Kemp Technologies LM-3600 IPv4 and IPv6 Performance Report

A Broadband-Testing Report
By Steve Broadhead, Founder & Director, BB-T

First published April 2012 (V1.0)

Published by Broadband-Testing
A division of Connexio-Informatica 2007, Andorra

Tel : +376 633010
E-mail : info@broadband-testing.co.uk
Internet : HTTP://www.broadband-testing.co.uk

TABLE OF CONTENTS

BROADBAND-TESTING

Broadband-Testing is Europe’s foremost independent network testing facility and consultancy organisation for broadband and network infrastructure products.

Based in Andorra, Broadband-Testing provides extensive test demo facilities. From this base, Broadband-Testing provides a range of specialist IT, networking and development services to vendors and end-user organisations throughout Europe, SEAP and the United States.

Broadband-Testing is an associate of the following:

Limbo Creatives (bespoke software development)

Broadband-Testing Laboratories are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

Broadband-Testing Laboratories operates an **Approval** scheme which enables products to be short-listed for purchase by end-users, based on their successful approval.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at HTTP://www.broadband-testing.co.uk

Broadband-Testing Consultancy Services offers a range of network consultancy services including network design, strategy planning, Internet connectivity and product development assistance.

EXECUTIVE SUMMARY

- Load-Balancing and network optimisation were originally designed - and priced - for high-end enterprise and service provider usage but, increasingly, the Small-Medium Business (SMB) market has the same needs for optimised application delivery and reliability.

- With its LoadMaster range of Load-Balancers/Application Delivery Controllers, Kemp Technologies is looking to offer that same level of "enterprise" features to the SMB and hosting company communities at an affordable price. Not only is this a good deal for the SMBs but it also brings the optimisation technology more in line with current server pricing, so it makes sense all-round.

- The impact of IPv6 may not be immediate in many parts of the world currently, but the need to be IPv6 "ready" is clearly a requirement. There is also a concern in many quarters that moving to IPv6 will result in potentially significant performance loss, especially when working at Layers 4-7. From its perspective, Kemp already supports IPv6 at Layer 7 on its appliances, so we put it to the test, using a test bed based on Spirent's Testcenter traffic generator, running a series of http-based tests, first with IPv4, then with IPv6.

- Taking Kemp's "mid-range" LM-3600 appliance we were able to show its ability to run at Gigabit line speed at both Layer 4 and Layer 7 with both IPv4 and IPv6.
- Using a multi-port test setup, we were able to validate Kemp's claims (marketing figures) of supporting 2.9Gbps throughput at Layer 7 with both IPv4 and IPv6.
- Throughout the testing we were able to show that moving to IPv6 had no impact on Layer 7 performance whatsoever and that CPU utilisation on the LM-3600 was never saturated, peaking at around 65% utilisation. This has excellent implications for both the scalability of the appliance, and its ability to handle complex Layer 7 operations at maximum performance levels.

INTRODUCTION: OPTIMISING SMB NETWORKS, BIG BUTS AND THE ONSET OF IPV6...

There's nothing new about Load-Balancing or Application Delivery Control (L-B/ADC).

The requirement - and solutions - have been around since the start of the .com movement and are every bit as relevant today as they were then. However, historically, the benefits of load-balancing and application acceleration, while clearly advantageous to any user who accesses data and applications stored in a data centre or server farm (wherever that may be) – and who doesn't – have been restricted to Enterprise and above. L-Bs were expensive - many still are - and designed to optimise expensive arrays of servers and their data. But - and this is a 'big but' - server prices, on a bang per buck basis, have fallen enormously over the past decade, so you can now get big power for small money.

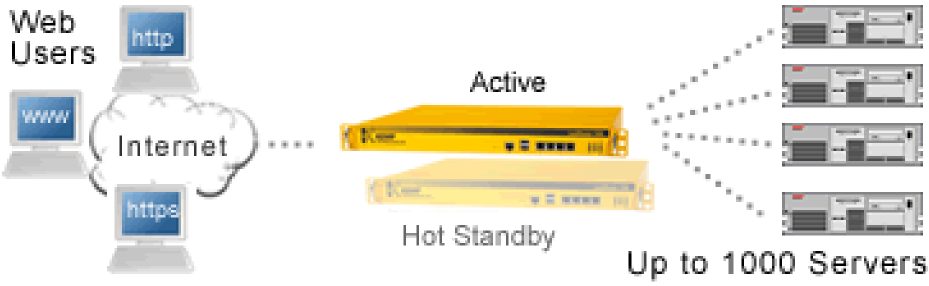

Figure 1 – A Modern Load-Balancer Deployment

What this means is that smaller businesses are able to buy into serious server horsepower and, remember, small-medium businesses (SMBs) account for around 93% of all IT business users. Moreover, their IT requirements are every bit as critical as those of a multi-national, to them at least. It therefore makes sense that SMBs, too, want to optimise their networks or, ensure that their hosting company is optimising their servers at their remote location, cloudy or otherwise. But - another 'big but' - is it realistic to pay tens of thousands of dollars for optimising technology when the servers themselves are costing a fraction of that?

Of course not. Which is where Kemp Technology comes into play. The company has created a range of L-B/ADC appliances which are designed to be cost-effective for the SMB market and lower-tier service providers/hosters. The idea is to play the 90:90 rule; offer the customer 90% of the features of the top-end competitors - ones that they use 90% of the time - for a fraction of the price of buying into the '100:100' products. It just makes good old-fashioned sense - give the customer what they want at a price they can afford - server and network optimisation for the masses.

So does this mean short-changing the users in terms of functionality and performance? Not at all. In its most basic form, an L-B/ADC directs application users to the best available server. It also manages online/offline server behaviour, re-routing traffic as required, while distributing user/application requests in as optimal a fashion as possible. This is the case whether the L-B/ADC is £3,000 or £30,000.

However, with its LoadMaster range of appliances, Kemp is looking to go beyond this, offering additional services such as SSL acceleration, offloading encryption/decryption processes from the application servers, and thereby avoiding the huge performance loss on servers associated with - and proved by Broadband-Testing in the past - SSL traffic.

Now there's an additional 'fly in the ointment' to contend with - the onset of IPv6. With the last of the public IPv4 address allocation now long gone and the Far East already deploying IPv6 big time, the reality is that we do all need to start thinking about moving from the "4" to the "6", albeit gradually in most cases. And with LTE around the corner in the mobile world, that being pure IP-based, how many new IP addresses will suddenly be demanded? And where are they going to get allocated from? Another worry - and one that is totally relevant to this report - is the potential performance hit when moving from IPv4 to IPv6.

This is something that vendors don't openly talk about - funny, that! - but we understand that up to 50% performance loss has been noted on internal testing by some vendors, when moving from IPv4 to IPv6 with a sub-10% loss generally seen as more than acceptable. Kemp Technologies - whose product is the focus of this report - is fully aware of the potential IPv6 performance problem and claims to have resolved any issues in this respect.

While it's realistic to say that, in many parts of the world at least, IPv6 take-up will be gradual, not being 'IPv6 ready' is no excuse. Axel Pawlik, MD of RIPE NCC, an independent, not-for-profit membership organisation that supports the infrastructure of the Internet in Europe, the Middle East and parts of Central Asia with a heavy focus on IPv6 deployment, sees IPv6 catching many businesses out, if they don't prepare for the migration. He believes that unless businesses start to act now to safeguard their networks, the future expansion of the Internet could be compromised. Designed to account for the future growth of the Internet, the pool of IPv6 addresses contains 340 trillion, trillion, trillion unique addresses. This huge number of addresses is expected to accommodate the predicted growth and innovation of the Internet and Internet-related services over the coming years.

In other words, it's coming and you'd better be ready for it, and ready to optimise it. Fundamentals such as NAT and DNS take on a new life - and meaning - in 64-bit form with IPv6. Everything is bigger and more complex, which usually means "performance hit".

In the test section of this report we compared IPv4 with IPv6 performance on the featured device - the Kemp LoadMaster 3600 - to see if Kemp is truly "IPv6 ready". But first, let us look at exactly what Kemp has to offer, products wise.

KEMP TECHNOLOGIES: LOADMASTER PRODUCT OVERVIEW

Kemp provides a range of both physical and virtual LoadMaster L-B/ADC products.

The appliance range runs from the LM-2200 offering 4xGigabit Ethernet ports, around a gigabit of throughput and 200 SSL transactions per second (claimed) through to the LM-5500 with 18xGigabit Ethernet (and optional 2x10GbE) ports, a claimed 6Gbps maximum throughput and up to 10,000 SSL transactions per second. In the middle of the range is the LM-3600, the focus of our testing here, supporting a claimed 3.4Gbps of L4 and 2.9Gbps of L7 throughput (from 8xGigabit Ethernet ports) and 5,000 SSL transactions per second.

Regardless of the appliance model, all are designed to offer the following features and functionality:

- Layers 4-7 Load Balancing (with several L-B methodology options)
- Content Switching
- Server Persistence
- SSL Offload/Acceleration
- Windows Terminal Services load balancing and persistence with Session Directory integration
- Application Front-end (Caching, Compression and IPS security)
- Industry leading price/performance value

Figure 2 – Kemp LM-3600

The LM-3600 under test here is described by Kemp as an advanced application delivery controller with Layer 7 content switching and integrated ASIC-based SSL acceleration. The hardware itself is based around an Intel Quad Core processor with 4GB of memory. It is designed to intelligently and efficiently distribute user traffic and offload and accelerate Layer 7 applications such as SSL security and content, to optimize web and application servers, ensuring users get the best experience possible. By this, Kemp defines the LM-3600 as:

- Providing 99.999% high-availability of application servers and removes SLB as single point of failure.

- Guaranteeing user requests will be directed to only "available" servers AND "available" applications.
- Ensuring that users maintain continuous connections with the specific server where their transactional data is available – even if the IP address changes during session.
- Enabling site administrators to optimise server traffic according to content type (images, multi-media, apps).
- Optimising server performance and user experience for encrypted application content.
- Reducing latency associated with internal network while further optimising performance over existing ISP links.
- Helping thwart application-level threats, even with SSL-encrypted traffic.

Management is via a fully configurable web-based user interface, while SSH and HTTPS remote access is also available. The configuration is wizard-based and includes an online help assistant. The management interface includes real-time performance and availability displays.

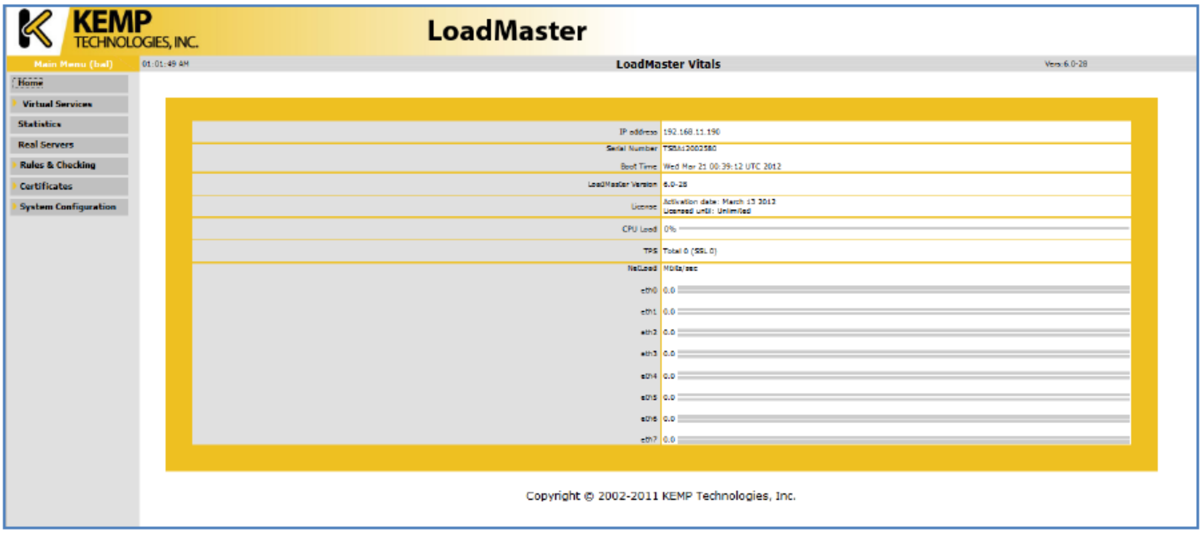

Figure 3 – Kemp LM-3600 Management Interface

Kemp also offers a range of Virtual LoadMasters which install and runs as a hardened, 'Guest' operating OS/Application on a dedicated virtual machine.

They provide the same feature set as the LoadMaster appliance including L4 load balancing, L7 content switching, SSL Offload, Server and Application Health Checking, IP and L7 Persistence, Caching, Compression, IPS etc. They also support stateful Active/Hot-standby configuration between two VLMs for redundancy and high-availability and are controlled by the same web-based management interfaces as the hardware-appliance versions of the LoadMaster.

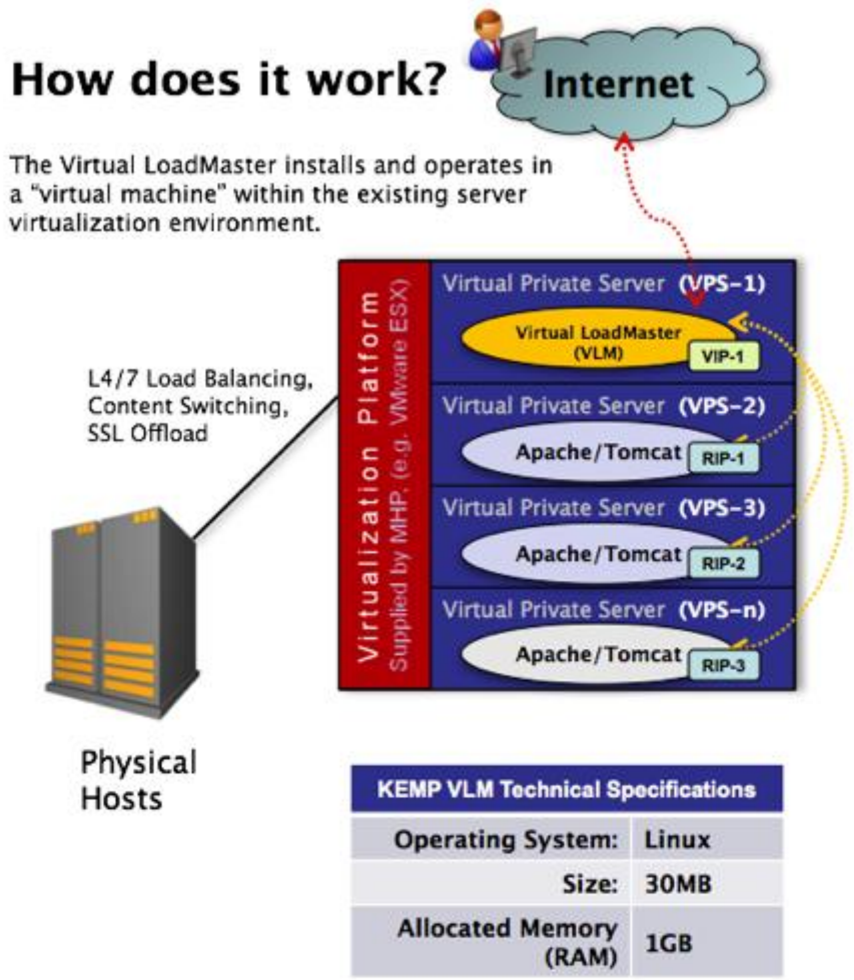

Figure 4 – How The Virtual LoadMaster Works

LOADMASTER 3600: PUT TO THE TEST

The Test Bed

To put the LM-3600 to the test, we created a test bed using Spirent's Testcenter, to generate web traffic with both IPv4 and IPv6.

Figure 5 – Spirent Testcenter

This was loaded with a Spirent CPU-5002 module. This has 8 cores and 8 gigabit ports. Internal architecture wise, one core is assigned to each port in order to maximise performance. We also used an HP 6600 Ethernet switch in order to link our performance and management networks.

Using the eight Spirent ports with the LM-3600 we created a 4-in, 4-out configuration, with clients on one side and servers on the other (both provided by the Testcenter). For the record, IP address ranges used were as follows:

IPv4 - Clients

- 192.168.11.2-180
- 192.168.12.2-180
- 192.168.13.2-180
- 192.168.14.2-180

IPv4 - Servers

- 10.12.0.50-51
- 10.13.0.50-51
- 10.14.0.50-51
- 10.15.0.50-51

IPv6 - Clients

FC8E:DECE:D8EC:DEAD::AA-FC8E:DECE:D8EC:DEAD::FF
FC8E:DECE:D8EA:DEAD::AA-FC8E:DECE:D8EA:DEAD::FF
FC8E:DECE:D8EB:DEAD::AA-FC8E:DECE:D8EB:DEAD::FF
FC8E:DECE:D8ED:DEAD::AA-FC8E:DECE:D8ED:DEAD::FF

IPv6 Servers

FC8E:DECE:D8EC:BEEF::3-FC8E:DECE:D8EC:BEEF::4
FC8E:DECE:D8EA:BEEF::3-FC8E:DECE:D8EA:BEEF::4
FC8E:DECE:D8EB:BEEF::3-FC8E:DECE:D8EB:BEEF::4
FC8E:DECE:D8ED:BEEF::3-FC8E:DECE:D8ED:BEEF::4

On the LM-3600 we created a basic configuration with simple L4 and L7 rules to use during the testing. Throughout the testing we did not adjust the configuration, so this was a real-world test scenario, in line with how an IT department would deploy the LoadMaster, rather than reconfiguring optimally for each test as is sometimes the case; hardly real-world in that scenario...

Since the focus was on throughput, we created http traffic using a 100KB transaction file. This would naturally limit the number of transactions per second achievable compared with using a much smaller file size, but did enable us to reach maximum potential throughput as quickly as possible, while still using a realistic file size. In all cases we were looking to push the device under test until we saw packet loss occur, so results are based on zero or minimal packet loss.

The Throughput Tests

IPv4 Layer 4: Gigabit Connection

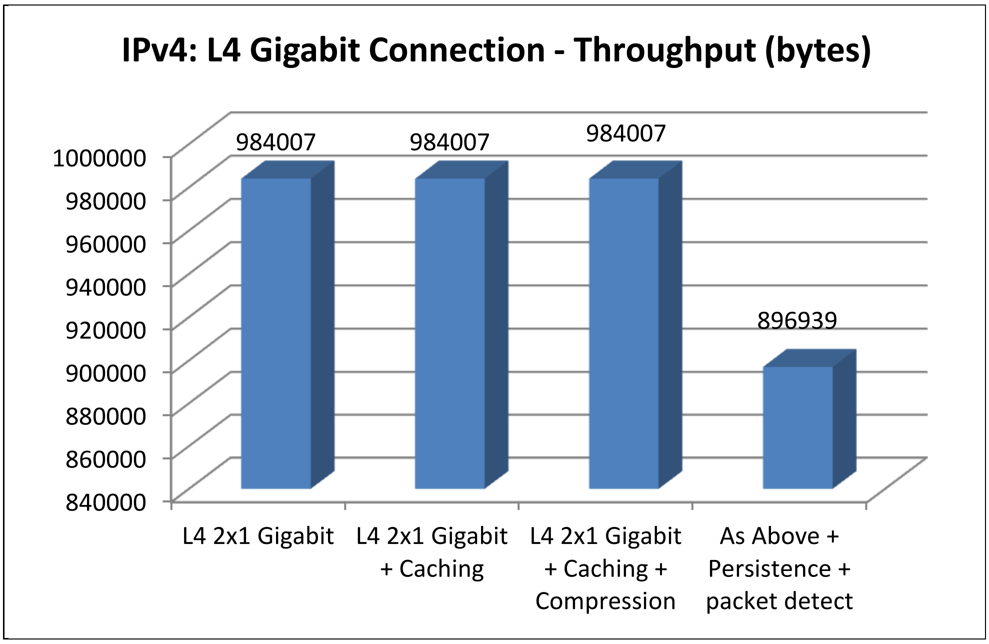

Figure 6 – IPv4 L4 Gigabit Test

Our first test was with a gigabit connection to clients and servers, using IPv4 at Layer 4, to see how close to gigabit line rate we could achieve. We started with a basic Layer 4 test with a simple configuration on the LM-3600, then repeated the test adding functionality each iteration - caching, then compression, then persistence and rogue packet detection - the latter creating significant CPU load on the appliance.

We found that we could achieve what is effectively line rate on the first three iterations, seeing CPU utilisation range from 6% to 18% during those three tests. On the fourth test, we also saw CPU utilisation "max" at around 18% but could only achieve around 900Mbps with this configuration - still an excellent result.

IPv4 Layer 7: Gigabit Connection

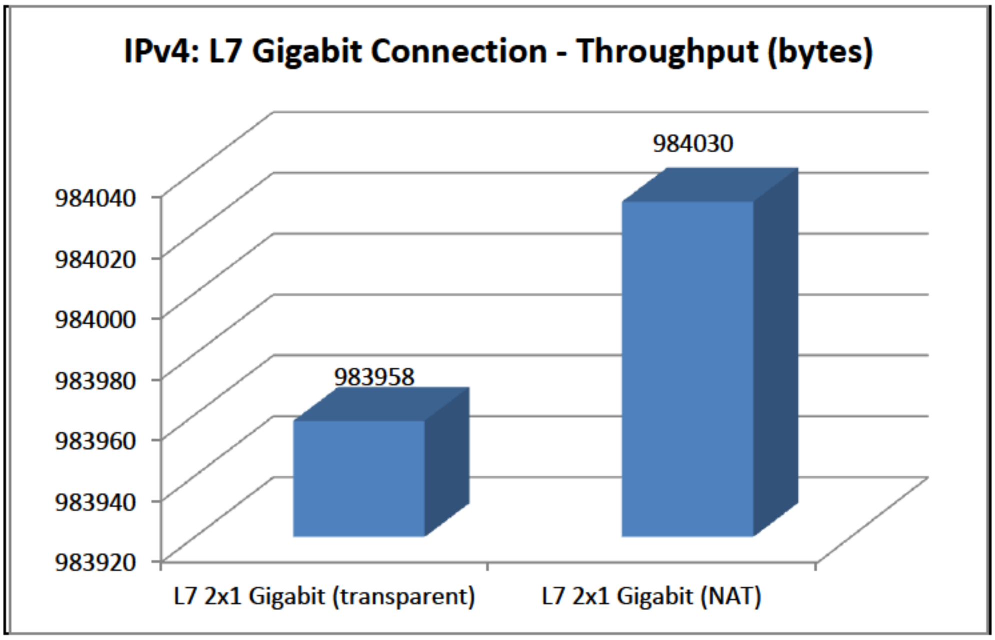

Figure 7 – IPv7 L4 Gigabit Test

We repeated the gigabit testing at Layer 7 with IPv4; firstly with transparent (native) mode and secondly with NAT enabled (non-transparent). We saw almost exactly the same performance in each case - around 980Mbps - again, effectively line rate allowing for the vagaries of the Ethernet standard. CPU utilisation peaked at 19%, so clearly lots of headspace left to use on the LM-3600.

IPv4 Layer 4: 4 x Gigabit (bidirectional) Connections

For the multi-port test (2 x 4 gigabit Ethernet connection in each direction (client/server) we again compared - with IPv4 at Layer 4 - performance using a basic configuration against a more advance configuration with caching and compression enabled. Again we found marginal difference in performance, in this case exceeding 2.6Gbps with CPU utilisation increasing from 31% (test 1) to 56% (test 2), so still plenty of headroom left.

We believe, given more time and more iterations, we could have increased the throughput closer to 3Gbps, making more use of the spare CPU (see next test).

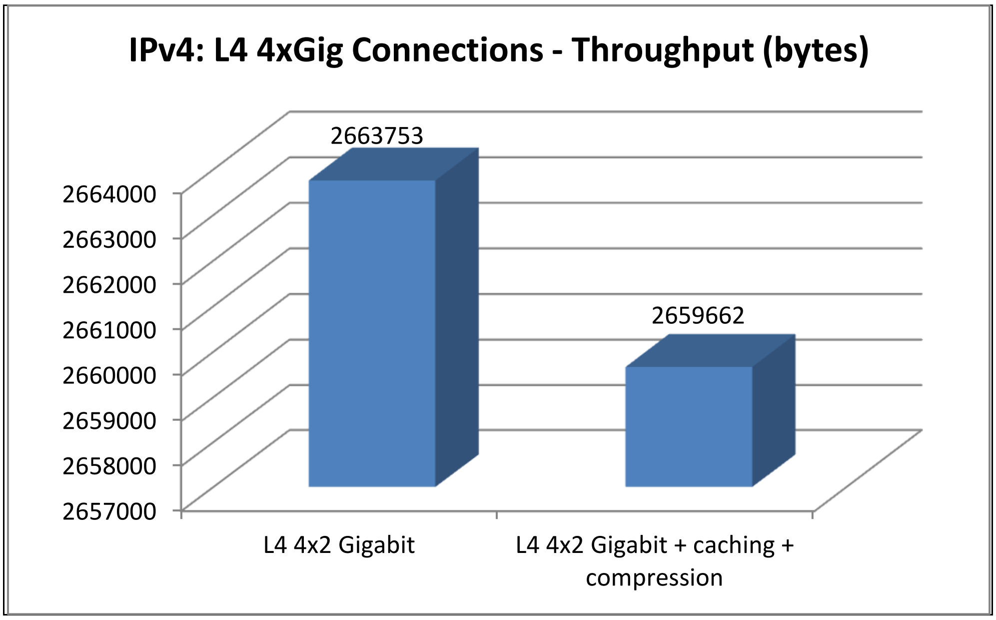

Figure 8 – IPv4 L4: 4xGigabit (bidirectional) Test

IPv4 Layer 7: 4 x Gigabit (bidirectional) Connections

We repeated the multi-port test with IPv4 at Layer 7, pushing until we did start to see some packet loss in this case, but at this point we were almost at 3Gbps in transparent mode, thereby exceeding Kemp's own published figures for Layer 7 performance with the LM-3600.

This fell slightly to around 2.7Gbps with NAT enabled, but still an excellent result. CPU utilisation "max'd" at around 65% during this series of tests - so still lots of headroom.

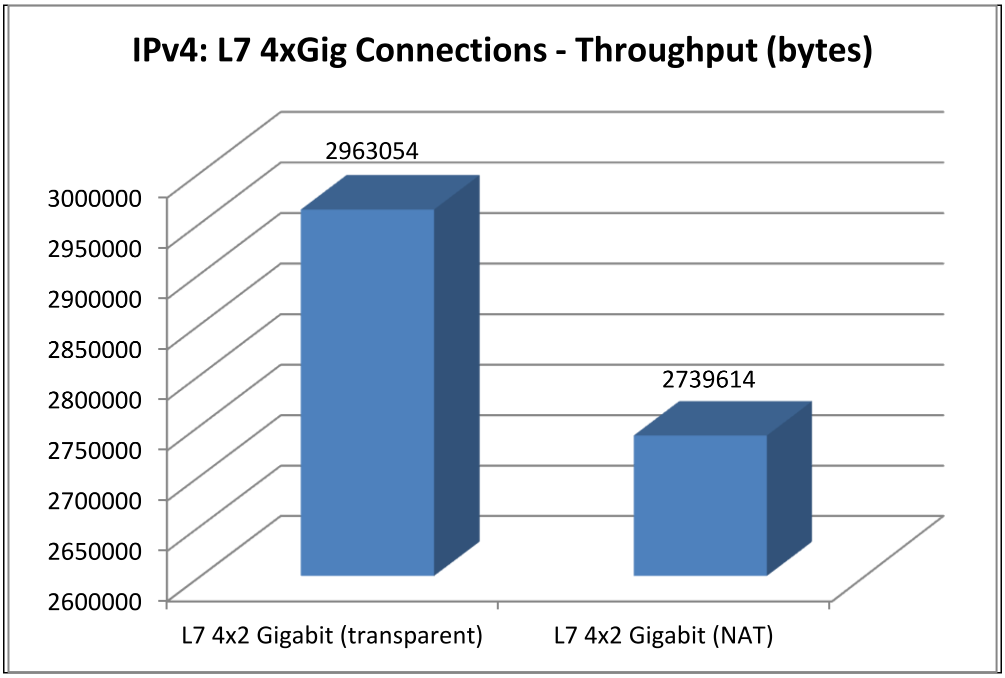

Figure 9 – IPv4 L7: 4xGigabit (bidirectional) Test

IPv6 Layer 7: Gigabit Connection

Moving to IPv6, we focused on Layer 7 (Layer 4 is currently not supported with IPv6 by Kemp) and reran and single gigabit and multi-port gigabit tests.

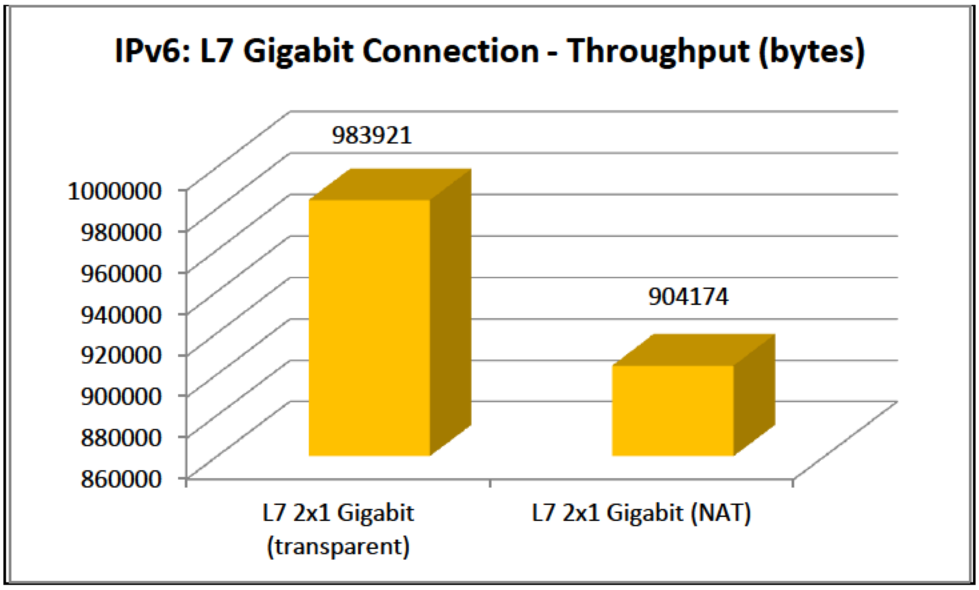

Figure 10 – IPv6 L7: Gigabit Test

Moving to IPv6 we note that we achieve all but line rate at Layer 7 in transparent mode and around 900Mbps with NAT enabled - an excellent result again. CPU utilisation peaked at 19%. In multi-port test mode we also see excellent results, once again exceeding Kemp's published performance figure for L7 throughput for IPv4 - but with IPv6 in this instance! Again, with NAT enabled, performance fell marginally to around 2.85Gbps - still superb considering we are working with IPv6 here. CPU utilisation peaked at 67%, again leaving plenty of overhead for other operations.

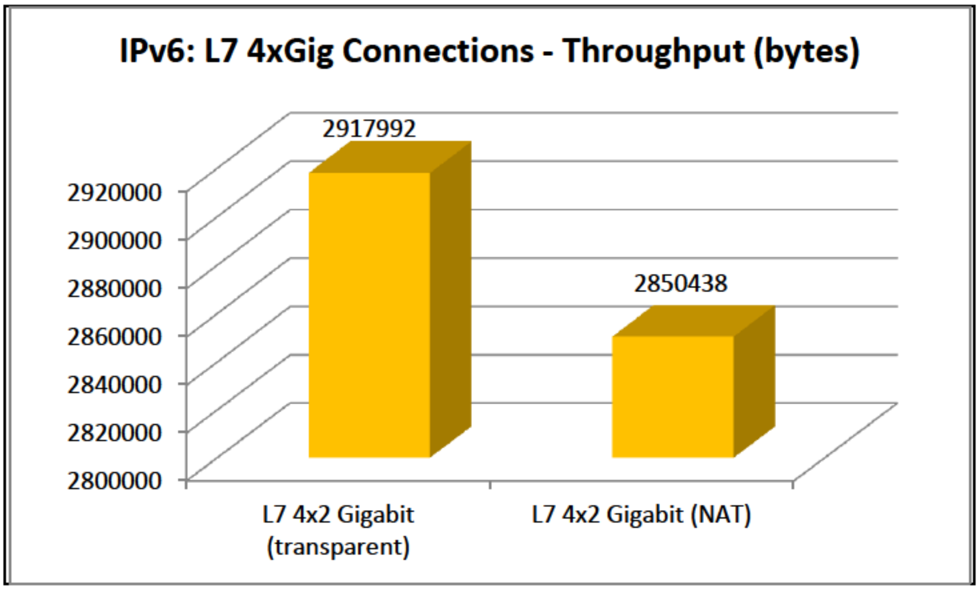

Figure 11 – IPv6 L7: 4xGigabit (bidirectional) Test

Comparing IPv4 Performance With IPv6

If we now look at Layer 7 performance with IPv4 and IPv6 side by side, we see that there is effectively zero performance degradation when moving from the former to the latter.

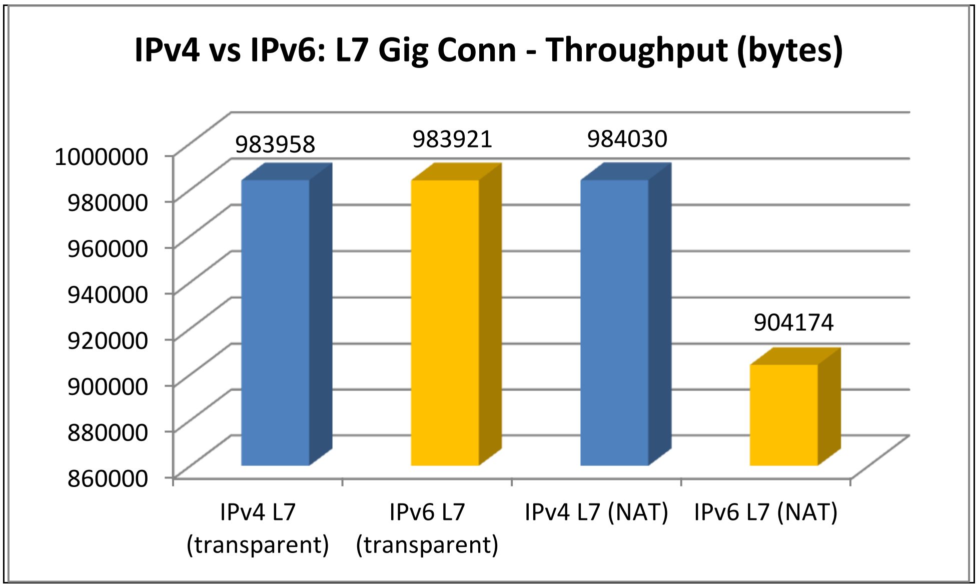

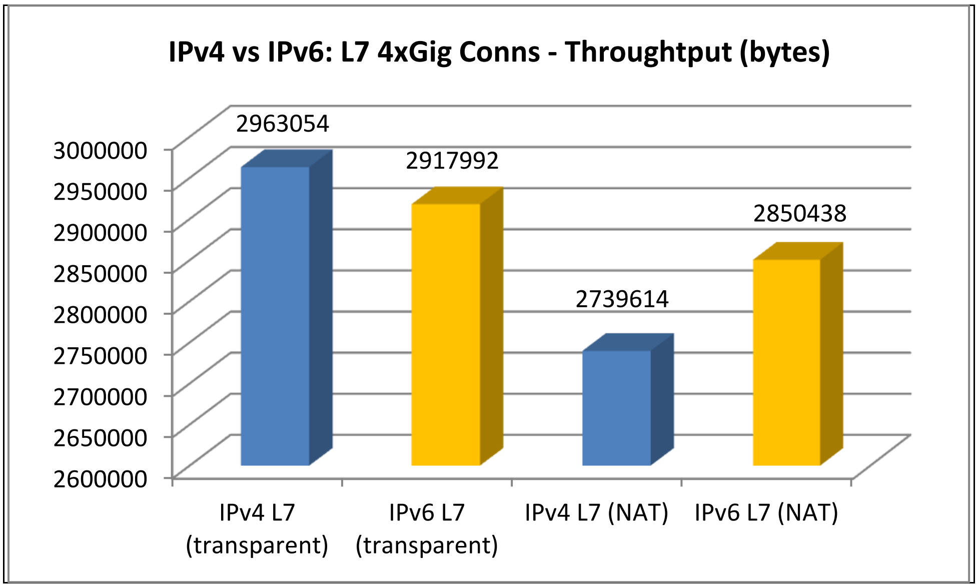

Figure 12 – Comparing IPv4 With IPv6 Performance

In some instances we see that IPv6 performance was actually better, test-on-test than with IPv4, with only the gigabit test with NAT enabled showing any remotely significant relative performance loss - around 80Mbps - less than 10% in other words.

SUMMARY & CONCLUSIONS

What Kemp is trying to do here is provide a very high level of network optimisation at an affordable (to pretty well any company) price.

Not only can we see from this test that it is done so without cutting corners, but the performance using IPv6, as well as IPv4, was consistently excellent. There is a concern - raised by many big-name vendors - that moving to IPv6 will result in potentially significant performance loss, especially when working at Layers 4-7, but we saw no significant degradation in performance here whatsoever.

Using a multi-port test setup, we were able to validate Kemp's claims (marketing figures) of supporting 2.9Gbps throughput at Layer 7 with both IPv4 and IPv6. While being able to show that moving to IPv6 had no impact on Layer 7 performance whatsoever, it is important to point out that that CPU utilisation on the LM-3600 was never saturated, peaking at around 65%. This has excellent implications for both the scalability of the appliance, and its ability to handle complex Layer 7 operations at maximum performance levels.

Overall, we therefore recommend anyone looking at Load-Balancer/Application Delivery Controller products to make sure they evaluate the Kemp LoadMaster range. Moreover, those companies who didn't believe they could afford this kind of technology should definitely think again...

Broadband Testing Approved