

Understanding Per-VC Queuing Options on the 4xOC3 ATM Line Card

Document ID: 12140

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Understanding SAR Memory

Recommended Values for SAR Pools

Understanding Buffer Memory

Recommended Values for Transmit Queue Limits

Understanding Delay

Use CAR to Reduce Delay

Low-Speed VCs

OAM Cells and Queuing

Summary

Related Information

Introduction

This technical note provides tips to maximize per-VC performance on the quad OC-3 (QOC-3) ATM line card for the Gigabit Switch Router (GSR).

The Quad OC-3 (QOC-3) is designed for service providers that require unspecified bit rate (UBR) PVCs with little or no traffic shaping to limit the peak cell rate (PCR) to less than the interface line rate. UBR is one of five ATM service categories defined in the Traffic Management Specification 4.0 of the ATM Forum. UBR is intended for non-real-time applications that do not require any maximum bound on the transfer delay or on the cell loss ratio. UBR allows for a high degree of statistical multiplexing and does not reserve any minimum bandwidth for each virtual circuit (VC). Typically, the VCs use the bandwidth up to the configured PCR, when available.

With any ATM interface, it is important to ensure that each VC has fair access to packet buffers. Per-VC queuing is designed to provide a dedicated set of packet buffers to each VC to ensure that one overused and congested VC does not impact performance on other (non-congested) VCs.

The QOC-3 does not offer per-VC queuing with multiple VCs. By default, the QOC-3 shares Segmentation and Reassembly (SAR) memory to all VCs on an interface and a maximum level of statistical buffer sharing of the 65535 transmit buffers. This technical note discusses how to use the SAR pools and tx-queue-limit features to implement subtotal limits on both SAR and transmit buffer memory, respectively, to ensure maximum performance for all VCs.

Note: The 1xOC12 and 4xOC12 ATM line cards do not support configurable SAR pools.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Understanding SAR Memory

Each line card on the GSR is composed of several key hardware components, such as physical memory modules. Figure 1 illustrates the architecture of the QOC-3 line card.

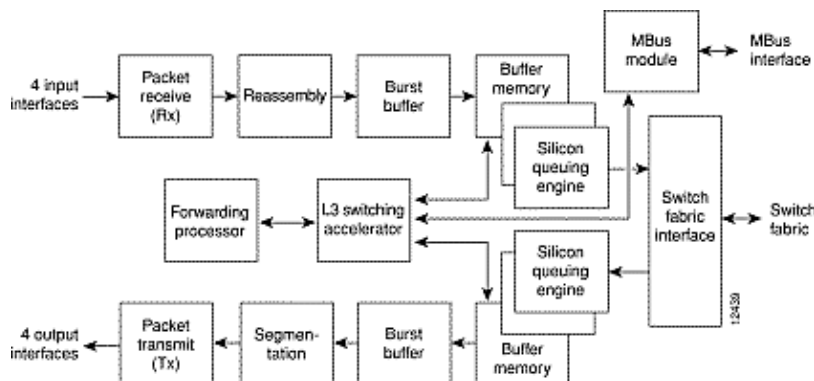


Figure 1: Diagram of the architecture of the Quad OC-3 line card.

The four interfaces on the QOC-3 share a single transmit SAR and 2 MB of SAR memory (SRAM) used to store and segment IP packets encapsulated with an ATM Adaptation Layer 5 (AL) trailer. This memory is divided evenly among the four interfaces. The system gives each interface 512 kB (1/4 of 2 MB). The GSR documentation refers to the SAR pools as output buffer memory.

This 512 kb is known as a SAR pool. The term "pool" refers to an isolated set of packet buffers. By default, the system assigns all VCs on an interface to a single pool (pool 0).

If an ATM interface can support more than one VC, we strongly recommend that you configure a non-default value of two or four SAR pools to implement further buffer isolation among the VCs. Specify four SAR pools in the configuration of the router to give 128 kb of SAR memory to each pool.

Use these three steps to configure SAR pools on your QOC-3 line card:

1. Execute the **sar txpools** command and specify two or four pools. We highly recommend 4 pools to reduce potential latency, even if only 1 PVC is configured.

```
GSR-1(config)#int atm 7/1
GSR-1(config-if)#sar txpools ?
 0 TxPools
 2 TxPools
 4 TxPools
GSR-1(config-if)#sar txpools 4
```

2. Create a VC and assign it to a pool with the **sar vc** command.

```
GSR-1(config-if)#pvc 1/1
GSR-1(config-if-atm-vc)#exit
GSR-1(config-if)#sar vc ?
  <0-65535> VC Identification Number | vpi
GSR-1(config-if)#sar vc 1 ?
  <0-2047> VCI
  txpool   Transmit Pool Number

GSR-1(config-if)#sar vc 1 1 ?
  txpool   Transmit Pool Number

GSR-1(config-if)#sar vc 1 1 txpool ?
  <0-3>    Transmit Pool Identification

GSR-1(config-if)#sar vc 1 1 txpool 2
```

3. Confirm your configuration with the **show sar txpools** command.

```
GSR-1#show sar txpools int atm 7/1
ATM 7/1: SAR TxPool 0: TxPool Size 128kb
        No VCs assigned
ATM 7/1: SAR TxPool 1: TxPool Size 128kb
        No VCs assigned

ATM 7/1: SAR TxPool 2: TxPool Size 128kb
        AAL /      Peak   Avg. Burst      Out      Out
        VCD VPI  VCI Type Encap   Kbps   Kbps Cells   Pkts     Bytes
        1   1    1  PVC AAL5-SNAP    0      0   0         0         0

ATM 7/1: SAR TxPool 3: TxPool Size 128kb
        No VCs assigned
```

Note: All switched virtual circuits (SVCs) must only use SAR pool 0. This behavior cannot be changed.

The **show controller atm 1 all** command provides additional information on the SAR pools, which includes how SAR pools use a counter, a threshold, and a hysteresis value for each pool. When the threshold value for a pool is exceeded, the QOC-3 asserts back pressure for that pool. The buffer memory ASIC (BMA), which controls the 16 MB to 64 MB of transmit buffer memory (see next section), maintains a queue of 16 packets that are sent regardless of any back pressure. In a worst-case condition, if all 16 packets are destined to a pool that is over the threshold, the "Buffers Used" value is $494 + 16 = 510$. Under rare conditions, you can see a higher number. This value does not cause a problem since the buffers are allocated from a general pool, and the per-pool thresholds are simply a rationing mechanism. If the sum of the pool thresholds exceeds the buffers, the system asserts a general back pressure to prevent an overrun of the total SAR buffers.

```
GSR-1#execute-on slot 7 show controller ATM 1 all
===== Line Card (Slot 7) =====
TX SAR (Patch 3.1.3) is Operational;
RX SAR (Patch 3.1.3) is Operational;

Interface Configuration Mode:
      ATM clock internal; STS-3c

[output omitted]

General: Quota = 8168  Usage = 3      Hysteresis = 120

PoolId  Buffers  Buffers  Hysteresis
```

	Per Pool	Used	Value
0	506	3	6
1	506	0	6
2	506	0	6
3	506	0	6
4	472	0	40
5	472	0	40
6	472	0	40
7	472	0	40
8	472	0	40
9	472	0	40
10	472	0	40
11	472	0	40
12	472	0	40
13	472	0	40
14	472	0	40
15	472	0	40

Note: The number 494 of 256-byte buffers. There are 2MB of SRAM connected to the Batman (on 1xOC12 ATM) or Batman2 (on 4xOC3 ATM) ASIC. On the Batman2, which is used on the 4xOC3 ATM line card, there is the ability to configure up to four SAR pools per interface for a total of up to 16 SAR pools per card. This causes the two MB of SRAM to be divided into approximately 128 KB of SRAM per SAR pool (or 512 256-byte buffers). This number is approximate because the backpressure signal from the BATMAN2 to the TxBMA does not take effect immediately (that is, packets in flight still arrive into the BATMAN2 even though the back pressure signal is on). Instead of setting the per-SAR-pool backpressure limit to 512, we set it to 494 so that back pressure takes effect at approximately 512 buffers.

Recommended Values for SAR Pools

Since the problems associated with buffer sharing are only an issue within periods of congestion, we can group VCs by levels of usage and place those with similar oversubscription levels in the same buffer queue.

For example, this configuration allows for 17 VBR-nrt shaped VCs and a number of UBR VCs on a single port:

- Transmit Pool One – Place two highly oversubscribed-shaped VCs in this pool.
- Transmit Pool Two – Place five moderately oversubscribed-shaped VCs in this pool.
- Transmit Pool Three – Place 10 lightly oversubscribed-shaped VCs in this pool.
- Transmit Pool Four – Place all best effort (UBR) VCs in this pool.

An alternate configuration is to place VCs with the same shaping parameters in the same SAR pool, particularly if there are big disparities between the traffic-shaping values. For example, there is an order of magnitude difference between a VC shaped to 1 MB and a VC shaped to 20 MB. If you configure both VCs on the same interface, configure them in different SAR pools.

Understanding Buffer Memory

Each line card has two silicon queuing engines: receive and transmit. The receive engine moves packets from the burst buffer to the switch fabric, and the transmit engine moves packets from the switch fabric to the transmit interface. In this way, the silicon queuing engine controls the placement of IP packets in buffer memory, as well as their removal from buffer memory.

By default, the QOC-3 comes with 32 MB of buffer memory. The 32 MB is divided into 16 MB of receive buffers and 16 MB of transmit buffers. Use the **show diag** command to view the amount of transmit buffer

memory currently installed on your line card. The QOC-3 supports up to 64 MB of receive buffers and 64 MB of transmit buffers.

This sample output from the **show diag** command was generated on a QOC-3 line card with 32 MB of receive buffers and 32 MB of transmit buffers.

```
router# show diag
SLOT 9 (RP/LC 9 ): 4 port ATM Over SONET OC-3c/STM-1 Single Mode
  MAIN: type 39, 00-0000-00 rev 71 dev 16777215
        HW config: 0x00      SW key: FF-FF-FF
  PCA:  73-3033-02 rev 71 ver 2
        HW version 1.0  S/N CAB02180072
  MBUS: MBUS Agent (1) 73-2146-07 rev B0 dev 0
        HW version 1.2  S/N CAB020800D7
        Test hist: 0xFF  RMA#: FF-FF-FF  RMA hist: 0xFF
  DIAG: Test count: 0xFFFFFFFF  Test results: 0xFFFFFFFF
  MBUS Agent Software version 01.35 (RAM) (ROM version is 01.33)
  Using CAN Bus A
  ROM Monitor version 00.0D
  Fabric Downloader version used 00.0D (ROM version is 00.0D)
  Board is analyzed
  Board State is Line Card Enabled (IOS  RUN )
  Insertion time: 00:00:11 (1d14h ago)
  DRAM size: 33554432 bytes
FrFab SDRAM size: 33554432 bytes
ToFab SDRAM size: 33554432 bytes
```

The GSR refers to the receive memory as "ToFab SDRAM" and to the transmit memory as "FrFab SDRAM". You also can use the **show controller frfab queue** command to view the transmit buffer memory. As illustrated, the QOC-3 uses a buffer-carving algorithm that creates four packet buffer pools referred to as "4 non-IPC free queues." In the example, the system created 26190 buffers for packets that are 80 bytes or less and 15919 buffers for packets that are 608 bytes or less. Overall, the system created 65535 packet buffers, as shown in the rightmost column titled "LenThresh."

```
GSR-1#execute-on slot 7 show controller frfab queue
===== Line Card (Slot 7) =====
Carve information for FrFab buffers
SDRAM size: 33554432 bytes, address: 20000000, carve base: 2009D100
! -- 32 MB SDRAM for transmit packet buffers.
32911104 bytes carve size, 0 SDRAM bank(s), 0 bytes SDRAM pagesize, 1 carve(s)
max buffer data size 4544 bytes, min buffer data size 80 bytes
51452/51452 buffers specified/carved
32907680/32907680 bytes sum buffer sizes specified/carved
  Qnum   Head   Tail           #Qelem LenThresh
  ----   ----   ----           -
4 non-IPC free queues:

    26190/26190 (buffers specified/carved), 50.90%, 80 byte data size
    1         9513   9512           26190 65535

    15919/15919 (buffers specified/carved), 30.93%, 608 byte data size
    2        26291  42209           15919 65535

    7703/7703 (buffers specified/carved), 14.97%, 1568 byte data size
    3        42210  49912            7703 65535

    1540/1540 (buffers specified/carved), 2.99%, 4544 byte data size
    4        49913  51452            1540 65535

IPC Queue:
  100/100 (buffers specified/carved), 0.19%, 4112 byte data size
```

30	32	31	100	65535
Raw Queue:				
31	0	37	0	65535
Interface Queues:				
0	0	9512	0	65535
	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
1	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
2	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
3	0	0	0	65535
	0	0	0	65535
	0	0	0	65535
	0	0	0	65535

Note: If you are not attached to the line card, you must start your command with **execute-on slot #**.

The output of the **show controllers frfab queue** command displays four queues under each interface. This is the fixed number of queues the system creates per interface.

By default, the system allows the VCs of one SAR pool to occupy all 65535 packet buffers on the line card. The use of a shared-memory design provides the benefit of complete sharing with a maximum level of statistical buffer sharing.

The transmit packet buffers are used when the QOC-3 starts to queue excess packets destined for an oversubscribed VC. First, the VC uses the assigned memory in its SAR pool. Next, it falls back to the transmit buffer memory. In a system with full sharing memory, a slowly draining VC with a relatively high output rate consumes much or all of the transmit buffer memory and leaves less memory for all the other VCs on all the other interfaces of the line card. It is important to understand that a single VC can affect VCs on all other interfaces of a line card unless you also implement transmit queue limits.

Transmit queue limits configure a hard threshold of the number of packets that can be queued for transmission on an interface. Once the queue reaches the threshold, any new packets switched to the port from the fabric are dropped. If an interface is configured with multiple SAR pools, the transmit queue limit is divided equally among the pools.

In this example, we limit the transmit queue size on interface ATM 7/0 to 4000 buffers with the **TX-queue-limit 4000** command.

```
GSR-1(config)#int atm 7/0
GSR-1(config-if)#sar txpools 4
GSR-1(config-if)#TX-queue-limit ?
<1-65535> Number (queue limit)

GSR-1(config-if)#TX-queue-limit 4000
GSR-1(config-if)#end
```

The queue limit for ALL four pools on interface atm7/0 is set to 4000. The tx queue limit is in units of buffers, so when the specified number of buffers are occupied by a TX queue (queue reaches the limit), traffic is dropped until the number of buffers falls below the limit. Note that the system equally divides the configured number of buffers (4000, in this case) among the configured number of SAR pools (four, in this

case) on the interface. In this example, each SAR pool is given buffers. This can be seen with the **show controller frfab queue** command.

```
GSR-1#execute-on slot 7 show controller frfab queue
===== Line Card (Slot 7) =====
Carve information for FrFab buffers
  SDRAM size: 33554432 bytes, address: 20000000, carve base: 2009D100
  32911104 bytes carve size, 0 SDRAM bank(s), 0 bytes SDRAM pagesize, 1 carve(s)
  max buffer data size 4544 bytes, min buffer data size 80 bytes
  51452/51452 buffers specified/carved
  32907680/32907680 bytes sum buffer sizes specified/carved

      Qnum      Head      Tail          #Qelem  LenThresh
      ----      -
4 non-IPC free queues:

26190/26190 (buffers specified/carved), 50.90%, 80 byte data size
  1          9807      9806          26190   65535

15919/15919 (buffers specified/carved), 30.93%, 608 byte data size
  2          26291     42209          15919   65535

7703/7703 (buffers specified/carved), 14.97%, 1568 byte data size
  3          42210     49912           7703   65535

1540/1540 (buffers specified/carved), 2.99%, 4544 byte data size
  4          49913     51452           1540   65535

IPC Queue:
  100/100 (buffers specified/carved), 0.19%, 4112 byte data size
  30         1         100            100    65535

Raw Queue:
  31         0         6              0      65535

Interface Queues:
  0          0         9806           0      1000
             0         0              0      1000
             0         0              0      1000
             0         0              0      1000

! -- Each SAR pool is assigned 1/4 of the configured TX-queue-limit (4000/4 = 1000)

  1          0         0              0      65535
             0         0              0      65535
             0         0              0      65535
             0         0              0      65535
  2          0         0              0      65535
             0         0              0      65535
             0         0              0      65535
             0         0              0      65535
  3          0         0              0      65535
             0         0              0      65535
             0         0              0      65535
             0         0              0      65535
```

If the configured number of SAR pools were 2, the system equally divides the 4000 buffers between the two SAR pools on the interface. Each SAR pool is given 2000 buffers.

```
GSR-1(config)#int atm 7/0
GSR-1(config-if)#sar txpools 2
GSR-1(config-if)#TX-queue-limit 4000
GSR-1#execute-on slot 7 show controller frfab queue
-----command output truncated-----
```

```

Interface Queues:
  0      0      0      0      2000
        0      0      0      2000
        0      0      0      2000
        0      0      0      2000
  1      0      0      0      65535
        0      0      0      65535
        0      0      0      65535
        0      0      0      65535
  2      0      0      0      65535
        0      0      0      65535
        0      0      0      65535
        0      0      0      65535
  3      0      0      0      65535
        0      0      0      65535
        0      0      0      65535
        0      0      0      65535

```

Keep these points about transmit queue limits in mind:

- Transmit queue limits refer to the number of packets and are independent of packet sizes. In other words, the configured value is in units of buffers.
- Transmit queue limits apply to non-IPC free queues only.
- IPC and raw queues are not affected by queue limits.

Recommended Values for Transmit Queue Limits

This section provides an example set of recommended TX-queue-limit sizes for the QOC-3 transmit buffer memory. It assumes that you are need to configure 16 VCs. Set the TX-queue-limit such that a queue never contains more than 1/16th of the total buffers.

1. Execute **show controller frfab queue** and determine the total number of buffers carved.

```

LC-Slot5#sho controller frfab queue
Carve information for FrFab buffers
  SDRAM size: 33554432 bytes, address: 20000000, carve base: 2009D100
  32911104 bytes carve size,  0 SDRAM bank(s), 0 bytes SDRAM pagesize, 2
  carve(s)
  max buffer data size 4544 bytes, min buffer data size 80 bytes
  51452/51452 buffers specified/carved
  32907680/32907680 bytes sum buffer sizes specified/carved
    Qnum      Head      Tail          #Qelem  LenThresh
    ----      ----      ----          -
  4 non-IPC free queues:

    26190/26190 (buffers specified/carved), 50.90%, 80 byte data size
    1          9972      9971          26190   65535

    15919/15919 (buffers specified/carved), 30.93%, 608 byte data size
    2          38043     38042         15919   65535

    7703/7703 (buffers specified/carved), 14.97%, 1568 byte data size
    3          42210     49912         7703    65535

    1540/1540 (buffers specified/carved), 2.99%, 4544 byte data size
    4          49913     51452         1540    65535

```

2. Assume that you test with 300-byte packets. In pool 2, the system created a total of 15919 buffers that can hold up to 608 bytes. The system uses these buffers for your 300-byte packets.

3. If you test 16 VCs simultaneously with the same packet size, your TX-queue-limit must be configured as $15919/16 = 994$. A value of 994 assures an equitable distribution of the appropriate size buffers. In other words, if you test with a fixed packet size that you can control, ensure that no queue contains more than 1/16th of the buffers that are used by that packet size.

Understanding Delay

Queuing always introduces a trade-off: it does not allow packets to be dropped, but it introduces delay. Packets experience delay on any networking device when that device needs to queue packets.

In other words, queuing happens when you present traffic at a rate greater than the configured traffic shaping values. On the QOC-3 line card, if the system constantly offers more than the sustained cell rate (SCR) of one vbr-nrt VC, the excess packets for this VC are queued in the SAR pools memory as well as in the buffer memory.

On the QOC-3, the configuration of one VC in a SAR pool provides the advantage of all the pool buffers to the VC, but introduces the disadvantage of significant latency if the pool is consistently maximized.

This is an example:

1. Configure a shaped VBR-nrt VC with a sustained cell rate (SCR) of three Mbps.
2. Configure four SAR pools on an ATM interface. As seen above, the system assigns 494 SAR buffers to each pool. Each buffer can hold 252 bytes of data. $494 \text{ buffers} * 252 \text{ bytes of data} = 124488 \text{ bytes of data}$.
3. Three Mbps equates to a transmission rate of about 7075 cells/sec with this formula: $3\text{Mbps} / 8 / 53$. Where eight is in units of bits per Byte and 53 is in units of bytes per cell.
4. 124488 bytes of data is roughly 2600 cells or 1/3 of a second worth of data.
5. If you fill the 494 SAR buffers, you can expect to see three seconds of delay before the last cell can be sent.

In other words, the maximum queuing delay is the sum of this:

- Time it takes for a packet to traverse the transmit buffer memory (65535 buffers by default, or a lower value as configured with the **TX-queue-limit** command.)
- Time it takes for a packet to traverse the SAR pool (about 500 buffers with four pools per interface.)

Use CAR to Reduce Delay

An oversubscribed VC quickly fills its allocated SAR pool memory. On a VBR-nrt PVC, excess traffic above the configured shaping values is buffered. If you configure PCR=SCR and MBS=1 on a VBR-nrt PVC, it effectively disables VBR-nrt bursting, but any traffic load beyond PCR/SCR is queued. The system continues to buffer the excess packets until all of the 64KB transmit buffers are exhausted or until the system reaches the configured TX-queue-limit for the VC. In these conditions, you can avoid queuing delay if you implement traffic policing in addition to shaping. In this configuration, shaping limits the output rate of the VC, and polices drops any excess.

Cisco IOS[®] traffic policing uses one of the following commands:

- **rate-limit** – Original command. Known as Committed Access Rate (CAR).
- **police** *bps burst-normal burst-max conform-action exceed-action action violate-action action* – Newer command with the Cisco modular QoS CLI syntax.

These are the advantages and disadvantages of Cisco IOS traffic policing when applied to ATM interfaces:

- Advantages:
 - ◆ Does not add delay since it drops all excess traffic without buffering.
 - ◆ Works on either input or output interfaces or subinterfaces.
 - ◆ Implements optional packet marking when it sets the IP precedence value of packets classified as conforming to or exceeding a particular rate limit.
- Disadvantages:
 - ◆ Simply limits rate and does not guarantee bandwidth. Drops rather than queues excess traffic that exceeds the maximum committed rate (plus the burst allowance).
 - ◆ Propagates bursts. It does no smoothing or shaping of traffic.

We recommend that you implement output CAR if you need to configure more than one VC in a SAR pool. In this configuration, you cannot depend on VBR traffic shaping on the QOC-3 line card to prevent head-of-line blocking between VCs in the same pool. Output CAR must be used to prevent each SAR pool from filling.

The trade-off involved with the application of CAR is a decrease in latency versus lower aggregate throughput. This trade-off is a result of the involvement of the CPU on both the receive side, where it performs L3 switching on inbound traffic, as well as on the transmit side where it enforces CAR rules. The resultant performance bottleneck is the CPU, which is limited in the total number of packets that it can process from both directions. The limit includes inbound traffic to be switched and any outbound traffic subject to CAR rules.

Low-Speed VCs

The QOC-3 line card is not recommended for use with slow VCs. When you use one or more VCs that are shaped to a PCR below one Mbps, configure output CAR rather than native ATM traffic shaping.

As an example, this UBR PVC has been assigned a PCR of 822 kbps.

```
interface ATM5/2.1111 point-to-point
 ip address 210.220.191.173 255.255.255.252
 no ip directed-broadcast
 no atm enable-ilmi-trap
 PVC 1/111
  ubr 822
```

We recommend replacing the **ubr** command with a CAR **rate-limit** command, as shown.

```
Interface ATM5/2.1111 point-to-point
 ip address 210.220.191.173 255.255.255.252
 no ip directed-broadcast
 rate-limit output 816000 5000 5000 conform-action transmit exceed-action drop
 no atm enable-ilmi-trap
 PVC 1/111
```

The value of 816000 matches the PCR of 822. For the burst bytes, we recommend you use a value that is greater than or equal to the MTU, which is 4470 by default on ATM interfaces.

We suggest this CAR solution for 4xOC-3 ATM cards:

- Assign all VCs to the UBR service category.
- Use output CAR with the maximum sustainable bit rate set to the maximum service level delivered on a particular VC.

- Configure burst values with this formulas. Tune the burst values to find the ideal trade-off between an increase in drops and slight increases in latency.

Normal Burst Value (bytes) = PCR in BPS / (8 bits/byte) * 1.5 seconds

Extended Burst Value (bytes) = Normal Burst Value * 2

- Ensure that the sum of the rates of the CAR rate limits do not exceed the aggregate bandwidth of the main interface on which they are configured.

If you configure more VCs in a single SAR pool, it reduces the chances of fair access to packet buffers. Output CAR sets a bound limit on latency but introduces a performance penalty, particularly with small packets. (ACLs and other features that use the CPU also affect throughput.)

OAM Cells and Queuing

It is important to understand the impact that queuing has on time-sensitive processes like Operations, Administration, and Maintenance (OAM), and even routing protocols like EIGRP.

OAM PVC management is designed to detect a transmission problem in the end-to-end path of a PVC. Without this feature, a PVC remains up on the end-devices. Routing table entries that point to the next-hop interface across the PVC remain in the table, and packets are lost until reconvergence is completed. Refer to Using OAM for PVC Management.

If the OAM loopback cells used to monitor PVC status experience long queuing delays and do not return to the source within the necessary interval, the remote router brings down the PVC when the OAM timers expire even though there is no problem in the end to end path.

As a workaround, increase the OAM retry intervals with this command:

```
Router(config-if-atm-vc)#oam retry up-count down-count retry-frequency
```

- Use the **up-count** argument to specify the number of consecutive end-to-end F5 OAM loopback cell responses that must be received in order to change a PVC connection state to up.
- Use the **down-count** argument to specify the number of consecutive end-to-end F5 OAM loopback cell responses that are not received in order to tear down a PVC.
- Use the **retry-frequency** argument to specify the frequency (in seconds) that end-to-end F5 OAM loopback cells should be transmitted when a change in UP/DOWN state is being verified. For example, if a PVC is up and a loopback cell response is not received after the frequency (in seconds) specified with the **oam-pvc** command, then loopback cells are sent at the retry-frequency to verify whether or not the PVC is down.
- By default, end-to-end F5 OAM loopback cell generation is turned off for each PVC. A PVC is determined as down when the router does not receive a loopback reply after a configured number of retries to send end-to-end F5 OAM loopback cells.

Summary

This illustration and bullet points summarize queuing on the QOC-3 ATM line card for the GSR. It is important to understand that you must configure both SAR pools and tx-queue-limits.

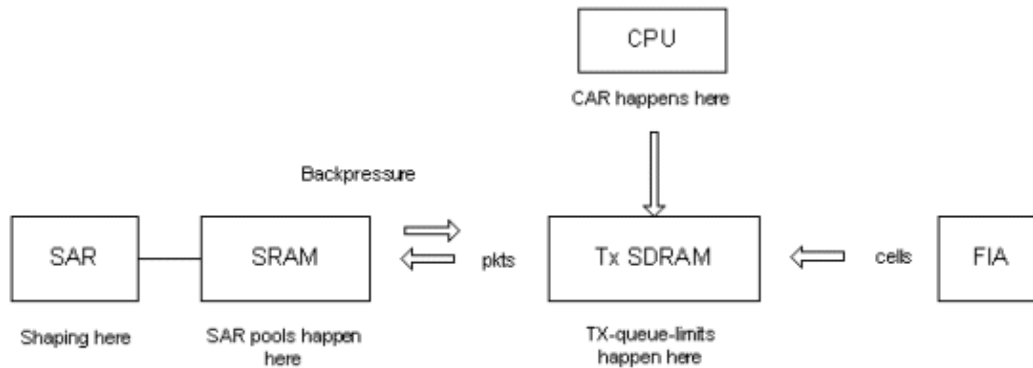


Figure 2: Queuing on the QOC-3 ATM line card for the GSR.

The SAR performs traffic shaping. When a VC offers more traffic than its configured shaping values, the excess packets are queued.

The SAR pools in SRAM buffer the excess packets. Configuring the recommended four SAR pools effectively isolates the one or more VCs assigned to a pool.

- A single oversubscribed VC can occupy most or all of the buffers of its assigned SAR pool. When this happens, the SAR exerts a back pressure signal to the transmit BMA. The line card then stores excess packets in the transmit buffer memory (SDRAM).

The buffer-carving algorithm creates four pools of various sizes in transmit buffer memory. Use the **show controller frfab queue** command to view these buffers.

By default, every VC can use all 65535 packet buffers carved on the line card. Use the **TX-queue-limit** command to limit the number of buffers that each pool can acquire. Transmit queue limits isolate the SAR pools (and VCs) from each other in the BMA memory. Without these limits, one VC can affect the entire line card, and the benefits of configuring SAR pools is effectively reduced. Packets for one VC can become stuck behind packets for another VC and effectively be prevented from using free memory in the SAR pool.

Configure rate-limiting on the outbound interface to prevent the router from queuing any excess traffic. Rate-limiting forces every packet to be routed from the transmit buffer memory to the CPU, then back to buffer memory before proceeding to the SAR. This architecture imposes limits on packet forwarding performance with CAR.

Related Information

- [Using OAM for PVC Management](#)
 - [Understanding the VBR-nrt Service Category and Traffic Shaping for ATM VCs](#)
 - [Understanding the UBR Service Category for ATM Virtual Circuits](#)
 - [Understanding the Maximum Number of Active Virtual Circuits on Cisco ATM Router Interfaces](#)
 - [Configuring VP Shaping on the 4xOC3 ATM Line Card for the GSR](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

