Lucent Technologies
Bell Labs Innovations



WHITE PAPER

Frame Relay Congestion Control

Using Congestion Detection and Notification Techniques as Service Differentiation Tools in a Competitive Market

February 2001

Table of Contents

1. Introduction 1

2. Understanding Congestion 2

Detecting Congestion 2

Notifying Upstream and Downstream Equipment of Network Congestion 3

Tracking Network Congestion 4

3. Using Congestion Control Techniques to Support QoS-Based Service Portfolios 4

Priority Frame 4

Congestion Control Framework 5

Implementing Congestion Marking in a Service Provider Network 7

4. Controlling Congestion in Frame Relay Networks That Rely on an ATM Backbone 9

5. Conclusion 9

Appendix: The Lucent Solution 11

GX 550 Multiservice WAN Switch 11

CBX 500 Multiservice WAN Switch 11

B-STDX 8000/9000 Multiservice WAN Switches 11

©2001 Lucent Technologies, Inc.

1. Introduction

Mechanisms that inform customers of frame relay network congestion can appear to service providers to be a losing proposition all around. After all, what good can come of alerting customers to congested service provider networks? Why not simply ignore temporary bouts of network congestion and hope the customer doesn't notice the performance hit until after the congestion clears?

The fact is that network congestion is a brilliant business opportunity disguised as a performance problem. If handled correctly, service providers can transform transient congestion conditions into increased customer satisfaction and by extension, reduced customer churn and greater revenues. Service providers can use congestion detection and notification techniques to:

- Build credibility with their customers
- Transform conversations about network performance from ones filled with queries and grievances to more productive, fact-based conversations focused on solutions
- Provide proof of service delivery
- Use network resources more efficiently by recovering faster from transient congestion conditions

Now that one-size-fits-all frame relay services are being displaced by a range of services tailored to the specific needs of a particular network, credibility with customers counts like never before. Today's enterprise network managers have access to network monitoring tools that give them the information they need to verify independently whether their service provider is delivering on their service agreement. This equipment gives enterprise LAN/WAN managers who purchase frame relay services high-level summaries of traffic throughput over the frame relay networks. Enterprise managers can easily detect disparities between the offered load to the frame relay network and the delivered throughput. And should that throughput fall short of the offered load, these managers are primed to challenge the service provider.

Congestion detection and notification mechanisms are a forthright approach to informing customers about actual network conditions. Service providers can use them to short circuit the enterprise manager's process of first watching for network performance disparities and then challenging the provider. Instead of debating whether there was congestion, service providers and customers can exactly characterize the congestion, understand its causes, and take appropriate steps. For example, the customer might decide to upgrade to a frame relay service that offers a higher level of service.

Once a service provider has earned credibility with a customer, that customer may consider the absence of congestion notifications in a functioning notification system as partial proof of the level of service delivered—making a functioning congestion notification service a critical component of a premium frame relay service.

Finally, congestion detection and notification systems can help the service provider use their available resources more effectively by promoting cooperation between the transmitting customer premise equipment (CPE) and the service provider network. This is particularly important in the lower end of the service spectrum, which targets applications that can tolerate occasional congestion. For such applications, the higher costs of premium services do not make economic sense. Still, a reliable and functioning congestion notification system allows cooperation between the network and the application traffic flow and enables rapid recovery from transient congestion conditions. Most frame relay customer premises equipment (CPE) now supports transmission rate modification features that enable them to respond intelligently to congestion conditions. For example, many IP routers connected over frame relay can throttle

©2001 Lucent Technologies, Inc.

down their transmission rate when they receive word of congestion. This quick response enables the service provider network to recover more quickly from congestion, thereby improving throughput and network performance for all service customers.

Whether addressing the high end of the market, the low end, or any point in between, congestion notification is a critical component of any service differentiation mix. Lucent provides a variety of Enhanced Frame Solutions, including Priority Frame™, that allow service providers to precisely control the quality of service (QoS) of every connection—whether that connection is a virtually lossless frame connection with minimum delay or simple best effort frame relay service with a cut rate pricing plan. This paper describes the role congestion detection and notification plays in supporting Priority Frame-based service portfolios.

2. Understanding Congestion

Before considering the tools used to manage congestion in a frame relay network, let's take a closer look at congestion itself. In a nutshell, network congestion means that there is a possibility that an output buffer leading to a transmission port could become too full of traffic to transmit all of its data. This situation creates the potential for the network to discard some traffic.

Note that simply having congestion in the network does not mean that any traffic will be dropped. Networks experience different levels of congestion depending on traffic volume. The likelihood that traffic will be discarded increases as the congestion worsens. In its least severe form, congestion is simply an indication that the network has the potential to become overburdened. Although lightly congested networks rarely need to discard traffic, if the traffic volume in that network remains constant or increases, it could increase the network congestion to a point that forces the network to discard traffic.

A practical way of thinking about network congestion is that during periods of congestion, the buffer holding traffic to be transmitted out a particular port could fill to the point that its volume exceeds a congestion threshold. This could trigger the network to take action to control the congestion. These actions range from requesting upstream equipment to slow its transmission rate to actually discarding traffic. The action taken depends upon the severity of the congestion and the Quality of Service (QoS) class of the affected traffic.

This paper focuses on the congestion detection and notification techniques available to network operators to minimize network congestion. For more detailed information on the full range of frame relay congestion control techniques, including buffer allocation, rate enforcement, congestion notification, discard priorities, and closed-loop congestion control mechanisms, see the *Enhanced Frame Solutions: Using Lucent Priority Frame™ to Support a Tiered Quality of Service Structure for Frame Relay* application note, available at www.lucent.com/ins.

Detecting Congestion

The precise means used to detect congestion in an output buffer varies according to the class of the virtual circuit. In the simplest case, the output buffer continues to accept traffic until it completely fills the buffer, at which point the network discards all incoming traffic until the congestion clears. Technically speaking, the network never formally identifies congestion in this scenario—it simply discards the excess traffic.

In the more complicated scenarios, the output buffer uses a series of thresholds to determine the level of congestion at that port. The network uses this information in conjunction with the traffic's QoS class to determine what action to take to control the congestion. Variable frame rate (VFR) virtual circuits, for example, are allowed to burst their transmission rates within agreed upon limits. If several circuits sharing the same output buffer experience simultaneous bursts in traffic, that output buffer could become congested. Therefore, VFR circuits require more involved congestion mechanisms that take into account the circuit's traffic contract, the level of

©2001 Lucent Technologies, Inc.

congestion in the output buffer, and the affected traffic's QoS requirements. These mechanisms will be discussed in Section 3.

Notifying Upstream and Downstream Equipment of Network Congestion

Frame relay networks employ a straightforward congestion notification process. Congestion notification takes two forms—one for each direction in which notification occurs. Both forms involve setting a bit in the frame relay header.

- To indicate congestion in the forward direction, the network sets the forward explicit congestion notification (FECN) bit in the frame relay header to one (1). This informs the receiving CPE that its received flow on the virtual circuit is experiencing congestion.
- To indicate congestion in the backwards direction, the network sets the backwards explicit congestion notification (BECN) bit in the frame relay header to one (1). This informs the transmitting CPE that its transmitted flow on the virtual circuit is experiencing congestion. Some CPE respond to BECN bits by throttling down transmission rates for a short period, then resuming full speed transmission when they stop receiving BECN bits.

Figure 1 clarifies the distinction between the FECN and BECN notification systems.

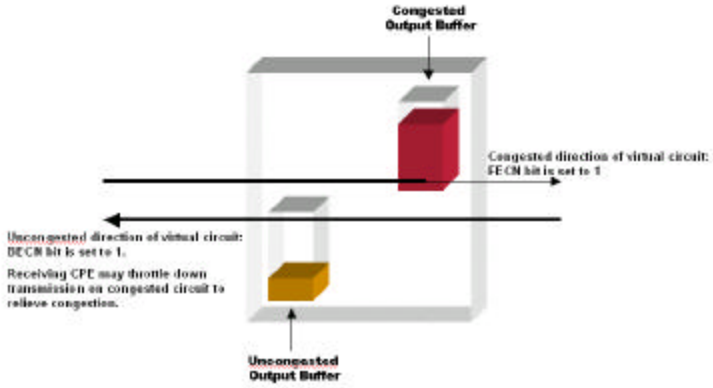

Figure 1—Forward and backward congestion notification in a frame relay network.

If the CPE is able to respond to receipt of BECN by throttling down transmission for a time, the congestion state in this switch might soon resemble Figure 2.

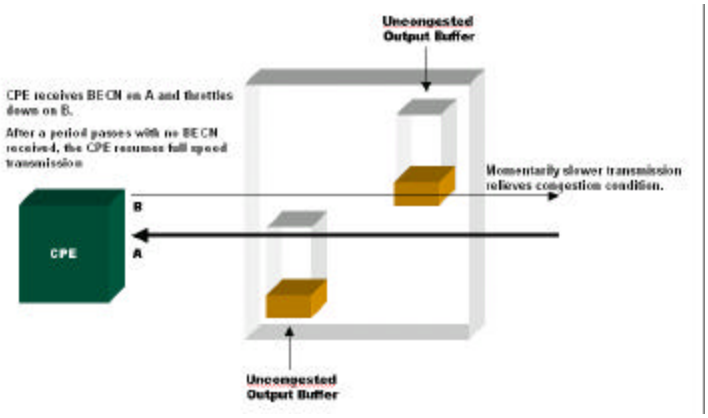

Figure 2—Service providers can use BECN to communicate with CPE and alleviate congestion conditions.

©2001 Lucent Technologies, Inc.

Although FECN and BECN are an established part of the frame relay standards, the standards do not define the manner in which switches detect congestion or the ways in which CPE respond to congestion notifications.

Tracking Network Congestion

Some enterprise managers actively track congestion conditions in the frame relay networks they purchase from service providers. One easy way for enterprise network managers to do that is to have their CPE routers count the FECN and BECN bits received from the service provider network. Many CPE routers can tally the FECN and BECN bits and record the results of that tally in their SNMP management information base (MIB). A suite of new network monitoring tools can periodically query the MIBs of these enterprise routers. These queries may occur several times an hour. The network monitoring tools then automatically generate reports that correlate the quantities and times of FECN and BECN bits with other observable network events, such as declines in the percentage of delivered frame relay frames. For many enterprises, checking up on their service provider has never been easier.

3. Using Congestion Control Techniques to Support QoS-Based Service Portfolios

Priority Frame

Lucent leads the market in providing true differentiated QoS for frame relay services. Part of the Lucent Enhanced Frame Solutions portfolio, Priority Frame allows providers to assign specific parameters including throughput, delay, and frame loss to various frame relay traffic types, creating a set of distinct QoS classes, including Variable Frame Rate-real-time (VFR-rt), Variable Frame Rate-non-real-time (VFR-nrt), and Unspecified Frame Rate (UFR) (see Figure 3). Traffic is then routed through the WAN according to its service class.

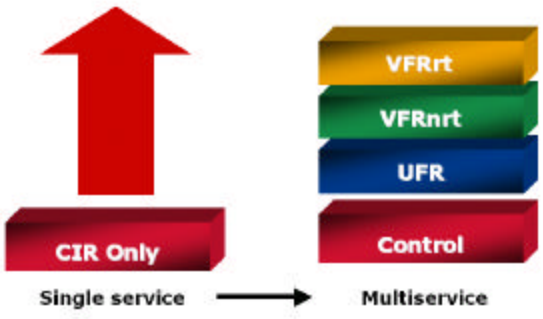

Figure 3— Lucent Priority Frame transforms the single service approach of traditional frame relay networks into a tiered service class structure capable of supporting multiservice portfolios on a single frame relay infrastructure.

Priority Frame does not simply guarantee bandwidth for services, it guarantees to deliver a specific set of information within a specific period of time. Service providers use its three frame relay service classes to distinguish between applications that can tolerate delay and frame loss, such as e-mail, residential internet access, and web browsing, and those with little to no tolerance for delay or frame loss, such as voice, video, and interactive business applications.

Priority Frame has several components, including:

- QoS objectives
- Routing

©2001 Lucent Technologies, Inc.

- Buffer allocation
- Transmission scheduling
- Congestion detection

All of the components used in Priority Frame are discussed in *Enhanced Frame Solutions: Using Lucent Priority Frame™ to Support a Tiered Quality of Service Structure for Frame Relay* white paper, available at www.lucent.com/ins. The remainder of this paper zeroes in on how FECN/BECN contributes to Priority Frame's ability to deliver absolute QoS guarantees for each service class.

Congestion Control Framework

Figure 4 below depicts the path frame relay traffic takes through a Lucent frame relay capable switch. At the egress point from the switch, each QoS class has a buffer used to temporarily hold frames received from the switching fabric while they are scheduled for transmission onto the line. Thus, each frame relay transmission port has four buffers, through which the frames of the appropriate frame relay QoS classes pass. As the figure shows, congestion detection in a Priority Frame network is performed separately for each QoS class at the output buffer just prior to the frame's transmission out of the network.

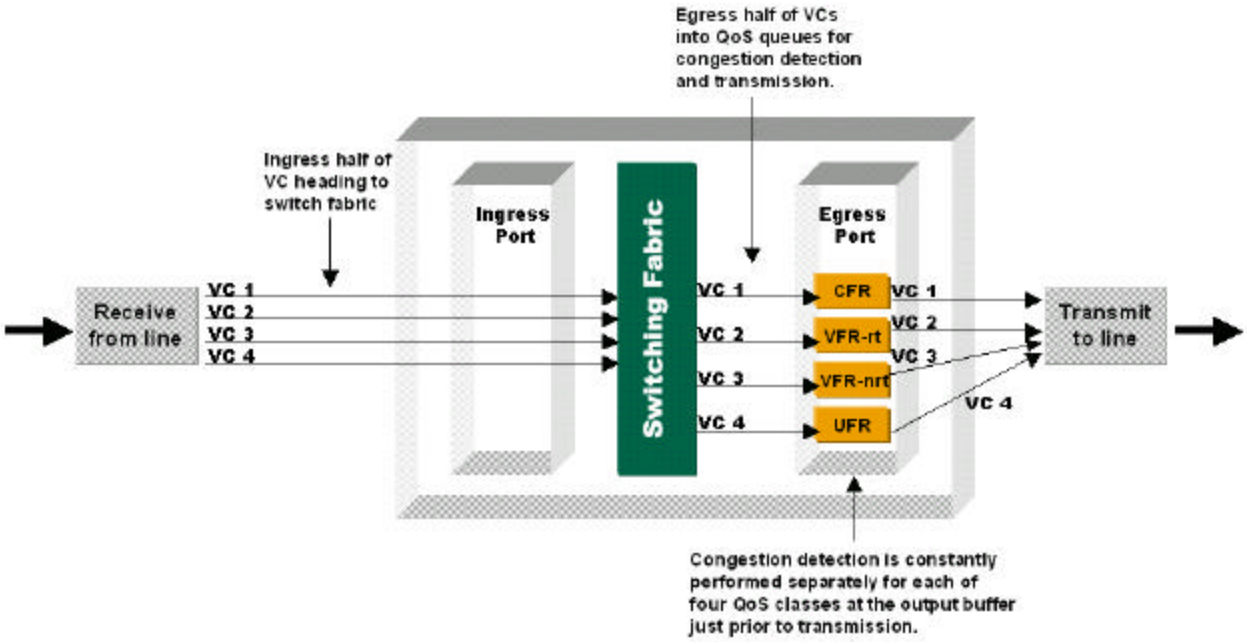

Figure 4—Congestion detection occurs constantly in output QoS buffers.

Two fundamental principles govern how the network reacts to congestion once it is detected in the output buffer:

- Punish non-conforming virtual circuits first
- Manage remaining traffic according to its QoS class

These rules are described below.

Punish Traffic from Non-Conforming Virtual Circuits First

Every virtual circuit is established with a traffic contract that specifies in mathematical terms the circuit parameters, such as its allowed transmission rate and the amount of allowed bursting on the circuit. The network checks all incoming frame relay circuits for conformance with their

©2001 Lucent Technologies, Inc.

traffic contract at the user-to-network interface (UNI), the circuits' ingress point into the frame relay network. Circuits that transmit too fast or otherwise violate their contract are said to be *non-conforming*. The network marks them as such so that all switches in the network will be able to distinguish non-conforming circuits from those circuits that comply with their contracts.

If congestion reaches a point that requires congestion notification or actual frame discarding, network switches will use this marking to determine which frames are eligible for discard. As we will see, frames from non-conforming circuits are marked and discarded first, regardless of their QoS class. This is a reasonable early step for service providers to take to relieve network congestion because non-conforming circuits are by definition in violation of their governing traffic contracts.

Manage Congested Traffic According to Its QoS Class

In addition to differentiating between conforming and non-conforming circuits, Lucent Priority Frame enables service providers to differentiate between the various frame relay QoS classes when determining what actions to take to relieve congestion in their networks. The process is described below.

Handling Constant Frame Rate Traffic

The Constant Frame Rate (CFR) QoS class has the same characteristics as ATM's Constant Bit Rate (CBR) service. Service providers define CFR services to provide a guaranteed frame rate, very low frame loss, very low frame delay, and very low frame delay variation.

Output queues for CFR circuits have no congestion detection mechanisms. This is non-intuitive on its face because CFR/CBR circuits have the highest QoS level of any service class. However, service providers only establish CFR circuits when there is sufficient free buffer space available to meet the high requirements of CFR traffic. When service providers establish CFR circuits for their customers, they set aside enough buffers on every switch through which the circuit passes to guarantee that their network will be able to handle the expected traffic flows. In addition, the network always services CFR/CBR queues first for transmission. Any traffic from other frame relay QoS classes that needs to be transmitted out a shared port is always transmitted after CFR/CBR traffic. These steps eliminate any possibility that CFR circuits that comply with their traffic contracts will experience network congestion.

If congestion occurs on CFR/CBR circuits, it must be the result of a non-conforming circuit. Frames from non-conforming CFR circuits are simply discarded as necessary without congestion notification. This is not so much congestion detection as a case of circuits violating their traffic contracts and paying the penalty (frame discard).

Implementing Congestion Detection for the Remaining Priority Frame QoS Classes

Congestion detection and notification for the remaining frame relay QoS classes is a bit more complicated. To understand why, one first needs to understand the classes themselves. Figure 5 summarizes the Priority Frame QoS classes.

Priority Frame Class	Maps to ATM QoS Class	Characteristics
Variable Frame Rate-real time (VFR-rt)	Variable Bit Rate-real time (VBR-rt)	• Minimum frame rate • Guaranteed frame rate • Ability to burst the frame rate within limits • Very low frame loss • Very low frame delay • Very low frame delay variation

©2001 Lucent Technologies, Inc.

Variable Frame Rate-non-real time (VFR-nrt)	Variable Bit Rate-non-real time (VBR-nrt)	• Minimum frame rate • Guaranteed frame rate • Ability to burst the frame rate within limits • Higher tolerance for frame loss, frame delay, and frame delay variation
Unspecified Frame Rate (UFR)	Unspecified Bit Rate (UBR)	No guarantees for frame rate, frame loss, or frame delay. Strictly "best-effort" service class.

Figure 5—The Priority Frame QoS service classes map to ATM QoS classes.

All three of these QoS classes are subject to unpredictable bursts of traffic. The two VFR classes are allowed to burst above the guaranteed rate for a specified amount of time. However, because the UFR class has no maximum frame rate set, it can burst all it wants (although the bursty traffic may never be delivered). Because all circuits of the same service class being transmitted out the same port share the same QoS class buffer, simultaneous bursts can cause transient congestion. For example, because all VFR-rt circuits on a given port share the same VFR-rt buffer, bursting in one or more VFR-rt circuits can create congestion for the other VFR-rt circuits using the port. Similar issues face VFR-nrt and UFR circuits.

The network applies two congestion thresholds to the VFR-rt, VFR-nrt, and UFR service class output buffers to respond to network congestion. VFR-rt queues use an instantaneous measurement to determine their current queue levels. VFR-nrt and UFR queues use an average measurement to determine their current queue levels. See Figure 6 below.

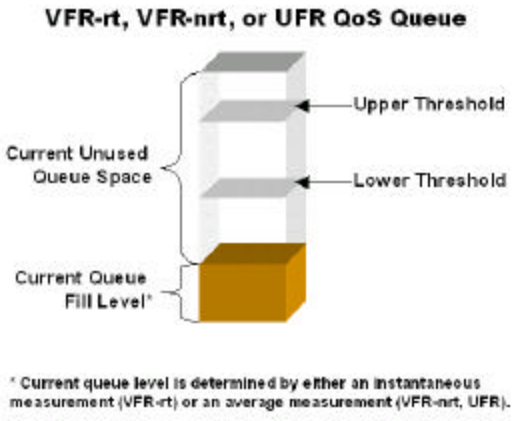

Figure 6—Two congestion thresholds exist in VFR-rt, VFR-nrt, and UFR queues.

Implementing Congestion Marking in a Service Provider Network

As explained above, the CFR queue is not subject to congestion marking.

For all other service classes, congestion marking is performed first on non-conforming virtual circuits, then on all circuits. When traffic in a VFR-rt, VFR-nrt, or UFR output queue passes the first threshold, the switch will mark non-conforming virtual circuits in that QoS class in both directions with FECN and BECN. If congestion in a particular QoS queue exceeds the second threshold, the switch will mark all virtual circuits in that QoS class in both directions with FECN and BECN. The thresholds ensure that traffic from non-conforming circuits is punished first if action is required to relieve congestion in the network.

©2001 Lucent Technologies, Inc.

Figures 7 through 9 illustrate this process.

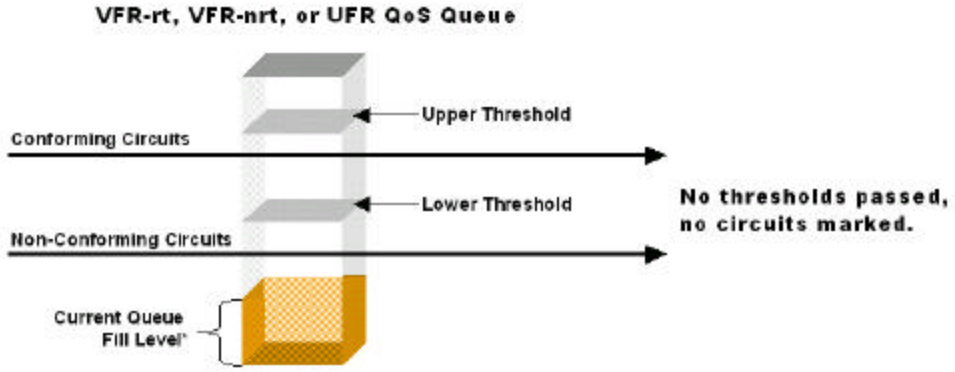

Figure 7—FECN/BECN markings are determined by queue saturation. When no thresholds are passed in the queue, no circuits are marked.

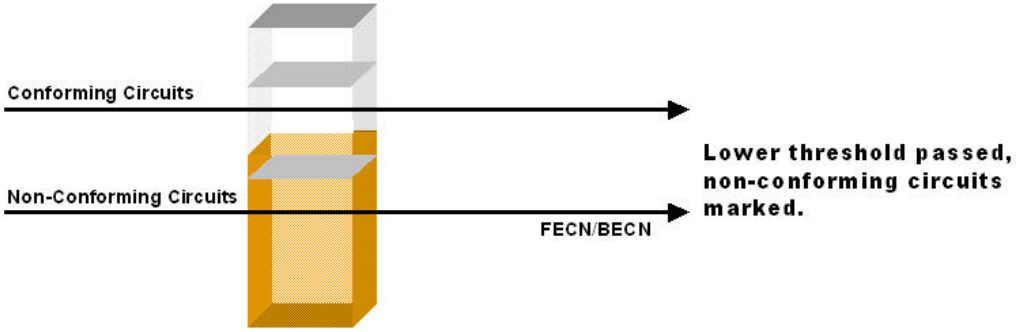

Figure 8—When volume in the VFR-rt, VFR-nrt, or UFR queue exceeds the lower threshold, the switch marks non-conforming circuits in both directions.

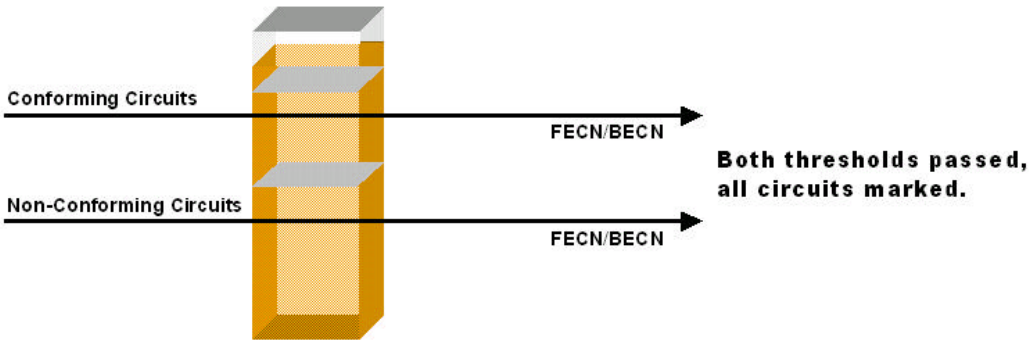

Figure 9—When volume in the VFR-rt, VFR-nrt, or UFR queue exceeds both thresholds, the switch marks all circuits in both directions.

©2001 Lucent Technologies, Inc.

(Note: When the network detects congestion, it sends notifications in both directions of a virtual circuit. That is, the network sets the FECN bit in the forward direction of the virtual circuit, and the BECN bit in the backward direction of the virtual circuit.)

4. Controlling Congestion in Frame Relay Networks That Rely on an ATM Backbone

End-to-end frame relay networks are becoming a thing of the past. Service providers have outgrown the capacity of their frame relay-only networks and are replacing them with multiservice networks that typically run ATM in the core and frame relay and IP at the edges. A fairly straightforward mechanism ensures that the frame relay congestion notification features described so far continue to operate in such multiservice networks.

When frame relay traffic is transmitted over an ATM network, it is first segmented into ATM cells using AAL 5 encapsulation. The new AAL5 cells contain fields used for congestion notification.

The AAL 5 cells travel on an ATM virtual circuit that meets the QoS requirements of the frame relay service. For example, incoming frame relay traffic marked for the frame CFR QoS class will travel over an ATM CBR circuit to ensure that its QoS requirements are met. Similarly, the network will route incoming frame relay data marked as frame VFR-rt over an ATM VBR-rt circuit, incoming frame VFR-nrt traffic over an ATM VBR-nrt circuit, and incoming frame UFR traffic over an ATM UBR circuit. Pairing frame relay QoS classes with their ATM QoS class equivalents ensures that the network will meet the QoS requirements for frame relay services that travel through a multiservice core.

If any ATM output buffer becomes congested, the network marks the congestion notifications in the AAL 5 header. At the ATM-to-frame relay border, the AAL 5 cells are reassembled into frame relay frames. The congestion markings contained in the AAL 5 header are written into the frame relay header. As a result, congestion occurring anywhere in the network, whether at the frame relay edge or the multiservice core, is communicated to the CPE, which can then respond appropriately.

5. Conclusion

Managing congestion in a service provider environment requires a congestion detection and notification system capable of accounting for each individual circuit's traffic contract and QoS requirements as well as the level of congestion in the output buffer. Lucent Priority Frame works in conjunction with the FECN/BECN marking systems to provide each circuit with fair access to network resources and to ensure that the network meets QoS requirements for premium services.

By implementing Priority Frame and FECN/BECN in their Lucent multiservice networks, service providers can transform network congestion from a network performance problem to a business opportunity that allows them to:

- Increase customer loyalty and reduce customer churn by earning their customer's trust
- Transform chronic congestion conditions into an opportunity to encourage customers to upgrade their level of network service
- Support service level agreements
- Use network resources more efficiently by cooperating with CPE to recover faster from transient congestion conditions

©2001 Lucent Technologies, Inc.

In spite of this, many frame relay services continue to operate without satisfactory congestion notification systems. This creates an opportunity for far-sighted service providers to differentiate their networks from their average competitor. Lucent's full fledged frame relay congestion notification for both pure frame relay networks and multiservice networks built with the Lucent GX 550™, CBX 500™, and B-STDX 8000/9000™ Multiservice WAN switches empowers service providers to address every market niche and to succeed in an ever more competitive marketplace.

©2001 Lucent Technologies, Inc.

Appendix: The Lucent Solution

The frame relay congestion detection and notification system described in this white paper is available as a standard feature on the Lucent GX 550, CBX 500, and B-STDX 9000 Multiservice WAN switches.

GX 550 Multiservice WAN Switch

The GX 550 Multiservice WAN Switch is a scalable, high-capacity switching system that provides the capacity, performance and fan-out capabilities required for multiservice networks. Specifically designed for the network core, the 25/32 Gbps GX 550 offers the scalability, availability, performance, and network management capabilities required in large service provider networks. The carrier-class GX 550 offers industry-leading port densities for OC-3/STM-1 and OC-12/STM-4 interfaces, industry-first OC-48/STM-16 high-speed trunking, lossless fabric switchover, port redundancy using 1+1 automatic protection switching (APS), line card redundancy, holdover timing, and extensive port and circuit capabilities. The robust GX 550 supports an efficient mix of all traffic classes on a single network, enabling service providers to use their GX 550 core to transport services such as voice, ATM access, frame relay, private line, MPLS, and IP on a single interface.

CBX 500 Multiservice WAN Switch

Service providers can deploy the CBX 500 multiservice WAN switch at the edge or in the core of their multiservice networks, depending upon their network size and growth projections. Service providers use the CBX 500's intelligence, reliability, and multiservice port density to create new revenue sources while scaling, simplifying, and consolidating converged public networks. Deployed at the network core or edge, this 5 Gbps switch delivers ATM, frame relay, and IP MPLS carrier-class networking to support voice, video, and data services.

B-STDX 8000/9000 Multiservice WAN Switches

Service providers have deployed the B-STDX 8000/9000 multiservice WAN switches for years at the core of their frame relay networks. Today, the B-STDX 9000 is naturally evolving to play a leading role at the access point to the next generation network. In addition to offering the broadest array of frame relay interfaces, standards support, and value-added features available today, the 1.2 Gbps B-STDX switch offers forward looking technologies such as IP Navigator™ MPLS and ATM-to-frame relay interworking modules for high-speed trunking and low-speed converged access. The B-STDX family offers a low-cost point of entry for smaller network deployments that do not require the bandwidth and raw horsepower of the CBX 500 or GX 550 products. As these networks grow, the B-STDX naturally moves to the network edge, where its breadth of interfaces allows service providers to continue to offer a variety of access options to their customers.

The Lucent GX 550, CBX 500, and B-STDX 9000 multiservice WAN switches offer the high availability, reliability, scalability, end-to-end SVC management, robust network management, and outstanding network performance required by carrier-class network environments. When used as a unit, the Lucent multiservice switching family can handle ATM, frame relay, and IP MPLS traffic, while providing trunking interfaces to the optical core.

©2001 Lucent Technologies, Inc.