

Virtual Trunking and Traffic Shaping on BPX 8600 Series

Introduction

With the widespread availability of public ATM services, service provider or enterprise customers deploying ATM wide-area networks now have an alternative means of interconnecting BPX switches—Virtual Trunking, as opposed to using more expensive and less flexible leased-line services or building their own transmission facility. To build a meshed network while only connecting to the public ATM service network using one single physical port, customers need the ability to “logically bundle” their network trunks. The BPX[®] (and IGX[™]) switch feature that supports this is called Virtual Trunking. The “many-to-one” virtual trunk-to-port relationship produces a one-to-many “fanout” connectivity effect, achieving phenomenal savings on equipment and recurring service costs. Virtual Trunking on the broadband switch module (BXM) is fully implemented on the BPX switch in Release 9.2. In Releases 9.1 and 8.4, users can use a hardware-based wraparound solution to implement Virtual Trunking.

Using public ATM services, customers are moving toward E3/T3, OC-3/STM-1, and even higher trunk speeds—speeds that were previously not economically viable. For service provider customers who already have a core ATM network infrastructure with no Cisco products in place, Virtual Trunking also provides the ability to expand the network using a Cisco solution while making use of the existing network for transport in the core. Using Virtual Trunking over an ATM network with no Cisco products instead of connecting to such a network via ATM Network-to-Network Interface (NNI) not only preserves the advanced BPX/IGX networking and traffic management features but also allows all the BPX and IGX switches to be managed as a single network.

All the connections on a virtual trunk are mapped to a single virtual path connection (VPC) with a public network-assigned virtual path identifier (VPI). The header format of cells may match the ATM-User-Network Interface (UNI) or ATM-NNI format because the interface to the public ATM network could be a UNI or NNI port. A software-configured virtual channel identifier (VCI) value that is unique for the cells of a VC is passed transparently across the public ATM network. When cells are received at the other end of the network, another BPX or IGX switch maps this VPI/VCI back to the original cell header.

The VI (virtual interface) and per-virtual circuit/virtual path (VC/VP) traffic-shaping features on the BXM card are key differentiators of the BXM Virtual Trunking feature. Each physical or virtual trunk is mapped to a VI.

The VI traffic shaping shapes the aggregate traffic of a virtual trunk on a multi-QoS basis so that the aggregate traffic is conformant to the traffic descriptor and characterization expected by the network. The VI traffic shaping reduces the cell delay variation (CDV) of a virtual trunk to the minimum, and the user can choose the least tariffed permanent virtual circuit (PVC) service for the same peak cell rate (PCR) while meeting the more stringent cell delay variation tolerance (CDVT) requirement. 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 per-VC/VP traffic shaping shapes each individual VC or VP by scheduling the cells using the WFQ (Weighted Fair Queuing) technique to ensure appropriate conformance dictated by the service category. For example, for a constant-bit-rate (CBR) connection, the VC/VP traffic shaping schedules the cells according to the negotiated PCR, and most of the CDV that may have accumulated in the path is canceled. The per-VC/VP traffic-shaping function is implemented on the BXM card as a virtual source (VS) function. It supports ATM Forum standard available-bit-rate (ABR) VS/VD behavior. Through a software upgrade, the VC/VP traffic-shaping function could also support new and different shaping protocols based on generation and processing of resource management (RM) cells. Such flexibility makes deploying future new services fast and easy.

The combination of VI and per-VC/VP traffic shaping not only enables multi-QoS traffic shaping for a virtual trunk, but also facilitates hierarchical traffic shaping on a UNI port with simple, future software upgrades. Hierarchical traffic shaping on a UNI port enables a BPX service provider customer to offer virtual-channel-connection (VCC) services to remote subscribers by using the public ATM VPC services as the transport.

Traffic Shaping for Virtual Trunking

The service class that the VPC uses for transporting virtual trunking traffic across the public ATM network affects the type of traffic that can be serviced. To maintain QoS guarantees, appropriate configuration of the traffic classes supported by a VPC type is required. Below is the recommended configuration for possible combinations of ATM classes of service in a virtual trunk that can be transported by different VPC types in a public ATM network.

| VPC Type | Component VCC Type | | | | |
|----------|--------------------|--------|---------|-----|-----|
| | ATM | | | | |
| | 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 |

The following is the recommended configuration for possible combinations of Fast Packet-based traffic (from IGX and MGX™ 8220 service modules) carrying high-priority, time-stamped, non-time-stamped, voice, bursty data A, or bursty data B VCCs in a virtual trunk that can be transported by different VPC types in a public ATM network.

| VPC Type | Component VCC Type | | | | | |
|----------|--------------------|--------------|------------------|-------|---------------|---------------|
| | Fast Packet | | | | | |
| | High Priority | Time-Stamped | Non-Time-Stamped | Voice | Bursty Data A | Bursty Data B |
| CBR | x | x | x | x | x | x |
| RT-VBR | 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 for the VPC at the entrance of the public ATM network.

Traffic shaping for a virtual trunk can be performed in two ways:

Single-QoS Traffic Shaping

VCCs of a virtual trunk are shaped through a single first-in/first-out (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 class-of-service (CoS) queues before being aggregated onto a VPC.

Figure 2 Multi-QoS Traffic Shaping



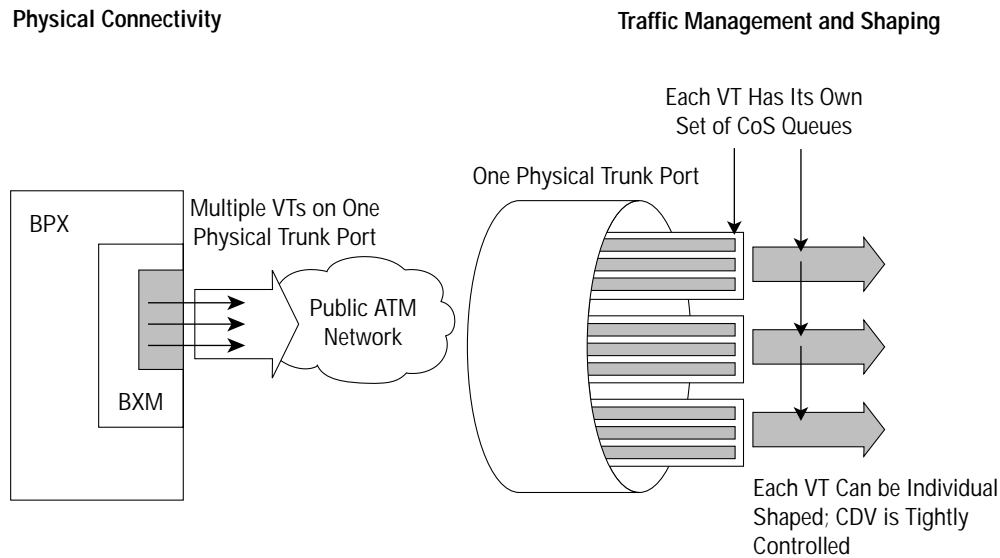
Both single-QoS and multi-QoS traffic shaping ensure the traffic on a VPC is compliant with the traffic contract for PCR. Multiple-QoS traffic shaping can further ensure that the QoS of each connection is respected while being shaped. Single-QoS traffic shaping is applicable if the component VCCs are all real-time type (for example, CBR, real-time variable bit rate [RT-VBR], voice). Otherwise, multi-QoS traffic shaping is required.

Release 9.2 BXM Virtual Trunking

The Virtual Trunking capability has already been developed for the BPX broadband network interface (BNI) trunk card. This capability is now extended to the BPX BXM and IGX universal-switching-module (UXM) cards in Release 9.2.

Release 9.2 supports up to 31 virtual trunks on a single BXM physical trunk port. The Release 9.2 Virtual Trunking feature includes per-virtual-trunk traffic shaping, enabling a tightly controlled CDV for each virtual trunk.

Figure 3 Virtual Trunking with Multi-QoS Traffic Shaping and Tightly Controlled CDV in Release 9.2



The Virtual Trunking supported on the BXM card has the following features and limitations:

1. The maximum number of virtual trunks per card is 31 for BXM.
2. Each virtual trunk has a set of 16 CoS (class-of-service) queues.
3. The maximum number of logical (physical and virtual) trunks per BPX node is 64.
4. Valid VPC types are CBR, VBR, and ABR.
5. ATM Forum Standard UNI and NNI VPI and VCI ranges are supported.
6. A VP cannot be routed over a virtual trunk. The routing algorithm excludes all virtual trunks from the routing topology for VP connections. The reason for this restriction is that by definition a VPC cannot be routed over another VPC.
7. A virtual trunk cannot be used as a feeder trunk.
8. Virtual Trunking supports partial Integrated Local Management Interface (ILMI) (including receiving ILMI traps on VPC status changes in the cloud and periodic cloud VPC ILMI status querying) over the UNI/NNI between a virtual trunk and a foreign switch.
9. Virtual Trunking supports F4/F5 flows on a virtual trunk as follows:
Alarm indication signal/remote defect identification Operation, Administration, and Maintenance cell (AIS/RDI OAM) flows:
 F5 (VCC) OAM flows are supported for end-to-end connections through a virtual trunk.
 F4 (VPC) OAM flows are not supported on virtual trunks.
OAM loopback:
 F5 (VCC) OAM flows are supported for end-to-end connections through a virtual trunk.
 F4 (VPC) OAM flows are not supported on virtual trunks.
10. The virtual trunk rate can be reconfigured while the trunk is in use, but reroute of connections will be triggered if the new rate results in too few resources to support the current load of connections.

Interim BXM "Wraparound" Virtual Trunking Solution for Releases 8.4 and 9.1

Release 8.4

Release 8.4 does not support Virtual Trunking on the BXM card, but an interim "wraparound" solution can be used to provide Virtual Trunking and per-virtual-trunk traffic shaping, as described below.

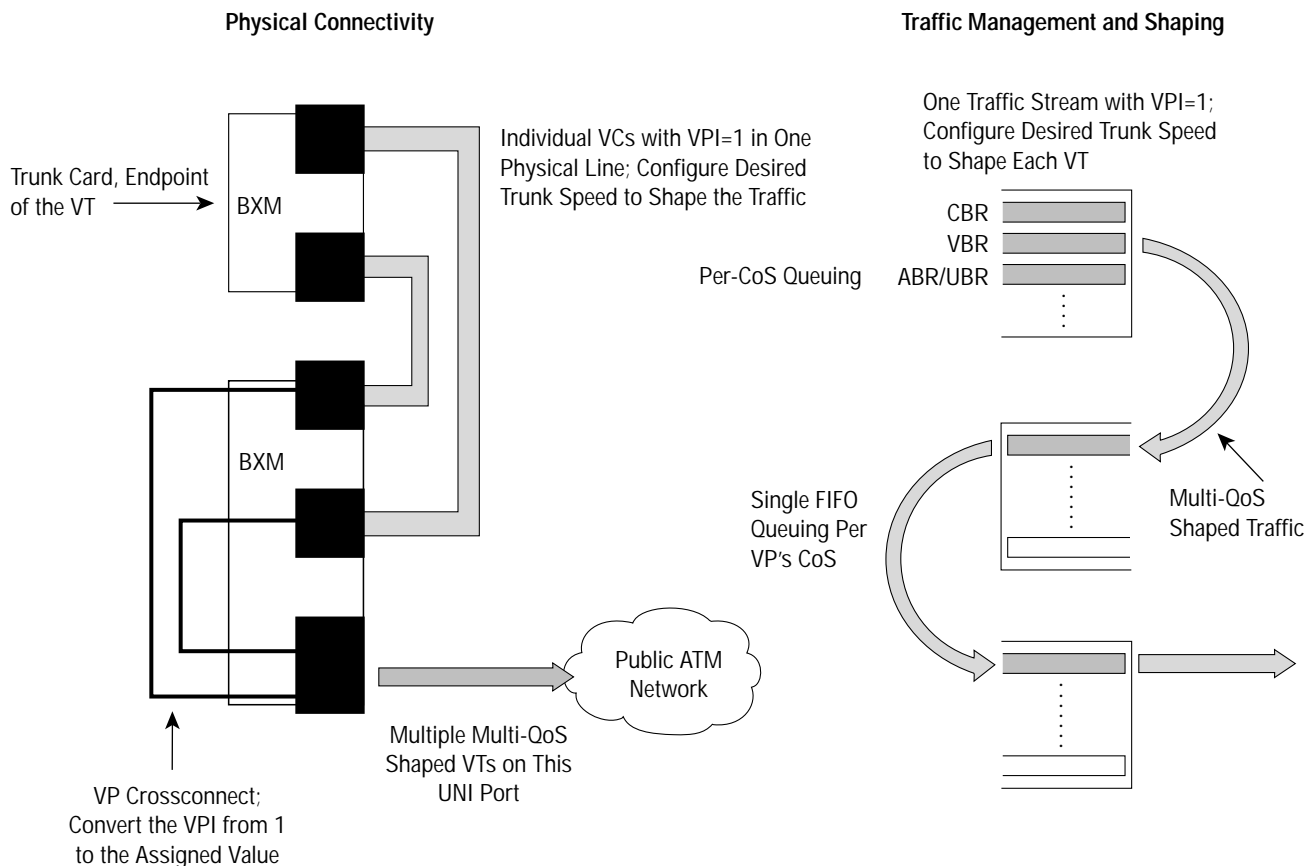
1. Use one BXM card as a trunk card (endpoint of virtual trunk). Each trunk port supports one virtual trunk. Cells transmitted from a trunk port all have VPI = 1.
2. Use the `cnftrk` command to set the Transmit Trunk Rate parameter to the desired shaping rate at the trunk port¹. Multi-QoS traffic shaping is supported. Cells are queued in multiple CoS queues at the trunk port.
3. Use one or more BXM cards to set up VP connections to convert the VPI from one to the network assigned VPI value.

One physical wraparound interconnect between two physical ports is required for each virtual trunk. The CDV is between 50 usec and 3 msec, depending on the bandwidth of the VP.

The number of physical ports available for wraparound interconnection limits the number of virtual trunks that can be supported on the UNI port to the public ATM network. This wraparound solution, although consuming extra numbers of ports, enables savings on recurring cost for accessing UNI ports on public ATM switches.

The following figure illustrates how Virtual Trunking with multi-QoS traffic shaping is supported in Release 8.4.

Figure 4 Virtual Trunking with Multi-QoS Traffic Shaping in Release 8.4



1. NOTE: The standard user command `cnftrk` requires deleting the trunk first to change the Transmit Trunk Rate. When the trunk is first set up, make sure the Transmit Trunk Rate is set to the highest rate ever needed in the future.

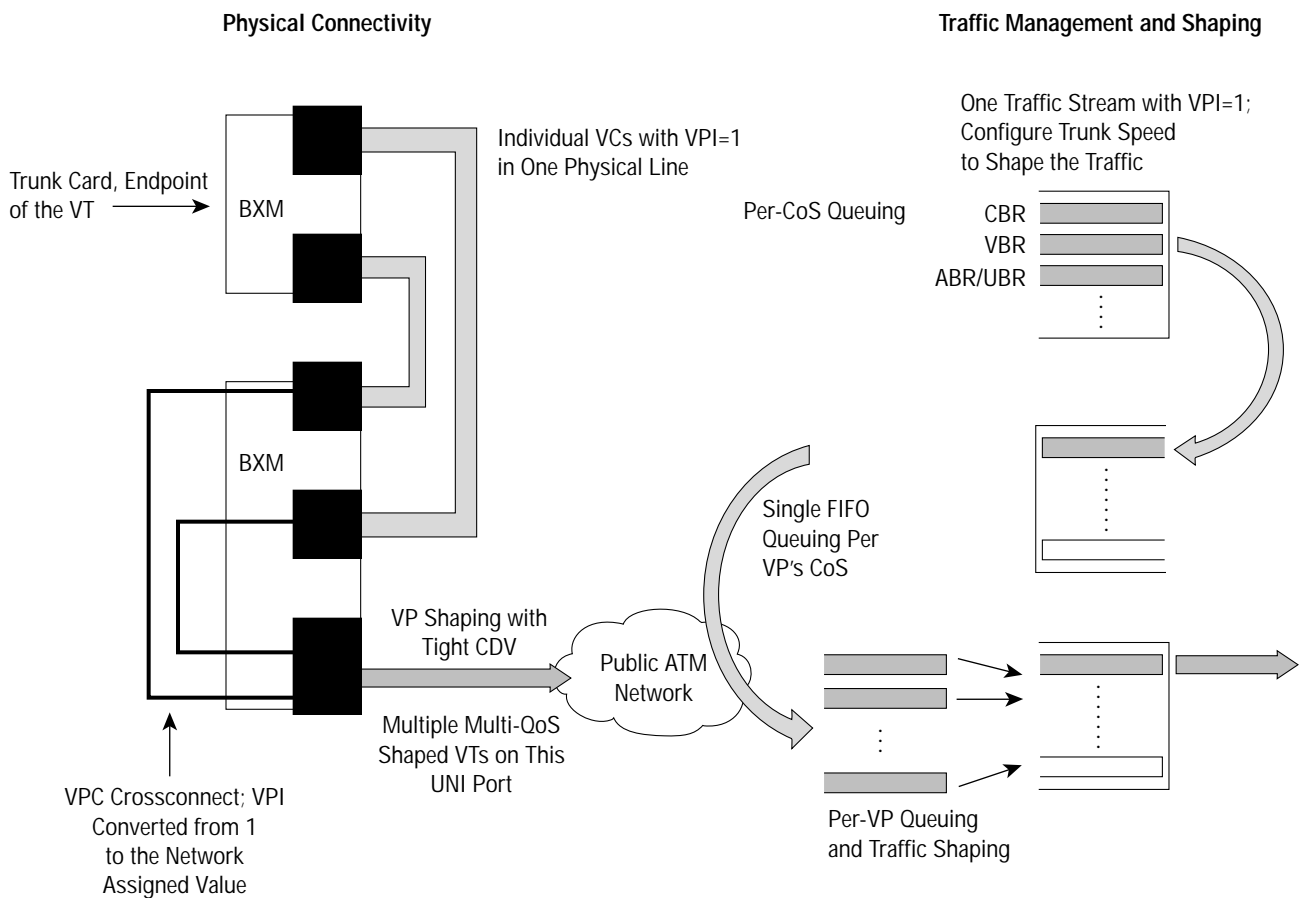
Note: When using wraparound solution for Virtual Trunking, the node number must be at least 32. The BPX switch uses VPI = 1 and VCI = node_number + 1 as the control channel on trunks between nodes. For a BPX switch with node number less than 32, the control channel could be mistaken for OAM or other supervisory cells in the public ATM network. VCI range from 0 to 15 is reserved by ITU-T, and VCI range from 16 to 31 is reserved by ATM Forum.

Release 9.1

Release 9.1 does not support Virtual Trunking on the BXM card. As in Release 8.4, an interim wraparound solution is required to provide Virtual Trunking and per-virtual-trunk traffic shaping as described in the previous section.

Release 9.1 supports per-VC/VP traffic shaping on UNI ports. Such a feature can be used to reduce the CDV of a virtual trunk. The following figure illustrates how this can be accomplished. The tariff of a public ATM PVC service for a PCR typically depends on CDVT. Reducing CDV on a virtual trunk allows the user to leverage the least tarified ATM PVC service for the same PCR.

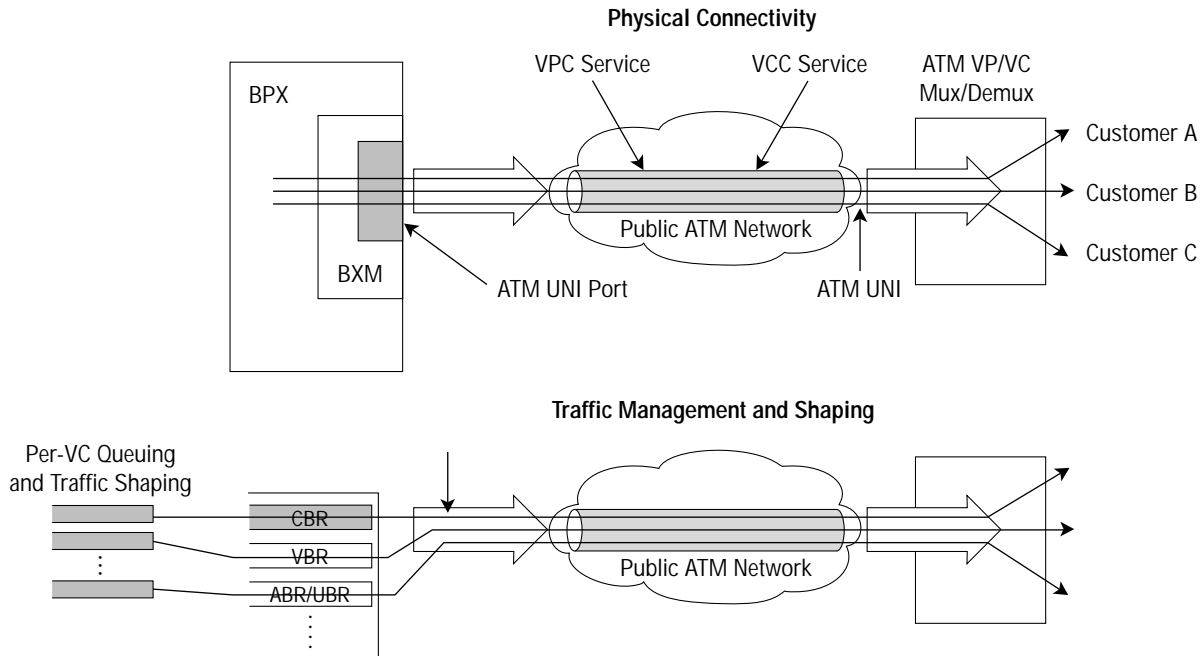
Figure 5 Virtual Trunking with Multi-QoS Traffic Shaping and Tight CDV in Release 9.1



UNI Port Hierarchical Traffic Shaping

A future software upgrade would allow the BPX card to take advantage of the hardware built-in VI and per-VC/VP traffic-shaping capabilities to enable hierarchical traffic shaping on a UNI port. Hierarchical traffic shaping on a UNI port enables a BPX service provider customer to use public ATM network VPC service as an alternative transport for providing VCC services from a UNI port to remote subscriber customers, as shown in the following figure.

Figure 6 Hierarchical Traffic Shaping Enables VCC Service Offerings to Remote Customers Using Public ATM VPC Service as Transport

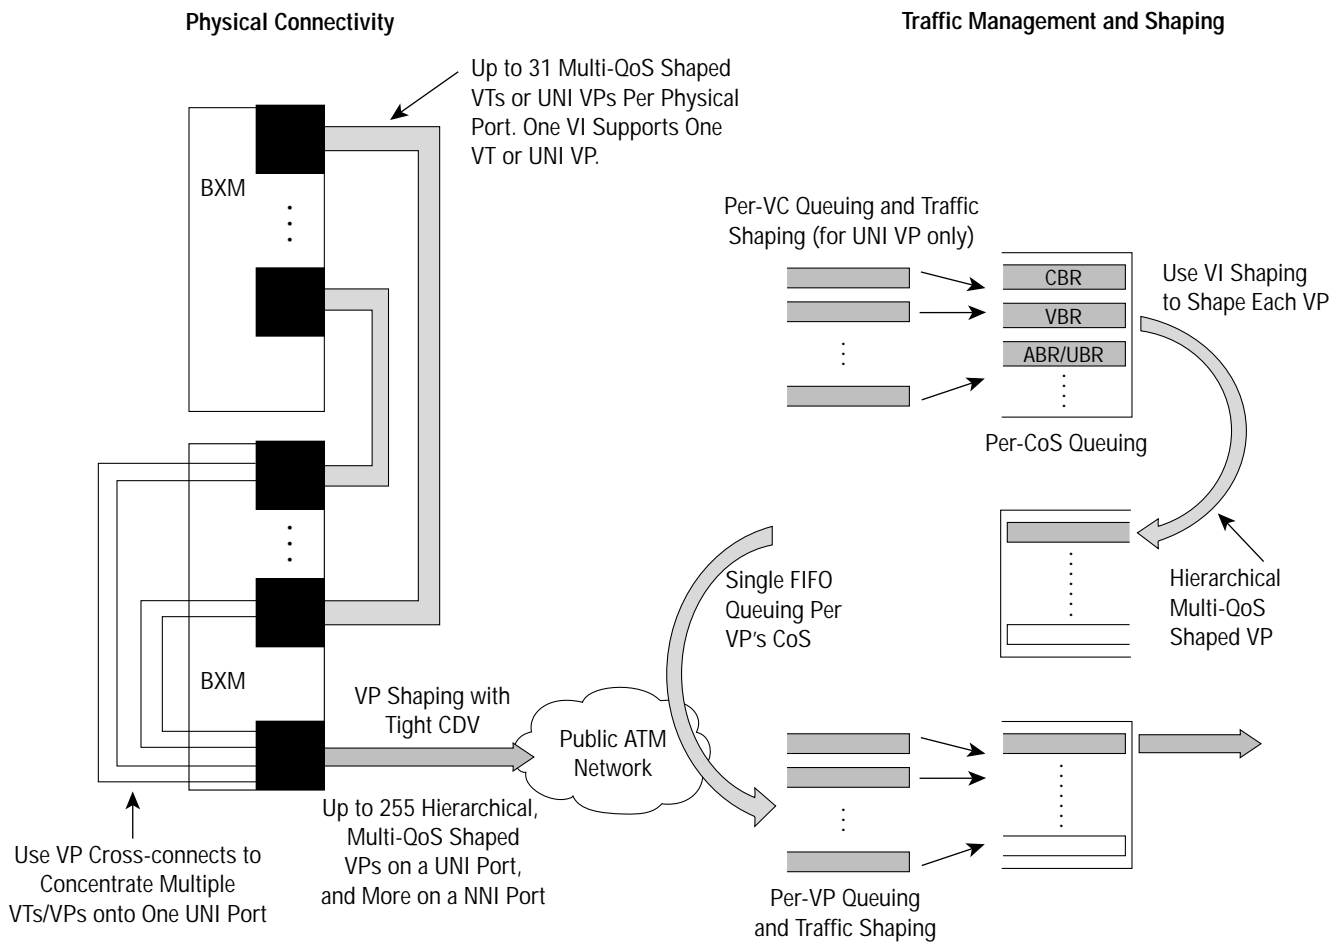


When a service provider uses public ATM network VPC services to provide VCC services to remote customers, hierarchical traffic shaping at both VP and VC levels is required. Traffic shaping at the VP level is for passing VPC aggregate traffic safely through the public ATM network. Traffic shaping at the VC level for each individual VCC is for fair bandwidth sharing and for matching the receiving speeds of the remote customer CPE equipment. Even if receiving speeds are not an issue, the CPE equipment may not have sufficient buffer to handle the “uncontrolled” CDV of the incoming cells. With VC traffic shaping, most of the CDV that may have accumulated in the path can be canceled.

Wraparound Solution for Supporting 255 and More Hierarchically Shaped VPs on a BXM Physical Port

When the UNI port hierarchical traffic shaping feature becomes available, up to 31 shapable VPs can be supported on each UNI port. Such a feature, when combined with the wraparound solution, can further increase the number of hierarchically shaped VPs on a UNI port to more than 31. Similarly, when the Virtual Trunking 9.2 feature is used in combination with a wraparound solution, the number of virtual trunks per port can be increased to more than 31.

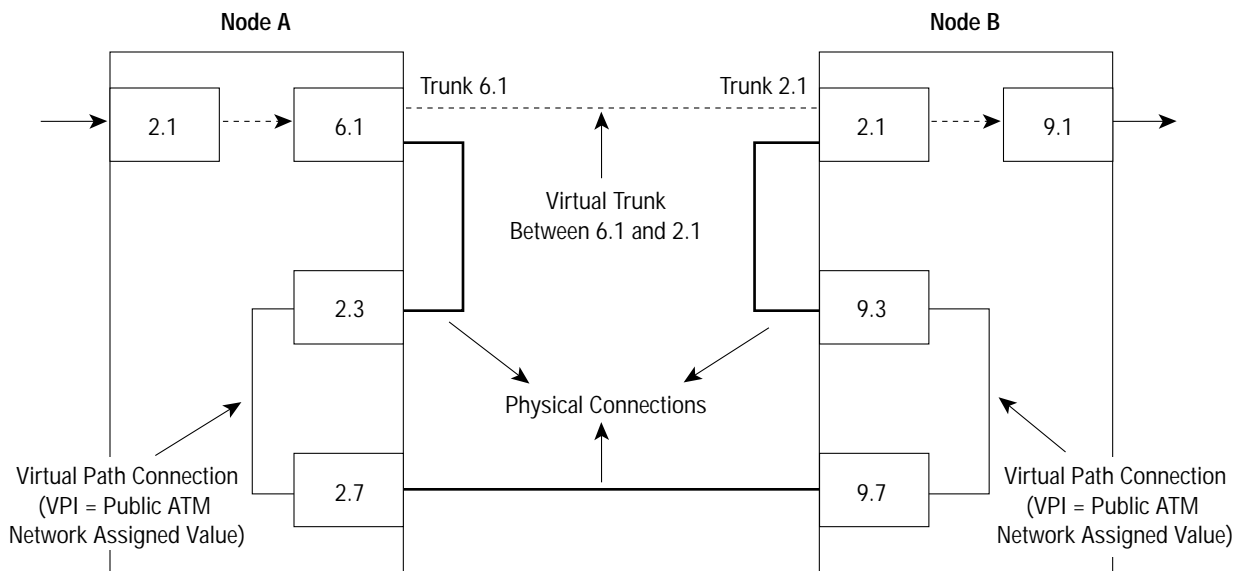
Figure 7 Support up to 255 Hierarchical, Multi-QoS Shaped VPs on a UNI Port, and More on a NNI Port



Appendix: Wraparound Setup Procedure

The example in the following figure will be used in this section to illustrate the virtual trunk wraparound setup. On each BPX node, there are two BXM cards; 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. Because Virtual Trunking is not supported on BXM until Release 9.2, the output traffic from the trunk cards does not necessarily use the same 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 cards is wrapped around the virtual path connections before sending to the public ATM network.

Figure 8 Virtual Trunk Wraparound Setup Example



BXM 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).

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

3. Turn on VP Shaping on Ports 2.7 and 9.7 at Node A and Node B, respectively.

```
cnfport 2.7 N N N N y 0 0 0 0
```

```
cnfport 9.7 N N N N y 0 0 0 0
```

4. Add a CBR virtual path connection with PCR = `trunk_speed`² 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 one, 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 *
```

5. Similarly 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 * * * .....
```

8. Add the trunk

```
Node A: addtrk 6.1
```

2. `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 rates of Trunk 6.1 on Node A and Trunk 2.1 on Node B.



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