Concordia University College of Alberta

Master of Information Systems Security Management (MISSM) Program

7128 Ada Boulevard, Edmonton, AB

Canada T5B 4E4

Study of Network Instability in VRRP and HSRP sub second timer implementation

By

Ekemezie, Emmanuel Ikechukwu

A research paper submitted in partial fulfillment of the requirements for the degree of

Master of Information Systems Security Management

Date: April 2012

Research advisors:

Dr. Pavol Zavarsky, Director of Research and Associate Professor, MISSM

Ron Ruhl, Director and Associate Professor, MISSM

Dr. Dale Lindskog, Assistant Professor, MISSM

Study of network Instability in VRRP and HSRP sub second timer implementation

By

Ekemezie, Emmanuel Ikechukwu
emmaekemezie@yahoo.com

Research advisors:

Pavol Zavarsky, Director of Research and Associate Professor, MISSM

Ron Ruhl, Associate Professor, MISSM

Dale Lindskog, Assistant Professor, MISSM

Reviews Committee:

Pavol Zavarsky, Associate Professor, MISSM

Ron Ruhl, Associate Professor, MISSM

Dale Lindskog, Assistant Professor, MISSM

The author acknowledges the significant contributions to the work by Academic Advisors and Review Committee Members and gives the right to Concordia University College to reproduce the work for the Concordia Library, Concordia Websites and Concordia MISSM classes.

Study of Network Instability in VRRP and HSRP sub-second timer implementation

A research paper submitted to Concordia University College of Alberta in conformity with the requirements for the degree of Master of Information Systems Security Management.

Author: **Ekemezie, Emmanuel Ikechukwu;** April 2012.

Research Advisors: Pavol Zavarsky, Ron Ruhl, Dale Lindskog

emmaekemezie@yahoo.com, {pavol.zavarsky, ron.ruhl, dale.lindskog}@concordia.ab.ca

Abstract – Network instability is a big concern when implementing sub-second timers. Implementing First Hop Redundancy Protocol (FHRP) sub-second advertisement interval can increase the potential for network instability. In this paper we studied the impact of sub-second timer implementation on network stability for fast failure detection and failover of two FHRP - VRRP (Virtual Router Redundancy Protocol) and HSRP (Hot Standby Redundancy Protocol). The goal of our research was achieved by simulating a production network, and implementing sub-second timers under induced process switched network traffic while varying timer values- increasing values set progressively from 50msec through 999msec- and varying the target of the induced network traffic between different interfaces on the active/master and backup/standby routers. From the results of our experiments, we conclude that sub-second timer implementation can adversely impact network stability given two conditions. When either the input queue of the interface on the backup/standby router used for keepalive message reception overflows and/or the output queue of the interface on active/master interface used for keepalive message forwarding overflows.

Keywords – Redundancy; high availability; FHRP; keepalive message; sub-second timers; network stability; HSRP;VRRP.

1.0 INTRODUCTION

Network instability caused by sub-second timer implementation is as a result of continuous change in state of FHRP. High availability with sub-second failover is a requirement in today's enterprise. Real-time business applications and transactions have stringent QoS requirements for packet loss, latency and delay. To mitigate the impact of link or node failure, modern networks are built with redundant components- links, chassis, processors etc. However, even with these measures, the default gateway still constitutes a single point of failure for nodes that lack the means to dynamically track the liveliness of the gateway. Virtual Switching System (VSS) is a Cisco proprietary mechanism that makes it possible for two Cisco Catalyst 6500 series switches to function as one logical device. It differs from HSRP in functionality and features provided. To illustrate, the VSS mechanism is implemented only in Cisco Catalyst 6500 series switches [1] and more importantly their proprietary nature makes interoperability impossible.

FHRP are designed to eliminate single point of failure by providing gateway redundancy by using multiple routers in a group to create a virtual router. FHRP ensure that the default gateway dynamically failover to another router if the gateway or link to the gateway becomes unavailable. VRRP [2][3][4] and HSRP [5][6][7] are two popular protocols that can be used to achieve this goal. Each of these protocols uses keepalive messages to monitor the health of the active/master router and to determine when it has failed though they use different terminologies to reference these mechanisms. VRRP uses advertisement, master down interval, master and backup while HSRP uses hello, hold timer, active and standby terminologies.

Gateway Load Balancing Protocol (GLBP) is another FHRP. But it adds load balancing automatically by allowing distribution of traffic flows between all the routers in a virtual router. It is a Cisco proprietary technology like HSRP [8]. By implementing FHRP, a number of routers can be designated as backup/standby routers in the event that the master/active router fails. When a master /active router fails, FHRP dynamically transfers the packet-forwarding responsibility to a backup/standby router. FHRP Virtual Router Group (VRG) enables a set of routers to be grouped as a virtual router that answers to a virtual IP address.

In a production environment, backup routers do not seat idle waiting for the master router to fail. They are normally configured in a manner that enables load sharing across all members in the group. Therefore, a router can function as a master in one virtual router group and as a backup in another virtual group. FHRP uses a mechanism which allows hosts to keep a single IP address for the default gateway, and maps this IP address to a well-known virtual MAC address. FHRP provides redundancy without user intervention or additional configuration at the end hosts.

Initially when redundancy protocols were created the need for sub-second failover was minimal because the industrial demand for real-time business applications and transactions were almost nonexistent. To meet the requirements created by the demand of these applications and transactions, newer versions of FHRP protocols are designed to support sub-second timers. HSRP version 2, VRRP version 2 and 3 include the ability to implement sub second timers in their protocol design in order to achieve fast failover. Though sub-second timers can help to achieve fast failover it can also increase the potential for network instability. Instability could occur due to processing requirements within the router preventing the processing of advertisement or loading conditions on the network preventing reception of these advertisements [9]. In a production network, network congestion or router overload is a real possibility. There are numerous conditions which can force a router to process switch a large volume of traffic. Some of these include excessive broadcast traffic, denial of service attacks against the router or large amounts of traffic from directly connected hosts with bad routes. Traffic such as these requires processing by the CPU, leading to high utilization. This condition can cause a delay in processing advertisements by the router.

1.1 Router interface

The network Interface of a router is used for receiving and forwarding of packets or traffic by the router. It has structures input/output queue used to store packets temporarily. Architecturally, modern routers have the ability to process switch or fast switch a packet.

Process switching requires a router to always look into the routing table whenever a packet is received before making a routing decision. This process is computationally expensive as it requires multiple CPU cycles for each lookup. Traffic that causes a router to forward the traffic through the same interface it was received on, is always process switched by the router.

Fast switching is a switching technology that enables a router to look into the routing table once and create a route cache for first packet received going out an interface. Subsequent packets belonging to the same flow is then forwarded by the router looking into the route cache instead of the routing table again.

The Cisco 3640 router used in this experiment is a shared memory router. In a shared memory router, a packet is held in the input queue corresponding to the ingress interface. Next, an output interface for the packet is possibly selected (perhaps after the process switching or fast switching task completes). Then the packet is held in the output queue of the egress interface, until it is finally sent to the output interface hardware [10].

The rest of this paper is organized as follows: Section 2 presents related work; section 3 presents experimental methodology and result; while section 4 presents discussions; section 5 presents research limitations; finally, section 6 concludes the paper.

2.0 RELATED WORK

There are only a few previous studies out there on the performance of FHRP redundancy protocol. Jen-Hao Kuo et al in their paper [11] evaluated the performance of VRRP and a derivative they developed -enhanced VRRP. The enhanced VRRP has the same basic features as VRRP but with the addition of dynamic load balancing. RFC 5798 describes a possible consequence of sub second timers implementation. Advertisements may not be received by backup routers because of packet loss or delay within the advertisement interval [12]. A study by Cisco mentioned that a random, momentary loss of data communication between FHRP peers is the most common problem that results in network instability, FHRP state changes are often due to High CPU Utilization or excessive network traffic; Cisco recommends setting the timer to a value no less than one second [13]. No independent previous work has studied the effect of sub-second timer implementation in the LAN environment.

3.0. EXPERIMENTAL METHODOLOGY AND RESULT

To study how VRRP and HSRP subsecond timer implementation impacts network stability, the setup described below was used to simulate a production network with two virtual groups. The simulated network consisted of three Cisco 3640(R4700) routers with processor speed of 100MHz, each with 2 fastethernet ports and one Cisco Catalyst 2950 switch with 24 fastethernet ports used to interconnect all the other components as depicted in figure 1 below. The routers were assigned the following roles; determined by their configured priorities depending on the virtual group (VG) they belong: Router 3 with the highest priority was Active/Master for VG 1 and standby/backup for VG 2; router 1 with the next highest priority was the standby/backup for VG 1 and master/active for VG 2 router.

Router 2 was backup/speak for both VG 1 and VG 2. Monitoring and traffic generator tools were installed on two Windows computers. Connectivity within the network was verified using ping and trace route utilities.

NetFlow version 9 was configured on the routers to export traffic flow information from each router. The exported traffic flow information was collected with the PRTG network monitoring tool server (192.168.20.5) to present the data in a graphical format. Colasoft capsa 7 was also used to monitor traffic on each router interface by sniffing the LAN interface of the server that is connected to the test bed network. Cisco IOS debug tools was used to collect the CPU utilization of the routers, wireshark was used to monitor the communication between VG members.

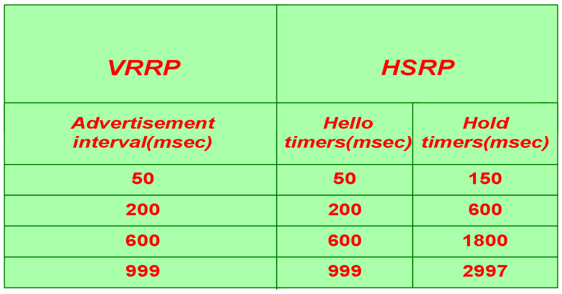

VRRP	HSRP	
Advertisement interval(msec)	Hello timers(msec)	Hold timers(msec)
50	50	150
200	200	600
600	600	1800
999	999	2997

Table 1: FHRP sub-second timers used for experiment testing

3.1 Experimental Procedure

The experiments were conducted using traffic that requires process switching like excessive broadcast traffic in order to study the impact of subsecond timer implementation on network stability. We designed the experiments to explore four scenarios. An explanation of our observations is provided after each scenario this then leads to the next experimental scenario.

First, sub-second timers where implemented with no traffic on the network in order to observe the impact of sub-second timers under these conditions. Second, sub-second timers were implemented with network traffic injected into the test bed network targeting the master/active router to simulate a DoS attack. Thirdly, sub-second timers were implemented with network traffic injected into the network targeting the backup router representing a DoS attack. Finally, sub-second timers were implemented with injected network traffic targeting a host in a different subnet from the network representing large amounts of traffic from directly connected hosts with bad routes. Each of these scenarios is detailed in the sections following below.

i) Sub-second implementation with no induced traffic on the network

The setup depicted in Figure 1 below was used to investigate the impact on network stability of VRRP/HSRP sub-second timer implementation with no other traffic on the network. This setup was used to determine the sub-second timer values that result in network instability and network resource usage. The values in table 1 lists the various timers used for this experiment. Both protocols were evaluated; VRRP first and subsequently HSRP.

When configuring VRRP, the advertisement interval is the only value that needs to be manually configured, the master down interval is automatically calculated from that. A continuous ping session was run from 192.168.20.5 to 192.168.50.3 (loopback address of a router on a different subnet); the essence was to enable observation of traffic flow from one subnet to the other that required use of the default gateway (virtual IP address 192.168.20.200). Hence, whenever there was a disruption in the FHRP operation it was observed as timeouts in the continuous ping session. The impact on network stability and traffic flow was observed using monitoring tools. Each timer value listed under the VRRP column in table 1 was configured and tested. Then VRRP was shut down on the three routers and HSRP was enabled. Thereafter the listed timer values under the HSRP column in table 1 were configured and instability was investigated. In configuring HSRP timer, hello and hold timers were configured because HSRP does not have the capability of calculating its own hold timer (master down interval timer in VRRP). The same procedure followed during the study of VRRP was used to observe and capture data for study.

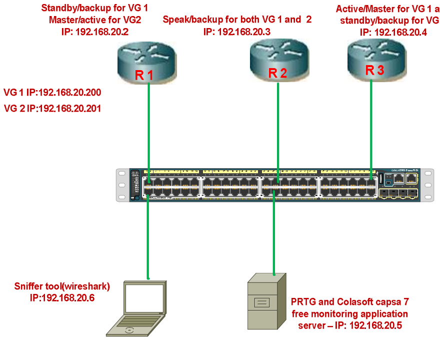

Figure 1: Test bed for sub second impact on network connectivity and traffic

Result

When VRRP was configured, we observed from the above described experimental procedure that the continuous stream of advertisement messages sent by master router increased the rate of traffic passing through the router and entire test bed network. We found out that the backup routers traffic remain constant as various timers value where configured, the reason being that they do not advertise keepalive message. Only the Master router influenced the total traffic on network because it advertises keepalive messages to backup routers. Also the smaller the timer value parameter the more traffic that is sent from the master router into the network.

During testing with HSRP sub-second timers, we observed that both the active and standby routers influenced the total traffic on the network because both advertise hello messages to prove they are alive. Router 2 retained the traffic level it had during the VRRP sub-second implementation, the reason being that it does not advertise any keepalive message but rather it listens to hello messages and renews its hold timer. For HSRP 15msec is the least configurable timer value while VRRP has 50 msec as the least configurable value in Cisco's implementation of these FHRP. When 15msec was configured we noticed standby router attempted transition to the active state but received a hello while still in transition and reverted back to standby state. When 50msec was configured there were no network instability observed, the same applied with the remaining timer values listed in table 1 above.

We noted the output from the show interface command for each router as shown in figure 2. This provides data about input and output packet rates, number of packet drops in the input and output queues. Our objective was to verify that there were no drops or errors in either the input or output queues at this point. From the foregoing we established a baseline; we obtained data about the input and output packet rates, errors and drops in the input and output queues when there is no traffic on the network. In figure 2 below, the following parameters can be observed; queuing strategy in use on this interface is first in first out (fifo) the first packet received is forwarded first. 5 minute input and output rate means average number of bits and packets transmitted per second in the last 5 minutes. Packet input and output describes the amount of error free packet received. Output and input queue displays the maximum size of the queue, and the number of packets dropped due to a full queue. Input error is the errors encounter during packet processing. To clarify, our emphasis is on the errors and drops parameter.

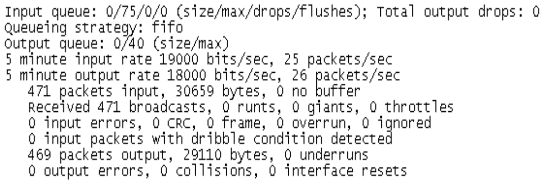

```
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 19000 bits/sec, 25 packets/sec
5 minute output rate 18000 bits/sec, 26 packets/sec
   471 packets input, 30659 bytes, 0 no buffer
   Received 471 broadcasts, 0 runts, 0 giants, 0 throttles
   0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
   0 input packets with dribble condition detected
   469 packets output, 29110 bytes, 0 underruns
   0 output errors, 0 collisions, 0 interface resets
```

Figure 2: Output show interface command without induced traffic.

Subsequently we varied traffic load on the network to study the effect of this on the test bed. We did this by inducing traffic in the network directed towards the master/active routers. The procedures used and results are outlined in the sections following below

ii) Sub-second timer implementation with a DoS attack targeting master/active Router

The setup illustrated in Figure 3 below was used to study network instability when master/active router is under DoS attack. A host machine was used to simulate a DoS attack by injecting packets with a destination IP address of 192.168.20.4; fastethernet0/0 of the master. This was the interface used to send keepalive messages. The timer values listed in table 1 were configured and the network observed for instability. The DoS attacker shown in Figure 3 was then reconnected to target IP address 192.168.40.2; interface fastethernet 0/1 of the master. This interface was not used for sending keepalive messages. The timer values listed in table 1 were implemented and the network again observed for instability.

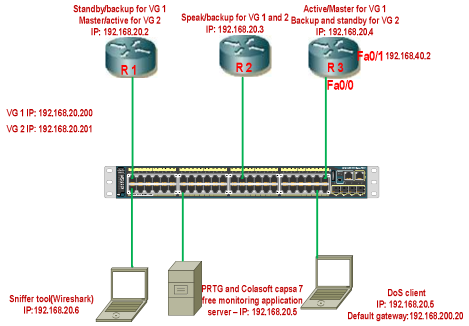

Figure 3: Test bed for induced network traffic test on active/master router

Result

As described above, two cases were studied when the master/active was targeted by a DoS attack. The essence was to determine the impact of high CPU utilization and the state of the interface input/output queue on network stability when implementing sub second timers. In the first case, the interface used for traffic forwarding and sending of keepalives was the target of the DoS attack. In the second case, the other interface that was not active in the relevant of VRRP/HSRP group was the target.

In the first case, the DoS traffic gradually increased the amount of traffic being processed by the active/master router to the extent that its CPU utilization approached 100%. Irregularities were observed in the intervals between keepalives sent by the master/active router. However no network instability was observed even when the CPU utilization hit the 100% mark; the same irregularities were observed in the intervals between keepalives sent by the master/active router in second case as well. Likewise no instability was observed in the network.

These irregularities were probably due to the drain on the router's resources-cpu, memory- caused by the processing required by the induced network traffic targeted at the router. Since the variance in the intervals between keepalives was smaller than the master_down_interval, these did not cause the standby/backup routers to transit to master/active state.

After establishing that high CPU utilization alone does not lead to network instability, we changed target of the DoS to the standby/backup router, which is the scenario discussed next.

iii) Sub-second timer implementation with DoS client targeted Standby/Backup Router

The setup shown in figure 4 was used to study network stability when a DoS attack is directed at backup/standby router. The DoS attack had a destination IP address of 192.168.20.2(fa0/0), which was the interface on the Standby/Backup used for keepalive message reception. The timer values listed in table 1 were configured. Then the setup was reconnected such that DoS traffic targeted destination IP address of 192.168.60.8. This was fastethernet 0/1 of the standby/backup, which was a different interface than the one that was used for keepalive message reception. Timers listed in table 1 were configured and network stability observed.

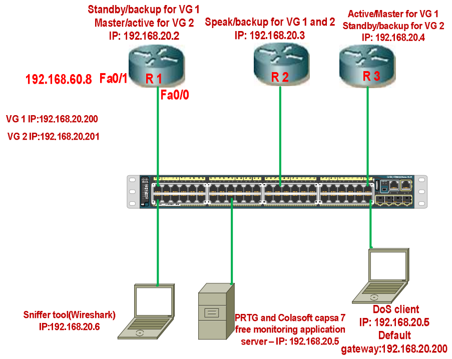

Figure 4: Test bed for induced network traffic test on backup/standby

Result

As described in the preceding section, in this scenario, two cases were explored. In the first case, the DoS attack targeted the interface used for keepalive messages reception. While in the second case, the DoS client targeted a different interface than that used for keepalive messages reception.

When the interface used for keepalive message reception was the target of the DoS attack, the backup/standby began transition to master/active state even before the processing caused by DoS tarffic increases CPU utilization to 30%.

On the other hand, when the interface targeted by DoS attack was different from that used for keepalive message reception, even when the CPU hit 100% utilization no instability was observed. This suggests that CPU utilization is not the cause of instability contrary to a study by Cisco stating that often CPU utilization is the cause of FHRP instability as mentioned in section 2 above(related research)[13].

Since the master/active router was not the target of the DoS attack; the intervals between keepalive remained constant.

PRTG network monitoring tool was used to observe network instability when the timers values listed in table 1 were configured progressing from 50msec to 999msec. Figure 5 below displays timer values and volume of traffic on the network when instability was observed with each configured sub-second timer. The figure demonstrates that the smaller the timer setting, the smaller the volume of traffic observed on network that causes instability.

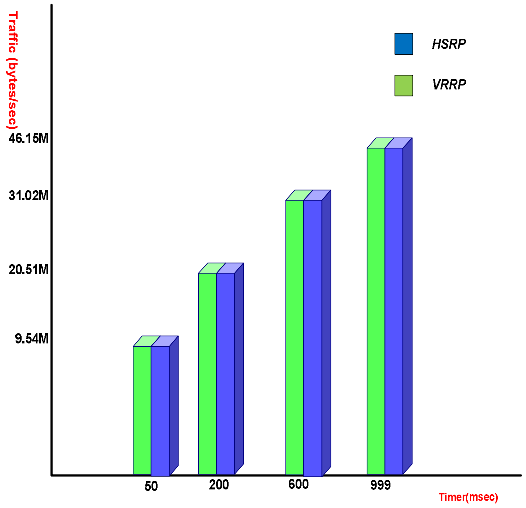

Figure 5: Traffic level when instability occurs at various VRRP/HSRP timer settings

The output from the Cisco IOS command "show interface" for fa0/0 suggests that the instability was caused by input queue packet error and drop as shown below. Bearing in mind that the output of the show command parameters below has been defined in previous section above. Input drops and errors encountered indicating that this queue overflowed. For our purposes though, the exact figure of the errors or drops is not really important but it provides data confirming that errors and drops were encountered on the input queue of the interface.

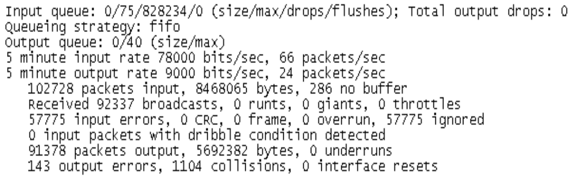

```
Input queue: 0/75/828234/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 78000 bits/sec, 66 packets/sec
5 minute output rate 9000 bits/sec, 24 packets/sec
   102728 packets input, 8468065 bytes, 286 no buffer
   Received 92337 broadcasts, 0 runts, 0 giants, 0 throttles
   57775 input errors, 0 CRC, 0 frame, 0 overrun, 57775 ignored
   0 input packets with dribble condition detected
   91378 packets output, 5692382 bytes, 0 underruns
   143 output errors, 1104 collisions, 0 interface resets
```

Figure 6: Drops on Fa0/0 of the backup/standby router

Given the above results, we decided to design another experiment that will simulate a host on a network that has a bad route which forwards traffic to the master/active router as described in next scenario.

iv) Sub-second timer implementation with a host on a network that has a bad route

The setup shown in figure 7 was used to study sub-second timer implementation impact on network stability when a host on a network has a bad route which will prompt it to forward traffic via a master/active router to another host on different subnet than the subnet where the host network traffic generator is located. A secondary IP address was configured on active/master interface fa0/0(192.168.21.4); a host was connected to the subnet and assigned an IP address of 192.168.21.6. The host was configured to use the secondary IP as its default gateway. A host with a bad route was used to send traffic to host 192.168.21.6. The objective was to force traffic through the output queue of interface fa0/0 of the active/master.

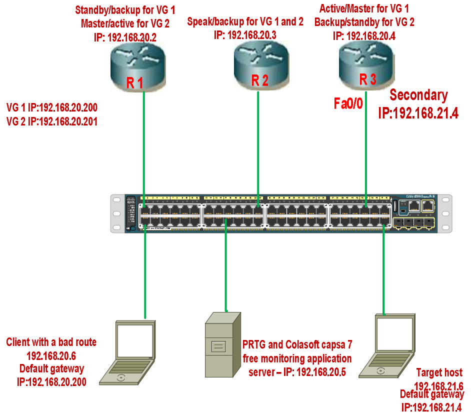

Figure 7: Test bed for induced network traffic test on a host on different subnet.

Result

As described above, the host with a bad route (192.168.20.6) induced traffic targeted to a host (192.168.21.6) in a different subnet. The default gateway on the targeted host was configured to be the secondary IP configured on the interface used for keepalive message sending. Network Instability was observed due to the fact that for forwarded traffic to get to the targeted host, it has to pass through the default gateway which uses output queue of interface fa0/0 for forwarding traffic to the targeted host. Since Fa0/0 output queue is the same output queue used for processing the keepalive message, there is a possibility of a delay or packet lost which will prevent the keepalive message from getting to the back/standby before hold timer expires.

The "show interface" command output verifies that there were errors and drops on the output queue of fa0/0 which caused the instability observed. This is because keepalives must be received within the master_down_interval otherwise the backup/standby routers will transit to become the master.

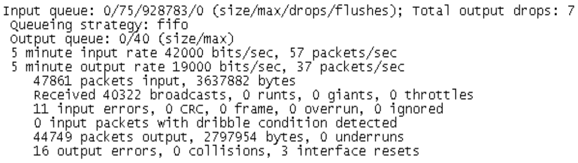

```
 Input queue: 0/75/928783/0 (size/max/drops/flushes); Total output drops: 7
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 42000 bits/sec, 57 packets/sec
  5 minute output rate 19000 bits/sec, 37 packets/sec
     47861 packets input, 3637882 bytes
     Received 40322 broadcasts, 0 runts, 0 giants, 0 throttles
     11 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     44749 packets output, 2797954 bytes, 0 underruns
     16 output errors, 0 collisions, 3 interface resets
```

Figure 8: Drops and errors on fa0/0 of master/active router

4.0 DISCUSSION

The above experimental study indicates that instability occurs when either the output queue on interface used for keepalive message processing or the input queue on the interface used for keepliave reception overflows. Furthermore, router generally uses different queues to store and process incoming and outgoing packets. In the first case where the DoS attack targeted f0/0 the master/active router interface which was used for keepalive message transmission, keepalives were stored before processing in the output queue while DoS traffic was stored before processing in the input queue. Hence keepalive messages did not have much traffic to contend with in the output queue. In the second case where backup/standby router was targeted by the DoS attack , the received keepalives were stored and processed using input queue so the probability of encountering packet delay or loss is much greater due to queue overflow or traffic contention.

In the last experiment, heavy traffic was directed towards a host on another subnet due to a bad route. This meant that traffic had to flow through the output queue of active/master router; the same queue used to store keepalive message on interface Fa0/0 before getting to the targeted host. The contention between the keeaplive messages and the induced traffic on the output queue interface can cause keepalive message delay or packet loss. Reception of keepalive messages is highly time sensitive, so once the packet experiences delay that surpasses the master_down_interval/hold timer, backup/standby router will start transiting to master/active.

5.0 RESEARCH LIMITATION

There are numerous and varied implementations of these protocols by different vendors in production today. Given the diversity of these implementations it is not practical to experiment with every possible implementation. Therefore, the specific conditions under which instability occurred, and even the specific causes identified applied to our environmental setup.

Due to time constraint we could not explore QoS that involves traffic prioritization recommended by Cisco. Although we know that QoS does not remove the packet drops or error encountered, but rather just determines what type of traffic to be dropped.

6.0 CONCLUSION

In this paper, we studied the impact of sub-second timer implementation on network stability. From analysis of data collected, we observed that when the master/active router was subjected to DoS attack directed through the input queue on the interface used for keepalive message transmission, sub-second timer implementation does not lead to instability in the network. This can be seen from the fact that even when CPU utilization on the master/active topped 100%, the irregularities in the intervals between keepalives sent by the router were smaller than the master_down_interval. Hence, the standby/backup routers never began the transition to master/active state. But when heavy traffic that was caused by a host with bad route flowed through the output queue of the master/active router used for keepalive message, network instability was observed.

When a backup/standby router was subjected to a DoS attack which sent traffic through the input queue of the same interface used for keepalive reception, instability was observed even when the CPU utilization was less than 30%. This is because for the backup/standby routers, same interface queue is used for reception of keepalive messages as for other traffic. Making it more likely that keepalives might be lost or delayed longer than the master_down_interval timer. This shows that when implementing subsecond timers instability is more likely to occur whenever there is traffic congestion on output queue of the interface used for keepalive message transmission or input queue of the interface used for keepalive message reception.

FHRP is being implemented on router interface and that interface is what will be used as a default route for host on network so it was design in such a way that it cannot run on a separate interface different from interface used as a default route. Due to the high sensitivity of FHRP keepalive message to delay and packet loss, one way to improve on FHRP design is to separate the keepalive message from other network traffic through allocating a separate interface for keepalive message transmission and reception or design it to be running on router global mode.

ACKNOWLEDGEMENT

I am indebted to my research advisors and reviews committee members who work tirelessly even during holiday's period to ensure that this work comes to an accomplishment. Also grateful to God almighty for giving me the strength, knowledge and power from the beginning of this work to the end.

REFERENCES

[1] Cisco Press, "Cisco Catalyst 6500 Series Virtual Switching System (VSS) 1440," 2010. Available: http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps9336/white_paper_c11_429338.pdf

[2] Akhil Srivastava, "Virtual Router Redundancy Protocol," White Paper, 2009. Available: http://www.techmahindra.com/Documents/WhitePaper/TechM_WhitePaper_VRRP.pdf Accessed on 05/6/11.

[3] [VRRP] R. Hinden, Ed., "Virtual Router Redundancy Protocol," RFC 3768, April 2004.

[4] ImageStream Internet Solutions, "Virtual Router Redundancy Protocol," White Paper, June 9, 2004. Available:http://www.imagestream.com/VRRP_WhitePaper.PDF. Accessed on 05/6/11

[5] J. Ranta, "Router Redundancy and Scalability Using Clustering," Seminar on Internetworking, Spring 2004, eds. A. Ylä-Jääski, N. Kasinskaja, [Online] Available:http://www.tml.hut.fi/Studies/T-110.551/2004/papers/Ranta.pdf,June 2004.

[6] [HSRP] T. Li, B. Cole, P. Morton and D. Li, "Cisco Hot Standby Router Protocol (HSRP)," RFC 2281, March 1998.

[7] Cisco systems, "Hot Standby Router Protocol version 2", Available: http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gthsrpv2.html

[8] Cisco Press, "Campus Network for High Availability Design Guide," 2008. Available:http://www.cisco.com/application/pdf/en/us/guest/netsol/ns431/c649/ccmigration_09186a008093b876.pdf

[9] R. Hott," Timer Enhancements to Reduce Failover Times for the Virtual Router Redundancy Protocol forIPv4," March 6, 2006

[10] Cisco Press, "Applying Cisco Troubleshooting Tools**,**" 2001**.** Available**:** http://www.ciscopress.com/articles/article.asp?p=25296&seqNum=3

[11] Jen-Hao Kuo, Siong-Ui Te, Pang-Ting Liao, Chun-Ying Huang, Pan-Lung Tsai, Chin-Laung Lei, Sy-Yen Kuo,
Yennun Huang, and Zsehong Tsai P. Padala, "An Evaluation of the Virtual Router Redundancy Protocol Extension withLoad Balancing" Dec. 2005

[12] S. Nadas,"Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6,"RFC 5798, March 2010

[13] Cisco Press, "Understanding and Troubleshooting HSRP Problems in Catalyst Switch Networks," May 05, 2009. Available: http://www.cisco.com/image/gif/paws/10583/62.pdf

[14] Nsasoft traffic emulator, available: http://www.nsauditor.com/network_tools/network_traffic_generator.html

[15] PRTG,Paessler Corporation available: http://www.paessler.com/manuals/prtg8/toplists.htm

[16] Solarwinds Inc, available: http://www.solarwinds.com/products/toolsets/engineer.aspx

[17] Colasoft capsa 7, available: http://www.colasoft.com/capsa/