

# Virtual Trunking “Wrap-Around” Solution and Traffic Shaping on Cisco IGX 8400

## Introduction

With widespread availability of public ATM services, customers deploying ATM wide-area networks (WANs) now have an alternate means of interconnecting IGX™ switches. Using public ATM services may prove to be economically attractive compared to using leased lines for interconnectivity. This requires the ability to support multiple virtual trunks on a single, physical ATM trunk interface, which is fully implemented on the IGX switch in Release 9.2. However, due to lack of virtual trunking support in Release 9.1, in the interim customers can use the hardware-based “wrap-around” solution to implement virtual trunking. This interim wrap-around solution requires a pair of UXM modules, one of which is configured as the trunk module and the other a UNI port module. The UXM UNI module simply switches virtual path connections (VPC) from the physical trunks and multiplexes them onto a single UNI port that is accessing the public ATM network. Corresponding to each physical trunk, the UNI port uses VPIs provided by the public network.

The virtual interface (VI) and VC/VP traffic shaping functionality is an important product differentiation for the IGX switch in enabling key customer applications over public ATM networks. The VI traffic shaping shapes the aggregate traffic of a VI on a multi-quality of service (QoS) basis so that the aggregate traffic conforms to the traffic description and characterization expected by the public ATM network that the VI is interfacing with. Multi-QoS based traffic shaping ensures that the QoS of each connection is respected while the aggregate traffic is being shaped. In particular, the multi-QoS based traffic shaping ensures that the real-time characteristics of real-time traffic are protected. The VC/VP traffic shaping shapes each individual VC or VP by scheduling the cells using the weighted fair queuing (WFQ) technique to ensure appropriate conformance dictated by the service category. For example, for a CBR connection, the VC/VP traffic shaping schedules the cells according to the negotiated peak cell rate (PCR) and cancels most of the cell delay variation (CDV) that may have accumulated in the path.

This application note describes the implementation and configuration of traffic shaping (VI and VC/VP) and virtual trunking using the wrap-around method supported on the IGX Release 9.1.08 and UXM firmware (FW) AAE. Further, it also documents the CDV performance with and without VP traffic shaping for a given configuration.

## Traffic Shaping for Virtual Trunking

IGX switches can be interconnected making use of the public ATM network using the virtual trunking feature. The virtual trunking capability allows definition of multiple logical trunk interfaces within a single, physical trunk interface. The “many-to-one” virtual trunk to port relationship further produces a one-to-many “fanout” connectivity effect, achieving phenomenal equipment cost savings. All the connections on a virtual trunk are mapped to a single VPC with a public network assigned virtual path identifier (VPI). The header format of cells may match the ATM user-to-network interface (UNI) or ATM network-to-network interface

(NNI) format since the interface to the public ATM network could be a UNI or NNI port. A software-configured VCI value that is unique for the cell is passed transparently across the public ATM network. When the other end of the network receives the cells, another IGX maps this VPI/VCI back to a correct cell header.

The virtual trunking capability has been developed on the UXM in Release 9.2. With the use of a wrap-around port interconnection solution explained later, the virtual trunking feature can be emulated in Release 9.1.08. UXM FW AAE is required to implement traffic shaping

The service class that the VPC uses for transporting virtual trunk traffic across the public ATM network affects the type of traffic that can be carried. To maintain QoS guarantees, appropriate configuration of the traffic classes supported by a VPC type is required. The following is the recommended configuration for an ATM port module (UXM) carrying VCCs of type CBR, RT-VBR, NRT-VBR, ABR, or UBR to be transported by corresponding VPC on the trunk module (BTM, ALM/B, UXM).

Table 1 ATM Type VCCs

VPC Type on ATM Trunk Module (BTM,ALM/B, UXM)	Component VCC Type on ATM Port Module (UXM)				
	CBR	RT-VBR	NRT-VBR	ABR	UBR
CBR	X	X	X	X	X
RT-VBR		X	X	X	X
NRT-VBR			X	X	X
ABR			X	X	X
UBR					X

Table 2 shows the recommended configuration for a fast packet-based port module (LDM, HDM, UVM, CVM, FRM, UFM) carrying VCCs of type high-priority, timestamped, nontimestamped, voice, bursty data A, or bursty data B to be transported by corresponding VPC on the trunk module (BTM, ALM/B, UXM).

Table 2 Fast Packet Type VCCs

VPC Type on ATM Trunk Module (BTM, ALM/B, UXM)	Component VCC Type on Fast Packet Based Module (LDM, HDM, UVM, CVM, FRM, UFM)					
	High priority	Timestamped	Nontimestamped	Voice	Bursty Data A	Bursty Data B
CBR	X	X	X	X	X	X
RT-VBR	X	X	X	X	X	X
NRT-VBR	X	X			X	X
ABR	X	X			X	X
UBR	X	X			X	X

As each virtual trunk interfaces with a public ATM network VPC, traffic shaping is required to ensure that the VPC traffic conforms to the traffic policing parameters (CDV) for the VPC at the entrance of the public ATM network. The public ATM network may drop the nonconformant traffic.

There are two ways of performing traffic shaping for a virtual trunk:


- Single-QoS traffic shaping—VCCs of a virtual trunk are shaped through a single FIFO queue before being aggregated onto a VPC

Figure 1 Single-QoS Traffic Shaping



- Multi-QoS traffic shaping—VCCs of a virtual trunk are shaped through a set of CoS queues before being aggregated onto a VPC

Figure 2 Multi-QoS Traffic Shaping



Both single-QoS and multi-QoS traffic shaping ensure the traffic on a VPC is compliant with the public ATM service traffic contract. Multiple-QoS traffic shaping can further ensure that the QoS of each connection is respected while being shaped. Single-QoS traffic shaping is applicable if the components VCCs are all of real-time type (such as CBR, non-timestamped, voice), otherwise multi-QoS traffic shaping is required.

### Virtual Trunking Wrap-Around Solution

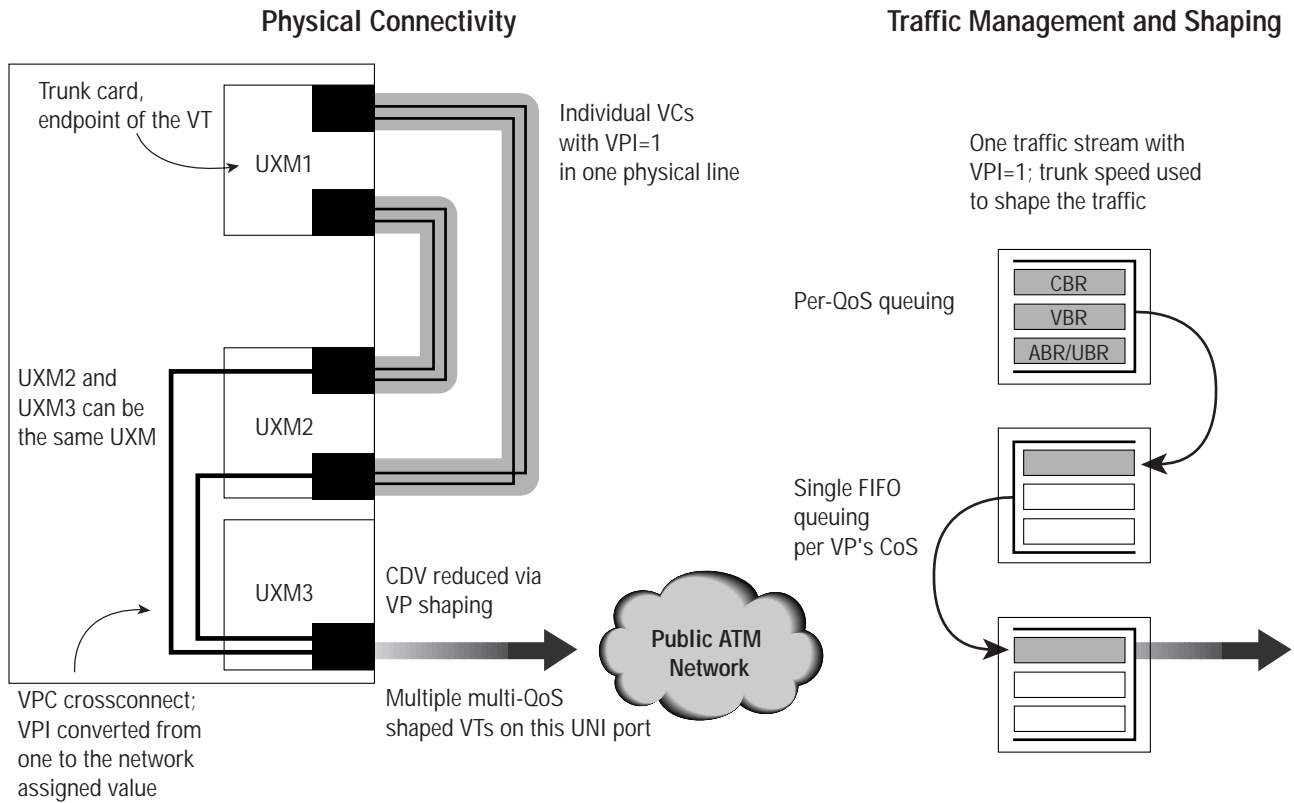
The UXM modules do not support the virtual trunking feature in Release 9.1. However, an interim wrap-around solution can be used to provide virtual trunking. Release 9.1.08 and UXM FW AAE supports VC/VP traffic shaping on the UXM to reduce the CDV of a virtual trunk to comply with any public ATM service. One physical interconnect between two physical ports is required for each virtual trunk.

The number of physical ports available for wrap-around interconnection limits the number of virtual trunks to the public ATM network that the UNI port can support. This wrap-around solution, although it consumes a large number of ports, enables savings on recurring costs for accessing UNI ports on public ATM switches. The solution is described below.

1. Use one UXM module as a trunk module (endpoint of virtual trunk). Each physical trunk port supports one virtual trunk. Cells transmitted from a trunk port are all with VPI=1.
2. Use VI shaping at the trunk port<sup>1</sup> and VP shaping at the UNI port to reduce the CDV entering the public network. Multi-QoS traffic shaping is supported and cells are queued in multiple CoS queues at the trunk port.
3. Use one or more UXM modules configured as a port module to set up VP connections to convert the VPI from one to the network assigned VPI value. The ports on the second and third UXM modules (UXM2 and UXM3 port modules) shown in Figure 3 can be on the same UXM port module.

<sup>1</sup> The standard user command `cnfrk` requires deleting the trunk first to change the RCV and XMT trunk rate. If the user requests changing the VI shaping rate without affecting the service, a workaround via the `scp/xcp` command is needed. This workaround is beyond the scope of this document.


Figure 3 Virtual Trunking with Multi-QoS Traffic Shaping and Reduced CDV



### Traffic Shaping for Supporting VCC Services over a Public VPC

Public ATM network VPC service can also be used to cost-effectively provide virtual circuit connection (VCC) services from a UNI port over a public ATM network VPC to remote subscribers as shown in Figure 4.

Figure 4 VCC Service over a Public VPC



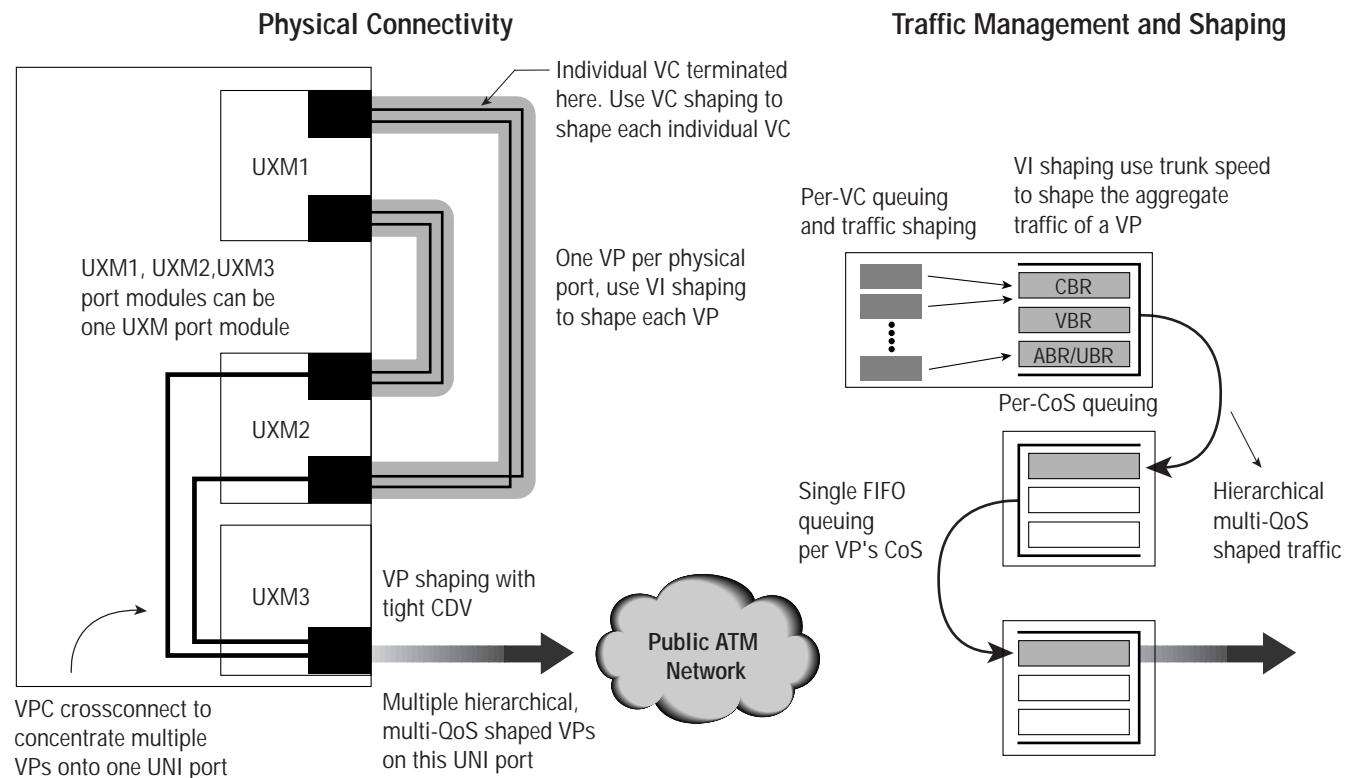
For a virtual trunk, the individual user connections are carried on a VPC and transparently transported across the public ATM network. For virtual trunking, multi-QoS traffic shaping is required and per-VC shaping is not needed. However, when a public ATM network VPC is used to carry multiple VCCs to remote customers, hierarchical traffic shaping at both VP and VC levels is required. Traffic shaping at VP level is required for passing VPC aggregate traffic through the public ATM network in a conformant fashion. Traffic shaping at VC level for each individual VCC is needed for fair bandwidth sharing and for matching the receiving speeds of the CPE equipment. Even if receiving speeds is not an issue, the CPE equipment may not have sufficient buffer to handle the large CDV of the incoming cells. Using VC traffic shaping cancels most of the CDV that may have accumulated in the path.

Release 9.1.08 and UXM FW AAE supports hierarchical traffic shaping at both VC and VP level. A wrap-around solution similar to that used in virtual trunking is required. Traffic shaping at VC level is accomplished by turning on per-VC shaping on the first UXM1 ports. VP shaping is performed via VI shaping on the first UXM. Additionally, VP shaping can be used at the UNI port on the last UXM card to control the CDV tightly.

The first UXM is configured as a port card in contrast to a trunk card in the wrap-around solution for virtual trunk. One physical wrap-around interconnect between two physical ports is required for each VP. As shown in Figure 5, UXM1, UXM2, and UXM3 port modules can be configured on a single UXM port module. The number of physical ports available for wrap-around interconnection thus limits the number of VPs per UNI port and number of UXM port modules required.

The following Figure 5 illustrates hierarchical traffic shaping support.

Figure 5 Support of Hierarchical, Multi-QoS Traffic Shaping



## Features and Limitations

Virtual trunking supported on the UXM modules has the following features and limitations:

- The maximum number of virtual trunks per card is 15 for UXM
- Each virtual trunk has a set of 16 class-of-service (CoS) queues
- The maximum number of logical (physical and virtual) trunks per IGX node is 32
- Valid VPC types are CBR, VBR, and ABR
- ATM Forum Standard UNI and NNI VPI and VCI ranges are supported
- A VPC is not allowed to be routed over a virtual trunk. The routing algorithm excludes all virtual trunks from the routing topology. The reason for this restriction is that by definition a VPC cannot be routed over another VPC
- A virtual trunk cannot be used as a feeder trunk
- ILMI between a virtual trunk and a foreign switch is supported
- F4/F5 OAM flows over virtual trunks are supported only when the flows originate from the cloud
- Node numbers 0 to 31 on the IGX router cannot be used
- The public ATM network will supply the clock

- The following F4/F5 flows are supported:

AIS/RDI OAM Flows:

- F5 (VCC) OAM flows are supported for end-to-end connections through a virtual trunk
- F4 (VPC) OAM flows not supported via virtual trunk
- F4 flows not supported between the cloud and the virtual trunk

OAM (Test Delay) Loopback

- F5 (VCC) OAM flows are supported for end-to-end connections through a virtual trunk
- F4 (VPC) OAM flows not supported via VT

- Y-cable redundancy is supported in the wrap-around virtual trunking solution provided there is no change in trunk speed or trunk speed is changed using the standard cnfrk command. In order to implement Y-cable redundancy, additional interfaces would be required. For instance as shown in Figure 6, for each trunk say 6.1 an additional trunk interface is required (6.2), a corresponding additional port interface is required (2.4), and for all the trunks an additional port interface (2.8) is required to implement Y-cable redundancy. Hence 6.1 and 6.2 would be connected by a Y-cable, similarly with 2.3, 2.4, 2.7, and 2.8.

**Note:** Changing the trunk speed using cnfrk command is service-affecting since it requires deleting the trunk first.

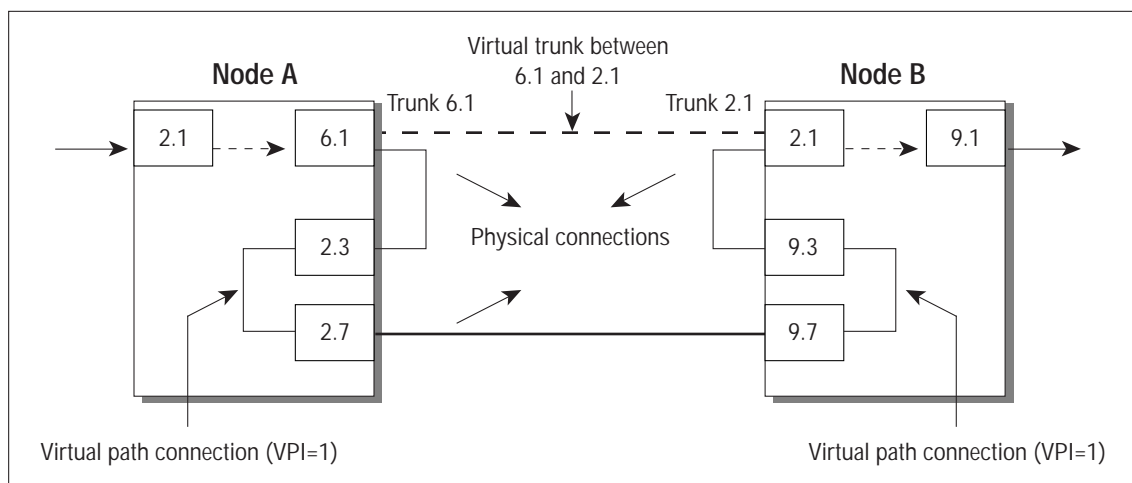
- When using a public ATM service, in order for virtual trunking using a VP service to be used, it is very important that the public ATM service provider configure the below besides the type of VP service (CBR, VBR, and so on.), rate of VP connection, and CDVT.

- Carrier should enable ILMI traps on the service provider switch port so that reroutes will be faster
- Carrier must configure a full VCI range of 16 bits on the VPC provided to the customer

### Virtual Trunk Wrap-Around Setup Procedure

The following Figure 6 illustrates the virtual trunk wrap-around solution. On each node, there are two UXM modules; one is used as a trunk card and the other as a port card. Note that the trunk card and the port card do not need to be on the same node as shown in the figure. Trunk 6.1 on Node A and Trunk 2.1 on Node B are the two ends of the virtual trunk routed through the public ATM cloud. Since virtual trunking is not supported on UXM until Release 9.2, the output traffic from the trunk modules does not necessarily use the same virtual path identifier (VPI) specified by the network, but instead uses VPI =1. Therefore, a virtual path connection is added to convert the VPI to the value specified by the public ATM network. Virtual path connections are added between port 2.3 and port 2.7 on Node A and between port 9.3 and port 9.7 on Node B. Traffic from the trunk modules is wrapped around the virtual path connections before it is sent to the public ATM network.

Figure 6 UXM Setup



## UXM Configuration

1. Bring up three ports on card 2 at Node A (ports 2.1, 2.3, and 2.7 are used in this example)

```
upln 2.1
upln 2.3
upln 2.7
upport 2.1
upport 2.3
upport 2.7
```

2. Bring up three ports on card 9 at Node B (ports 9.1, 9.3, and 9.7 are used in this example). Ensure that the line and trunk configuration on both ends have the same parameters (such as line coding, payload scramble, and so on)

```
upln 9.1
upln 9.3
upln 9.7
upport 9.1
upport 9.3
upport 9.7
```

3. Turn on VP shaping on port 2.7 and 9.7 at Node A and Node B respectively

```
cnfln 2.7 ..... y
cnfln 9.7 ..... y
```

4. Add a CBR virtual path connection with PCR = trunk\_speed<sup>2</sup> in cells per second, and CDV = 10000 usec between ports 2.3 and 2.7 on Node A. The VPI at port 2.3 should be set to 1, and the VPI at port 2.7 should be the VPI that the public ATM network specifies. Assume that the public ATM network specifies VPI = 99 in this example, the addcon command should be

```
addcon 2.3.1 * Node A 2.7.99 * cbr trunk_speed * 10000 *
```

**Note:** In the case of T-1/E-1 trunk interfaces, when adding a low trunk rate (150cps, 301cps) the CDV may have to be configured at a high value, around 400000 usec, in order to avoid comm failure on the trunk.

<sup>2</sup>Trunk\_speed is the rate of the virtual trunk. It should be the same as the PCR (in cells per second) of the virtual path connections 2.3.1.\* at Node A and 9.3.1\* at Node B, and the trunk 6.1 on Node A and trunk 2.1 on Node B.

5. Similar as in Step 5, add a CBR virtual path connection with PCR = trunk\_speed, and CDVT = 10000 usec between ports 9.3 and 9.7 on Node B.

```
addcon 9.3.1 * Node A 9.7.99 * cbr trunk_speed * 10000 *
```

6. Up trunk 6.1 on Node A and trunk 2.1 on Node B

```
Node A: uptrk 6.1
Node B: uptrk 2.1
```

7. Configure the trunk speed on trunk 6.1 at Node A and trunk 2.1 at Node B

```
Node A: cnftrk 6.1 <trunk_speed> <trunk_speed> * * * ....
Node B: cnftrk 2.1 <trunk_speed> <trunk_speed> * * * .....
```

The cnftrk command on UXM requests for both transmission rate and receiving rate. Make sure that the transmission rate agrees with the receiving rate on the other end of the trunk.

8. Add a trunk between 6.1 on Node A and 2.1 on Node B

```
Node A: addtrk 6.1
```

In the above procedure, configure a virtual trunk of 2 Mbps using trunk\_speed=4716cps.

## Changing the Virtual Trunk Rate

This section describes the procedure to change the virtual trunk rate.

**Note:** Changing the trunk speed using this procedure causes service interruption since the trunk has to be deleted first. Deleting the trunk may cause connections on that trunk to be rerouted or even deleted if deleting the trunk splits the network.

### Procedure

1. Delete the trunk. Note that this may cause connections routed through the virtual trunk being rerouted or even deleted from the network.  
Node A: *deltrk 6.1*  
or  
Node B: *deltrk 2.1*
2. Use *cnftrk* to change the trunk rate  
Node A: *cnftrk 6.1 <new\_trunk\_rate> <new\_trunk\_rate> \* \* \* .....*  
Node B: *cnftrk 2.1 <new\_trunk\_rate> <new\_trunk\_rate> \* \* \* .....*
3. Use *cnfcon* to change the PCR of the virtual path connections  
Node A: *cnfcon 2.3.1.\* new\_trunk\_rate \* \* \**  
Node B: *cnfcon 9.3.1.\* new\_trunk\_rate \* \* \**
4. Add the virtual trunk back  
Node A: *addtrk 6.1*
5. If by deleting the trunk, connections were lost, then add the user connections again

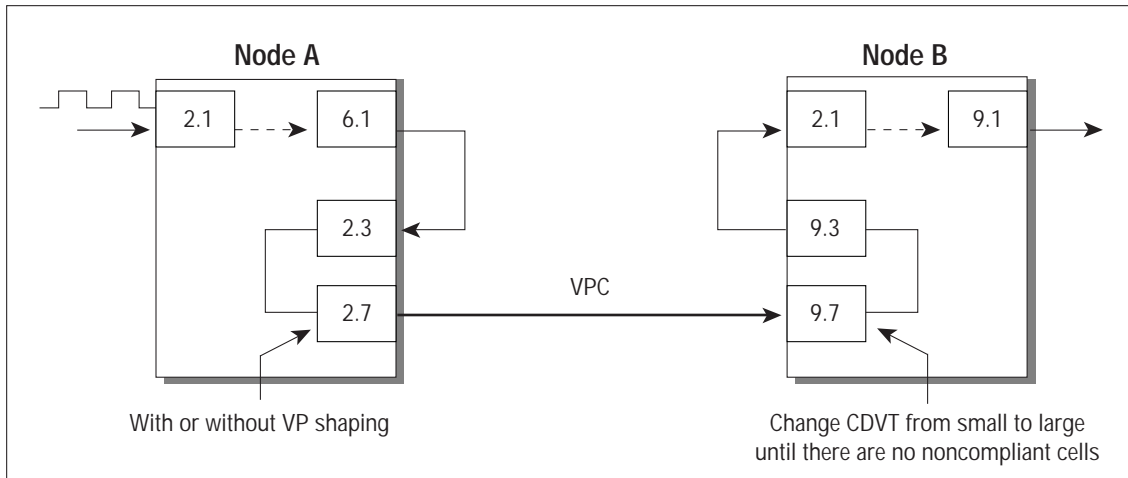
The trunk between Node A trunk 6.1 and Node B 2.1 has been set up. If the trunk rate has to be increased from 2 Mbps to 3 Mbps, follow the above procedure using *new\_trunk\_rate=7075cps*.

## Measured Cell Delay Variation (CDV)

The measured worst-case CDV values for the output traffic from the IGX router with or without VP shaping are presented in this section assuming that the carrier switch represented by Node B is a Cisco WAN ATM switch. Note that the CDV of the output traffic should comply with the CDVT that the public ATM network requires, or else cells may be discarded at the edge of the public ATM network.

The CDV of the output traffic from port 2.7 (in Figure 7) on IGX node A is measured using the following method. Change the CDVT setting for the VPC at node B until there are no non-compliant cells. The reading of CDVT thus indicates the CDV of the output traffic of the VPC. With multiple virtual trunks in the same physical trunk and per VP shaping enabled, the CDV of each virtual trunk may not be affected. The CDV is not affected since cells from a virtual trunk are queued in an individual queue and shaped on a per-VP basis.

Figure 7 CDV Measurement Method



The set up procedure for the CDV measurement is as follows:

1. Add a data connection between port 2.1 of Node A and port 9.1 of Node B.  
`addcon 2.1.vpi.vci NodeB 9.1.vpi.vciubr trunk_speed * * d d n`
2. Change the CDVT of 9.7.99.\* of Node B  
`cnfcon 9.7.99.* <trunk_speed> * <cdvt_to_measure> *`
3. Use dspchstats to check if any noncompliant cells are discarded at port 9.7 on Node B.  
`dspchstats 9.7.99.* 1`

The following tables summarize the measured CDV values for an IGX router with a worst-case input traffic. The worst-case input traffic pattern is shown in following Figure 8.

Figure 8 Traffic Pattern

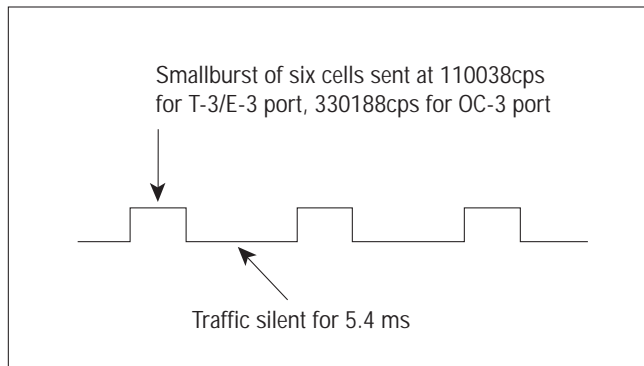


Table 3 UXM CDV Measurement on T-1 Port

## T-1 port

Virtual Trunk Rate	CDV (in microseconds)	
	VP Shaping Off	VP Shaping On
64 Kbps (150 cps)	6500	< 500
128 Kbps (301 cps)	10000	< 600
256 Kbps (603 cps)	10500	< 500
0.5 Mbps (1179 cps)	10500	< 700
1 Mbps (2358 cps)	8000	< 500
1.536 Mbps (3622 cps)	20	< 30

Table 4 UXM CDV Measurement for E-1 Port

## E-1 port

Virtual Trunk Rate	CDV (in microseconds)	
	VP Shaping Off	VP Shaping On
64 Kbps (150 cps)	1100	< 500
128 Kbps (301 cps)	10000	< 450
256 Kbps (603 cps)	5500	< 500
0.5 Mbps (1179 cps)	11000	< 400
1 Mbps (2358 cps)	2400	< 400
1.92 Mbps (4528 cps)	250	< 200

Table 5 UXM CDV Measurement for T-3 Port

## T-3 port

Virtual Trunk Rate	CDV (in microseconds)	
	VP Shaping Off	VP Shaping On
0.5 Mbps (1179 cps)	10000	< 30
1 Mbps (2358 cps)	5000	< 30
2 Mbps (4716 cps)	3000	< 30
3 Mbps (7075 cps)	1600	< 30
4 Mbps (9433 cps)	1300	< 30
5 Mbps (11792 cps)	800	< 30
6 Mbps (14150 cps)	800	< 30
7 Mbps (16509 cps)	600	< 30
8 Mbps (18867 cps)	550	< 30
9 Mbps (21226 cps)	475	< 30
10 Mbps (23584 cps)	475	< 30
20 Mbps (47169 cps)	140	< 30
30 Mbps (70754 cps)	50	< 30
40 Mbps (94339 cps)	5	< 30

Table 6 UXM CDV Measurement for E-3 Port

## E-3 port

Virtual Trunk Rate	CDV (in microseconds)	
	VP Shaping Off	VP Shaping On
0.5 Mbps (1179 cps)	11000	< 30
1 Mbps (2358 cps)	5300	< 30
2 Mbps (4716 cps)	2600	< 30
3 Mbps (7075 cps)	1500	< 30
4 Mbps (9433 cps)	1200	< 30
5 Mbps (11792 cps)	1100	< 30
6 Mbps (14150 cps)	700	< 30
7 Mbps (16509 cps)	600	< 30
8 Mbps (18867 cps)	500	< 30
9 Mbps (21226 cps)	450	< 30
10 Mbps (23584 cps)	400	< 30
20 Mbps (47169 cps)	80	< 30
30 Mbps (70754 cps)	18	< 30

Table 7 UXM CDV Measurement for OC-3 Port

## OC-3 port

Virtual Trunk Rate	CDV (in microseconds)	
	VP Shaping Off	VP Shaping On
0.5 Mbps (1179 cps)	< 6500	< 50
1 Mbps (2358 cps)	< 2200	< 50
2 Mbps (4716 cps)	< 700	< 50
3 Mbps (7075 cps)	< 450	< 50
4 Mbps (9433 cps)	< 350	< 50
5 Mbps (11792 cps)	< 300	< 50
6 Mbps (14150 cps)	< 250	< 50
7 Mbps (16509 cps)	< 250	< 50
8 Mbps (18867 cps)	< 200	< 50
9 Mbps (21226 cps)	< 200	< 50
10 Mbps (23584 cps)	< 150	< 50
20 Mbps (47169 cps)	< 100	< 50
30 Mbps (70754 cps)	< 100	< 50
40 Mbps (94339 cps)	< 50	< 50
50 Mbps (117924 cps)	< 50	< 50
70 Mbps (165094 cps)	< 200	< 50
100 Mbps (235849 cps)	< 150	< 50
130 Mbps (306603 cps)	< 100	< 50
155 Mbps (365566 cps)	< 100	< 50

For more information, please contact Mihir Maniar at [mmaniar@cisco.com](mailto:mmaniar@cisco.com).



Corporate Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

European Headquarters  
Cisco Systems Europe s.a.r.l.  
Parc Evolic, Batiment L1/L2  
16 Avenue du Quebec  
Villebon, BP 706  
91961 Courtaboeuf Cedex  
France  
<http://www-europe.cisco.com>  
Tel: 33 1 69 18 61 00  
Fax: 33 1 69 28 83 26

Americas  
Headquarters  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-7660  
Fax: 408 527-0883

Asia Headquarters  
Nihon Cisco Systems K.K.  
Fuji Building, 9th Floor  
3-2-3 Marunouchi  
Chiyoda-ku, Tokyo 100  
Japan  
<http://www.cisco.com>  
Tel: 81 3 5219 6250  
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the  
**Cisco Connection Online Web site at <http://www.cisco.com/offices>.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE  
Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia  
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore  
Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela