



Delivering Predictable Host Integration Services

Delivering Predictable Host Integration Services

Consolidation of different network traffic types onto the same infrastructure means that sophisticated, interactive, mission-critical applications share resources with non-mission-critical data, bandwidth-hungry applications (such as multicast video), and delay-sensitive applications (such as voice). Each application type uses the network differently, has different priority within the corporation, and is often targeted to a subset of the employees. Furthermore, today's tiered network designs include points where media speeds converge, such as where a gigabit Ethernet backbone link feeds a 10-Mbps wiring closet link.

In these environments quality of service (QoS) is required to move traffic according to the needs of the business and prevent congestion and subsequent session loss. QoS is defined as those mechanisms that give network managers the ability to control the mix of bandwidth, delay, jitter, and packet loss in the network. QoS achieves these goals by using tools to manage network congestion, shape network traffic, use expensive WAN links more efficiently, and set traffic policies end to end from the application across the network to the client.

This paper describes a set of features and tests that can be used to provide QoS from an IBM S/390® Parallel Enterprise Server™ through a Cisco-routed network. These tests were a joint Cisco - IBM effort, and both companies have published this white paper.

QoS Technologies

Providing end-to-end QoS requires a toolbox of QoS technologies:

- *QoS signaling* requests a level of service from the network. Applications, clients, and network servers can tag packets to identify the level of priority required for that packet.
- *Congestion management* tools use queuing algorithms to sort traffic and determine a method of priority to schedule traffic to output links based on QoS signaling, which is set by the application, client, or network server.
- *Traffic shaping and policing* tools manage traffic and congestion between the QoS domains—for example, where bandwidth speeds change.
- *Link efficiency* tools improve service-level efficiency and predictability. Examples include traffic compression for all traffic types and silence suppression for voice traffic.
- *Congestion avoidance* tools monitor network traffic loads and take actions to avoid congestion at common network bottleneck points. Often this includes discarding packets to avoid congestion.

White Paper

- *Bandwidth schedulers* allot bandwidth for each queue or session.

QoS is not a device feature or a link-layer feature set. QoS is an end-to-end system architecture that stretches from the application to the end user. The right QoS solution includes a variety of technologies that interoperate to give you scalable, media-independent services across your LAN and WAN. Correctly implemented, QoS architectures can play a significant role in lowering costs, enabling end-to-end service-level guarantees, and providing optimal application performance.

As a first step in implementing QoS, packets must be classified (for example, as high priority, low delay or low priority, moderate delay). Although a human must make the actual classification decision, that decision can be implemented by the application, the server in which the application resides, or in a networking device such as a router.

Implementing QoS classification can be as simple as an application setting the IP precedence bits in all the packets sent. Alternatively, a policy manager running in a server can set IP precedence as appropriate for the applications running in that server, the time of day, the individual end user, or any of a variety of conditions that may be known only to that server—overriding that which was set by the application if necessary. Traffic can also be classified in the network itself, according to network-wide policy based on congestion or multiservice traffic requirements. Regardless of whether the application, server, or network chooses classification, the packet is marked to provide QoS signaling. Without packet marking, every router in the network would have to incur additional overhead to make a policy decision.

It is common for a network to identify priority traffic based on a physical IP address or TCP port number. With the advent of intranets and multiple Web-based applications on a single server, you need to differentiate application flows that have the same physical address. Application-aware QoS distinguishes mission-critical Web traffic from standard Web traffic and preferentially expedites one flow over the other. Sophisticated QoS mechanisms inspect packet payloads at the connection level to identify application subflows for greater control. For instance, it might be desirable to give priority to an SAP financial transaction over an SAP print job. As you can see from the discussion above, where QoS signaling occurs is dependent on a variety of factors. At different times and under different conditions the placement of QoS signaling may vary. Establishing an overall QoS policy will determine when and where the decision is made.

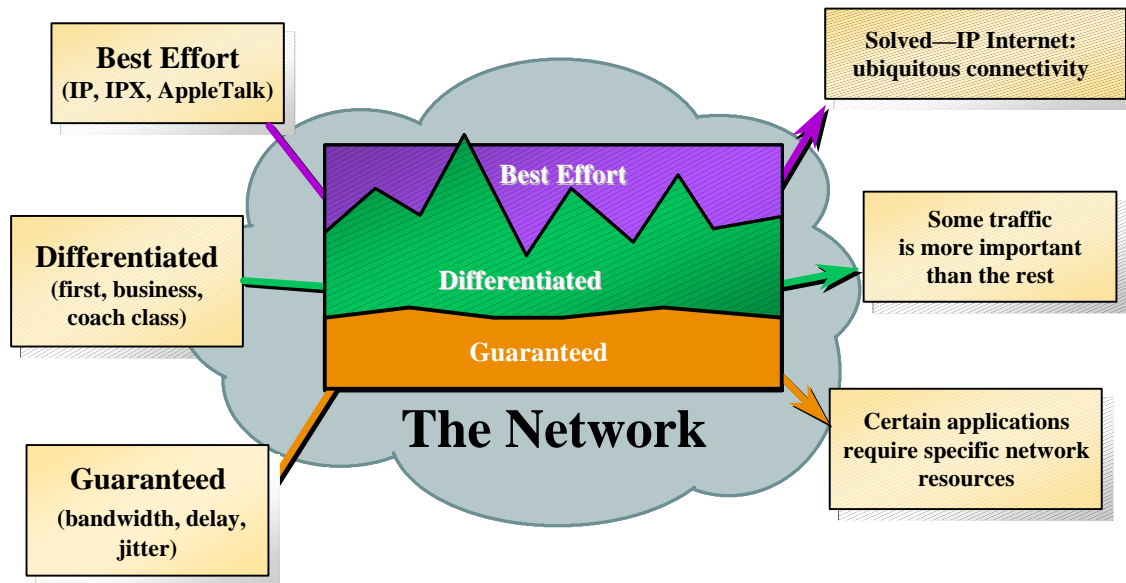
QoS Service Levels

Service levels refer to the actual end-to-end QoS capabilities, meaning the ability of a network to deliver service needed by specific network traffic from end to end or edge-to-edge. The QoS services differ in their level of “QoS strictness,” which describes how tightly the service is bound by specific bandwidth, delay, jitter, and loss characteristics.

White Paper

There are three basic levels of end-to-end QoS service that can be provided across a heterogeneous network, as shown in Figure 1: best effort, differentiated, and guaranteed service.

Figure 1: End-to-End QoS Service Levels



- *Best effort service* is better known as “lack of QoS” and provides basic queuing during congestion with first-in, first-out (FIFO) packet delivery on the link. There are no priorities or guarantees. Examples of this type of traffic include low-priority e-mail and bulk data transfers.
- *Differentiated service* is also called “qualitative QoS.” In this case some traffic is treated better than the rest (faster handling, more bandwidth, lower loss rate). This service, however, gives statistical preference rather than hard and fast guarantees. Examples of this type of traffic could be mission-critical interactive applications.
- *Guaranteed service* is also called “quantitative QoS.” In this case network resources are reserved for specific traffic. This service guarantees that adequate bandwidth is available for applications. This type of service is for delay-sensitive traffic, such as voice and video.

Many networks employ all three types of service, depending on the application requirements. If the network is not carrying voice and video traffic, differentiated service is most likely adequate. This service must be carefully analyzed because, in most cases, deploying guaranteed service is more costly in bandwidth and tools than deploying differentiated services.

This paper deals with a set of differentiated service tests performed in a data environment, using specific standards and vendor-unique products to provide the required QoS.

QoS Standards

While individual vendors may implement certain aspects of QoS architectures using a proprietary solution, developing an end-to-end architecture requires a set of standards that are consistently implemented among vendors. If a server vendor signals QoS requirements in a way that is not understood by the congestion management platform, effective QoS cannot be delivered.

There are two key standards for QoS signaling when providing differentiated services in an IP environment: differentiated services and IP precedence settings. Both specifications are defined in a request for comment (RFC) from the Internet Engineering Task Force (IETF). Both standards provide a way to identify QoS in an IP environment using a field in the IP header.

The IP precedence settings were defined in the original IP protocol specification in September 1981 as part of RFC 791. Within the IP header an eight-bit type of service (ToS) was defined. This field was defined as a setting of parameters used to guide the selection of service parameters when transmitting a packet across the network. The field is shown in Figure 2.

Figure 2: IP Precedence Field

0	1	2	3	4	5	6	7
Precedence			Delay	Throughput	Reliability	Reserved	

The values for the delay, throughput, and reliability fields are identified as follows:

- For delay, bit three is set to normal delay ('0'b) or low delay ('1'b).
- For throughput, bit four is set to normal throughput ('0'b) or high throughput ('1'b).
- For reliability, bit five is set to normal reliability ('0'b) or high reliability ('1'b).


Precedence mappings signal QoS to congestion control technologies. With b"111" as the highest priority traffic and b"000" as lowest priority, the receiving network devices then know how to put the traffic in queues and release the traffic from the queues into the network. The precedence mappings are defined as follows:

- '111'b is network control
- '110'b is internetwork control
- '101'b is CRITIC/ECP
- '100'b is flash override
- '011'b is flash
- '010'b is immediate
- '001'b is priority
- '000'b is routine

White Paper

After RFC 791, RFC 2472 was submitted to the IETF in December 1998 and defines the differentiated services field, which overlays the ToS field in IPV4 and the traffic class field in IPV6. The purpose of defining this new standard was to “enable scalable service discrimination in the Internet without the need for per-flow state and signaling at every hop.” Like the IP precedence bits, when the differentiated services bits are set as part of QoS signaling, the network devices can determine how packets are forwarded based on these settings (congestion control). The differentiated services field is shown in Figure 3.

Figure 3: Differentiated Services Field in the IP Header



DSCP: Differentiated services code point

CU: Currently unused

The six DSCP bits are used to select per-hop behavior (PHB)—the queuing and prioritization applied at a differentiated services-compliant node. The CU field corresponds to the reserved field in the ToS header. Using a six-bit field instead of a three-bit field enables a maximum of 64 classes of service instead of the eight classes provided by the IP precedence field. The six-bit field provides more granularity in prioritizing traffic, although not all of the 64 options have completed the standards process at this time. An initial 32 code points will be identified ('xxxxx0'b) for prioritization, while the remaining 32 code points are designated as experimental, local, or waiting for future standardization. At this time, only the first eight code points ('xxx000'b) have been identified as corresponding to the original eight IP precedence settings.

Both standards ensure that QoS signaling produces the desired result when the packet is queued in the network.

Cisco and IBM QoS Testing

Cisco and IBM conducted a joint test with IBM S/390 Enterprise Servers, Cisco networking routers and switches, and IBM Host Integration desktop products to demonstrate the benefits of prioritizing IP traffic end-to-end. Prioritizing IP traffic provides the predictable IP performance required by most service-level agreements and meets the demands of today's e-business applications. The ability of the S/390 Enterprise Server to assign application traffic priority coupled with the ability of Cisco routers and switches to enforce QoS throughout the network ensures users receive predictable response times even when operating in mixed-traffic environments.

IBM and Cisco developed their prioritization techniques based on the open IETF standards discussed earlier in this document. The testing validated that the S/390 QoS signaling and the Cisco congestion management tools were indeed compatible and that

White Paper

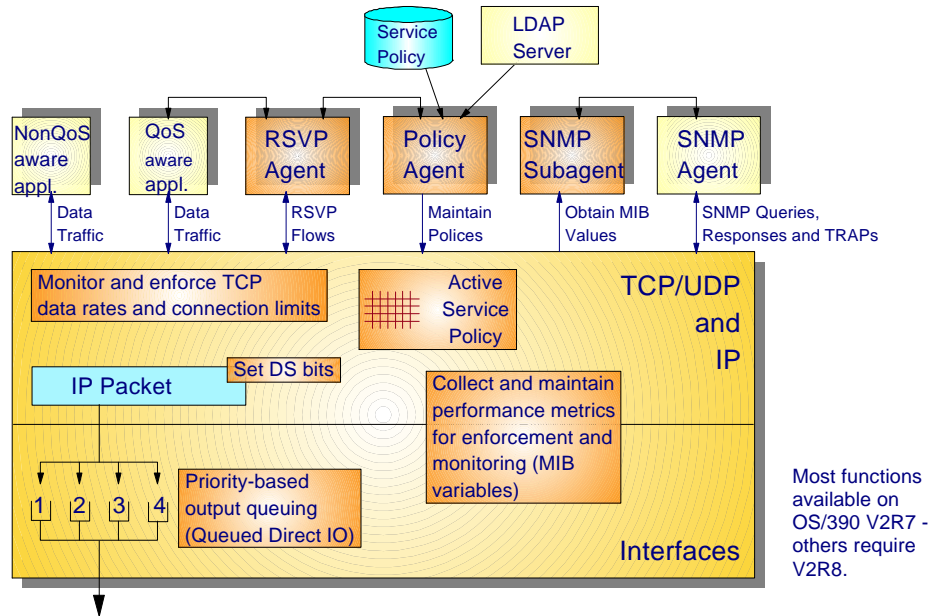
the benefits of consistent response time and effective bandwidth utilization could be realized. The key components of the test were the IBM Policy Agent for the Operating System for S/390® (OS/390®), which provided the QoS signaling, and the Cisco Weighted Fair Queuing (WFQ) algorithm in the Cisco IOS software, which provided congestion management.

IBM Policy Agent

Not all applications can establish priority, and a priority established by an application can conflict with corporate policy. OS/390 Version 2 Release 7 (V2R7) provides a Policy Agent that enables the system programmer to overrule or establish priority for any application. As the source (and destination) of application traffic, a server has full knowledge of all data flows and is therefore effective and efficient in setting QoS service levels. System administrators can provide QoS signaling at the application source to interface with congestion management technologies in the network. This approach improves network performance by eliminating the need to classify the traffic explicitly at each WAN interface in the core or backbone network, avoiding per-packet processing. Defining and setting policy at the source also avoids problems with client and application identification when security encryption hides the original IP header. The S/390 Enterprise Server enables IP security mechanisms, such as IPSec, to be used end to end—from the application server to the end client.

IBM Policy Agent QoS definitions are extensive and can be specified based on application type, individual user, user group, time of day, and day of week. Policy Agent service policy rules and categories are stored either in an LDAP directory or in a local configuration file. These policy rules classify S/390 traffic for a *connection* that generally consists of the application type (TCP/UDP port numbers), host addresses, address groups (source/destination IP addresses), IP protocol ID (TCP/UDP), and route information (inbound/outbound interface). They also contain information about their period of validity. The integration of the Policy Agent into OS/390 is shown in Figure 4.

Figure 4: QoS Components in OS/390



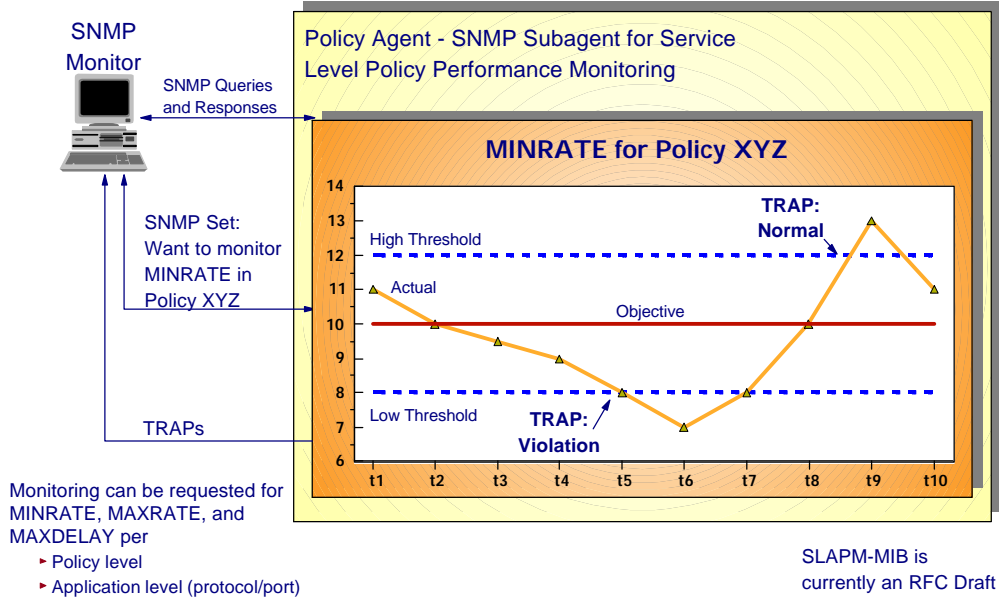
For each policy rule, a corresponding service category defines the appropriate QoS. A service category generally contains the priority of the traffic, the minimum and maximum TCP connection throughput, and the number of TCP connections allowed at any given time. The S/390 provides this per-connection bandwidth management function through TCP window manipulation. The per-connection level of granularity provided by the S/390 can be used to “fine-tune” the aggregate bandwidth management provided by the router.

In addition, network interface priority output queuing could be applied when Gigabit Ethernet or other Queued Direct Input Output (QDIO)-capable S/390 adapters are used. The IP precedence value in the IP header is used by the S/390 to route packets to appropriate priority queues within the QDIO adapter for transmission out into the network. Network devices can then use these values to enforce traffic prioritization as the packets traverse the network to their destination.

A S/390 Service Level Agreement Management Information Base (SLA MIB) enables performance monitoring of the service policy. Available in OS/390 V2R8, the SLA MIB monitors MINRATE, MAXRATE, and MAXDELAY at the policy level or at the application (protocol/port) level, as shown in Figure 5. Performance thresholds can be set to send traps to notify the Simple Network Management Protocol (SNMP) manager of deviations. This SLA MIB can also be effectively utilized to gather accounting and billing data.

White Paper

Figure 5: SLA Performance Monitoring - MIB



Policy Agent enhancements available with OS/390 V2R8 improve real-time QoS enforcement. In V2R8, Policy Agent automatically updates policy information whenever it is changed. The Policy Agent monitors changes to policy definitions by re-reading the local configuration file and by polling the defined LDAP server where the policy may be stored.

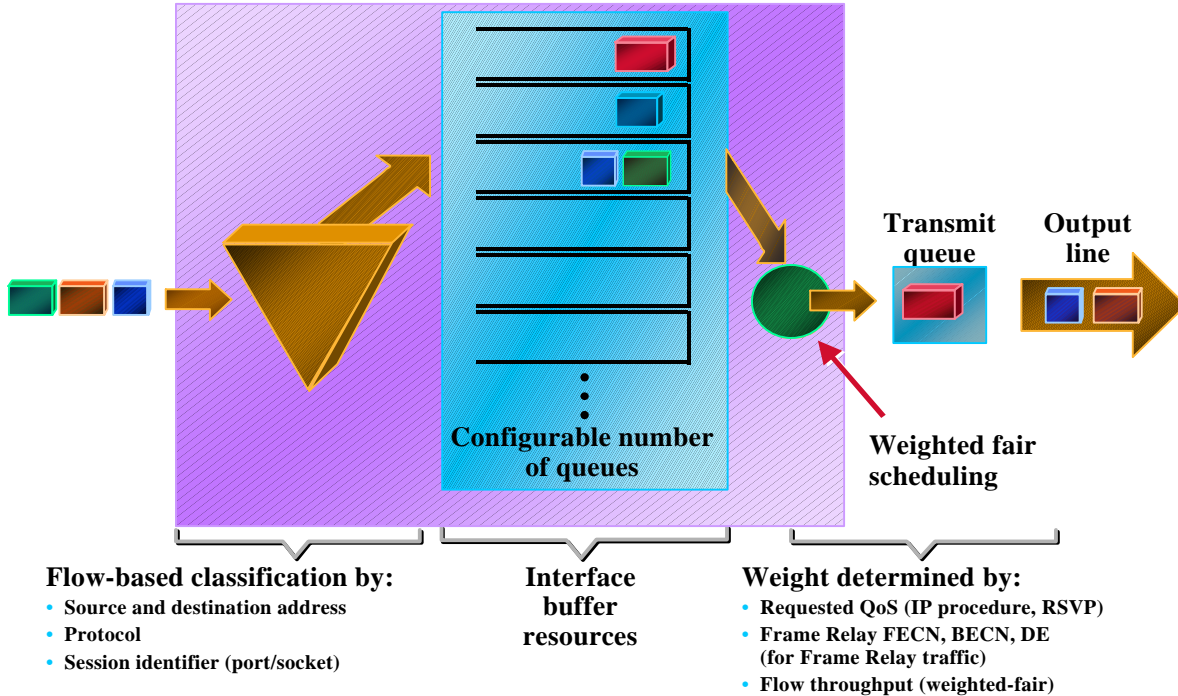
Through its support of standard-based differentiated services and IP precedence setting, the S/390 Policy Agent provides versatile application-level control of policy settings and granular per-connection bandwidth management. Used in conjunction with the traffic prioritization and enforcement functions provided by Cisco network devices, a powerful end-to-end QoS network solution can be implemented to provide more predictable access to mission-critical applications residing on the S/390 Enterprise Server.

Policy Agent was included as a software component of OS/390 V2R7 and was used for this test. Relevant portions of the configuration for OS/390 and the Policy Agent are shown in appendixes B, C, and D.

Cisco WFQ

WFQ is one of the premier Cisco queuing techniques. It is a flow-based queuing algorithm that simultaneously schedules interactive traffic to the front of the queue to reduce response time and fairly shares the remaining bandwidth between high-bandwidth flows. WFQ ensures that queues do not starve for bandwidth and that traffic gets predictable service. Low-volume traffic, which is the majority of traffic, receives preferential service. High-volume traffic shares the remaining capacity proportionally, as shown in Figure 6.

Figure 6: WFQ



WFQ minimizes configuration effort and adapts automatically to changing network traffic conditions. It is the default queuing mode on most serial interfaces configured to run at or below T1/E1 speeds. WFQ efficiently uses whatever bandwidth is available to forward traffic from lower priority flows if no traffic from higher priority flows is present.

WFQ is IP precedence and differentiated services-aware, which means that it is able to detect higher priority packets from the QoS signaling fields and can schedule them faster, providing superior response time for this traffic. Using the IP precedence field, for example, it has values between 0 (the default) and 7. As the precedence value increases, the algorithm allocates more bandwidth to that connection to make sure that it gets served more quickly when congestion occurs. WFQ assigns a weight to each flow, which determines the transmission order for queued packets. In this scheme, lower weights are served first. IP precedence serves as a divisor to this weighting factor. For instance, traffic with an IP precedence field value of 7 gets a lower weight than traffic with an IP precedence field value of 3 and, thus, has priority in the transmission order.

Configuration for the router using WFQ is shown in Appendix A.

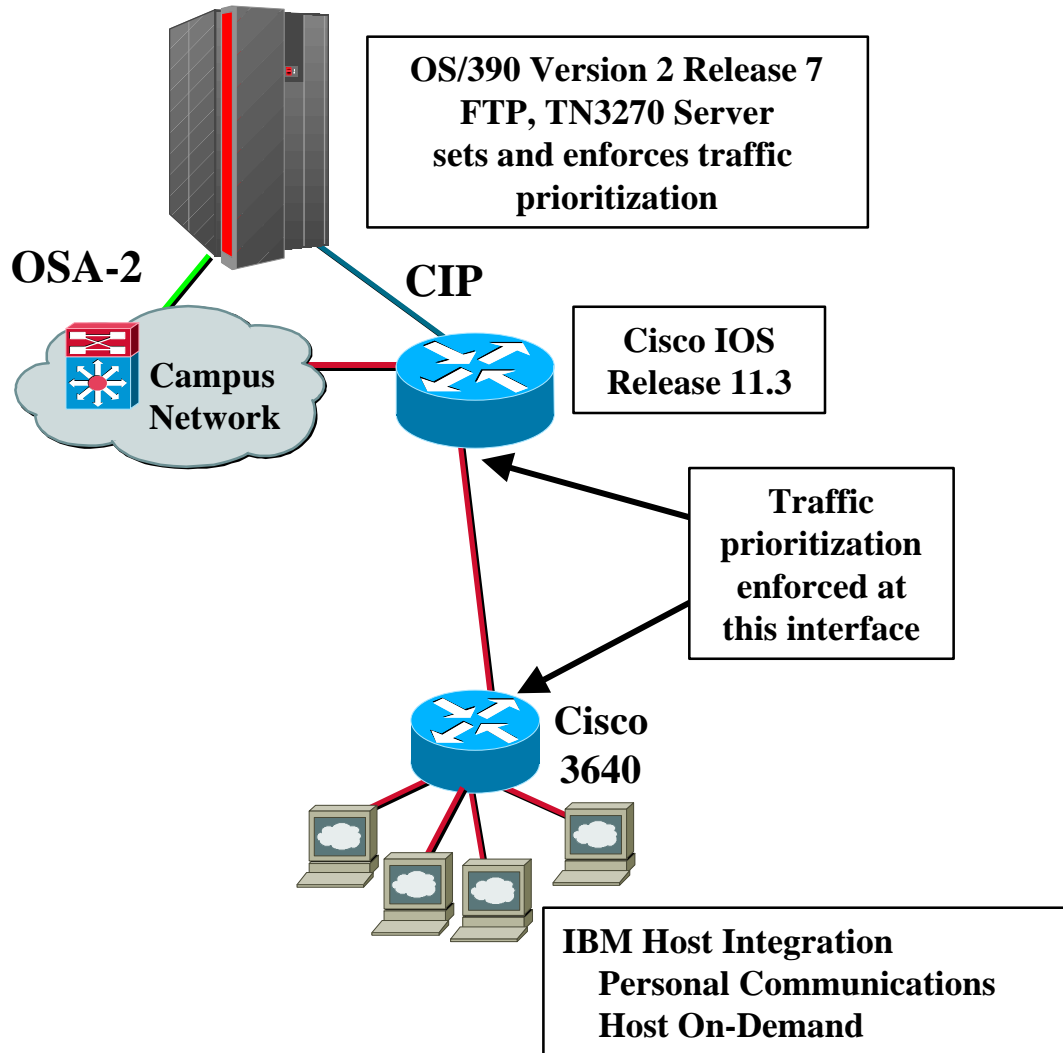
Test Configuration

When designing and setting up this test, the objective was to demonstrate the ability to prioritize mission-critical S/390 traffic through a Cisco routed network when there is severe WAN congestion. The test configuration is shown in Figure 7. The key elements

White Paper

are in the software of the S/390 Enterprise Server and the Cisco 7513 and 3640 routers. OS/390 V2R7 contains the Policy Agent that sets the IP precedence bits for traffic prioritization. In the Cisco routers, Cisco IOS Release 11.3 contains the WFQ prioritization solution for congestion management.

Figure 7: QoS Test Configuration



A WAN bandwidth of 256 kbps was chosen as a compromise between high- and low-speed links, but similar results can be achieved at any bandwidth. File transfers between a File Transfer Protocol (FTP) server and client were used to create congestion on the WAN link. The network consisted of a pair of routers. For the purposes of this test, we presumed that there was a very heavy bias toward traffic flowing from the Cisco 7513 router, representing a headquarters or data center location, toward the Cisco 3640 router, representing a medium- to large-branch location.

White Paper

After the file transfers created network congestion, higher priority interactive traffic was added on the network link. The TN3270e server provided by Communications Server for OS/390 Version 2 Release 7 and IBM Personal Communications Version 4 Release 3 and Host-on-Demand Version 4 access products generated the interactive traffic. Message sizes were set at 1400 bytes out from the data center and 100 bytes in from the branch. These values were chosen to represent a typical TN3270 transaction profile. Chariot MVS Endpoint Software Version 3 Release 3 from Ganymede Software Inc. provided additional TN3270e clients and accurate response time measurements.

Attachment to the S/390 Enterprise Server was provided by the S/390 Open Systems Adapter Version 2 (OSA-2) and the Cisco Channel Interface Processor (CIP) over an Enterprise Systems Connection (ESCON®) channel.

Test Scenarios

Test parameters were varied, and multiple scenarios were run. The S/390 Policy Agent set TN3270 priority based on the differentiated services and IP precedence bit definitions. (Because there were only the two applications, varying TN3270 priority levels was not required. In a corporate network each application and user connection could have a different priority level.) TN3270 priority effects were measured with FIFO and WFQ in the Cisco routers. Measurements were taken for both the CIP/ESCON and OSA-2 Fast Ethernet data paths. Each test configuration was established and allowed to run for five minutes. TN3270 response times, TN3270 transaction rates, FTP throughputs, and FTP response times were measured. Sample console displays from the router demonstrate the traffic prioritization and are shown in Appendix E.

By far the most important metric for transaction processing systems is response time. Response time is a measure of how long terminal operators must wait, from the time that they press the key that initiates a transaction, until the results are returned to their screens. Sophisticated systems and technologies have been developed to ensure that the portion of response time spent performing the requested work on OS/390 host systems is minimized. It is equally important that the portion of response time that the information is traveling through the network is minimized. Of equal importance, transactions should complete within the same amount of time, independent of other loads on the network and host systems. In the tests an application simulation tool was used: Chariot from Ganymede Software, which is designed to eliminate any extraneous host activity, so that only network time is measured. The types and amount of traffic generated by Chariot are shown in Appendix F.

Eight tests were run and the parameters and results of each are shown in Appendix G:

- Differentiated services priority settings, FIFO queuing, CIP
- Differentiated services priority settings, FIFO queuing, OSA-2
- IP precedence priority settings, FIFO queuing, CIP
- IP precedence priority settings, FIFO queuing, OSA-2
- Differentiated services priority settings, WFQ, CIP

White Paper

- Differentiated services priority settings, WFQ, OSA-2
- IP precedence priority settings, WFQ, CIP
- IP precedence priority settings, WFQ, OSA-2

The only changes to either of the routers from one test to the next was to turn on or off WFQ on the Cisco 7513 router and to lower the queue depth available in hardware buffers on the output interface of the Cisco 7513 router.


Test Results

The test demonstrated that establishing application priority and using WFQ in the network provided substantial improvements in response times and increased transaction rates. There was nearly a fivefold reduction in response times from an average of 4.7 seconds to an average of 0.95 seconds¹. A comparable fivefold increase in transactions accompanied the response time reduction. This improvement was with both the CIP and the OSA-2 data paths.

Figure 8 shows the results of the eight different tests. All of the tests were run with a constant load of background traffic made up of 30 simulated file transfers from the OS/390 side of the network toward the client side. The eight lines on the graph show the response time results for the interactive traffic only; the batch traffic is not included.

The four lines that track between three and seven seconds show the measured response time for the interactive transactions when forced to compete for network bandwidth on a FIFO basis. The four lines that cluster near one second show the measured response time for the same interactive transactions when they are marked for priority handling by the Policy Agent in OS/390 and given priority handling by the WFQ algorithm in the router.

Figure 8: Summary Graph of Response Times



¹ The performance data was jointly collected by IBM and Cisco in the dedicated system environment described in this white paper. Results obtained in other configurations or operating system environments may vary.

White Paper

The combination of Policy Agent and WFQ dramatically affected the response times, bringing the human wait times down from a frustrating three to seven seconds, to a far more acceptable one second. Less obvious, but equally important, the response times fell within a predictable range that ensured that the terminal operator always saw the same response time.

Summary

Building a consolidated IP network that supports a mix of application traffic requires a QoS strategy that ensures that mission-critical applications receive higher priority than less critical applications. A QoS strategy includes several components. The components tested here were QoS signaling, as set by the Policy Agent in the S/390 Enterprise Server, and congestion management, as demonstrated using the Cisco WFQ algorithm.

Combining S/390 application priority setting with Cisco networking provided consistent subsecond response time for clients such as those provided by the IBM Host-on-Demand and Personal Communications products. The combined capability of S/390 QoS signaling and Cisco congestion management provide IP networks with application-to-desktop performance that can meet the needs of the service-level requirements of today's e-business solutions.

White Paper

Appendixes

Appendix A: Cisco Router Configuration

The following configuration statements were used in the Cisco 7513 router, which provided the serial interface with traffic prioritization, the channel interface, and the OSA Fast Ethernet interface.

```
!  
! The following interface connects to a branch router via  
! 256 kbps PPP. The branch router configuration is not shown  
!  
interface Serial5/0/0  
  ip address 133.10.40.1 255.255.255.0  
  no ip directed-broadcast  
  encapsulation ppp  
  no ip route-cache optimum  
  bandwidth 256  
  load-interval 30  
  tx-queue-limit 1  
  fair-queue 150 4096 0  
  hold-queue 150 out  
!  
! This interface connects to the OS/390 Enterprise Server via  
! OSA Fast Ethernet.  
!  
interface FastEthernet5/1/0  
  ip address 133.10.30.99 255.255.255.0  
  no ip directed-broadcast  
  no ip route-cache optimum  
  ip policy route-map inter  
!  
! This channel interface connects to the OS/390 Enterprise Server via  
! ESCON.  
!  
interface Channel10/0  
  description Channel to MVS074  
  ip address 133.10.20.99 255.255.255.0  
  no ip directed-broadcast  
  ip ospf network point-to-point  
  ip ospf demand-circuit  
  ip policy route-map inter  
  no keepalive  
  claw D900 10 133.10.20.1 H9672A C7513 TCPIP TCPIP broadcast  
!  
!  
router ospf 7513  
  network 22.10.7.0 0.0.0.255 area 0.0.0.0  
  network 133.10.0.0 0.0.255.255 area 0.0.0.0  
!
```

White Paper

Appendix B: OS/390 TCP/IP Profiles

The following configuration shows the options specified as part of the TCP/IP definitions in OS/390:

```
;*****
;          CHANNEL DEVICE TO CONNECT TO C7513
;*****
;
DEVICE CIP2A CLAW E30 H9672A C7513 NONE 15 15 4096 4096
LINK CIP1 IP 0 CIP2A
START CIP2A
;
;*****
;          OSA Fast Ethernet Connection
;*****
;
DEVICE OSAETH    LCS    2E72
LINK LOSAETH    ETHERNET 0 OSAETH
START OSAETH
;
;*****
;          VIRTUAL IP ADDRESS FOR MVS074 IMAGE
;*****
;
DEVICE VIPA01    VIRTUAL 0
LINK  VIPA1     VIRTUAL 0  VIPA01
;
;*****
;  PROCS TO AUTOSTART
;*****
;
AUTOLOG
  FTPR7                ; FTP Server
  TNOEA                ; OE TELNET Server
  OMPROUTE             ; OE Routed Server
ENDAUTOLOG
;
;*****
;  WELL KNOWN PORTS
;*****
;
PORT
  20  TCP OMVS NOAUTOLOG           ; FTP server data port
  21  TCP OMVS                     ; FTP server
control port
  23  TCP OMVS                     ; OE TELNET
Server
  520 TCP OMPROUTE NOAUTOLOG       ; OE Routed Server
  520 UDP OMPROUTE NOAUTOLOG       ; OE Routed Server
;
;*****
;  HOME ADDRESSES FOR THIS STACK
;*****
```


White Paper

```
;
HOME
;
; ADDRESS LINK_NAME
; =====
133.10.20.1 CIP1 ; This goes to Cisco
CIP (7513)
133.10.30.1 LOSAETH ; OSA Ethernet to Cisco 7513
135.1.1.1 VIPA1 ; This is the VIPA
Address
;
;*****
; ASSORTED PARMS AND OPTIONS STATEMENTS
;*****
;
ASSORTEDPARMS
VARSUBNETTING
ENDASSORTEDPARMS
;
IPCONFIG DATAGRAMFWD SOURCEVIP A IGNOREREDIRECT MULTIPATH
;
```

Appendix C: OSA-2 Fast Ethernet Configuration

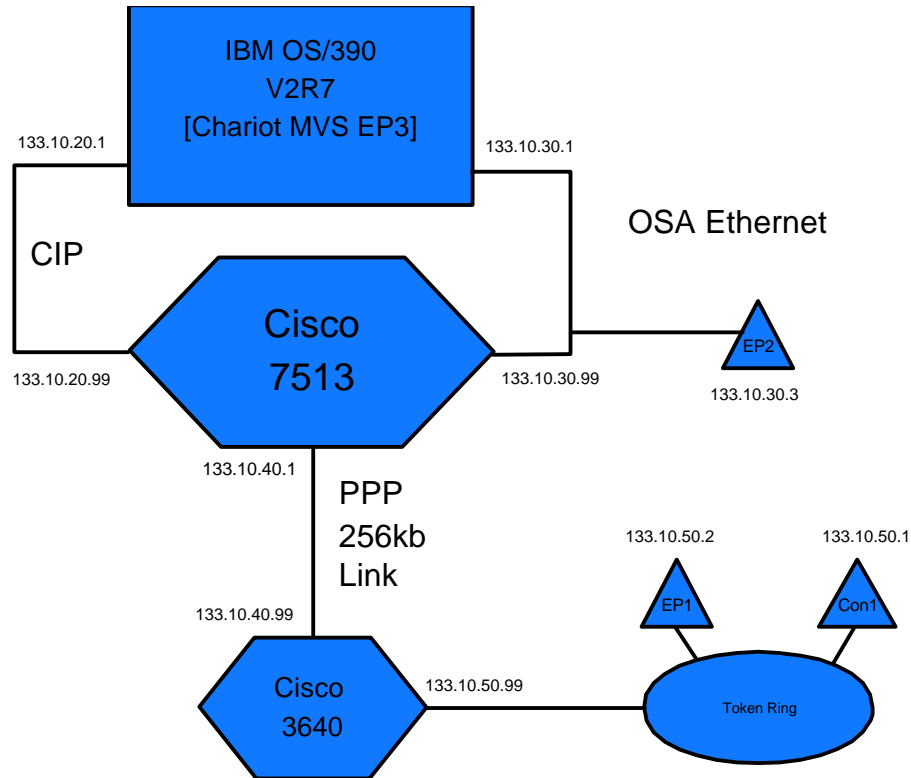
The following configuration for the OSA-2 adapters was used in the test:

```
CHPID PATH=40 ,TYPE=OSA ,SHARED
CHPID PATH=D8 ,TYPE=OSA ,SHARED
CNTLUNIT CUNUMBR=2E60 ,PATH=40 ,UNIT=OSA
CNTLUNIT CUNUMBR=2E70 ,PATH=D8 ,UNIT=OSA
IODEVICE
ADDRESS=( 2E60 , 12 ) ,CUNUMBR=2E60 ,UNIT=OSA ,STADET=Y ,      X
      UNITADD=00
IODEVICE
ADDRESS=( 2E6C , 2 ) ,CUNUMBR=2E60 ,UNIT=OSA ,STADET=Y ,      X
      UNITADD=FC
IODEVICE
ADDRESS=( 2E6E , 1 ) ,CUNUMBR=2E60 ,UNIT=OSAD ,STADET=Y ,      X
      UNITADD=FE
*
IODEVICE
ADDRESS=( 2E70 , 12 ) ,CUNUMBR=2E70 ,UNIT=OSA ,STADET=Y ,      X
      UNITADD=00
IODEVICE
ADDRESS=( 2E7C , 2 ) ,CUNUMBR=2E70 ,UNIT=OSA ,STADET=Y ,      X
      UNITADD=FC
IODEVICE
ADDRESS=( 2E7E , 1 ) ,CUNUMBR=2E70 ,UNIT=OSAD ,STADET=Y ,      X
      UNITADD=FE
*
```

Appendix D: OS/390 Policy Agent Profiles

Figure D-1 presents the configuration from an IP standpoint, providing the IP addresses for the various components. With the OS/390 Policy Server profiles, this figure demonstrates how QoS policies are defined in the OS/390 Policy Agent.

Figure D1: IP Configuration in the Data Center



The OS/390 Policy Server profiles were created for the traffic using the OSA-2 Fast Ethernet path. The SubnetAddr operand at the beginning of the profile, along with the Interfaces in the ServiceCategories definitions, match the IP address of the -2 interface (133.10.30.1) shown in Figure 13. There are three service categories defined based on destination address and IP precedence bit settings.

The categories are followed by three service policy rules that identify what traffic is assigned to which service category based on port number and destination address. While there are more options available for assigning categories, such as day of the week, these were not used in this particular test. The ports used for the interactive traffic in this particular test with OSA-2 Fast Ethernet were in the range of 700 to 719. The destination address for this particular connection is 133.10.50.2. Therefore, *rule1* in the set of service policy rules applied and dictated that the ToS bits be set based on service category *inter1*.

OS/390 Policy Server Profiles

```
#####
# CS for OS/390 v2r7
# SMP/E distribution path: /usr/lpp/tcpip/samples/IBM/EZAPAGCO
#####
#
# 5647-A01 (C) Copyright IBM Corp. 1998.
TcpImage TCP FLUSH
SetSubnetPrioTosMask
{
    SubnetAddr          133.10.30.1
    SubnetTosMask       11100000
    PriorityTosMapping  1 11100000
    PriorityTosMapping  1 11000000
    PriorityTosMapping  1 10100000
    PriorityTosMapping  2 10000000
    PriorityTosMapping  2 01100000
    PriorityTosMapping  2 01000000
    PriorityTosMapping  3 00100000
    PriorityTosMapping  3 00000000
}
ServiceCategories inter1
{
    PolicyScope      DataTraffic
    MaxRate          100000
    MinRate          10000
    Interface        133.10.30.1
    OutgoingTOS     11000000
}
ServiceCategories inter2
{
    PolicyScope      DataTraffic
    MaxRate          100000
    MinRate          10000
    Interface        133.10.30.1
    OutgoingTOS     01100000
}
ServiceCategories inter3
{
    PolicyScope      DataTraffic
    MaxRate          100000
    MinRate          10000
    Interface        133.10.30.1
    OutgoingTOS     00100000
}
ServicePolicyRules rule1
{
    PolicyScope      DataTraffic          #
    Direction        Both                 #
    Permission       Allowed              #
}
```

White Paper

```

    ProtocolNumber 6
    DaysOfWeekMask 1111111
    TimeOfDayRange 0:01-23:59
    SourceAddressRange 135.1.1.1
    DestinationAddressRange 133.10.50.2
    SourcePortRange 700 719
ports are in this range
    ServiceReference inter1a
    ServiceReference inter1b
}
ServicePolicyRules rule2
{
    PolicyScope DataTraffic
    Direction Outgoing
    Permission Allowed
    ProtocolNumber 6
    DaysOfWeekMask 1111111
    TimeOfDayRange 0:01-23:59
    SourceAddressRange 133.10.30.1
    DestinationAddressRange 133.10.50.2
    SourcePortRange 600 619
    ServiceReference inter2
}
ServicePolicyRules rule3
{
    PolicyScope DataTraffic
    Direction Outgoing
    Permission Allowed
    ProtocolNumber 6
    DaysOfWeekMask 1111111
    TimeOfDayRange 0:01-23:59
    SourceAddressRange 133.10.30.1
    DestinationAddressRange 133.10.50.2
    SourcePortRange 800 819
    ServiceReference inter3
}
# tcp
# weekdays
← 3270 Script
```

Appendix E: Cisco 7513 Router Console Displays

The messages shown in Figure E-1 were cut and pasted from the console display of the Cisco 7513 router while test traffic was being routed through it. It shows the view of the IP precedence and differentiated services field settings from within the router.

Examining the information provides the following details:

- The weights for the differentiated services (512) and IP precedence (585) traffic were set differently. This would provide different levels of service using WFQ. This, in fact, can be observed in the ToS field, where the ToS value for differentiated services (252) is higher than the ToS field for IP precedence (192). The correlation is that a higher weight results in a lower class of service.
- The queue depths demonstrate that there is congestion in the network. If the queue depths were zero it would imply that traffic was flowing freely and the congestion control algorithms were, in fact, not being executed. Congestion management is required only as queues build.
- The lengths of the messages and source address identify this as outbound interactive traffic from the S/390 to the client.

Figure E-1: Cisco 7513 Router Console Displays

While running Differentiated Services test:

```
(depth/weight/discards/tail drops/interleaves) 1/512/0/0/0  
Conversation 600, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x5DAC, ttl: 63,  
TOS: 252 prot: 6, source port 700, destination port 1282
```

```
(depth/weight/discards/tail drops/interleaves) 1/512/0/0/0  
Conversation 604, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x5E6B, ttl: 63,  
TOS: 252 prot: 6, source port 701, destination port 1285
```

```
(depth/weight/discards/tail drops/interleaves) 2/512/0/0/0  
Conversation 607, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x5ECA, ttl: 63,  
TOS: 252 prot: 6, source port 702, destination port 1287
```

While running IP Precedence test:

```
(depth/weight/discards/tail drops/interleaves) 2/585/0/0/0  
Conversation 722, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x6AEB, ttl: 63,  
TOS: 192 prot: 6, source port 700, destination port 1282
```

```
(depth/weight/discards/tail drops/interleaves) 2/585/0/0/0  
Conversation 725, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x6B42, ttl: 63,  
TOS: 192 prot: 6, source port 701, destination port 1285
```

```
(depth/weight/discards/tail drops/interleaves) 1/585/0/0/0  
Conversation 727, linktype: ip, length: 1444  
source: 133.10.30.1, destination: 133.10.50.2, id: 0x6B80, ttl: 63,  
TOS: 192 prot: 6, source port 702, destination port 1287
```

Appendix F: Chariot Test Tool Traffic Generation

The Chariot test tool from Ganymede was used to generate traffic, in addition to the clients from IBM. Chariot also analyzes the traffic. Information is included in this appendix to demonstrate the types and amount of traffic generated. Figure F-1 displays the scripts for generating the batch FTP traffic. Endpoint 1 is the FTP client, sending 100-byte files. Endpoint 2 is the FTP server, which is generating 10,000-byte files.

Figure F-1: Chariot Scripts for FTP Traffic

Endpoint 1 -----	Endpoint 2 -----
SLEEP initial_delay=0	
CONNECT_INITIATE port_number=AUTO	CONNECT_ACCEPT port_number=AUTO
LOOP number_of_timing_records=100 START_TIMER	LOOP number_of_timing_records=100
LOOP transactions_per_record=1	LOOP transactions_per_record=1
SEND size_of_record_to_send=100 size_of_record_to_send=100 send_datatype=NOCOMPRESS send_data_rate=UNLIMITED	RECEIVE size_of_record_to_send=100 size_of_record_to_send=100
RECEIVE file_size=10000 receive_buffer_size=DEFAULT	SEND file_size=10000 send_buffer_size=DEFAULT send_datatype=NOCOMPRESS send_data_rate=UNLIMITED
INCREMENT_TRANSACTION	
END_LOOP	END_LOOP
END_TIMER	
SLEEP transaction_delay=0	
END_LOOP	END_LOOP
DISCONNECT	DISCONNECT

Figure F-2 shows the scripts used to generate the interactive traffic. Endpoint 1 is the TN3270 client, sending 100-byte messages. Endpoint 2 is the TN3270 server, sending 1400-byte messages.

White Paper

Figure F-2: Chariot Scripts for Interactive Traffic

Endpoint 1 -----	Endpoint 2 -----
SLEEP initial_delay=0	
CONNECT_INITIATE port_number=700	CONNECT_ACCEPT port_number=700
LOOP number_of_timing_records=50	LOOP number_of_timing_records=50
START_TIMER	
LOOP transactions_per_record=1	LOOP transactions_per_record=1
SEND size_of_record_to_send=100 inquiry_send_buffer=DEFAULT send_datatype=NOCOMPRESS send_data_rate=UNLIMITED	RECEIVE size_of_record_to_send=100 inquiry_receive_buffer=DEFAULT
RECEIVE reply_size=1400 reply_receive_buffer=DEFAULT	SLEEP delay_before_responding=0
	SEND reply_size=1400 reply_send_buffer=DEFAULT send_datatype=NOCOMPRESS send_data_rate=UNLIMITED
INCREMENT_TRANSACTION	
END_LOOP	END_LOOP
END_TIMER	
SLEEP transaction_delay=0	
END_LOOP	END_LOOP
DISCONNECT	DISCONNECT

For each of the eight tests, three TCP/IP connections were established between the TN3270 clients and server to generate interactive traffic. Thirty TCP/IP connections were established to carry the FTP traffic.

Appendix G: Detailed Information from Individual Tests

The information provided in the body of the report reflects the summary of the eight tests. This appendix will provide the detailed information for each test.

Diffserv-FIFO-CIP Test

In this test, the differentiated services standard was used in the S/390 to set the ToS bits in the IP header. The routers used FIFO for traffic prioritization (that is, there was no traffic prioritization in the network). The S/390 was connected to the network via the CIP.

Table G-1 provides the totals for the traffic generated for the test.

Table G-1: Traffic Generated for Diffserv-FIFO-CIP Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
Diffserv-FIFO-CIP Total Traffic	955	955	95,500	7,933,200
Interactive Traffic	188	188	18,800	263,200
FTP Traffic	767	767	76,700	7,670,000

Table G-2 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types.

Table G-2: Response Time Averages for Diffserv-FIFO-CIP Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
Diffserv-FIFO-CIP Total Traffic	10.92308	1.13600	21.58100
Interactive Traffic	4.73353	1.13600	7.37000
FTP Traffic	11.54204	4.16100	21.58100

Figure G-1 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four FIFO examples present similar variations, based on which message gets to the queue first. In all examples you can see significant response time variations are based on the lack of network prioritization.

White Paper

Figure G-1: Response Time for Interactive and FTP Traffic for Diffserv-FIFO-CIP Test





Figure G-2 shows the throughput during the test. The upper line is throughput for the FTP traffic. It is fairly consistent until the queues are filled, congestion occurs, packets are dropped, and the TCP backoff mechanism is invoked, at which point the amount of traffic begins to significantly increase and decrease based on which message is released from the queue. It closely parallels the response time curve. The lower line is the throughput for the interactive traffic. There is little variation based on the scale of the graph and the small size of the messages. The small variations seen parallel the response time variations.

Figure G-2: Throughput for Diffserv-FIFO-



Diffserv-FIFO-OSA-2 Test

In this test, the differentiated services standard was used in the S/390 to set the ToS bits in the IP header. The routers used FIFO for traffic prioritization (that is, there was no traffic prioritization in the network). The S/390 was connected to the network via the OSA-2.

Table G-3 provides the totals for the traffic generated for the test.

Table G-3: Traffic Generated for Diffserv-FIFO-OSA-2 Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
Diffserv-FIFO-OSA-2 Total Traffic	949	949	94,900	7,899,000
Interactive Traffic	185	185	18,500	259,000
FTP Traffic	764	764	76,400	7,640,000

Table G-4 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types.

Table G-4: Response Time Averages for Diffserv-FIFO-OSA-2 Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
Diffserv-FIFO-OSA-2 Total Traffic	10.99561	0.16300	22.26000
Interactive Traffic	4.80483	0.16300	7.00900
FTP Traffic	11.61469	4.20500	22.26000

Figure G-3 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four FIFO examples present similar variations, based on which message gets to the queue first. The large variation in response time reflects the lack of traffic prioritization in the network.

White Paper

Figure G-3: Response Time for Interactive and FTP Traffic for Diffserv-FIFO-OSA-2 Test





Figure G-4 shows the throughput during the test. The upper line is throughput for the FTP traffic. It is fairly consistent until the queues are filled, congestion occurs, packets are dropped, and the TCP backoff mechanism is invoked, at which point the amount of traffic begins to significantly increase and decrease based on which message is released from the queue. It closely parallels the response time curve. The lower line is the throughput for the interactive traffic. There is little variation based on the scale of the graph and the small size of the messages. The small variations seen parallel the response time variations.

Figure G-4: Throughput During Diffserv-FIFO-OSA-2 Test



White Paper

IP Precedence-FIFO-CIP Test

In this test, the original standard for setting the IP precedence bits was used in the S/390 to set the ToS bits in the IP header. The routers used FIFO for traffic prioritization (that is, there was no traffic prioritization in the network). The S/390 was connected to the network via the CIP.

Table G-5 provides the totals for the traffic generated for the test.

Table G-5: Traffic Generated for IP Precedence-FIFO-CIP Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
IP Precedence-FIFO-CIP Total Traffic	944	944	94,400	7,814,600
Interactive Traffic	189	189	18,900	264,600
FTP Traffic	755	755	75,500	7,550,000

Table G-6 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types.

Table G-6: Response Time Averages for IP Precedence-FIFO-CIP Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
IP Precedence-FIFO-CIP Total Traffic	11.07209	0.08100	22.70200
Interactive Traffic	4.71790	0.08100	7.01700
FTP Traffic	11.70750	6.01400	22.70200

Figure G-5 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four FIFO examples present similar variations, based on which message gets to the queue first. The large variation in response time reflects the lack of traffic prioritization in the network.

White Paper

Figure G-5: Response Time for Interactive and FTP Traffic for IP Precedence-FIFO-CIP Test

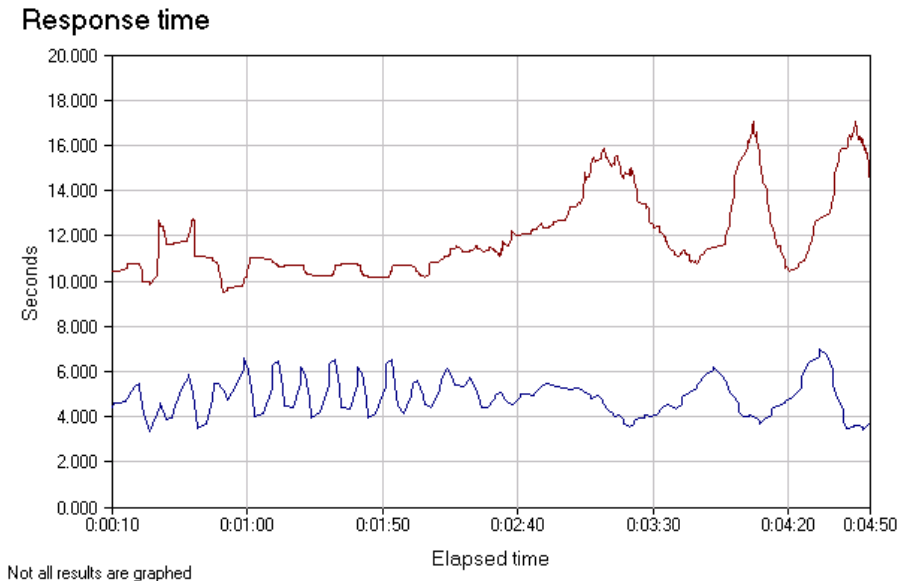



Figure G-6 shows the throughput during the test. The upper line is throughput for the FTP traffic. It is fairly consistent until the queues are filled, congestion occurs, packets are dropped, and the TCP backoff mechanism is invoked, at which point the amount of traffic begins to significantly increase and decrease based on which message is released from the queue. It closely parallels the response time curve. The lower line is the throughput for the interactive traffic. There is little variation based on the scale of the graph and the small size of the messages. The small variations seen parallel the response time variations.

Figure G-6: Throughput for IP Precedence-FIFO-CIP Test



White Paper

IP Precedence-FIFO-OSA-2 Test

In this test, the original standard for setting the IP precedence bits was used in the S/390 to set the ToS bits in the IP header. The routers used FIFO for traffic prioritization (that is, there was no traffic prioritization in the network). The S/390 was connected to the network via the OSA-2.

Table G-7 provides the totals for the traffic generated for the test.

Table G-7: Traffic Generated for IP Precedence-FIFO-OSA-2 Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
IP Precedence-FIFO-OSA-2 Total Traffic	936	936	93,600	7,734,600
Interactive Traffic	189	189	18,900	264,600
FTP Traffic	747	747	74,700	7,470,000

Table G-8 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types.


Table G-8: Response Time Averages for IP Precedence-FIFO-OSA-2 Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
IP Precedence-FIFO-OSA-2 Total Traffic	11.20120	0.06300	26.57100
Interactive Traffic	4.72481	0.06300	7.40800
FTP Traffic	11.84883	4.25100	26.57100

Figure G-7 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four FIFO examples present similar variations, based on which message gets to the queue first. The large variation in response time reflects the lack of traffic prioritization in the network.

White Paper


Figure G-7: Response Time for Interactive and FTP Traffic for IP Precedence-FIFO-OA-2 Test



Not all results are graphed

Figure G-8 shows the throughput during the test. The upper line is throughput for the FTP traffic. It is fairly consistent until the queues are filled, congestion occurs, packets are dropped, and the TCP backoff mechanism is invoked, at which point the amount of traffic begins to significantly increase and decrease based on which message is released from the queue. It closely parallels the response time curve. The lower line is the throughput for the interactive traffic. There is little variation based on the scale of the graph and the small size of the messages. The small variations seen parallel the response time variations.

Figure G-8: Throughput During IP Precedence-FIFO-OA-2 Test



Not all results are graphed

White Paper

Diffserv-WFQ-CIP Test

In this test, the differentiated services standard was used in the S/390 to set the ToS bits in the IP header. The routers used WFQ for traffic prioritization in the network. The S/390 was connected to the network via the CIP.

Table G-9 provides the totals for the traffic generated for the test. With WFQ for traffic prioritization the interactive traffic receives higher priority and, therefore, more bandwidth. Comparing these results with those achieved in the Diffserv-FIFO-CIP test, approximately 65 percent more transactions occurred using WFQ, and the number of interactive transactions increased approximately 476 percent while the number of FTP transactions decreased by approximately 13 percent.

Table G-9: Traffic Generated for Diffserv-WFQ-CIP Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
Diffserv-WFQ-CIP Total Traffic	1575	1575	157,500	8,053,000
Interactive Traffic	895	895	89,500	1,253,000
FTP Traffic	680	680	68,000	6,800,000

Table G-10 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types. Comparing these results to the results using FIFO, the average response time for the total traffic has increased significantly because the response time of the FTP traffic has increased and there were more individual FTP connections than interactive connections.

Comparing these results specifically to the Diffserv-FIFO-CIP test, the interactive traffic response time has been reduced from an average of 4.73353 seconds to 1.00288 seconds. The FTP traffic average response time has increased from 11.54204 seconds to 16.85668 seconds.

Table G-10: Response Time Averages for Diffserv-WFQ-CIP Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
Diffserv-WFQ-CIP Total Traffic	15.41543	0.36200	57.01600
Interactive Traffic	1.00288	0.36200	2.27300
FTP Traffic	16.85668	5.95700	57.01600

White Paper

Figure G-9 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four WFQ examples present similar variations.

The response time for the interactive traffic has decreased significantly and is much more consistent, with only minor variations from the average one-second response time. With the FTP traffic, response time gets worse and then improves toward the end of the measured period as TCP/IP windowing takes effect—response time improves, while throughput begins to decline.

Figure G-9: Response Time for Interactive and FTP Traffic for Diffserv-WFQ-CIP Test

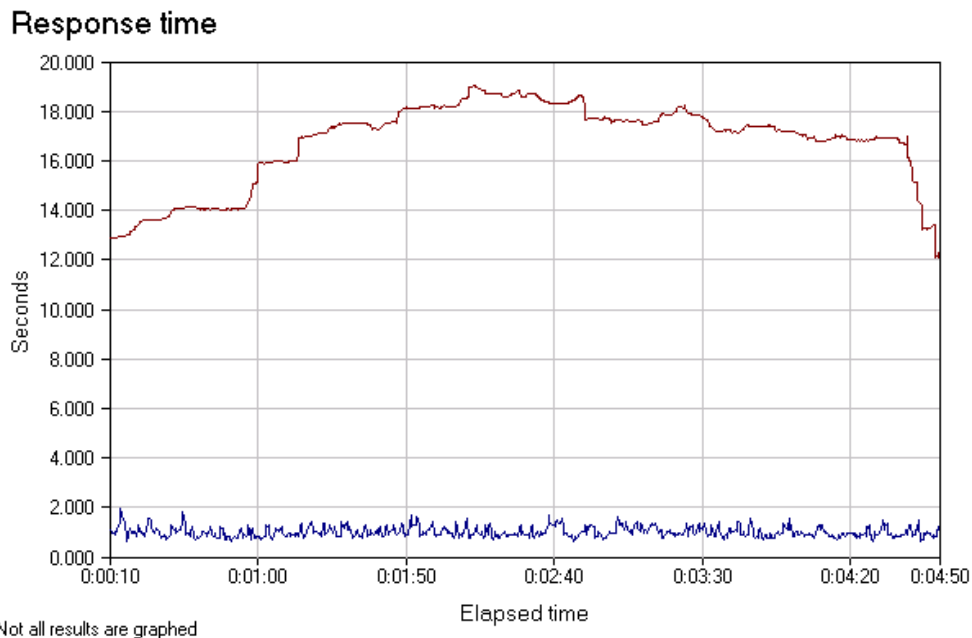



Figure G-10 shows the throughput during the test. The upper line is throughput for the FTP traffic. The throughput for the FTP traffic begins to decline as interactive traffic is added to the mix and continues to decline as TCP/IP windowing reduces the size of FTP windows being sent. The lower line is the throughput for the interactive traffic. It is significantly increased over the FIFO throughput and, most importantly, is consistent over the measured time period.

White Paper

Figure G-10: Throughput for Diffserv-WFQ-CIP Test



Not all results are graphed

Diffserv-WFQ-OSA-2 Test

In this test, the differentiated services standard was used in the S/390 to set the ToS bits in the IP header. The routers used WFQ for traffic prioritization in the network. The S/390 was connected to the network via the OSA-2.

Table G-11 provides the totals for the traffic generated for the test. With WFQ for traffic prioritization the interactive traffic receives higher priority and, therefore, more bandwidth. Comparing these results with those achieved in the Diffserv-FIFO-OSA-2 test, approximately 72 percent more transactions occurred using WFQ, and the number of interactive transactions increased approximately 528 percent while the number of FTP transactions decreased by approximately 15 percent.

Table 11: Traffic Generated for Diffserv-WFQ-OSA-2 Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
Diffserv-WFQ-OSA-2 Total Traffic	1628	1628	162,800	7,886,400
Interactive Traffic	976	976	97,600	1,366,400
FTP Traffic	652	652	65,200	6,520,000

Table G-12 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less

White Paper

interesting than the response time for the individual traffic types. Comparing these results to the results using FIFO, the average response time for the total traffic has increased significantly because the response time of the FTP traffic has increased and there were more individual FTP connections than interactive connections.

Comparing these results specifically to the Diffserv-FIFO-OSA-2 test, the interactive traffic response time has been reduced from an average of 4.80483 seconds to 0.91980 seconds. The FTP traffic average response time has increased from 11.61469 seconds to 17.04647 seconds.

Table G-12: Response Time Averages for Diffserv-WFQ-OSA-2 Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
Diffserv-WFQ-OSA-2 Total Traffic	15.58041	0.17300	46.92000
Interactive Traffic	0.91980	0.17300	2.12400
FTP Traffic	17.04647	5.47700	46.92000

Figure G-11 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four WFQ examples present similar variations.

The response time for the interactive traffic has decreased and is much more consistent, with only minor variations from the average nine-tenths-second response time. With the FTP traffic, response time gets worse and then improves toward the end of the measured period as TCP/IP windowing takes effect—response time improves, while throughput begins to decline.

Figure G-11: Response Time for Interactive and FTP Traffic for Diffserv-WFQ-OSA-2 Test





Figure G-12 shows the throughput during the test. The upper line is throughput for the FTP traffic. The throughput for the FTP traffic begins to decline as interactive traffic is added to the mix and continues to decline as TCP/IP windowing reduces the size of FTP windows being sent. The lower line is the throughput for the interactive traffic. It is significantly increased over the FIFO throughput and, most importantly, is consistent over the measured time period.

Figure G-12: Throughput for Diffserv-WFQ-OSA-2 Test



IP Precedence-WFQ-CIP Test

In this test, the original IP standard was used to set the IP precedence bits in the IP header in the S/390. The routers used WFQ for traffic prioritization in the network. The S/390 was connected to the network via the CIP.

Table G-13 provides the totals for the traffic generated for the test. The interactive traffic receives higher priority with WFQ and, therefore, more bandwidth. Comparing these results with those achieved in the IP Precedence-FIFO-CIP test, approximately 66 percent more transactions occurred using WFQ, and the number of interactive transactions increased approximately 471 percent while the number of FTP transactions decreased by approximately 11 percent.

Table G-13: Traffic Generated for IP Precedence-WFQ-CIP Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
IP Precedence-WFQ-CIP Total Traffic	1563	1563	156,300	7,976,000
Interactive Traffic	890	890	89,000	1,246,000
FTP Traffic	673	673	67,300	6,730,000

Table G-14 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less interesting than the response time for the individual traffic types. Comparing these results to the results using FIFO, the average response time for the total traffic has increased significantly because we have increased the response time of the FTP traffic and there were more individual FTP connections than interactive connections.

Comparing these results specifically to the IP Precedence-FIFO-CIP test, the interactive traffic response time has been reduced from an average of 4.71790 seconds to 1.00889 seconds. The FTP traffic average response time has increased from 11.70750 seconds to 17.18999 seconds.

Table G-14: Response Time Averages for IP Precedence-WFQ-CIP Test

Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
IP Precedence-WFQ-CIP Total Traffic	15.71898	0.37300	49.22000
Interactive Traffic	1.00889	0.37300	2.44500
FTP Traffic	17.18999	4.43400	49.22000

White Paper

Figure G-13 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four WFQ examples present similar variations.

The response time for the interactive traffic has decreased and is much more consistent, with only minor variations from the average one-second response time. With the FTP traffic response time gets worse and then improves toward the end of the measured period as TCP/IP windowing takes effect—response time improves, while throughput begins to decline.

Figure G-13: Response Time for Interactive and FTP Traffic for IP Precedence-WFQ-CIP Test


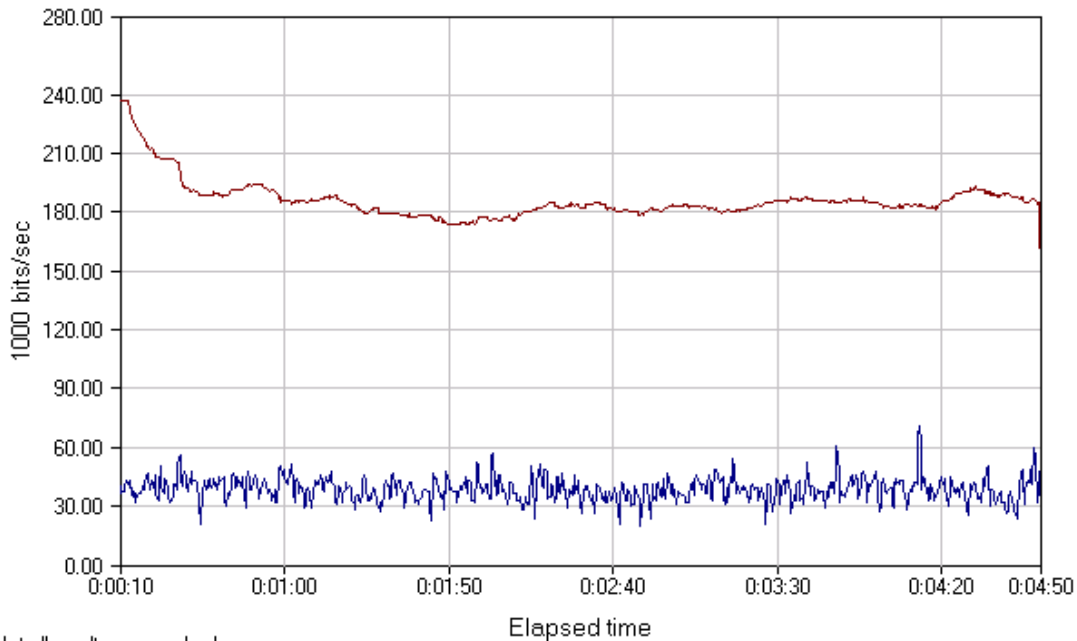


Figure G-14 shows the throughput during the test. The upper line is throughput for the FTP traffic. The throughput for the FTP traffic begins to decline as interactive traffic is added to the mix and continues to decline as TCP/IP windowing reduces the size of FTP windows being sent. The lower line is the throughput for the interactive traffic. It is significantly increased over the FIFO throughput and, most importantly, is consistent over the measured time period.

Figure G-14: Throughput for IP Precedence-WFQ-CIP Test

Throughput



Not all results are graphed

IP Precedence-WFQ-OSA-2 Test

In this test, the original IP standard was used to set the IP precedence bits in the IP header in the S/390. The routers used WFQ for traffic prioritization in the network. The S/390 was connected to the network via the OSA-2.

Table G-15 provides the totals for the traffic generated for the test. The interactive traffic receives higher priority and, therefore, more bandwidth. Comparing these results with those achieved in the IP Precedence-FIFO-OSA-2 test, approximately 71 percent more transactions occurred using WFQ, and the number of interactive transactions increased approximately 490 percent while the number of FTP transactions decreased by approximately 10 percent.

Table G-15: Traffic Generated for IP Precedence-WFQ-OSA-2 Test

Group/Pair	Number of Timing Records	Transaction Count	Bytes Sent by Endpoint 1	Bytes Received by Endpoint 1
IP Precedence-WFQ-OSA-2 Total Traffic	1563	1563	156,300	7,976,000
Interactive Traffic	890	890	89,000	1,246,000
FTP Traffic	673	673	67,300	6,730,000

Table G-16 provides the response time averages for the test. The response time average in the first row includes both interactive and FTP traffic and is, therefore, probably less

White Paper

interesting than the response time for the individual traffic types. Comparing these results to the results using FIFO, the average response time for the total traffic has increased significantly because the response time of the FTP traffic has increased and there were more individual FTP connections than interactive connections.

Comparing these results specifically to the IP Precedence-FIFO-OSA-2 test, the interactive traffic response time has been reduced from an average of 4.72481 seconds to 0.96818 seconds. The FTP traffic average response time has increased from 11.84883 seconds to 16.41391 seconds.

Table G-16: Response Time Averages for IP Precedence-WFQ-OSA-2 Test


Group/Pair	Response Time Average	Response Time Minimum	Response Time Maximum
IP Precedence-WFQ-OSA-2 Total Traffic	15.00975	0.12900	40.19800
Interactive Traffic	0.96818	0.12900	2.24700
FTP Traffic	16.41391	4.77000	40.19800

Figure G-15 shows a graph of the response time for the test. The first and last ten seconds are not shown because of the wide fluctuations during start-up and take-down. Two graphs are shown. The upper line is the response time for the FTP traffic. The lower line is the response time for the interactive traffic. All four WFQ examples present similar variations.

The response time for the interactive traffic has decreased and it is much more consistent, with only minor variations from the average one-second response time. With the FTP traffic, response time gets worse and then improves toward the end of the measured period as TCP/IP windowing takes effect—response time improves, while throughput begins to decline.

Figure G-15: Response Time for Interactive and FTP Traffic for IP Precedence-WFQ-OSA-2 Test

Response time




Not all results are graphed

Figure G-16 shows the throughput during the test. The upper line is throughput for the FTP traffic. The throughput for the FTP traffic begins to decline as interactive traffic is added to the mix and continues to decline as TCP/IP windowing reduces the size of FTP windows being sent. The lower line is the throughput for the interactive traffic. It is significantly increased over the FIFO throughput and, most importantly, is consistent over the measured time period.

Figure G-16: Throughput During Diffserv-WFQ-OSA-2 Test

Throughput



Not all results are graphed



©Copyright IBM Corporation 2000

IBM Corporation
Department AQYA
3039 Cornwallis Road
Research Triangle Park, NC 27709

Printed in the United States of America
3-00
All Rights Reserved

IBM, OS/390, S/390, and S/390 Parallel Enterprise Server are trademarks of International Business Machines Corporation in the United States, other countries or both.

Other company, product and service names may be trademarks or service marks of others.