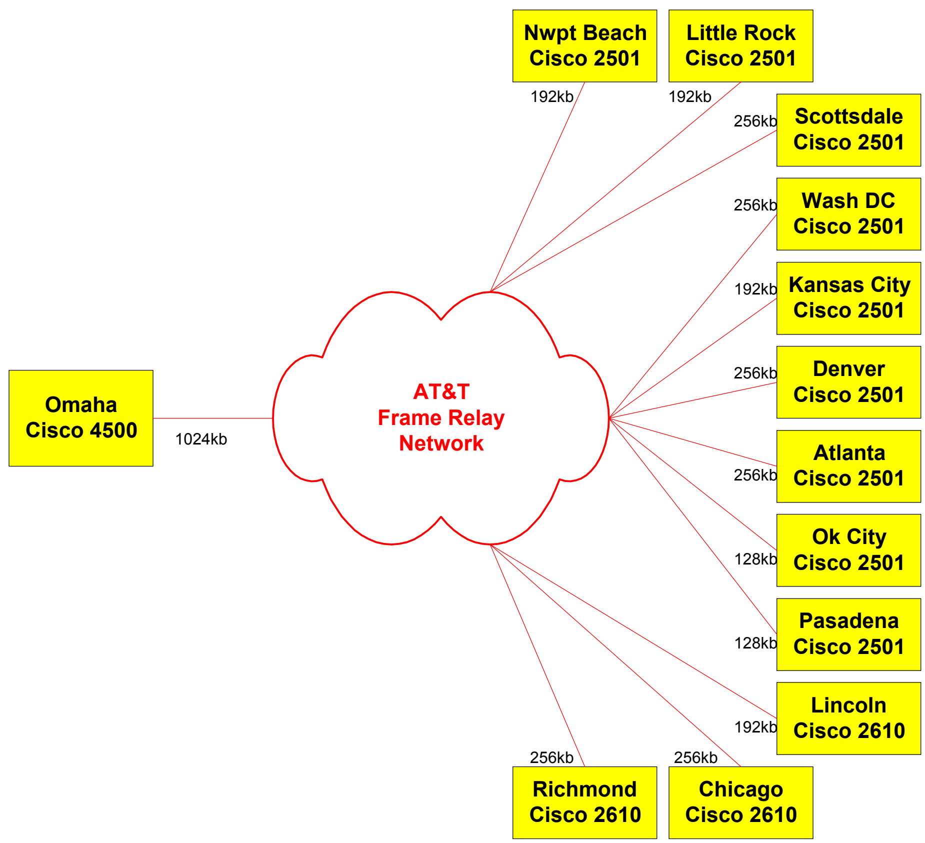


The Frame Relay network was observed to be functioning as engineered, however several changes are recommended to improve the quality of service as seen by Hallmark users.

The sum of all Port Speeds indicate the Frame Relay network has been over-subscribed by 2.5, while the sum of all CIR speeds indicate the network has been over-subscribed by 1.6. Although it is common practice in some organizations to over-subscribe, the decision to do so is based primarily on a good understanding of the exact data flows within the enterprise network. The following sections of this document will help provide a technical understanding of the effects within Hallmark, and the changes necessary to improve current performance limitations.

### *Frame Relay – Statistical Multiplexing*

Current Frame Relay standards address only Permanent Virtual Circuits (PVCs) that are administratively configured and managed by the carrier for each customer, as represented in Figure 6 by ***A***. The engineers responsible for designing a carrier's Frame Relay network provide bandwidth between their switches based on the sum of all customers' Committed Information Rate (CIR) values. In Figure 6, the bandwidth between FR Switch #2 and FR Switch #3 is the sum of company #1 and company #2 CIR values.

## Figure 6

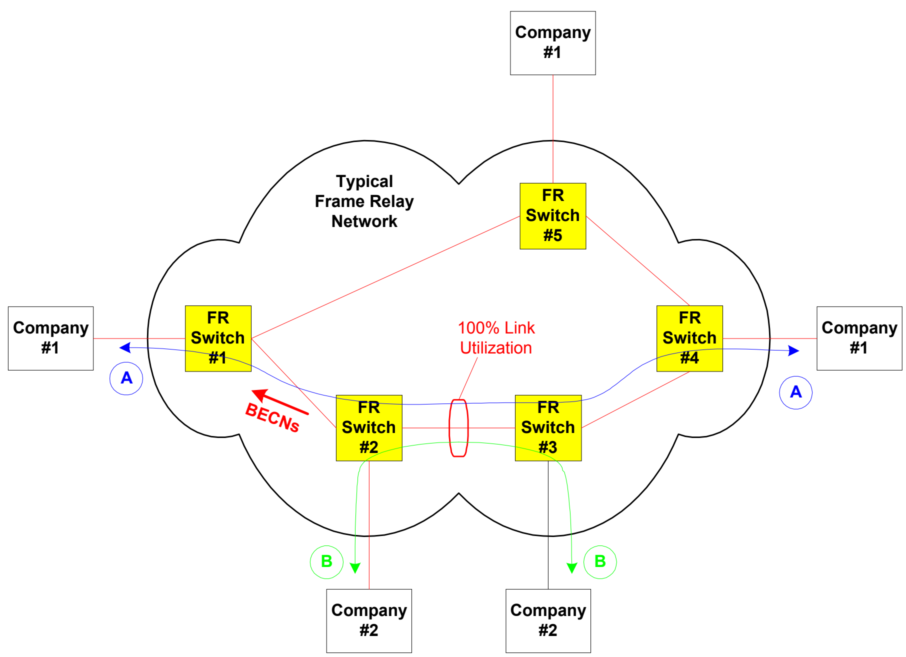


At any point in time, company #1 can send bursts of data in excess of their CIR values through path ***A***, however if company #2 is already consuming their allotted share of the bandwidth (path ***B***), many of company #1's packets will be *dropped* and Backward Explicit Congestion Notification (BECN) packets will be sent back to customer #1. It is the responsibility of customer #1's router to reduce the amount of data traffic presented, or, simply allow the end-stations (e.g., server and work station) to recognize missed packets through a long TCP/IP protocol timeout period. The protocol timeout method is highly visible by the end-user in the form of erratic and elongated network response times.

The following chart represents Hallmark's actual network traffic between Omaha and Denver over a four-hour period. The blue line indicates that all available bandwidth was consumed for each of four instances of sending a 57-megabyte file from Omaha to Denver.

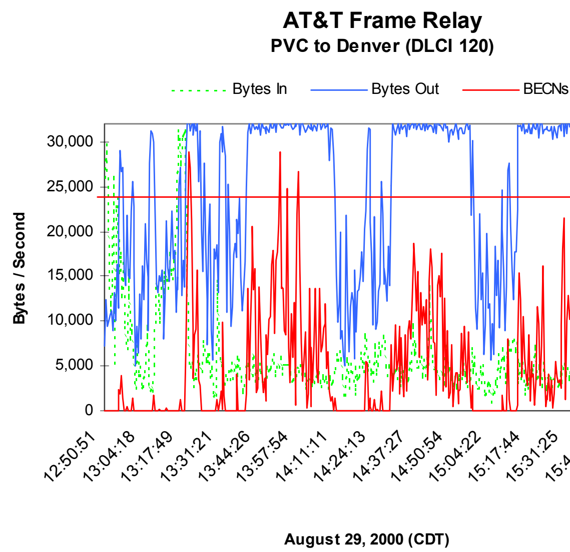


The red horizontal line indicates Hallmark's chosen CIR value for this circuit, and since the data traffic exceeded the CIR value, BECN packets (shown in red) were sent back to Omaha signifying network congestion. The Hallmark router did not slow the flow of traffic, but rather continued its transmissions to Denver. Undoubtedly, many packets where dropped by the Frame Relay network causing the Hallmark server and/or work stations to detect the dropped packets via the protocol timeout mechanism. It is important to note that traffic was able to burst to the full Denver Port Speed (256,000 bits per second divided by 8 bits per byte equals 32,000 bytes per second).

Note the quantity of BECN packets is significantly different between the 13:44, 14:50, 15:31 and 16:10 portions of the chart. The differences are directly related to congestion within the Frame Relay network, and have no relationship to Hallmark's chosen CIR values or port speeds.

Each of the four transfers shown in this chart required 30 minutes to transmit the 57 megabytes of data. While these transfers were occurring, any Hallmark user attempting to communicate across this path would have been severely impaired. Whether any Hallmark user or other server actually attempted is purely *statistical* in nature.

### *Frame Relay – Managing Bandwidth*

Figure 7, below, represents Hallmark's data flows through the Cisco routers. All traffic sent by Server #1 must be received by the router in an *Inbound Queue* (buffer), inspected by the router software to determine the path on which to send the data, and placed in an *Outbound Queue* (buffer) for transmission to the distant point. A single outbound queue (associated with the Frame Relay v.35 serial port) is shared by all PVCs.

**Figure 7**

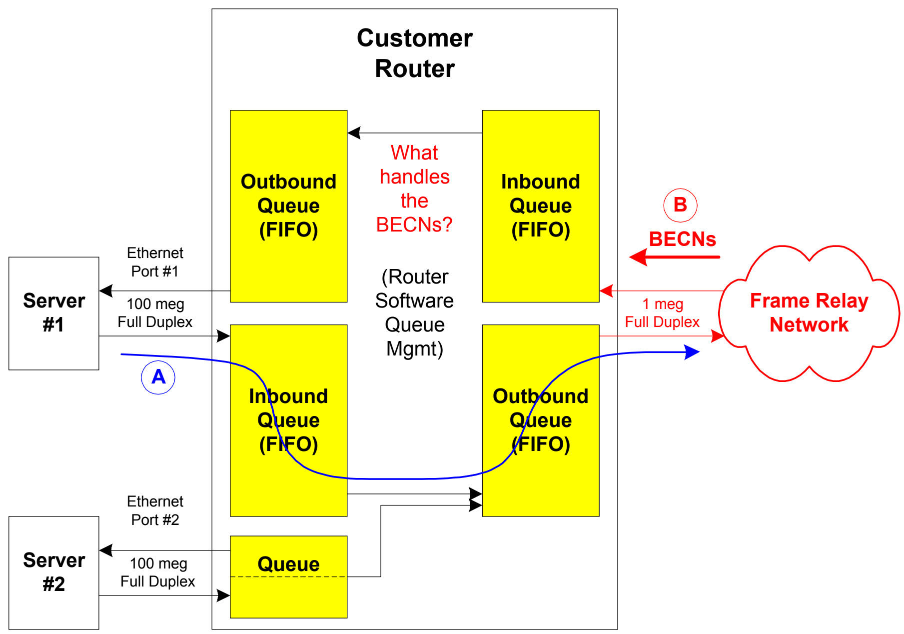


If the Outbound Queue happens to be full due to CIR *or* port speed constraints within the Frame Relay network, the inbound packets from the server are simply dropped on the floor, causing the end-stations (e.g., server and/or work station) to detect the discarded packets through the elongated protocol timeout process.

All router vendors, including Cisco, provide mechanisms for managing the Inbound and Outbound Queues. Which mechanism is implemented within a customer's network is *dependent* upon the customer's specific traffic mixture and specific network policies.

Figure 8, below, reflects the data flows as observed at the Omaha Cisco 4500 router. Two performance related issues are important to note from this chart.

First, the full speed of the Omaha 1,024,000 bit per second (or 128,000 bytes per second) serial link into the Frame Relay network is never achieved. Rather, only 50% to 70% of the available bandwidth is used.

## Figure 8

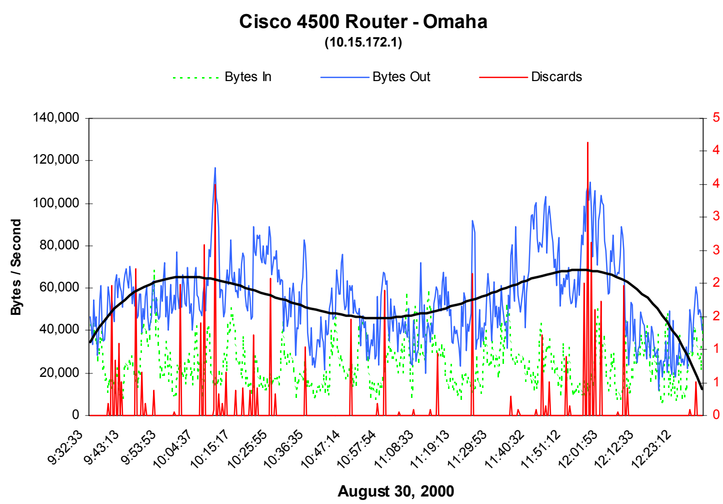


Second, the router's serial interface is discarding packets at the rate of one-to-four packets per second indicating the *Outbound Queue* (buffer) is full leaving the router with no choice other than to drop packets on the floor.

Both conditions are the direct result of Frame Relay congestion caused by inadequate CIR values, remote port speeds, and lack of a router queue management mechanism.

### *Managing End-to-End Network Bandwidth*

The intent of the Frame Relay protocol developers was for this protocol to be extended to each end-station (e.g., server and work station), where BECN packets received by the transmitting end-station would be truly interpreted by that end-station in such a manner as to slow the flow of data, thereby reducing congestion. There are no such implementations in the world today. Rather, the customer's Frame Relay router effectively encapsulates TCP/IP (and other protocols) into Frame Relay packets. Since there is no mechanism to translate Frame Relay BECN packets into TCP/IP flow control messages, the end-stations cannot be notified of congestion and therefore continue to send data at its current rate.

Frame Relay congestion management is handled in two ways: *congestion avoidance* or *congestion control*. Congestion avoidance is essentially the age-old method of purchasing the bandwidth needed to satisfy all end-station data flows (e.g., through more bandwidth at network problems). Small networks frequently use this technique as staffing for performance management is seldom supported from a business perspective.

Congestion control implies that *someone* understands the exact types of data flows occurring over the customer's network, and manages the data flows by implementing and monitoring some Outbound Queue management mechanism (as noted in Figure 7). Queue management implies that a set of Network Policies has been defined. The policies may be similar to:

- Network management protocols = High priority
- Telnet sessions = High priority
- OSPF/RIP routing protocol = High priority
- Netbios sessions over TCP/IP = Medium priority
- FTP file transfers = Low priority
- WEB surfing = Low priority

The individual responsible for implementing the queue management then defines High, Medium and Low priority queues within each router (possibly for each router interface) and assigns values (or time slices) to each queue instructing the router to give certain levels of priority to data contained within each queue.

The above sample describes one of several different queue management techniques available; however the technique that provides the greatest benefit for one customer's network may be substantially different than for another customer's network. Companies such as could not possibly be successful in constructing and managing these queues unless they understood the detailed data flows for each customer, the importance of each data flow (as documented in some network policy statement), and had the instrumentation in place to detect changes within the customer's network. (Note: a change could be that a customer purchased a client-server accounting package that happens to use TCP Port 1511 as its underlying communications transport mechanism.)

Figure 9, below, attempts to graphically depict organizational responsibilities in terms of the Open Systems Interconnection (OSI) model.

## Figure 9

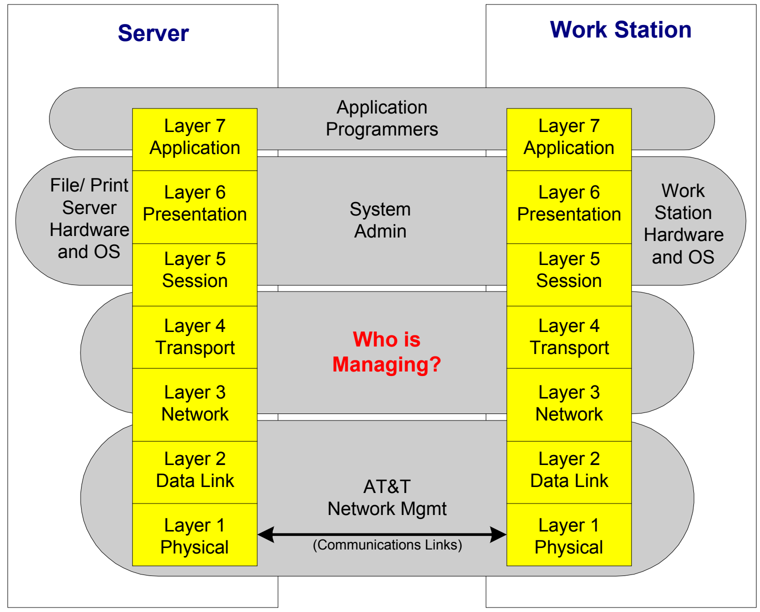


It is clear that Application Programmers are responsible for writing applications, that System Admin personnel manage the File/Print Server configurations and that  Network Management is monitoring Layer 1 (Physical), Layer 2 (Frame Relay), and a small portion of Layer 3 (TCP/IP). It is also clear that *no one* is managing the balance of Layer 3 and all of Layers 4 and 5. (Note: the only reason that  has any interest in Layer 3 is their network management system uses the Simple Network Management Protocol, SNMP, to remotely monitor the status of each Frame Relay interface. If some other Frame Relay network management protocol existed, they would have no interest in Layer 3 at all.)

Many telephone companies have sold network management services to customers without any reasonable definition or mutual understanding as to what those responsibilities mean. Traditionally, those services are limited to ensuring the Frame Relay network is active (Layer 1 and 2 on one serial port), the router is operational, and to some extent, upgrading router software to new releases. Since Hallmark support personnel have no router login capability, it should be clear from the above discussion that on one is managing Layer 1 and 2 for all other router interfaces, interface queues and discarded packets, data transmission priorities, selection of routes, IP route stability, as well as many other Layer 3 and above functions.

Hallmark management must reverse the router login roles that currently exist in order to be successful in supporting Hallmark's users in the future. Network Management has little, if any, need to log into these devices. Hallmark support personnel (or their agent) have a large need as demonstrated above.

As new technologies are deployed from a business perspective, Hallmark must have the ability to manage available bandwidth and the quality of service provided to Hallmark users. As servers and workstations become faster, bandwidth constraints will become more acute and more difficult to manage.

Hallmark may also wish to consider other network management vendors, including those from Qwest; however, the services to be provided should be well understood from a technical viewpoint to ensure the level of service expected can truly be provided.

## Frame Relay Usage Summary

The following table reflects the relative usage of the Frame network over an 80-minute sample interval:

| City | Bytes From | % | Bytes To | % | Total Traffic | % of Total |
|---|---|---|---|---|---|---|
| Scottsdale | 7,179,122 | 10.0% | 64,407,660 | 90.0% | 71,586,782 | 26.6% |
| Wash DC | 24,594,728 | 40.9% | 35,604,792 | 59.1% | 60,199,520 | 22.3% |
| Denver | 12,565,663 | 24.3% | 39,237,725 | 75.7% | 51,803,388 | 19.3% |
| Little Rock | 8,861,196 | 30.1% | 20,619,594 | 69.9% | 29,480,790 | 10.9% |
| Atlanta | 7,542,252 | 28.8% | 18,681,570 | 71.2% | 26,223,822 | 9.7% |
| Kansas City | 4,550,543 | 45.2% | 5,509,179 | 54.8% | 10,059,722 | 3.7% |
| Pasadena | 642,993 | 7.6% | 7,858,470 | 92.4% | 8,501,463 | 3.2% |
| Nwpt Beach | 1,212,663 | 18.7% | 5,283,860 | 81.3% | 6,496,523 | 2.4% |
| Richmond | 545,607 | 19.0% | 2,320,852 | 81.0% | 2,866,459 | 1.1% |
| Ok City | 236,098 | 12.4% | 1,673,011 | 87.6% | 1,909,109 | 0.7% |
| Lincoln | 159,574 | 48.4% | 170,451 | 51.6% | 330,025 | 0.1% |
| Chicago | NA | | NA | | NA | |
| Total | | 25.3% | | 74.7% | | 100% |

The table indicates that 74.7% of all Frame Relay traffic is outbound from Omaha to the remote offices, and 25.3% is inbound. Hallmark may find an economical advantage to price non-symmetrical Frame Relay CIR values. For example, Omaha-to-Atlanta may require a CIR of 128, while Atlanta-to-Omaha may only need a CIR of 32. Each Frame Relay provider has different CIR rate plans; non-symmetrical CIR values may not always be published.

The table also indicates that 79.1% of all network traffic is associated with Scottsdale, Washington DC, Denver and Little Rock. Managing the bandwidth to/from these four locations will have a substantial impact on the quality of service provided to the remaining eight locations. (The *Remote Site Internet Usage* section of this document indicates that 32.5% of all Frame Relay traffic is the result of these four location's use of the Internet. If this Internet traffic could be handled in another way, the quality of service provided to all remote locations would be significantly improved.)

## Remote Site Usage Detail

The following sections provide a short duration Baseline for each of the Frame Relay sites. Each PVC utilization chart has been arranged to depict the remote site's *Port Speed* as the upper x-axis limit (e.g., 24,000 bytes per second x 8 bits per byte = 192,000 bit per second port speed). This is the maximum speed at which data can possibly be transmitted. The red horizontal datum line reflects the *Committed Information Rate* (CIR) currently configured for each PVC.

### *Atlanta (PVC 110)*

The Atlanta PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 79 kb |
| **Peak Speed** | 243 kb |
| **Last PVC Change** | Oct 28, 1999 |
| **Last Remote Router Reboot** | July 26, 2000 |
| **BECNs since Aug 22, 1999** | 195,191 |

The PVC Utilization over a four-hour period is depicted in the following chart.

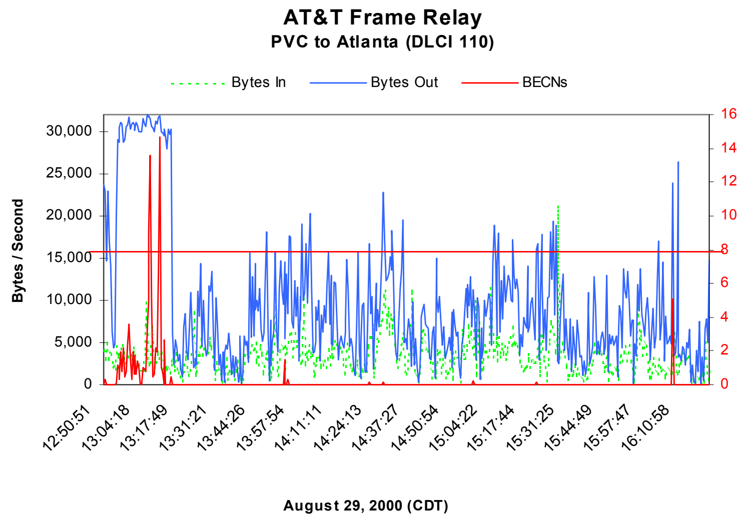


The following chart indicates that 56% of all traffic from Atlanta is associated with Internet usage.

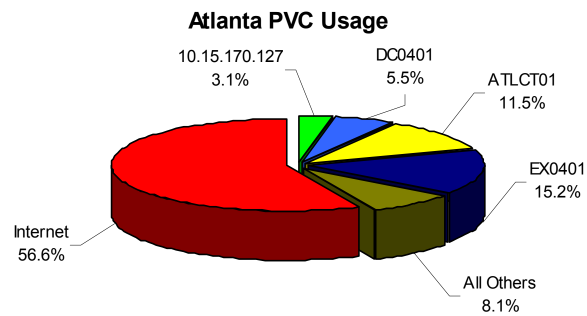


**Recommendations:**

1. The Atlanta PVC is properly sized for the traffic currently being presented.
2. If the Internet traffic could be routed via an Atlanta DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time. A 38kb or greater circuit would be required for the Internet traffic.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the Atlanta router.

### *Denver (PVC 120)*

The Denver PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 192 kb |
| **Average Actual Speed** | 190 kb |
| **Peak Speed** | 249 kb |
| **Last PVC Change** | Apl 28, 2000 |
| **Last Remote Router Reboot** | Aug 24, 1999 |
| **BECNs since Aug 22, 1999** | 4,040,441 |

The PVC Utilization over a four-hour period is depicted in the following chart.

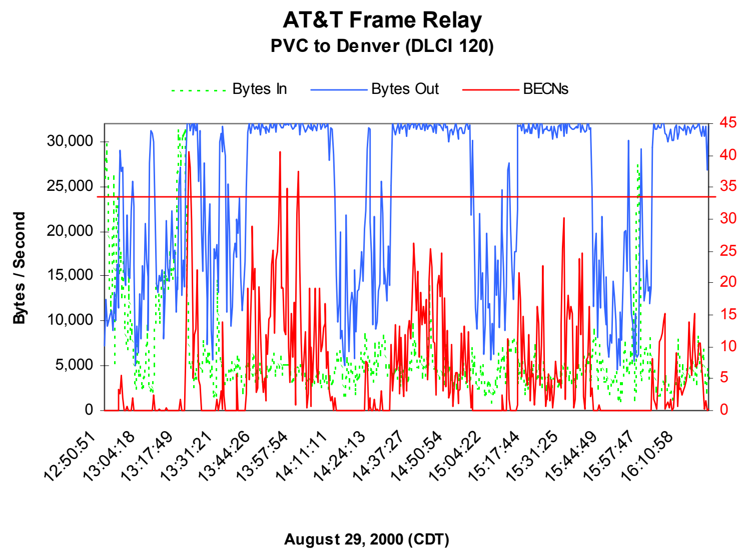


The following chart indicates that 52% of all traffic for Denver is associated with Internet usage.

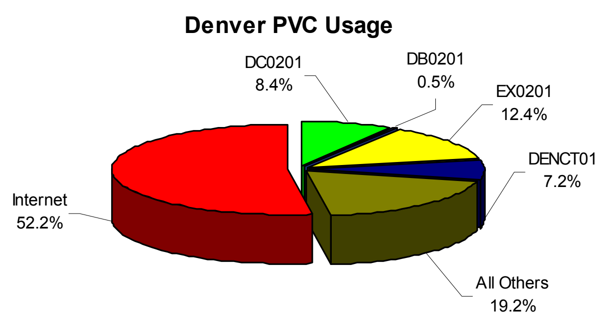

**Recommendations:**

1. The Denver PVC should be upgraded. A port speed of 512 kb with a CIR of 320 is recommended for current traffic.
2. If the Internet traffic could be routed via a Denver DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time. A 56kb or greater local circuit could handle the Internet traffic.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the Denver router.

### *Kansas City (PVC 130)*

The Kansas City PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 192 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 19 kb |
| **Peak Speed** | 154 kb |
| **Last PVC Change** | Aug 7, 2000 |
| **Last Remote Router Reboot** | Aug 4, 2000 |
| **BECNs since Aug 22, 1999** | 49,263 |

The PVC Utilization over a four-hour period is depicted in the following chart.

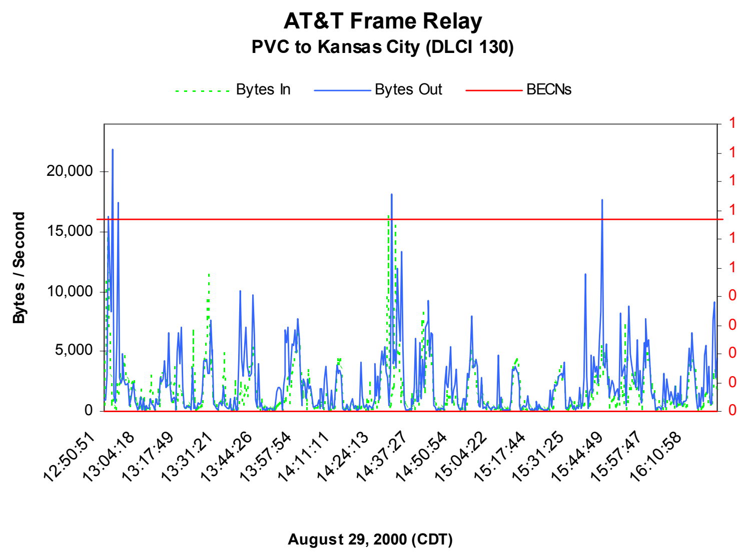


The following chart indicates that 15% of all traffic for Kansas City is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

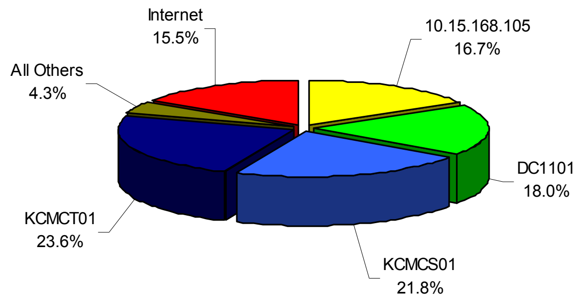


**Recommendations:**

1. The Kansas City CIR value can be reduced from the current 128 value, to 96 and still provide the same level of service.
2. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the router.

## *Little Rock (PVC 141)*

The Little Rock PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 192 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 33 kb |
| **Peak Speed** | 143 kb |
| **Last PVC Change** | Jul 6, 2000 |
| **Last Remote Router Reboot** | May 15, 2000 |
| **BECNs since Aug 22, 1999** | 190,804 |

The PVC Utilization over a four-hour period is depicted in the following chart.

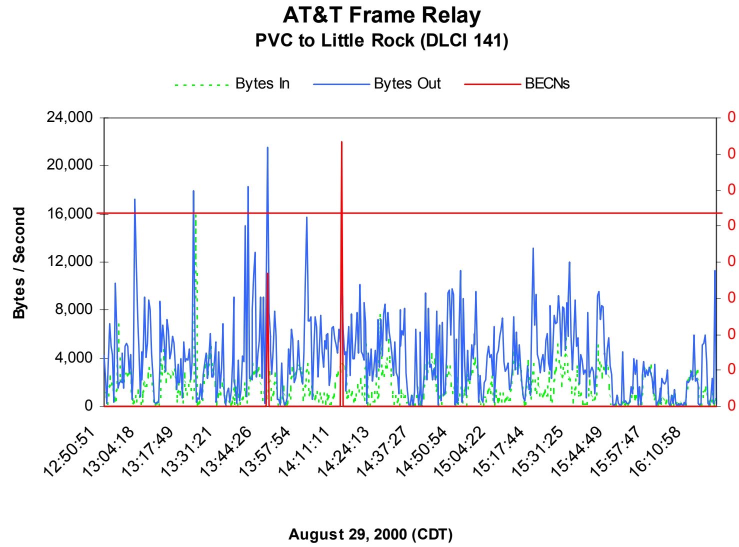


The following chart indicates that 70% of all traffic for Little Rock is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

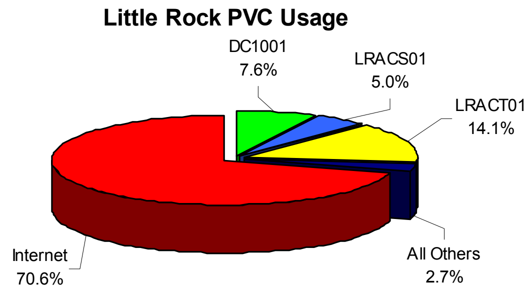


**Recommendations:**

1. The Little Rock CIR value could be reduced from the current 128, to a value of 96 and still provide the same level of service.
2. If the Internet traffic could be routed via a Little Rock DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time. A 45 kb or greater speed circuit would be required for the Internet traffic.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the Little Rock router.

### *Newport Beach (PVC 150)*

The Newport Beach PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 192 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 28 kb |
| **Peak Speed** | 155 kb |
| **Last PVC Change** | Aug 24, 2000 |
| **Last Remote Router Reboot** | Oct 28, 1999 |
| **BECNs since Aug 22, 1999** | 25,643 |

The PVC Utilization over a four-hour period is depicted in the following chart.

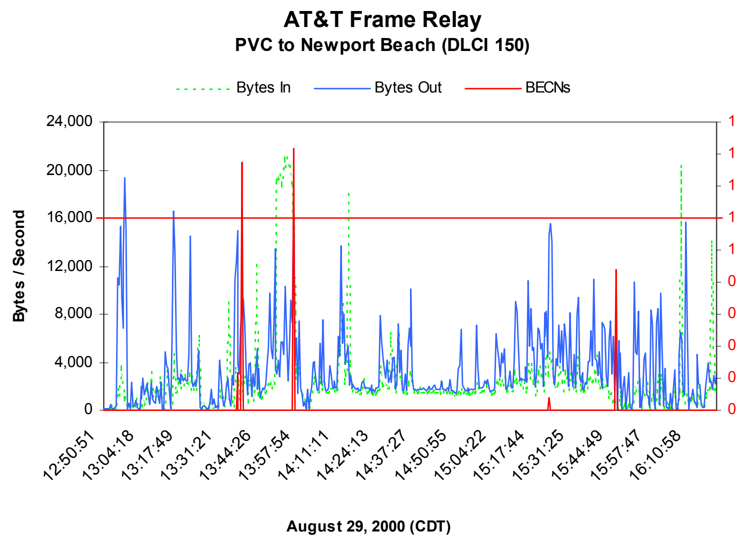


The following chart indicates that 41% of all traffic for Newport Beach is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

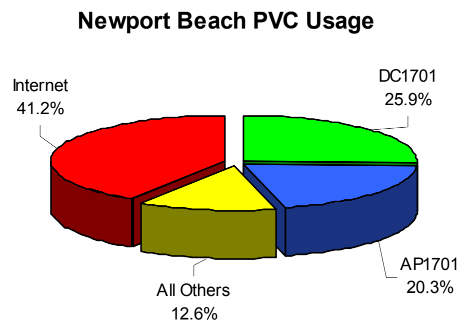


**Recommendations:**

1. The Newport Beach CIR value could be reduced from the current 128 value, to a 96 value and still provide the same level of service.

### *Oklahoma City (PVC 160)*

The Oklahoma City PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 128 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 12 kb |
| **Peak Speed** | 108 kb |
| **Last PVC Change** | Jun 26, 2000 |
| **Last Remote Router Reboot** | Aug 28, 1999 |
| **BECNs since Aug 22, 1999** | 20,941 |

The PVC Utilization over a four-hour period is depicted in the following chart.

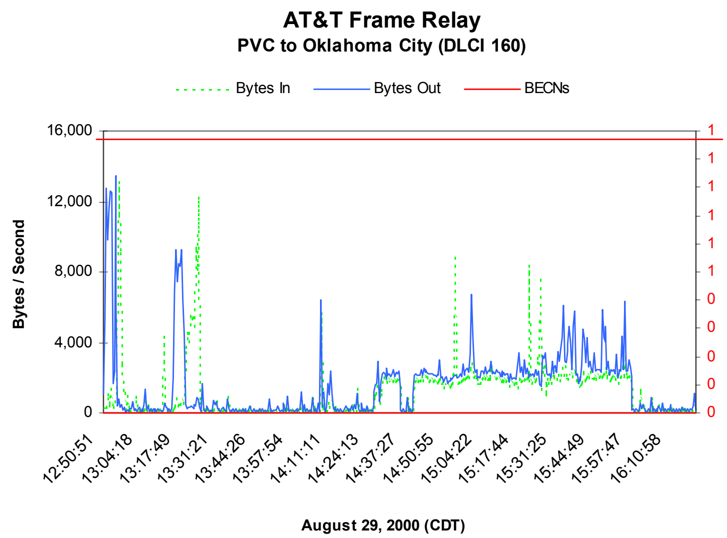


The following chart indicates that 0.5% of all traffic for Oklahoma City is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

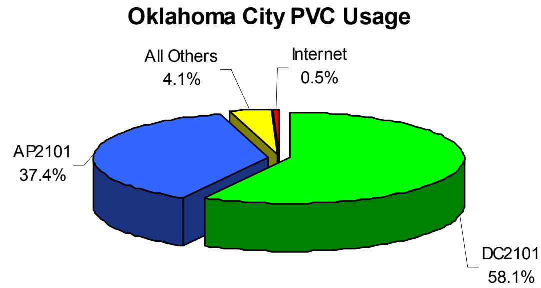


**Recommendations:**
1. The Oklahoma City CIR could be reduced from the current 128 value, to a 32 value and still provide the same level of service.
2. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the router.

## *Scottsdale (PVC 180)*

The Scottsdale PVC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 67 kb |
| **Peak Speed** | 250 kb |
| **Last PVC Change** | Aug 24, 2000 |
| **Last Remote Router Reboot** | Oct 23, 1999 |
| **BECNs since Aug 22, 1999** | 313,049 |

The PVC Utilization over a four-hour period is depicted in the following chart.

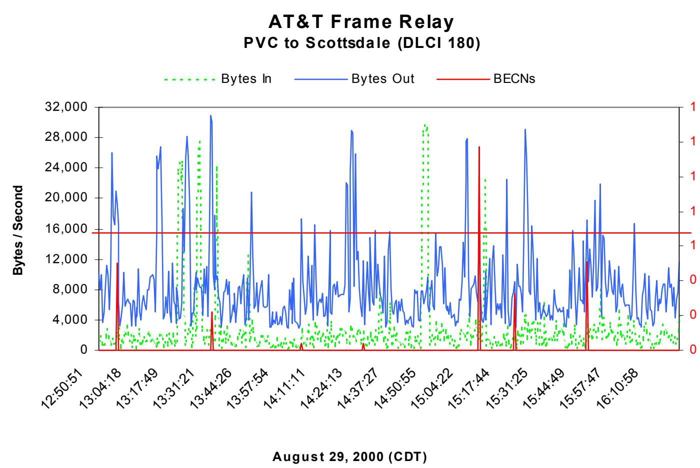

The following chart indicates that 19% of all traffic for Scottsdale is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

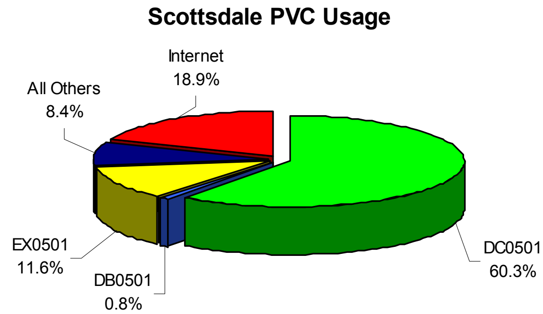


**Recommendations:**

1. The Scottsdale PVC is properly sized for the traffic currently being presented.
2. If the Internet traffic could be routed via a Scottsdale DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time. A 25kb or greater circuit would provide the required service.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the router.

### *Washington DC (PVC 200)*

The Washington DC traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 63 kb |
| **Peak Speed** | 187 kb |
| **Last PVC Change** | May 1, 2000 |
| **Last Remote Router Reboot** | Sep 4, 1999 |
| **BECNs since Aug 22, 1999** | 278,262 |

The PVC Utilization over a four-hour period is depicted in the following chart.

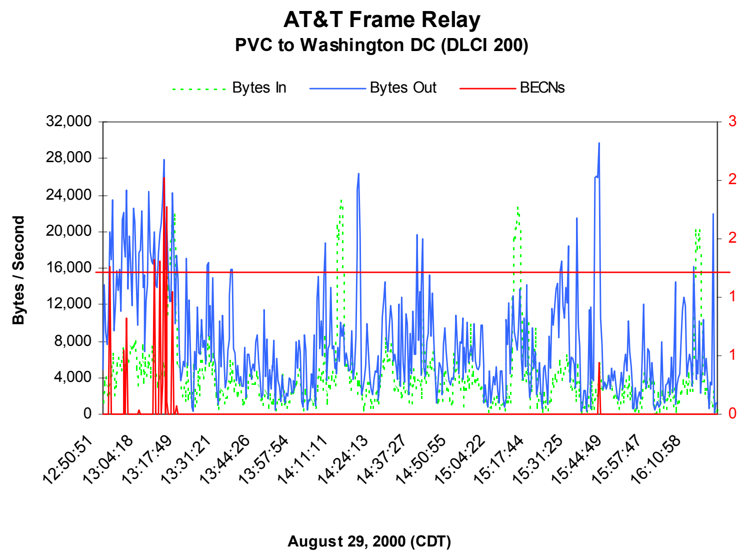


The following chart indicates that 44% of all traffic for Washington DC is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

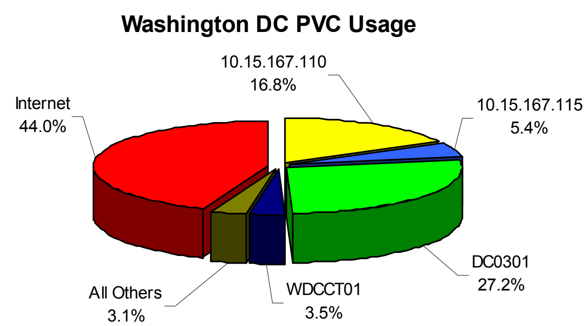


**Recommendations:**

1. The Washington DC PVC is properly sized for the traffic currently being presented.
2. If the Internet traffic could be routed via a Washington DC DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time. A 72kb or greater circuit would be required to handle the Internet traffic.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the router.

### *Lincoln (PVC 210)*

The Lincoln traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 192 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 10 kb |
| **Peak Speed** | 104 kb |
| **Last PVC Change** | Jun 16, 2000 |
| **Last Remote Router Reboot** | Jun 14, 2000 |
| **BECNs since Aug 22, 1999** | 27,554 |

The PVC Utilization over a four-hour period is depicted in the following chart.

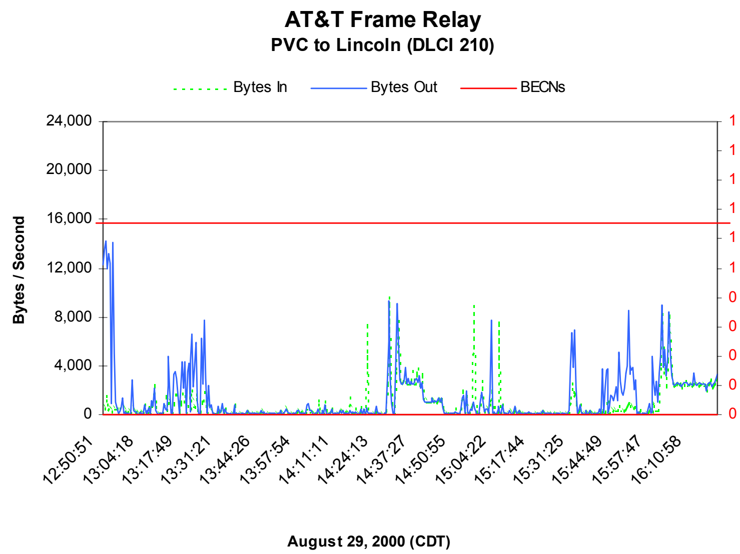


The following chart indicates that 0% of all traffic for Lincoln is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

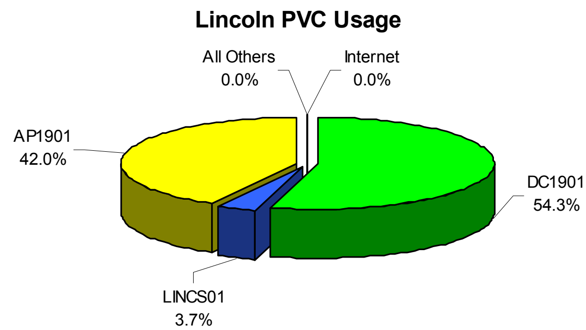

**Recommendations:**

1. The Lincoln CIR value could be reduced from the current 128, to a value of 64 and still provide the same level of service.

### *Pasadena (PVC 220)*

The Pasadena traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 128 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 23 kb |
| **Peak Speed** | 114 kb |
| **Last PVC Change** | Aug 24, 2000 |
| **Last Remote Router Reboot** | Jul 10, 2000 |
| **BECNs since Aug 22, 1999** | 124,688 |

The PVC Utilization over a four-hour period is depicted in the following chart.

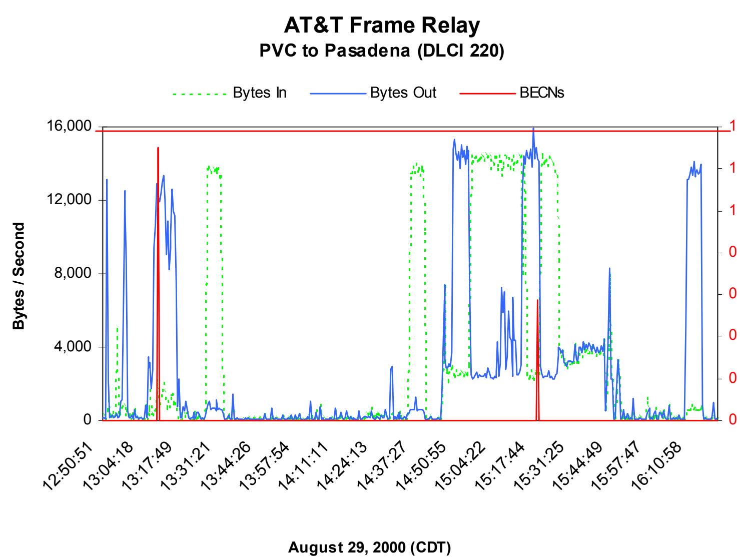


The following chart indicates that 5% of all traffic for Pasadena is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

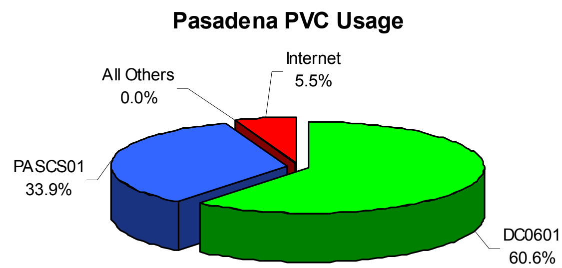


**Recommendations:**
1. The Pasadena PVC should be upgraded to a port speed of 192kb. The current CIR value of 128 is adequate for current traffic.

## *Richmond (PVC 230)*

The Richmond traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 30 kb |
| **Peak Speed** | 155 kb |
| **Last PVC Change** | Aug 20, 2000 |
| **Last Remote Router Reboot** | Feb 6, 2000 |
| **BECNs since Aug 22, 1999** | 40,239 |

The PVC Utilization over a four-hour period is depicted in the following chart.

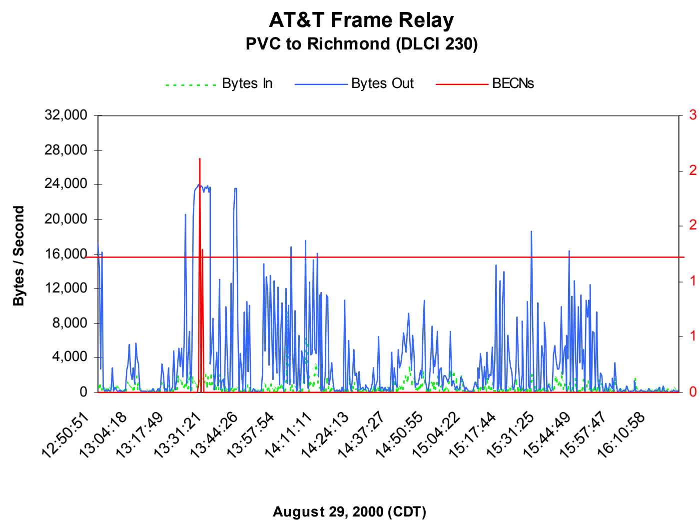

The following chart indicates that 16% of all traffic for Richmond is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

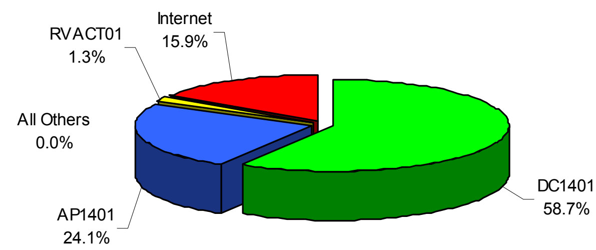


**Recommendations:**

1. The Richmond PVC should be downgraded. The port speed can be reduced from the current 256kb, to a speed of 192kb. The CIR value can also be reduced from the current 128, to a value of 96.

### *Chicago (PVC 240)*

The Chicago traffic characteristics are summarized as follows:

| | |
|---|---|
| **Current Port Speed** | 256 kb |
| **Current CIR** | 128 kb |
| **Average Actual Speed** | 177 kb |
| **Peak Speed** | 256 kb |
| **Last PVC Change** | Aug 17, 2000 |
| **Last Remote Router Reboot** | Sep 7, 2000 |
| **BECNs since Sep 7, 2000** | 1,353 |

The PVC Utilization over a one-hour period is depicted in the following chart. It should be noted that the Chicago site was activated on September 7th, therefore the traffic is reflected in this chart was not obtained in the same timeframes as all other locations.

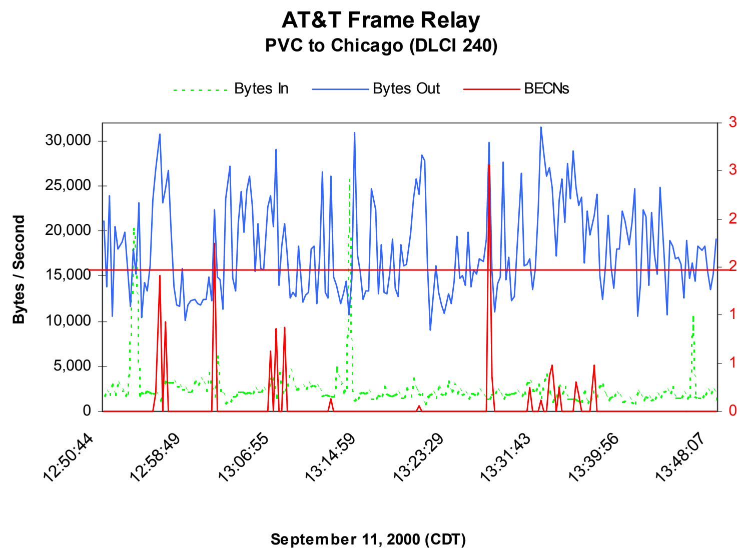


The following chart indicates that 5% of all traffic for Chicago is associated with Internet usage, while the balance is intranet interactions with Omaha devices.

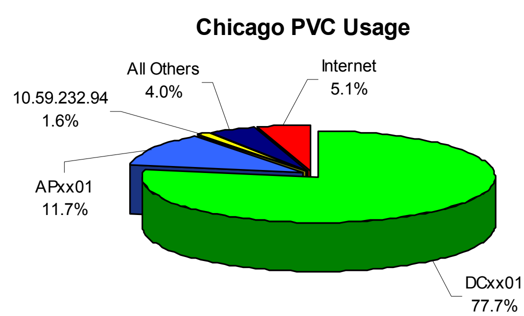


**Recommendations:**

1. The Chicago CIR value should be increased from the current 128, to a value of 192.
2. If the Internet traffic could be routed via a Chicago DSL circuit to a local Internet Service Provider, the existing PVC will provide adequate service for a longer period of time.
3. If Novell IPX traffic to/from other Hallmark locations is no longer required, the IPX routing definitions should be removed from the router.

## Remote Site Internet Usage

All remote Hallmark offices are currently configured to transport all Internet-destined traffic through the Omaha office as shown by ***A*** in the following diagram.

## Figure 10

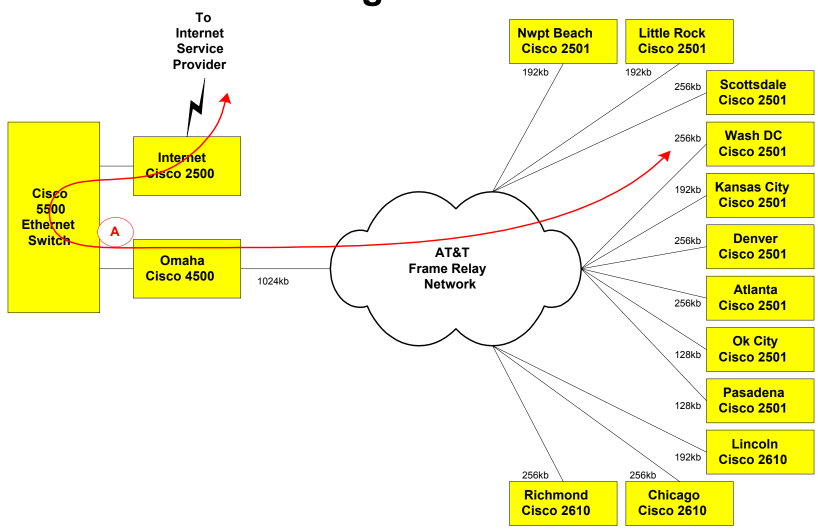


The amount of Internet traffic to/from each remote location is shown in the following table.

| City | Internet Traffic (Bytes) | % of Total Internet | % of Total FR Traffic |
|---|---|---|---|
| Washington DC | 42,641,699 | 27.9% | 10.9% |
| Denver | 35,484,653 | 23.2% | 9.0% |
| Little Rock | 26,764,802 | 17.5% | 6.8% |
| Atlanta | 22,823,515 | 14.9% | 5.8% |
| Scottsdale | 14,387,523 | 9.4% | 3.7% |
| Newport Beach | 4,432,534 | 2.9% | 1.1% |
| Kansas City | 2,692,497 | 1.8% | 0.7% |
| Richmond | 2,569,544 | 1.7% | 0.7% |
| Pasadena | 1,204,749 | 0.8% | 0.3% |
| Ok City | 19,739 | 0.0% | 0.0% |
| Lincoln | 0 | 0.0 | 0.0 |
| Chicago* | 5,906,956* | --* | --* |
| Total Internet | 153,021,255 | 100.0% | 39.0% |
| Total Intranet | 239,212,241 | -- | 61.0% |
| Total FR Traffic | 392,233,496 | -- | 100.0% |

(*Note: The Chicago location was added to the network on September 7, 2000. The network Baseline data for all other locations was obtained prior to the Chicago site activation. The volume of traffic shown in this table for Chicago is for reference purposes only (order of magnitude), and has not included in column totals.)

Overall, the Hallmark Frame Relay network consists of 39% Internet and 61% Intranet traffic. It should be noted the first five locations in this table generate 36% of all data traffic handled by the Frame Relay network (or 93% of all Internet traffic transported by the network).

Hallmark personnel should review the availability and cost of deploying local DSL solutions, and protect the Intranet by installing a single Firewall / Proxy Server at those locations. This server should only allow WEB and FTP services to pass to the Internet, and stop all incoming (from the Internet) regardless of service. Email and other corporate services should continue to be transported via the Frame Relay network to Omaha. Only a single ISP provided IP address would be required at each location.

## Frame Relay Recommendations

Three approaches will be discussed that resolve the Frame Relay network performance issues. The three approaches include:

- Single Network (Port Speed and CIR Changes Only)
- Single Network (Speed Changes, & Remote Internet Access)
- Dual Homed Network (Omaha & Denver, & Speed Changes)

Each of the approaches assume that router queue management has been or will be implemented.

### *Single Frame Relay Network (Speed Changes Only)*

This approach to resolving the Frame Relay network performance is the most conservative. The approach increases the Omaha serial port speed, three remote location port speeds, and adjusts the CIR values at eight locations to more close reflect current traffic levels.

| City | Current | | Recommended | |
|---|---|---|---|---|
| | Port Speed | CIR | Port Speed | CIR |
| Omaha | 1024 kb | | 1536 kb | |
| Scottsdale | 256 kb | 128 kb | 256 kb | 128 kb |
| Wash DC | 256 kb | 128 kb | 256 kb | 128 kb |
| Denver | 256 kb | 192 kb | 512 kb | 320 kb |
| Little Rock | 192 kb | 128 kb | 192 kb | 96 kb |
| Atlanta | 256 kb | 128 kb | 256 kb | 128 kb |
| Kansas City | 192 kb | 128 kb | 192 kb | 96 kb |
| Pasadena | 128 kb | 128 kb | 192 kb | 128 kb |
| Newport Beach | 192 kb | 128 kb | 192 kb | 96 kb |
| Richmond | 256 kb | 128 kb | 192 kb | 96 kb |
| Ok City | 128 kb | 128 kb | 128 kb | 32 kb |
| Lincoln | 192 kb | 128 kb | 192 kb | 64 kb |
| Chicago | 256 kb | 128 kb | 256 kb | 192 kb |

In addition to the speed changes, some form of outbound queue management must be implemented within each of the Cisco routers.

### *Single Frame Relay Network (Plus Remote Internet)*

This approach is slightly more aggressive and combines the speed changes mentioned above, with off-loading the Frame Relay network of Internet traffic at five remote locations as shown in the Figure 12, below.

## Figure 12

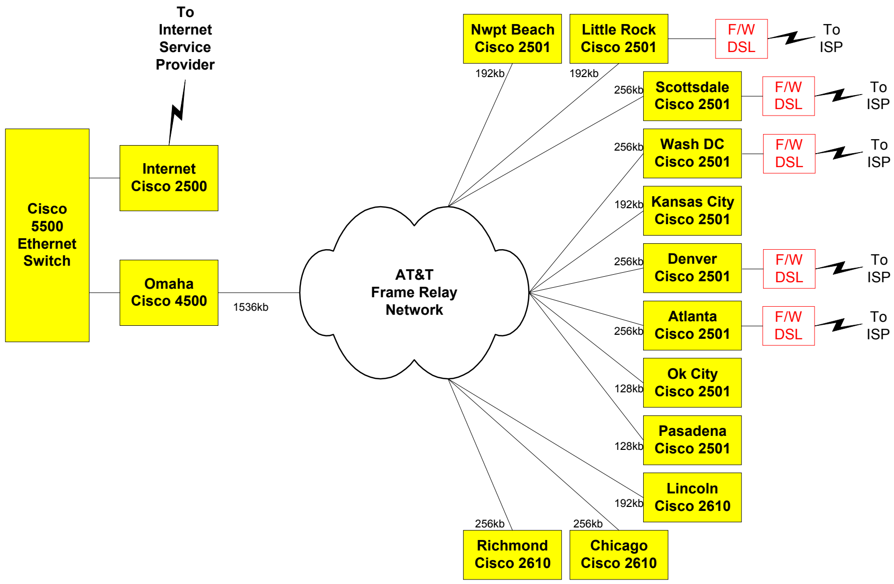


The port speeds and CIR speeds are the same as those described in the previous option.

| City | Current | | Recommended | |
|---|---|---|---|---|
| | Port Speed | CIR | Port Speed | CIR |
| Omaha | 1024 kb | | 1536 kb | |
| Scottsdale | 256 kb | 128 kb | 256 kb | 128 kb |
| Wash DC | 256 kb | 128 kb | 256 kb | 128 kb |
| Denver | 256 kb | 192 kb | 512 kb | 320 kb |
| Little Rock | 192 kb | 128 kb | 192 kb | 96 kb |
| Atlanta | 256 kb | 128 kb | 256 kb | 128 kb |
| Kansas City | 192 kb | 128 kb | 192 kb | 96 kb |
| Pasadena | 128 kb | 128 kb | 192 kb | 128 kb |
| Newport Beach | 192 kb | 128 kb | 192 kb | 96 kb |
| Richmond | 256 kb | 128 kb | 192 kb | 96 kb |
| Ok City | 128 kb | 128 kb | 128 kb | 32 kb |
| Lincoln | 192 kb | 128 kb | 192 kb | 64 kb |
| Chicago | 256 kb | 128 kb | 256 kb | 192 kb |

In addition to the speed changes, some form of outbound queue management must be implemented within each of the Cisco routers.

The primary benefit to this approach is off-loading 36% of the existing Frame Relay traffic. The disadvantage is the increased complexity of managing five remote Firewall/Proxy servers.

### *Dual Homed Frame Relay Network (Omaha-Denver)*

The dual homed Frame Relay approach is intended to split the Internet traffic between two locations (Omaha and Denver), and to

## Figure 13

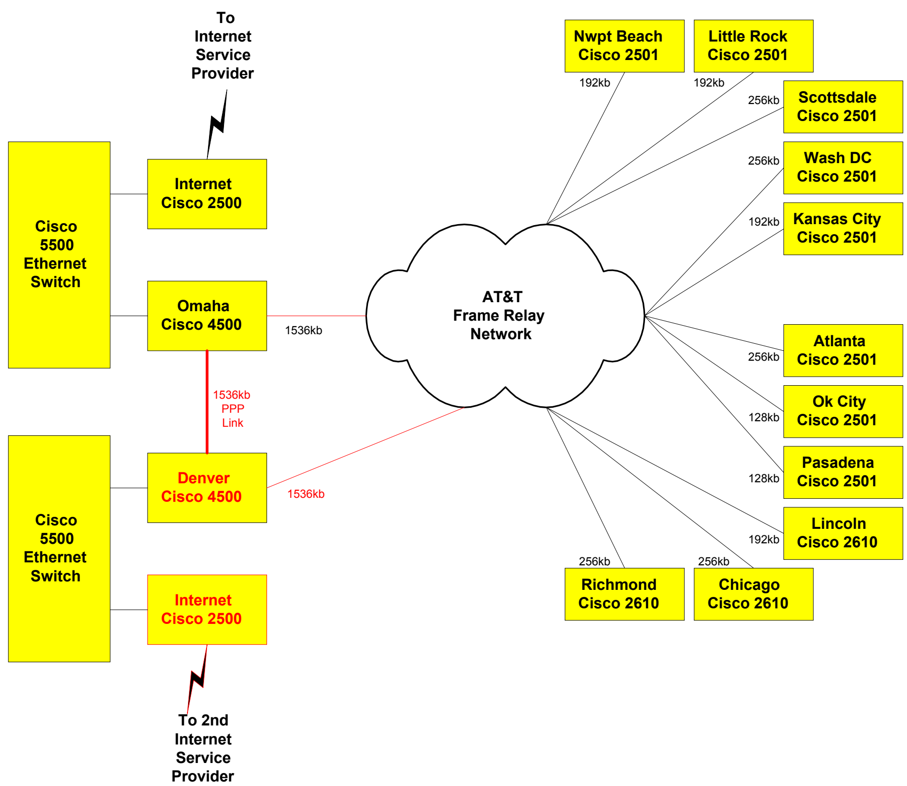


improve overall network availability and disaster recovery by creating two data paths from each remote location (one to Denver and one to Omaha). In addition, a dedicated Point-to-Point (PPP) link is recommended between Omaha and Denver, off-loading the traffic between these two cities from the Frame Relay network

As in all options, some form of outbound queue management must be implemented within each of the Cisco routers.

The port speeds and CIR speeds are the same as those described in the previous option with minor exceptions (e.g., Denver removed from the Frame Relay network).

| From Omaha to following: | Current | | Recommended | |
|---|---|---|---|---|
| | Port Speed | CIR | Port Speed | CIR |
| Omaha | 1024 kb | | 1536 kb | |
| Scottsdale | 256 kb | 128 kb | 256 kb | 128 kb |
| Wash DC | 256 kb | 128 kb | 256 kb | 128 kb |
| Denver | 256 kb | 192 kb | NA | NA |
| Little Rock | 192 kb | 128 kb | 192 kb | 96 kb |
| Atlanta | 256 kb | 128 kb | 256 kb | 128 kb |
| Kansas City | 192 kb | 128 kb | 192 kb | 96 kb |
| Pasadena | 128 kb | 128 kb | 192 kb | 128 kb |
| Newport Beach | 192 kb | 128 kb | 192 kb | 96 kb |
| Richmond | 256 kb | 128 kb | 192 kb | 96 kb |
| Ok City | 128 kb | 128 kb | 128 kb | 32 kb |
| Lincoln | 192 kb | 128 kb | 192 kb | 64 kb |
| Chicago | 256 kb | 128 kb | 256 kb | 192 kb |
| PPP to Denver | NA | NA | 512 kb | 0 |

The following table represents the *secondary* PVC CIR values created from each remote location to Denver.

| From Denver to following: | Current | | Recommended | |
|---|---|---|---|---|
| | Port Speed | CIR | Port Speed | CIR |
| Omaha | NA | NA | NA | NA |
| Scottsdale | 256 kb | 128 kb | 256 kb | 64 kb |
| Wash DC | 256 kb | 128 kb | 256 kb | 128 kb |
| Denver | 256 kb | 192 kb | NA | NA |
| Little Rock | 192 kb | 128 kb | 192 kb | 64 kb |
| Atlanta | 256 kb | 128 kb | 256 kb | 64 kb |
| Kansas City | 192 kb | 128 kb | 192 kb | 64 kb |
| Pasadena | 128 kb | 128 kb | 192 kb | 64 kb |
| Newport Beach | 192 kb | 128 kb | 192 kb | 64 kb |
| Richmond | 256 kb | 128 kb | 192 kb | 64 kb |
| Ok City | 128 kb | 128 kb | 128 kb | 64 kb |
| Lincoln | 192 kb | 128 kb | 192 kb | 32 kb |
| Chicago | 256 kb | 128 kb | 256 kb | 64 kb |
| PPP to Omaha | NA | NA | 512 kb | 0 |

The CIR values for the secondary paths represent arbitrary values intended to support Internet traffic, and provide connectivity to other locations (should either Denver or Omaha location fail) at some reduced level of service.

The primary advantage to this approach is the disaster recovery functions, and the splitting of Internet traffic from the remote locations to Omaha and Denver thus off-loading the PVCs to Omaha. The disadvantage is the cost of the additional network facilities, and, a slightly more complex network from a route management perspective.

Obviously, many other combinations exist that would provide the services expected by the Hallmark users.

# Appendices

Appendix A: Router Configuration Summary

# Appendix A: Router Configuration Summary

## *Omaha 4500 Router*

```
Device Summary (As of 9/11/00 )

Node: 10.15.172.1
SysName: Omaha[HQ]    Last Reset: 8/22/99 11:34:22 PM
Location:
Contact:
System Descr:
  Cisco Internetwork Operating System Software <crlf>IOS (tm) 45
  00 Software (C4500-DS-M), Version 11.2(17)P,  RELEASE SOFTWARE
   (fc1) <crlf>Copyright (c) 1986-1999 by cisco Systems, Inc.<cr
  lf>Compiled Tue 12-Jan-99 14:42 by pwade
Object Id: 1.3.6.1.4.1.9.1.14

Interface Summary:
 Interface: 1   Ethernet0  Status: up/up
   Ethernet   MAC:00E01E4E0B91  Speed: 10,000,000   MTU: 1500
   IP Addr: 10.15.171.1      255.255.255.0     Max Reas: 18,024

 Interface: 2   Ethernet1  Status: up/up
   Ethernet   MAC:00E01E4E0B94  Speed: 10,000,000   MTU: 1500
   IP Addr: 10.15.172.1      255.255.252.0     Max Reas: 18,024

 Interface: 3   Serial0  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,024,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     110  Active  9/11/00 10:30:22 AM     0
     120  Active  9/11/00 10:30:22 AM     0
     130  Active  9/11/00 10:30:22 AM     0
     141  Active  9/11/00 10:30:22 AM     0
     150  Active  9/11/00 10:30:22 AM     0
     160  Active  9/11/00 10:30:22 AM     0
     170  Inact   9/11/00 10:30:22 AM     0
     180  Active  9/11/00 10:30:22 AM     0
     190  Invalid 9/11/00 10:30:22 AM     0
     200  Active  9/11/00 10:30:22 AM     0
     210  Active  9/11/00 10:30:22 AM     0
     220  Active  9/11/00 10:30:22 AM     0
     230  Active  9/11/00 10:30:22 AM     0
     240  Active  9/11/00 10:30:22 AM     0
     896  Active  9/11/00 10:30:22 AM     0


 Interface: 4   Serial1  Status: down/down
   Propri-PPP-Serial   MAC: (na)  Speed: 1,544,000   MTU: 1500

 Interface: 5   Serial0.1  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.5      255.255.255.252   Max Reas: 18,024

 Interface: 6   Serial0.2  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.9      255.255.255.252   Max Reas: 18,024

 Interface: 7   Serial0.3  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.13     255.255.255.252   Max Reas: 18,024

 Interface: 8   Serial0.4  Status: up/up
```

```
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.17     255.255.255.252   Max Reas: 18,024

Interface: 9   Serial0.5  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.21     255.255.255.252   Max Reas: 18,024

Interface: 10   Serial0.6  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.25     255.255.255.252   Max Reas: 18,024

Interface: 11   Serial0.7  Status: up/down
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.41     255.255.255.252   Max Reas: 18,024

Interface: 12   Serial0.8  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.45     255.255.255.252   Max Reas: 18,024

Interface: 13   Serial0.9  Status: up/down
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.33     255.255.255.252   Max Reas: 18,024

Interface: 14   Serial0.10  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.37     255.255.255.252   Max Reas: 18,024

Interface: 15   Serial0.11  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.29     255.255.255.252   Max Reas: 18,024

Interface: 16   Serial0.12  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.49     255.255.255.252   Max Reas: 18,024

Interface: 17   Serial0.13  Status: up/up
  FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
    DLCI State   Last Chg                CIR  Destination
    ---- ------- ----------------------- ---- ---------------
    (none)

  IP Addr: 10.15.160.53     255.255.255.252   Max Reas: 18,024
```

```
 Interface: 18   Serial0.14  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,536,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.57     255.255.255.252   Max Reas: 18,024

 Interface: 19   Serial0.15  Status: up/up
   FrameRelay   MAC: (na)  Speed: 1,024,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.61     255.255.255.252   Max Reas: 18,024

Routes: 38
 (First 10000 routes only)
 Destination     Netmask         Next Hop         HopCnt    IF   Prot Type
 -----------     -------         --------         ------    ---- ---- -----
 0.0.0.0         0.0.0.0         10.15.172.2      0         0    loc  indir
 10.1.1.1        255.255.255.255 10.15.160.54     20        17   ospf indir
 10.15.160.4     255.255.255.252 10.15.160.5      0         5    loc  dir
 10.15.160.8     255.255.255.252 10.15.160.9      0         6    loc  dir
 10.15.160.12    255.255.255.252 10.15.160.13     0         7    loc  dir
 10.15.160.16    255.255.255.252 10.15.160.17     0         8    loc  dir
 10.15.160.20    255.255.255.252 10.15.160.21     0         9    loc  dir
 10.15.160.24    255.255.255.252 10.15.160.25     0         10   loc  dir
 10.15.160.28    255.255.255.252 10.15.160.29     0         15   loc  dir
 10.15.160.36    255.255.255.252 10.15.160.37     0         14   loc  dir
 10.15.160.44    255.255.255.252 10.15.160.45     0         12   loc  dir
 10.15.160.48    255.255.255.252 10.15.160.49     0         16   loc  dir
 10.15.160.52    255.255.255.252 10.15.160.53     0         17   loc  dir
 10.15.160.56    255.255.255.252 10.15.160.57     0         18   loc  dir
 10.15.160.60    255.255.255.252 10.15.160.61     0         19   loc  dir
 10.15.160.254   255.255.255.255 10.15.160.50     66        16   ospf indir
 10.15.161.0     255.255.255.0   10.15.160.14     20        7    ospf indir
 10.15.162.0     255.255.255.0   10.15.160.30     20        15   ospf indir
 10.15.164.0     255.255.255.0   10.15.160.38     20        14   ospf indir
 10.15.165.0     255.255.255.0   10.15.160.18     20        8    ospf indir
 10.15.167.0     255.255.255.0   10.15.160.46     20        12   ospf indir
 10.15.168.0     255.255.255.0   10.15.160.26     20        10   ospf indir
 10.15.169.0     255.255.255.0   10.15.160.10     20        6    ospf indir
 10.15.170.0     255.255.255.0   10.15.160.22     20        9    ospf indir
 10.15.171.0     255.255.255.0   10.15.171.1      0         1    loc  dir
 10.15.172.0     255.255.252.0   10.15.172.1      0         2    loc  dir
 10.59.224.0     255.255.255.0   10.15.160.54     20        17   ospf indir
 10.59.225.0     255.255.255.0   10.15.160.50     20        16   ospf indir
 10.59.226.0     255.255.255.0   10.15.160.58     20        18   ospf indir
 10.59.232.0     255.255.255.0   10.15.160.62     20        19   ospf indir
 10.100.1.20     255.255.255.255 10.15.172.7      1         2    rip  indir
 10.100.1.197    255.255.255.255 10.15.172.7      1         2    rip  indir
 10.109.76.1     255.255.255.255 10.15.172.7      1         2    rip  indir
 10.109.76.254   255.255.255.255 10.15.172.7      1         2    rip  indir
 138.12.0.0      255.255.0.0     10.15.172.7      1         2    rip  indir
 198.62.193.0    255.255.255.0   10.15.160.6      0         0    loc  indir
 207.25.176.0    255.255.255.0   10.15.172.7      1         2    rip  indir
 216.141.32.0    255.255.252.0   10.15.160.6      0         0    loc  indir

ARP Table Enties: 333
```

## *Denver 2501 Router*

```
Node: 10.15.160.10
SysName: Denver    Last Reset: 8/24/99 7:23:42 AM
Location:
Contact:
System Descr:
  Cisco Internetwork Operating System Software <crlf>IOS (tm) 25
  00 Software (C2500-DS-L), Version 11.2(17), RELEASE SOFTWARE (
  fc1)<crlf>Copyright (c) 1986-1999 by cisco Systems, Inc.<crlf>
  Compiled Mon 04-Jan-99 17:13 by ashah
Object Id: 1.3.6.1.4.1.9.1.17

Interface Summary:
 Interface: 1   Ethernet0  Status: up/up
   Ethernet   MAC:00E0B05B2295  Speed: 10,000,000   MTU: 1500
   IP Addr: 10.15.169.1      255.255.255.0     Max Reas: 18,024

 Interface: 2   Serial0  Status: up/up
   FrameRelay   MAC: (na)  Speed: 256,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     100  Active  6/28/00 3:45:34 AM      0


 Interface: 3   Serial1  Status: down/down
   Propri-PPP-Serial   MAC: (na)  Speed: 1,544,000   MTU: 1500

 Interface: 4   Serial0.1  Status: up/up
   FrameRelay   MAC: (na)  Speed: 256,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.10     255.255.255.252   Max Reas: 18,024

Routes: 32
 (First 10000 routes only)
 Destination     Netmask         Next Hop         HopCnt    IF   Prot Type
 -----------     -------         --------         ------    ---- ---- -----
 0.0.0.0         0.0.0.0         10.15.160.9      1         4    ospf indir
 10.1.1.1        255.255.255.255 10.15.160.9      20        4    ospf indir
 10.15.160.4     255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.8     255.255.255.252 10.15.160.10     0         4    loc  dir
 10.15.160.12    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.16    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.20    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.24    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.28    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.36    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.44    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.48    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.52    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.56    255.255.255.252 10.15.160.9      455       4    ospf indir
 10.15.160.60    255.255.255.252 10.15.160.9      487       4    ospf indir
 10.15.160.254   255.255.255.255 10.15.160.9      456       4    ospf indir
 10.15.161.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.162.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.164.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.165.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.167.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.168.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.169.0     255.255.255.0   10.15.169.1      0         1    loc  dir
 10.15.170.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.171.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.15.172.0     255.255.252.0   10.15.160.9      20        4    ospf indir
 10.59.224.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.59.225.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.59.226.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 10.59.232.0     255.255.255.0   10.15.160.9      20        4    ospf indir
 198.62.193.0    255.255.255.0   10.15.160.9      20        4    ospf indir
 216.141.32.0    255.255.252.0   10.15.160.9      20        4    ospf indir

ARP Table Enties: 129
```

## *Newport Beach 2501 Router*

```
Node: 10.15.160.14
SysName: NewportBeach    Last Reset: 10/28/99 11:20:37 PM
Location:
Contact:
System Descr:
  Cisco Internetwork Operating System Software <crlf>IOS (tm) 25
  00 Software (C2500-DS-L), Version 11.2(17), RELEASE SOFTWARE (
  fc1)<crlf>Copyright (c) 1986-1999 by cisco Systems, Inc.<crlf>
  Compiled Mon 04-Jan-99 17:13 by ashah
Object Id: 1.3.6.1.4.1.9.1.17

Interface Summary:
 Interface: 1   Ethernet0  Status: up/up
   Ethernet   MAC:00E0B05B2244  Speed: 10,000,000   MTU: 1500
   IP Addr: 10.15.161.1      255.255.255.0     Max Reas: 18,024

 Interface: 2   Serial0  Status: up/up
   FrameRelay   MAC: (na)  Speed: 192,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     100  Active  9/9/00 7:12:14 PM       0


 Interface: 3   Serial1  Status: down/down
   Propri-PPP-Serial   MAC: (na)  Speed: 1,544,000   MTU: 1500

 Interface: 4   Serial0.1  Status: up/up
   FrameRelay   MAC: (na)  Speed: 192,000   MTU: 1500
     DLCI State   Last Chg                CIR  Destination
     ---- ------- ----------------------- ---- ---------------
     (none)

   IP Addr: 10.15.160.14     255.255.255.252   Max Reas: 18,024

Routes: 32
 (First 10000 routes only)
 Destination     Netmask         Next Hop         HopCnt    IF   Prot Type
 -----------     -------         --------         ------    ---- ---- -----
 0.0.0.0         0.0.0.0         10.15.160.13     1         4    ospf indir
 10.1.1.1        255.255.255.255 10.15.160.13     20        4    ospf indir
 10.15.160.4     255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.8     255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.12    255.255.255.252 10.15.160.14     0         4    loc  dir
 10.15.160.16    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.20    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.24    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.28    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.36    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.44    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.48    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.52    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.56    255.255.255.252 10.15.160.13     585       4    ospf indir
 10.15.160.60    255.255.255.252 10.15.160.13     617       4    ospf indir
 10.15.160.254   255.255.255.255 10.15.160.13     586       4    ospf indir
 10.15.161.0     255.255.255.0   10.15.161.1      0         1    loc  dir
 10.15.162.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.164.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.165.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.167.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.168.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.169.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.170.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.171.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.15.172.0     255.255.252.0   10.15.160.13     20        4    ospf indir
 10.59.224.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.59.225.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.59.226.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 10.59.232.0     255.255.255.0   10.15.160.13     20        4    ospf indir
 198.62.193.0    255.255.255.0   10.15.160.13     20        4    ospf indir
 216.141.32.0    255.255.252.0   10.15.160.13     20        4    ospf indir

ARP Table Enties: 0
```