Measuring Forwarding Delay Across a Campus Network

OVERVIEW

In June 2006, I spent a morning with Mike Pennachi and Chris Greer of Network Protocol Specialists, measuring forwarding delay across our network. This document summarizes these measurements.

1. The more expensive the device, the greater its forwarding delay.
2. In-line taps do not contribute to forwarding delay.
3. Our network typically inserts ~20-80us of latency, inclusive of packet insertion type and exclusive of queuing delay.

METHODOLOGY

We employed a Finisar THG Analyzer to measure forwarding delay, placing the device under test between the two interfaces of the THG unit, producing a stream of 10,000 packets from one interface, capturing with the other interface, exporting the results into Excel, and calculating an average and standard deviation of the results. In all tests, the standard deviation was sufficiently small that we did not bother to record it.

Finisar advertises that the THG unit's clock is accurate to 10ns.

RESULTS

Lab Gear

Individual device, connected only to the THG unit, i.e. the only traffic flowing over the wire, in addition to our test traffic, is whatever BPDU, CDP, etc. frames the device under test produces by default.

Device	**Frame Size**	**Latency**
Datacomm copper aggregation tap using 100BaseTX ports	64 bytes	320ns
Datacomm in-line copper tap using 100BaseTX ports	64 bytes	0ns
NetGear 10/100Mb switching hub using 100BaseTX ports	64 bytes	330ns
Catalyst 4003 w/Sup I using 100BaseTX Ports	64 bytes	3170ns
Catalyst 4003 w/Sup I using 1000BaseSX Ports	1518 bytes	3170ns
Catalyst 4503 w/Sup V using 1000BaseSX ports	64 bytes	3300ns
Catalyst 4503 w/Sup V using 1000BaseSX ports	1518 bytes	7120ns

Production Gear

At this point, we moved onto our production network: ~7000 nodes spread across a campus network containing a dozen buildings.

For the Campus Network test, we sent pings from a 1000BaseTX attached PC in one building across our network to a 1000BaseSX attached PC in another building and inserted in-line taps between each PC and its local switch. The in-line taps attached to the two ports of the THG unit, employing in-house glass as needed to drag the remote tap's signal back to the THG unit.

The packets traversed the following electronics: Catalyst 4503 w/Sup V (access-layer: m1w-esx), Catalyst 6506 w/Sup 720 (building distribution layer: md-a-rtr), Catalyst 6506 w/Sup 720 (core layer: core-a-rtr), Catalyst 6506 w/Sup 720 (building distribution layer: ja-a-rtr), Catalyst 4503 w/Sup V (access-layer: j4-test-esx), plus about 500m of copper and glass cabling.

Device	**Frame Size**	**Latency**
Catalyst 6506 w/Sup 720 using 1000BaseSX ports	64 bytes	5000ns
Catalyst 6506 w/Sup 720 using 1000BaseSX ports	1518 bytes	7120ns
Campus Network	70 bytes	20,240ns

To sanity check the 20,240ns number, I added together the individual forwarding delay numbers acquired during our Lab Gear and Production Gear tests:

m1w-esx contributes	~3us
md-a-rtr contributes	~5us
core-a-rtr contributes	~5us
ja-a-rtr contributes	~5us
j4-test-esx contributes	~3us
500m of cable contributes	~3us
Total equals	~24us

That 24us number is pretty close to the 20,240ns number we recorded during the Campus Network.

LATENCY

Our campus network – consisting typically of five Ethernet switches plus ~500m of cabling between any two points – inserts 20,240ns ~= 20us of forwarding delay.

Given the store-and-forward nature of our electronics, and a typical width of five Ethernet switches, I calculate latency across this campus network by adding the cost of insertion delay[1] (for five switches) to the 20,240ns figure.

64 byte packets require 22,800ns ~= 23us to cross our network.
1518 byte packets require 80,960ns ~= 81us to cross our network.

---

[1] Insertion delay for 64 byte packet at Gigabit Ethernet = 512ns. Insertion delay for 1518 byte packet at Gigabit Ethernet = 12144ns.