

# TCP/IP Jumbo Frames Network Performance Evaluation on A Test-bed Infrastructure

Shaneel Narayan, Paula Raymond Lutui

*Department of Computing, Unitec Institute of Technology, Auckland, New Zealand*

---

## Abstract

Enhancing network performance has been studied by a number of researchers. The need to provide greater throughput on network infrastructure has been the key driver for these studies. The use of jumbo frames is considered one of the methodologies that can be employed to increase data throughput on networks. In this research undertaking, the authors implement jumbo frames on a test-bed implemented with Windows Server 2003/2008 networks and performance related metrics are measured for both IPv4 and IPv6 implementations. The results obtained in this empirical study shows that performance metrics values are different in various scenarios.

**Index Terms:** Jumbo frames, IPv4, IPv6, network performance evaluation, Windows Server 2008, Windows Server 2003.

© 2012 Published by MECS Publisher. Selection and/or peer review under responsibility of the Research Association of Modern Education and Computer Science

---

## 1. Introduction

The Internet continues to play a dominant part in global communication and its usage is always on an exponential incline. Social networking sites and the use of multimedia applications warrant high data transmission speed; to this effect software and hardware vendors continually develop products with increased communication capabilities. Not only that, the Internet Protocol itself is undergoing major architectural enhancements. Issues that were prevalent in IPv4 have been addressed by IPv6 however the maximum transmission unit (MTU) still remains unchanged at 1500Bytes. This is seen as one of the characteristics that limit throughput in IP based network infrastructures. The use of jumbo frames, which can offer maximum frame size of 9000Bytes, is a possible solution to this problem. Therefore there is a need to study the impact of using jumbo frames on network implementations. In [1], we tested performance of jumbo frames on Windows Server 2008 network, now we extend that work and compare jumbo frame performance by including Windows Server 2003 implementation.

Studies related to jumbo frames have also been undertaken by number of other researchers. In [2], Gigabit over Ethernet (GoE) is discussed and shown that jumbo frames are a viable candidate for event builder.

\* Corresponding author.

E-mail address: [snarayan@unitec.ac.nz](mailto:snarayan@unitec.ac.nz)

Experiences with offloading protocol to a programmable Gigabit Ethernet are presented in [3]. Here it is shown that network throughput almost doubles when jumbo frames are implemented. Jumbo frames and its impact on TCP packet processing is the subject of research in [4] and it is shown that jumbo frames can form a substantial component in data transfer. Performance under maximum-segment sizes is discussed in [5] and shown that segment size does affect performance. The impact of using jumbo frames in field-buses and multimedia applications to enhance performance is presented in [6]. Similar theme, where impact of using jumbo frames with multimedia applications, is presented in detail in [7]. Performance differences when jumbo frames are implemented on wireless network have been evaluated on test-bed in [8] and shown that performance enhancements can be achieved with larger frames. With optical networks, jumbo frames performance is analyzed in [8].

The rest of this paper is organized as follows. In II we outline the test-bed used in this research, and in III results are presented and discussed. This is followed by conclusions.

## 2. Experimental Setup

Computers with similar hardware (CPU: Intel Pentium C2D, RAM: 2GB, NIC: PCI Intel Pro 1000, Hard Drive: Seagate 160GB) were connected using cross-over cables and Microsoft Windows Server 2008 was installed on test computers on this test-bed (Figure 1). IPv4 as the communication protocol was configured first and data was collected. Later this was replaced with IPv6 ensuring that all other test-bed parameters remained the same. D-ITG 2.6.1d was the primary tool employed to evaluate performance of protocols on operating systems. D-ITG generates traffic at network/transport and application layer and sends it from sender to generator node and can measure performance related metrics. On completing Windows Server 2008 runs, Windows Server 2003 was installed and the experiment was rerun keeping all other variables same. To ensure high data accuracy, all tests were executed 20 times, and to get the maximum throughput for a given packet size, each run had duration of 30 seconds.

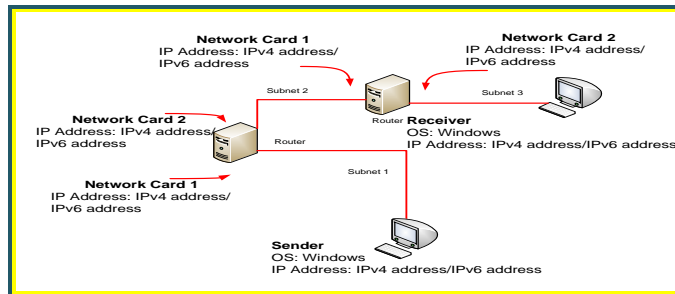


Fig1: Experimental Setup

## 3. Results and Discussion

We now present and discuss the findings of this research. In Figure 2, throughput values for TCP traffic type is presented when the network was implemented with jumbo frames. Here it is seen that there is significant difference between Windows Server 2003 IPv4 values and the other three scenarios. This difference of approximately 10% is constant for most packet sizes. The other three scenarios register almost similar values for all packet sizes. UDP throughput values (Figure 3) also shows that Windows Server 2003 IPv4 is a comparatively low performer than the other three scenarios. In this case, the difference is almost 25%. Like TCP, UDP values for the other three scenarios are almost constant for most packet sizes.

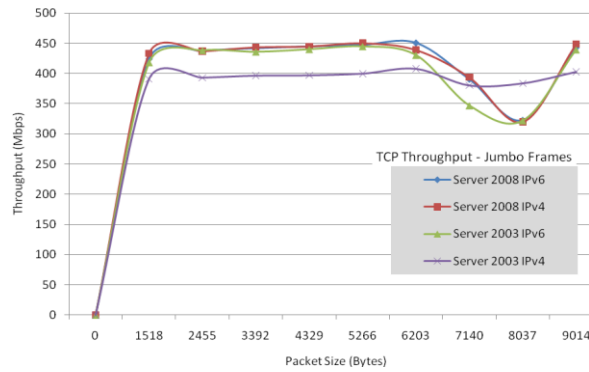


Fig 2: TCP Throughput – Jumbo Frames




Figure 3: UDP Throughput – Jumbo Frames

In Figures 4 and 5, throughput values for normal packet sizes (up to 1518Bytes) are presented for TCP and UDP respectively. Here it is seen that these throughput values are all lower than that obtained when the same network is implemented with jumbo frames. In case of TCP, average performance difference is almost 15% and for UDP the difference is approximately 25%. Also, it is evident that with normal frames, Windows Server 2003 IPv4 consistently gives low throughput values for TCP traffic types, however for UDP normal frames; Windows Server 2003 IPv6 gives the highest values for most packet sizes.

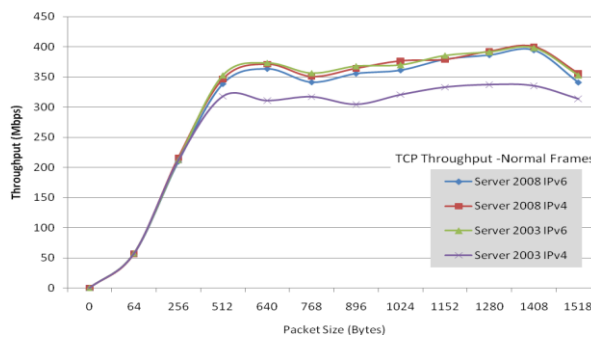


Fig 4: TCP Throughput – Normal Frames




Fig 5: UDP Throughput – Normal Frames

Delay values are presented next in Figures 6 and 7. For TCP traffic type, Windows Server 2008 IPv4 registers that highest delay of all scenarios (difference almost 80%). IPv6 delays are significantly lower that IPv4 for TCP, however for UDP this is not the case. Here there is a clear distinction between Server 2003 and 2008 with latter giving UDP delay values significantly higher than Server 2003. Also, all TCP delay values are less than 0.6ms while that of UDP is up to 14ms. In Figures 8 and 9, normal packet delay values are presented and it is seen that Windows Server 2008 values are mostly higher than that of the counterpart operating system for most packet sizes.

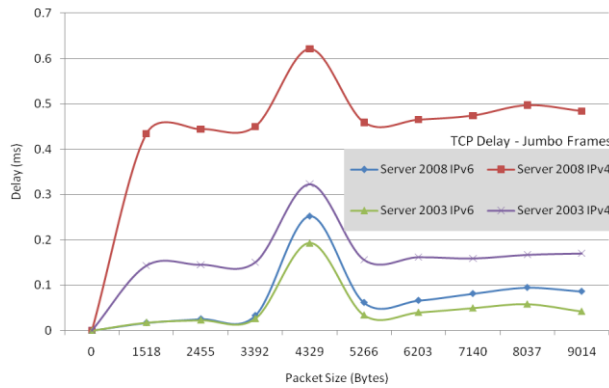


Fig 6: TCP Delay – Jumbo Frames

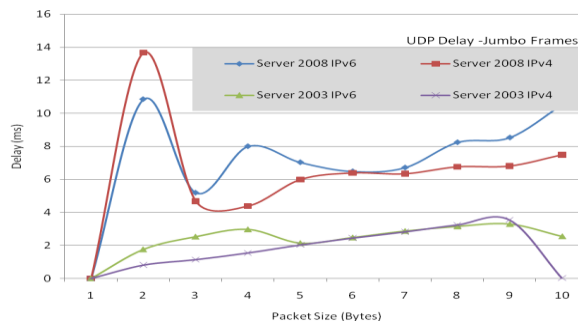


Fig 7: UDP Delay – Jumbo Frames




Fig 8: TCP Delay – Normal Packets




Fig 9: UDP Delay – Normal Packets

Jumbo frames jitter values are presented in Figures 10 and 11. For TCP traffic type, the difference between the scenarios is insignificant, however a few packet sizes give jitter values significantly higher than the rest. Like TCP, UDP also registers low values (highest value 0.014ms) but the trend is different to that of TCP. When these values are compared to normal frames (Figure 12 and 13), it is seen that most values are close to zero, thus there does not appear to be any major difference between jumbo and normal frame jitter values. Also, for normal frames, Server 2008 jitter values are slightly higher than that of Server 2003.

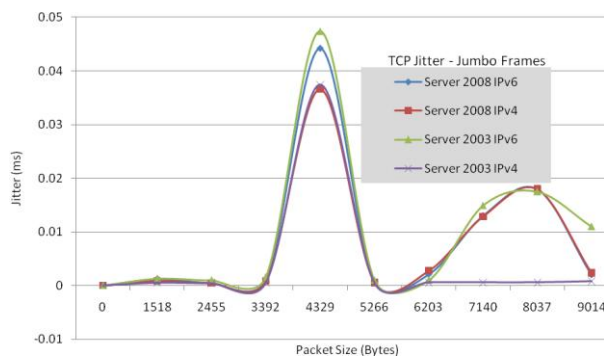


Fig 10: TCP Jitter – Jumbo Frames




Fig 11: UDP Jitter – Jumbo Frames




Fig 12: TCP Jitter – Normal Frames




Fig 13: UDP Jitter – Normal Frames

For UDP traffic type, packet dropped during data transmission was measured. These values are presented in Figures 14 and 15. Jumbo frames give highest drop rate (60%) for Windows Server 2003 IPv4 while the other three scenarios give comparable values averaging approximately 40%. Between IP versions, the difference is insignificant. When compared to normal frames (Figure 15), most values are comparable and no distinctive trends are evident. Average of almost 40% is seen and all values range between zero and 60%.




Fig 14: UDP Packet Dropped – Jumbo Frames




Fig 15: UDP Packet Dropped – Normal Frames

Finally, TCP CPU usage values are presented in Figures 16 and 17. In both the graphs, difference between the scenarios is minimal; however jumbo frames use much less CPU resources on the sending router than normal frames.




Fig 16: TCP CPU Usage Router 1 – Jumbo Frames




Fig 17: TCP CPU Usage – Normal Frames

#### 4. Conclusion

In this research an unbiased empirical analysis of performance of jumbo frames on TCP/IP networks with IPv4 and IPv6 were conducted on Windows Server 2003/2008 operating system. TCP and UDP traffic were generated and performance related metrics were measured. From this research, we can draw the following specific conclusions:

1. Jumbo frames give higher throughput than normal frames for both TCP and UDP traffic types.
2. IPv6 delay values are significantly lower than IPv4 for TCP traffic type.
3. Jitter values and UDP packet dropped values are comparable between jumbo and normal frames.

The authors anticipate extending this work to include more operating systems, especially Linux distributions, and different network scenarios.

#### References

- [1] S. Narayan and P.R. Lutui, "Test-bed Network Performance Evaluation of Jumbo-frames on IPv4/IPv6 TCP/IP Network Infrastructure" Proceedings of the IEEE 4<sup>th</sup> International Conference on Internet Multimedia Systems Architecture and Application (IMSAA), 2010.
- [2] Y. Yasu, Y. Nagasaka and A. Manabe, "Evaluation of Gigabit Ethernet with quality of service for event builder", Proceedings of the 11th IEEE Real Time Conference NPSS, Santa Fe, pp.433-436, 1999.
- [3] A. B. Maccabe, W. Zhu, J. Otto and R. Riesen, "Experience in offloading protocol processing to a programmable NIC," Proceedings of the IEEE International Conference on Cluster Computing", pp. 67- 74, 2002.
- [4] S. Makineni and R. Iyer, "Performance characterization of TCP/IP packet processing in commercial server workloads," Proceedings of the IEEE International Workshop on Workload Characterization WWC-6", pp. 33- 41, 27 Oct. 2003.
- [5] M. Katevenis and G. Passas, "Variable-size multipacket segments in buffered crossbar (CICQ) architectures," Proceedings of the IEEE International Conference on Communications", vol.2, pp. 999-1004, 16-20 May 2005.
- [6] J. Silvestre, V. Sempere and T. Albero, "Impact of the use of large frame sizes in fieldbuses for multimedia applications," Proceedings of the 10th IEEE Conference on Emerging Technologies and Factory Automation ETFA", vol.1, pp. 440-448, 19-22 Sept. 2005.
- [7] Y. Hsiao, M. Chen, K. Huang, Y. Chu and C. Yeh, "High speed UDP/IP ASIC design," Proceedings of the International Symposium on Intelligent Signal Processing and Communication Systems, ISPACS", pp.405-408, 7-9 Jan. 2009.
- [8] A.P. Iyer, G. Deshpande, E. Rozner, A. Bhartia and Q. Lili, "Fast Resilient Jumbo frames in wireless LANs," Proceedings of the 17th International Workshop on Quality of Service IWQoS", pp.1-9, 13-15 July 2009.
- [9] W.Yixin Wu; S. Kumar and P. Seung-Jong, "On Transport Protocol Performance Measurement over 10Gbps High Speed Optical Networks," Proceedings of 18<sup>th</sup> International Conference on Computer Communications and Networks, ICCCN", pp.1-6, 3-6, Aug. 2009