
ATM Forum Document Number: ATM_Forum/96-0519

Title: General Considerations for Frame-Level Performance Measurement of ATM Switches.

Abstract: This contributions discusses issues that apply to several performance metrics related to ATM switches. These issues include measurement configurations, traffic patterns, protocol layers, and measurement points. The first two issues were discussed in our February 1996 contribution and have been revised for this contribution. The last two points have been elaborated based on the feedback from the previous meetings.

Source:

Raj Jain, Bhavana Nagendra, and Gojko Babic
The Ohio State University

Raj Jain is now at Washington University in Saint Louis, jain@cse.wustl.edu <http://www.cse.wustl.edu/~jain/>

The presentation of this contribution at the ATM Forum is sponsored by NASA.

Date: April 1996, Anchorage, Alaska

Distribution: ATM Forum Technical Working Group Members (AF-TEST)

Notice: This contribution has been prepared to assist the ATM Forum. It is offered to the Forum as a basis for discussion and is not a binding proposal on the part of any of the contributing organizations. The statements are subject to change in form and content after further study. Specifically, the contributors reserve the right to add to, amend or modify the statements contained herein.

This contribution deals with the general considerations like traffic patterns, test configurations and applicability of metrics at the specified layers. The metrics namely throughput and latency are dealt with in a separate contribution [96-0520].

Both these contributions are enhancements of the ideas presented in our February 96 contribution [96-0180]. During February 96 meeting, a number of good suggestions were made. We have tried to update the presentations with those suggestions.

Since this is a new endeavor, it is limited in several respects.

This particular proposal concentrates on the data traffic (ABR and UBR service classes) since that is expected to be the bulk of traffic on ATM networks initially. Other service classes will be added later.

TRAFFIC PATTERNS:

There are two types of traffic based on application's response to network performance.

- a. Open loop traffic
- b. Closed loop traffic

Case a) With open loop traffic, the application does not reduce its load when the network performance degrades in terms of throughput or delay. Periodically occurring events generally lead to such traffic patterns. UDP traffic is an example of an open loop traffic.

Case b) With closed loop traffic, the application does slow down when the network response is slower. In many client-server applications, clients will not generate new requests if the previous requests have not been served. TCP, which is expected to be a big part of the ATM market at least initially, is an example of a closed loop application. If the network performance degrades and TCP packets are delayed excessively or lost, TCP will reduce its window and resulting load on the network.

The following figure shows some of the application layer protocols that run on TCP and UDP, respectively.

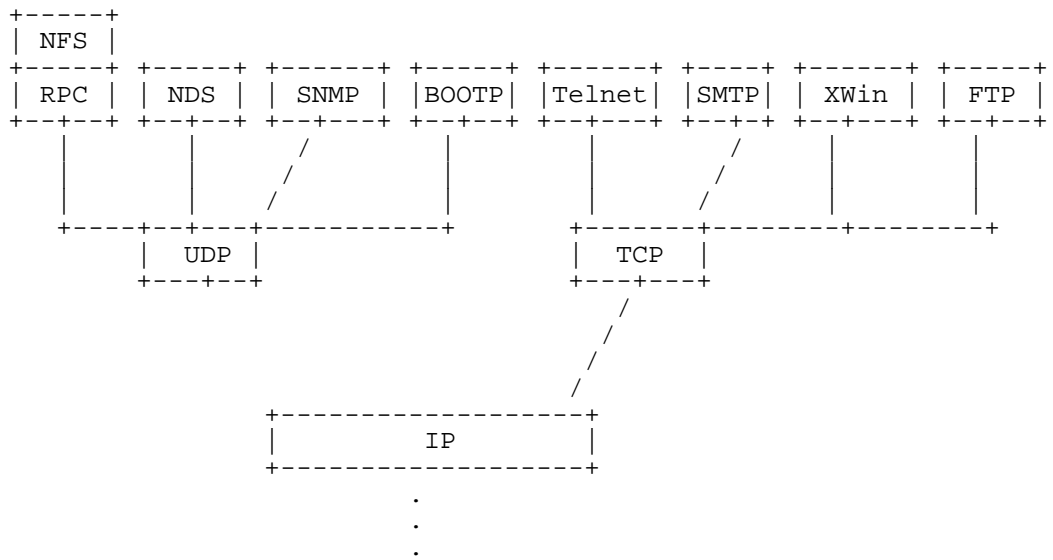


Figure 1 - The protocol stack above UDP and TCP

WHICH LAYER TO MEASURE THE PERFORMANCE?

The idea of coming up with some measuring points is to have as few measuring points as possible and still be able to characterize the network effectively. There are atleast four

layers, where frame-level performance can be measured: AAL, Datalink, Network, and Transport. At each layer, there may be several alternatives. For example, the performance can be measured at AAL5, LAN emulation, IP, or TCP.

The lowest layer where frame-level performance can be measured is the AAL5 layer. This gives us the performance of ATM transport alone without the influence of the higher layer. Throughput, latency and other metrics can be measured at this layer. However, we cannot compare technologies. For example, users interested in comparing ATM to traditional LAN technologies may be more interested in comparing LANE throughput with those at datalink layers of other technologies.

TCP is the reliable transport layer protocol. The reliability of TCP comes at a certain overhead. The performance at the TCP level needs to be studied to understand this tradeoff.

With UDP traffic, the performance is essentially the same as that with IP alone. By examining the passing of datagrams from IP to UDP, it can be seen that measuring at UDP is the same as measuring at the IP. UDP software does not execute as a separate process. Instead it consists of conventional procedures that the IP process executes to handle an incoming UDP datagram. These procedures examine the destination UDP protocol port number and use it to select an operating system queue (port) for the user datagram. The IP process deposits the UDP datagram on the appropriate port, where an application program can extract it. Thus, measuring UDP performance is not included in the minimal set.

As an aside, the above reasoning does not apply to TCP since most designs use a separate process to handle incoming TCP segments. Because of this, IP and TCP must use an interprocess communication mechanism to communicate. Hence, TCP layer is an important measuring point.




Figure 2 - Examples of measurement alternatives

As a result, as shown in Figure 2, above, performance could be measured at any of the five layers: AAL5, LANE, IP over RFC 1577, and TCP. The metrics (for example, throughput, latency, frame loss rate, back to back burst size and call collection time) are different when measured at various points in the protocol stack.

measure the output. This configuration represents an overloaded condition as N inputs are flowing into one switch and there is a single output. Such a condition would result in lower throughput, increased frame loss, lower back to back burst size, higher latency etc. Fairness can also be measured, i.e. if the switch discriminates among the sources.

For Configuration B, the traffic can be sent in the following four ways.

- i) Host_iL sends all of its traffic to Host_iR , $i = 1,2,\dots,N$.
- ii) Host_iL sends $1/N$ of its traffic to every Host_jR, $j = 1,2,\dots,N$, traffic, $i = 1,2,\dots,N$.
- iii) Same as i), but with bidirectional traffic.
- iv) Same as ii), but with bidirectional traffic.

N needs to be determined, overloading depends on the number of sources.

Increase load symmetrically on all ports and measure on the corresponding outputs and this configuration can be used to measure the fairness of the switch.

PERFORMANCE METRICS:

Table 1 lists a number of performance metrics and shows the measurement points to which they apply. As discussed above, there are four measurement points: AAL5-level metrics, LANE-level metrics, IP-level metrics, TCP-level metrics. The table does not list traffic management performance measures and network management performance measures. Those will be added later. Traffic management metrics measure ability of the switches to avoid overload and to efficiently and fairly resolve contention among various VCs when there is overload. Network management metrics are defined to aid characterization of the switch in responding to network management requests.

Output	AAL5	TCP	IP/LANE	IP/RFC1577
Peak Throughput	X	Note 1	X	X
Lossless Throughput	X	X	X	X
Latency	X	X	X	X
Frame loss	X	Note 2	X	X
Back to back Burst size	Note 3	X	X	X
Call establishment latency	X	Note 4	Note 4	Note 4
Max Call establishment rate	X	Note 4	Note 4	Note 4

Table 1 - Metrics and layers to which they apply.

The metrics throughput (both peak and lossless) and latency are defined in [96-0520], presented at this ATM Forum meeting. The rest of the metrics represented in the above table, frame loss, back to back burst size, and call establishment time will be addressed in the contributions in the future ATM Forum meetings.

The following notes justify why measuring a particular metric at a particular layer is not meaningful.

Note 1 - Peak throughput and lossless (zero loss) throughput for TCP are same since TCP is a reliable transport protocol. Any lost packets are retransmitted by TCP and so the TCP users get zero loss. Peak throughput, as defined in [96-0520] is the maximum throughput that can be achieved under lossy conditions. For TCP, this metric is, therefore, not applicable.

Note 2 - Extending the same argument as above leads to the conclusion that frame loss at the TCP layer is zero and so measuring frame loss at TCP is not useful.

Note 3 - The gap between frames in back to back traffic is very minimal. Traffic for back to back metric is sent in bursts. This metric is meaningful if the traffic is bursty and will get through at peak rate due to buffering.

Note 4 - Call establishment time is the time taken to setup a connection with the destination and so is related to the VCs. Hence it is applicable only at the AAL5 layer and the metric cannot be defined at the higher layers.

Call establishment performance is measured in two ways: number of calls that can be made per second (maximum rate) and time to make each call (latency).

REFERENCES:

[95-1347] Raj Jain, "Performance Benchmarking BOF," AF-ALL/95-1347, October 1995.

[95-1662] Raj Jain, Bhavana Nagendra, "Performance Benchmarking of ATM Switches", AF-TEST/95-1662, December 1995.

[96-0180] Raj Jain, Bhavana Nagendra, Gojko Babic, "Scope For ATM Forum's Performance Benchmarking Work Item," AF-TEST/96-0180, February 1996.

[96-0520] Raj Jain, Bhavana Nagendra, Gojko Babic, "Considerations for throughput and latency measurements of ATM Switches," AF-TEST/96-0520, April 1996.

Note: All our past ATM forum contributions and presentations are available on-line at <http://www.cse.wustl.edu/~jain/>