digital

GIGAswitch/FDDI System
IP Packet Switching Throughput Testing

Digital Equipment Corporation sponsored a series of IP packet throughput tests under the observation of Scott Bradner to determine the throughput of a DIGITAL GIGAswitch/FDDI System utilizing 34 full duplex ports. The goal was to provide a packet-per-second forwarding rate for the IP Packet Switching features provided with GIGAswitch/FDDI Systems. These features allow packets destined for different IP subnets than the source to be forwarded at a rate identical to the rate for packets forwarded within an IP subnet.

This is accomplished with an ARP server and other features within the GIGAswitch/FDDI System, which eliminate the need to broadcast ARP requests to the entire network.

The resulting throughput can be compared to the throughput of other layer 2 switches, as well as to the throughput of routers, which would otherwise be required to forward inter subnet traffic.

Summary of Results

The results demonstrate the ability of the GIGAswitch/FDDI to forward packets that would otherwise be routed, at more than 4 million pps with low latency.

Because of equipment constraints, a full 34 port test could not be performed. Instead 4 and 8 port tests were run to see if there is any scaling factor involved in using additional GIGAswitch/FDDI ports. Since the results for the 4 and 8 port tests were about the same (specifically the per port performance on the 8 port test was not less than the per port performance on the 4 port test), no scaling factor was seen, and therefore it is reasonable to extrapolate the observed performance to that of a fully configured switch.

Extrapolating to 34 full duplex ports yields a forwarding rate of greater than 4 million packets per second, for 64 byte packets.

Performance of the ARP server was also measured, to assure it does not present a bottleneck to predicted inter subnet performance. No such bottleneck exists.

Detail of Tests

Goals

1. To establish performance figures for throughput of 64 octet data packets on full-duplex links of the GIGAswitch/FDDI. "Throughput" is defined in accordance with rfc 1242, as "The maximum rate at which none of the offered frames are dropped by the device". The transmission period must be a minimum of 60 seconds of simultaneous transmission on all testers.

2. To demonstrate that full-duplex performance of all 34 ports in a fully loaded GIGAswitch/FDDI can be extrapolated by performing 4 and 8 port tests, and determining whether a scaling factor is required when utilizing a larger number of ports.

3. To demonstrate the ARP serving capabilities of the GIGAswitch/FDDI.

Throughput Tests

a. Setup and Configuration

The test devices used were W&G DA30 FDDI analyzers.

In all cases, full-duplex links were established between the DUT (Device Under Test, or “central” GIGAswitch/FDDI System), and up to three external GIGAswitch/FDDI Systems. The external GIGAswitch/FDDI Systems were each internally divided into 2 or 3 Logical Bridge Domains, with one link (Tn) to a DA30 transmitter, one link (Rn) to a DA30 receiver, and a Full-Duplex link (FDn) to the DUT in each domain. See the section on “Specific Test Data” for additional configuration information.

The utilization of GIGAswitch/FDDI Systems as the “feed devices” to the DUT was required due to the lack of support for full duplex FDDI on the W&G DA30 Test Devices.

For a description of the test methodology utilized, see the section entitled "Test Methodology Overview".

b. Measured Performance

The highest average lossless throughput demonstrated was 118,800 packets per second, per port, transmitted and received simultaneously on 8 full-duplex ports on the DUT. The peak range was 118,781 to 118,846 pps, per port.

c. Scalability

The per-port forwarding performance of 4 Full-Duplex ports was identical to that of 8 Full-Duplex ports, 118,800 pps. In both cases, increasing the data rate beyond that led to a decrease in throughput (but no packet loss) due to FDDI claiming.
Since no scaling factor was observed it is reasonable to assume that a fully populated GIGAswitch/FDDI can sustain the same 118,800 PPS per port, which would mean an aggregate rate of 4,039,200 64 byte packets per second.

ARP Tests

100 ARP requests for a non-existent IP address were sent from a DA30. These ARP requests were formatted to appear to have come from 100 separate IP addresses. This allowed the GIGAswitch/FDDI System to add these 100 IP addresses to its ARP database.

100 ARP requests were then generated for the above 100 IP addresses from a 2nd DA30.

By varying the frequency of ARP requests and noting the times of the responses, these tests demonstrated a sustained ARP response rate of approximately 2000 responses per second.

SPECIFIC TEST DATA

Location and Equipment used in Tests

The serial number of the switch under test was AS23601003. All modules in this switch, as well as the intermediate switches, were running the v3.1 release of firmware. The tests took place at the Digital Equipment Corporation, Littleton, Massachusetts facility, under the observation of Scott Bradner.

How the Switches Were Connected

In the tables below cables are identified as

- Rn - cable to receive port on tester n
- Tn - cable to transmit port on tester n
- FDn - cable between 2 GIGAswitch/FDDI full duplex ports, one on the DUT, the other in logical bridge domain n on one of the other switches.

Each table shows which cable is connected to each port, as well as the logical bridge domain of each port (except for GS #1, which only has a single logical bridge domain.

GS #1 (DUT)

Slot	Card	Port	Cable
6	fgl-4	6.1	FD6
		6.2	FD4
		6.3	FD3
		6.4	FD5
9	fgl-4	9.1	FD8
		9.2	FD1
		9.3	FD2
		9.4	FD7

GS #2

Slot	Card	Port	Cable	LB Domain
2	fgl-4	2.1	T1	1
		2.3	R1	1
		2.4	R2	2
3	fgl-4	3.1	FD1	1
		3.4	R3	3
9	fgl-4	9.1	FD2	2
		9.3	T2	2
11	fgl-4	11.1	FD3	3
		11.3	T3	3

GS #3

Slot	Card	Port	Cable	LB Domain
5	fgl-2	5.1	R4	4
		5.2	T4	4
6	fgl-2	6.1	FD4	4
		6.2	R5	5
10	fgl-2	10.1	T5	5
		10.2	FD5	5

GS #4

Slot	Card	Port	Cable	LB Domain
10	fgl-2	10.1	T6	6
		10.2	R6	6
11	fgl-4	11.1	T7	7
		11.2	R7	7
		11.3	FD7	7
		11.4	FD6	6
12	fgl-4	12.1	T8	8
		12.2	R8	8

12.4	FD8	8

NOTE: The 4 port configuration utilized only testers 3,4,5,6. The links to ports on the fgl-4 in slot 9 were not used in this test.

TEST METHODOLOGY OVERVIEW

"Fanout De-Multiplexed" Throughput Test Methodology for GIGAswitch/FDDI

This test requires 8 Windows-based DA30s, each configured with 2 FDDI analyzer modules. On each DA30, analyzer #1 utilized the FLEXmit application to continuously transmit 100 PDUs, each 64 octets in size. On DA30 #1, these PDUs alternated their Destination MAC addresses (DA) between 00-00-00-00-00-01, 00-00-00-00-00-02, 00-00-00-00-03, and 00-00-00-00-00-04. The 2nd DA30 continuously transmitted 100 PDUs to DAs which cycle through 00-00-00-00-00-02 through 00-00-00-00-00-05, the 3rd DA30 transmitted to 00-00-00-00-00-03 through 00-00-00-00-00-06, and so on. The source MAC address of these 100 PDUs must be unique for each DA30; a convention of 00-00-00-01-00-00 for the SA of the 100 PDUs transmitted by the first DA30, 00-00-00-02-00-00 for the second, etc., was used.

Prior to enabling traffic receipt, Analyzer #2 on each DA30 (the receiver) was configured to transmit 1 or more broadcast "learning" PDU's. The source MAC addresses of these PDU's was the MAC addresses which the GIGAswitch/FDDI would learn, and utilize to forward the test PDU's transmitted by the Analyzer #1's in the "threshold" part of this test. Thus, Analyzer #2 on the first DA30 used a SA of 00-00-00-00-00-01, AN2 on DA30 #2 transmitted using a SA of 00-00-00-00-00-02, and so on. Address age-out on the GIGAswitch/FDDI must be set to a large enough value to preclude the aging-out of the learned MAC address during the threshold testing. In our testing, a value of 30 minutes was utilized for age-out. In addition, the Spanning Tree Protocol was disabled, so that STP BPDU's would not affect the packet count.

Once the "Learning PDU's" had been transmitted by each Analyzer #2, the packet counters of all 16 analyzers were reset to zero, and all 8 transmitters were manually started, each transmitting their continuously repeating stream of 100 PDU's.

After all transmitting analyzers had been simultaneously transmitting for a duration of at least 60 seconds, they were all manually stopped. The packet counts of all AN1 analyzers were added together, and the packet counts of all AN2 analyzers were added together. "Lossless Throughput" could be gauged at the transmitted packet rate where the two classes of cumulative packet counts (all AN1 PDU's and all AN2 PDU's) are equal, indicating 0 packets lost.

Transmit Topologies

Each packet transmitter (link T1,T2,…,T8) continuously transmitted sets of 100 PDU's to 4 receivers (Rn). 25% of the 100 packets went to each of the transmitter's 4 destinations. This was achieved by alternating the MAC DA such that PDU's 1,5,9,13,…,97 were sent to the first receiver; PDU's 2,6,…,98 were sent to the second; PDU's 3,7,…,99 were sent to the third; and PDU's 4,8,…,100 were sent to the fourth receiver.

In this way each transmit port sent to 4 receive ports - not including itself. The receive port in a transmit port's Logical Bridge Domain was never in the destination set of the transmitter. This insured that all transmitted packets traversed the Device Under Test.

NOTE: When changing from 8 to 4 Full-Duplex ports in order to test scalability, it was necessary to modify the number of destination MAC addresses sent by each transmitter from 4 to 2, in order to assure that none of the transmitters was sending to a receiver in the same Logical Bridge Domain. For this reason, final testing at 118,800 FPS for both 4 and 8 ports was performed with multiplexed destination traffic consisting of 2 MAC Destination addresses, rather than 4, as per the above.