Fault Timing Analysis Test Report

GarrettCom
Ethernet at Its Best

Test name: Gigabit S-Ring (Power ON/OFF) with Tagbased Vlan

Test By: Rajeev at GCI / Fremont Date: 08/12/04

Redundancy Software used: _____ Tester's Contact Info, ph: 510-580-2783.

Release # Rel 2.9 beta email: rajeev@garrettcom.com

MNS-6K or other net mgt software used: MNS-6K. Rev #: Mig.10.935A.or

Media types are: (circle all that apply) (Fiber) (RJ-45) Speed: 10 (100) (1G)

Describe the ring or mesh members (mfg, model, etc): 6K magnum switched connected in an S-ring

Number of Ethernet units in the ring / mesh: Hubs: _____ Switches: 6. Other: _____

Describe nodes attached to the LAN, type, quantity, special features: M6K-5 running S-ring software or has forwarding gigabit port # 17 and Hostlight 25

Distances of media: 10 feet.

Topology, brief description: Gigabit S-Ring

Ring, mesh, or combination: Ring.

Other products in the LAN, h/w and s/w: FTA on M-6K-6 and Target connected to M-6K-1. M-6K-5 running S-ring software.

FTA Software running on: (describe the PC) Computer (P-4)

Operating System on the FTA computer: Windows 2000.

Ping Target is: mP62

Notes: (Tester's comments, other ring activity noted during the test, where was fault introduced, etc.)
Gigabit S-ring to tested for power failure on Mgr. mgr switch is cold rebooted to check the ring redundancy & fault timings measured as below.

Results: Fault Times (in seconds, with 3 decimal places)

Minimum	Average	Maximum	No. of incidents
0.264	0.280	0.297	20

GarrettCom Test 8/12/04
Ring Manager Fault Test

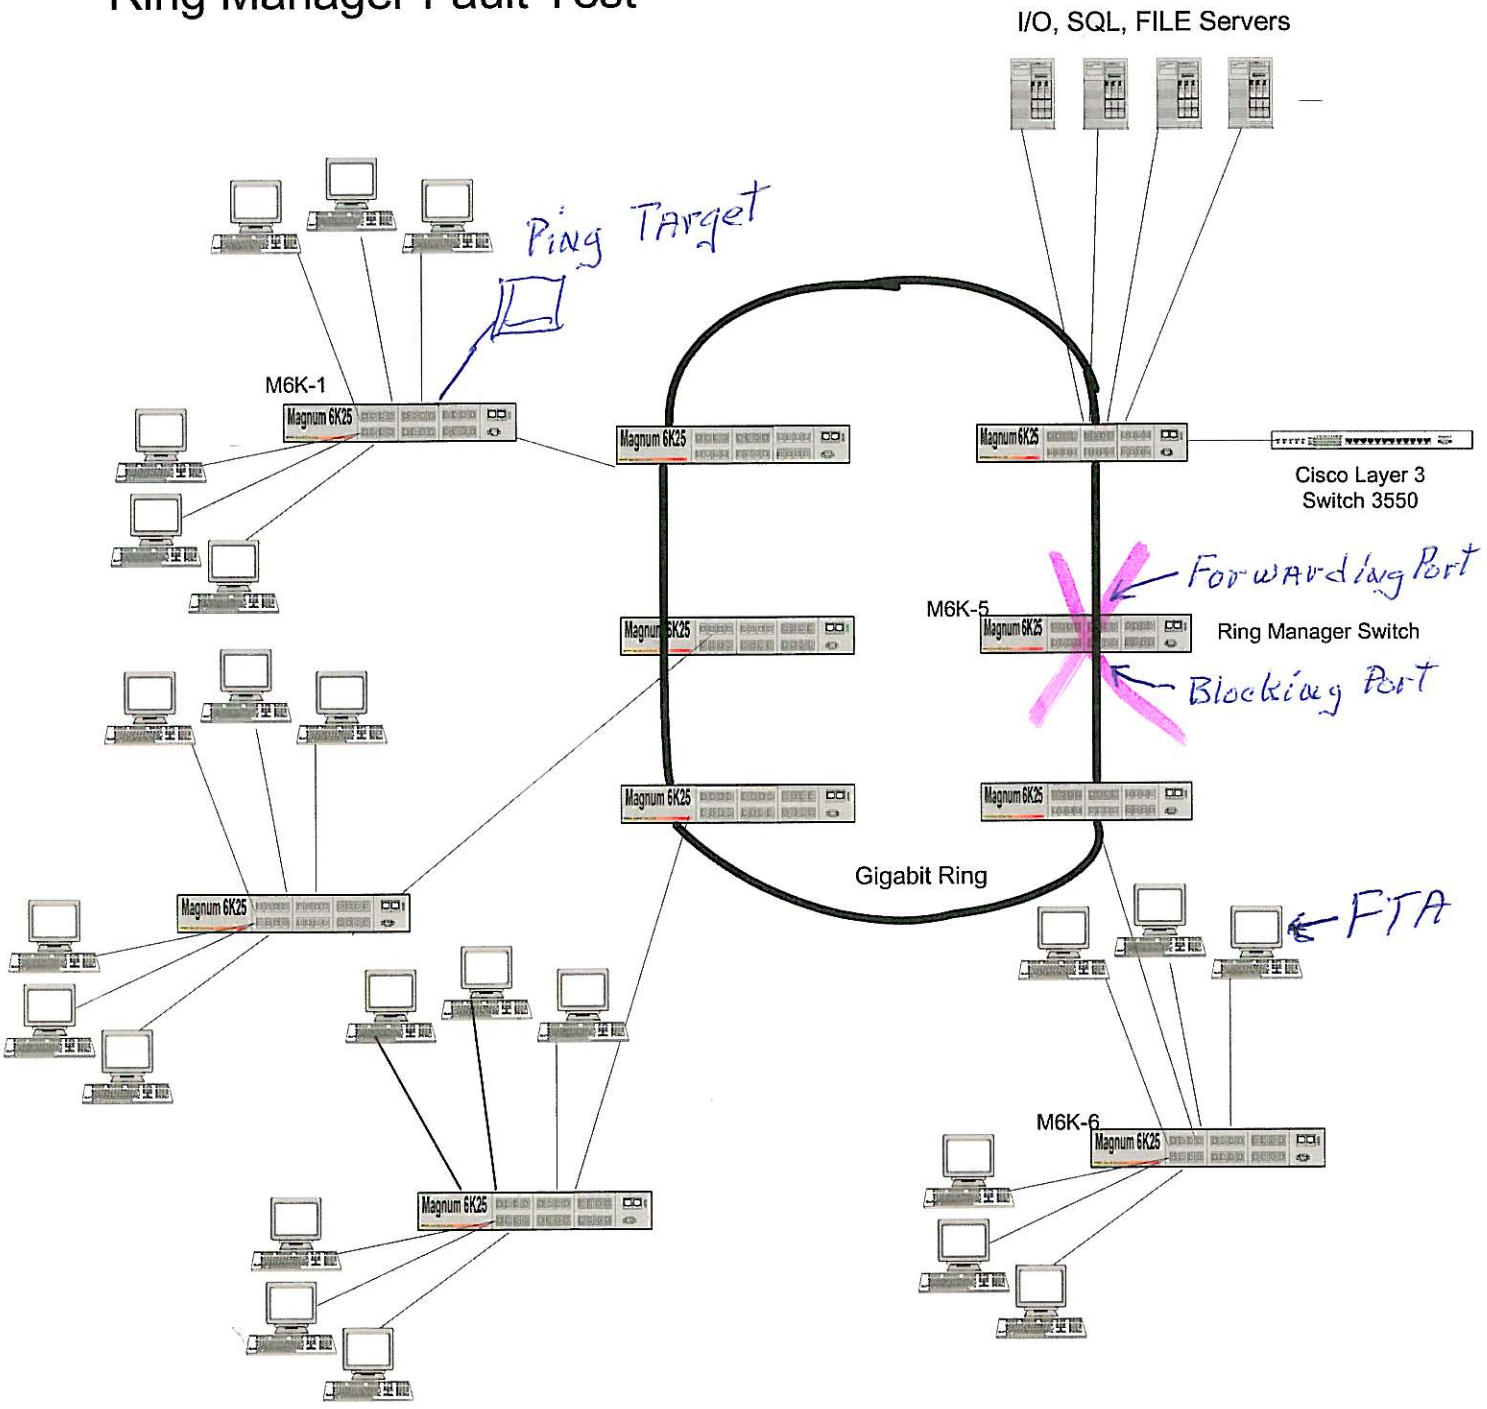

Test Objective: Show the effect of S-Ring Ring Manager Switch failure, a potential "Single Point of Failure."

Servers -HMI-IO, File servers, SQL server, VOIP

S-Ring running in Gigabit Ring

Test Set-Up: topology per diagram. With FTA running, pull the power plug on X, measure the effect on ring traffic.

Measuring S-Ring Fault Recovery with FTA in a Gigabit Ring with no Single Point of Failure

Here is FTA in action, with an example of its use from a recent lab test performed at GarrettCom. We were focusing specifically on the issue of single-point-of-failure with S-Ring, and how such a redundant LAN could be constructed and operated.

We operated the test LAN as a normal set-up, with S-Ring and LLL configured and operating on all of the Magnum 6K25 ring switches with two gig ports each. For the Ring Manager, we specifically chose one 6K25 (designated as MK6-5, see diagram attached) that had no other devices attached to it, so that the loss of this particular 6K25 would not take down any other operating equipment in the LAN.

With things running normally, we pulled the power plug on the Ring Manager Switch for this fault test, and measured the result with FTA. We repeated this test 20 times, as shown in the Test Report attached.

The result was a loss of traffic in the gig ring for about 1/4 second. This is exactly what the theory says should happen. Understanding this test result is part of understanding S-Ring and self-healing LANs generally. Referring to the diagram, you see that the Ring Manager Switch in normal operation has one ring port forwarding and one ring port blocking. This means that none of the ring traffic is actually passing through the Ring Manager Switch itself . . . it obviously cannot with one port blocking. The ring is actually a string, and the Ring Manager is at the end of the string. When the Ring Manager Switch fails, the traffic in the string is maintained and all of it passes through the same paths as it did before the Ring Manager failure. Nothing in any of the traffic paths has changed.

The fault time of about 1/4 second occurs because the ring switches are triggered to flush their buffers and re-learn all node positions on a link-loss. Recall that S-Ring and Link-Loss-Learn, upon a loss-of-link on a ring port anywhere in the ring, cause that switch to flush buffers and also propagate the fault signal down the line to other switches to do the same thing. In this test case, the switches on either side of the Ring Manager see a link-loss on a gig port because the Ring Manager Switch has died. They flush their buffers, propagate the fault signal, and re-learn all node positions. Since no node positions have changed in this particular test case, all ring member switches re-learn the same nodes set-up and resume operation. This consumes the 1/4 second of fault time.

If you need a refresher on the theory of Link-Loss-Learn, it is on our web site at http://www.garrettcom.com/techsupport/papers/lll_techbr.pdf.

This test result does not mean that customers should always use S-Ring in their self-healing gig rings, or any other rings for that matter. It only means that S-Ring is available, and we believe it should not be ruled out by a “no-single-point-of-failure” requirement. Both S-Ring and RSTP are available from GarrettCom, and customers now can have fault timing test data on each of them to help you make their choice for use in redundant LANs.

GarrettCom, September, 2004