LOAD BALANCING IN THE MODERN DATA CENTER WITH BARRACUDA LOAD BALANCER FDC T740

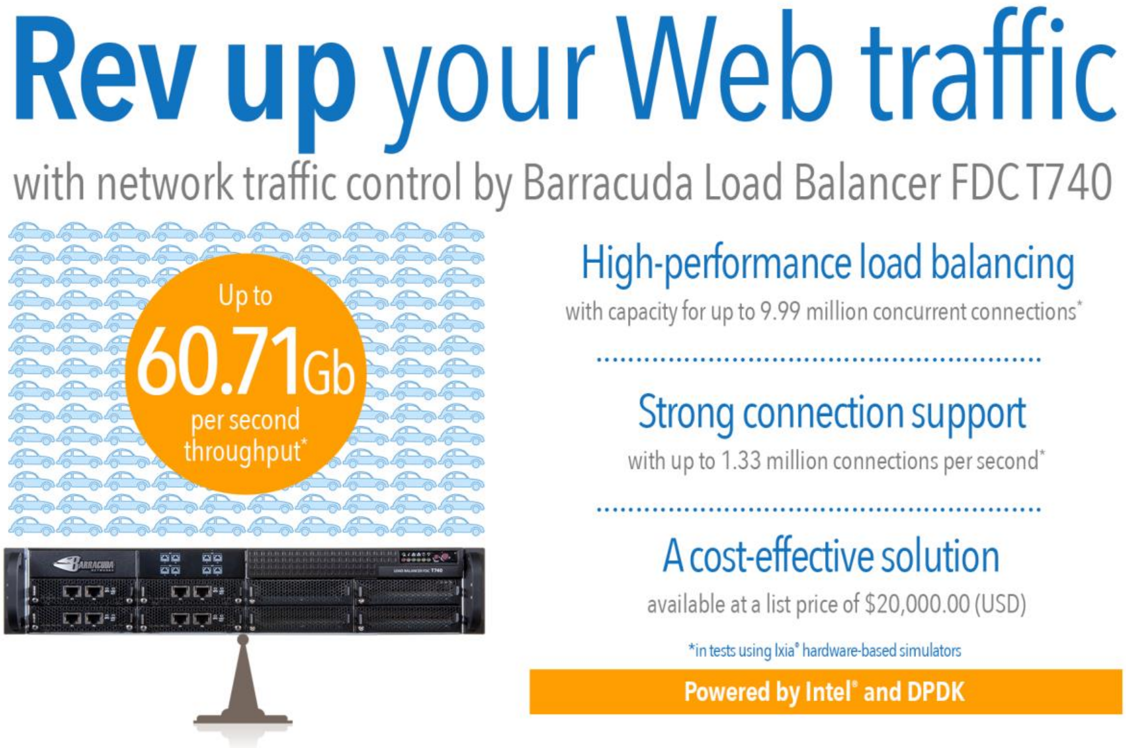

Balancing Web traffic across your front-end servers is key to a successful enterprise Web presence. Web traffic is much like a stadium parking lot full fans trying to leave at the end of the big game. Without someone directing, sending some cars to each available roadway, everyone gets stuck trying to get to the same place at the same time. Web traffic works much the same way. A server infrastructure needs a tool that divvies up user requests to each server so that one isn't unfairly overburdened while other servers cruise along comfortably.

These traffic directors, called load-balancers, vary in ability and price. In tests at the Principled Technologies data center using the Ixia IxLoad® Web-client/server simulator, we found that Barracuda Load Balancer FDC T740 provided a "big, open highway" for traffic to help you speed things up—in fact it handled up to 60.71 Gb/s of throughput. By leveraging the Data Plane Development Kit (DPDK), Barracuda Networks can offer this level of performance for just $20K.[1] Using a variety of requests, the Barracuda FDC T740 handled 9.99 million simultaneous Web connections, which translated to a good experience for users accessing even the most popular pages. The Barracuda FDC T740 also handled 1.33 million connections per second, which can help cover cancelled requests such as those from users changing their minds after clicking on a Web element, for example. With a load balancer capable of directing a high Web traffic load, users accessing your Web sites aren't stuck waiting, which can mean keeping their attention longer.

[1] List price quoted from Barracuda Networks, Inc. Includes no discounts or sales tax. For sales in the United States.

MARCH 2015
A PRINCIPLED TECHNOLOGIES REPORT
Commissioned by Barracuda Networks, Inc. and Intel Corp.

BARRACUDA LOAD BALANCER FDC INCREASES DATACENTER PERFORMANCE

Not just for Web
Load balancing isn't important to just Web infrastructures, but to other applications as well. In this study, we tested a simulated Web server environment, but all sorts of workloads that travel across a network benefit from having a traffic director that distributes work to servers. From databases to mail servers to analytics, load balancers can help you maximize your application delivery for greater capacity and reliability.

A user clicks on the checkout link to purchase a book from your online store. The wheels are turning, but nothing happens—the page is taking a long time to load. Frustrated, the customer opens another tab, types in your competitor's address, and orders from them instead. When a slow Web infrastructure frustrates countless visitors trying to make a purchase or simply look at your content, the lost revenue can really add up.

The Barracuda FDC seeks to remedy this problem by distributing network or application traffic to available servers in your environment. By balancing server traffic in the physical layer rather than the application virtualization layer, the Barracuda FDC T740 can handle many requests efficiently, which modern data centers require.

Based on the Data Plane Development Kit (DPDK), the Barracuda FDC makes the most of the packet processing on Intel Xeon® processors to provide fast application delivery to users. The DPDK implements a low overhead run-to-completion model for fast data plane performance and accesses network interface cards (NICs) via polling to eliminate the performance overhead of interrupt processing. While building high-speed networking appliances, system memory latencies typically pose a major bottleneck. DPDK implements techniques such as software pre-fetch and cache line alignments to overcome the performance impediments. This software architecture combined with Intel Data Direct Input Output (Intel DDIO) technology, which allows using part of CPU L3 cache as high-speed system memory, enables extremely high levels of system performance for packet processing workloads.

The Barracuda FDC can reduce expensive custom hardware costs by using these open-source technologies to offer fast load balancing for TCP and UDP traffic at a more affordable cost. The Barracuda FDC is part of Application Delivery Networking, which works to provide application availability, security, visibility, and acceleration—all crucial in the modern data center.

WHAT WE FOUND

The DPDK-based Barracuda FDC did its job directing Web traffic requests, handling an impressive amount of throughput, simultaneous users, and number of connections per second at an affordable cost. Read on for the details. To learn about how we tested, see Appendix A.

Let more users in with 40 Gb/s throughput

If a load balancer can handle lots of users or connections, it's less likely that there will be a big backup of users waiting to access your servers. Using asymmetrical traffic to simulate a real-world environment, we found that the Barracuda FDC handled 40.01 Gb/s to 60.71 Gb/s of throughput for Web-page payloads of 1KB to 1 MB (see Figure 1).

Page size	Throughput (Gb/s)
1 KB	40.01
4 KB	50.03
1 MB	60.71

Figure 1: The throughput the Barracuda FDC T740 achieved at various page sizes.

To understand the throughput numbers, consider the relationship between page size and the number of connected users. A 1KB page size meant less data per user request, but it also meant more users were able to connect, resulting in 40.1 Gb/s throughput. Larger page sizes with more data in each request resulted in higher throughput because fewer overall requests went through.

The DPDK-based Barracuda FDC worked quickly to process requests at all page sizes, which can translate into many benefits:

- A better experience for users, who don't have to wait on sluggish Web pages or other applications
- Increased user capacity, so more users can access your servers
- Possible savings in hardware and infrastructure costs, because you don't have to purchase more equipment to handle the load

More simultaneous connections = more users accessing your Web pages

Another way to look at load balancing performance is to see how many simultaneous users the appliance can handle. More users mean better reliability—in a popular Web server environment, for example, you may never have a page that gets millions of concurrent users, but having a high threshold means that the high numbers you do get can browse comfortably.

We found that the Barracuda FDC was able to handle just under 10 million simultaneous connections accessing a Web environment, which shows it can handle significant loads. (See Figure 2).

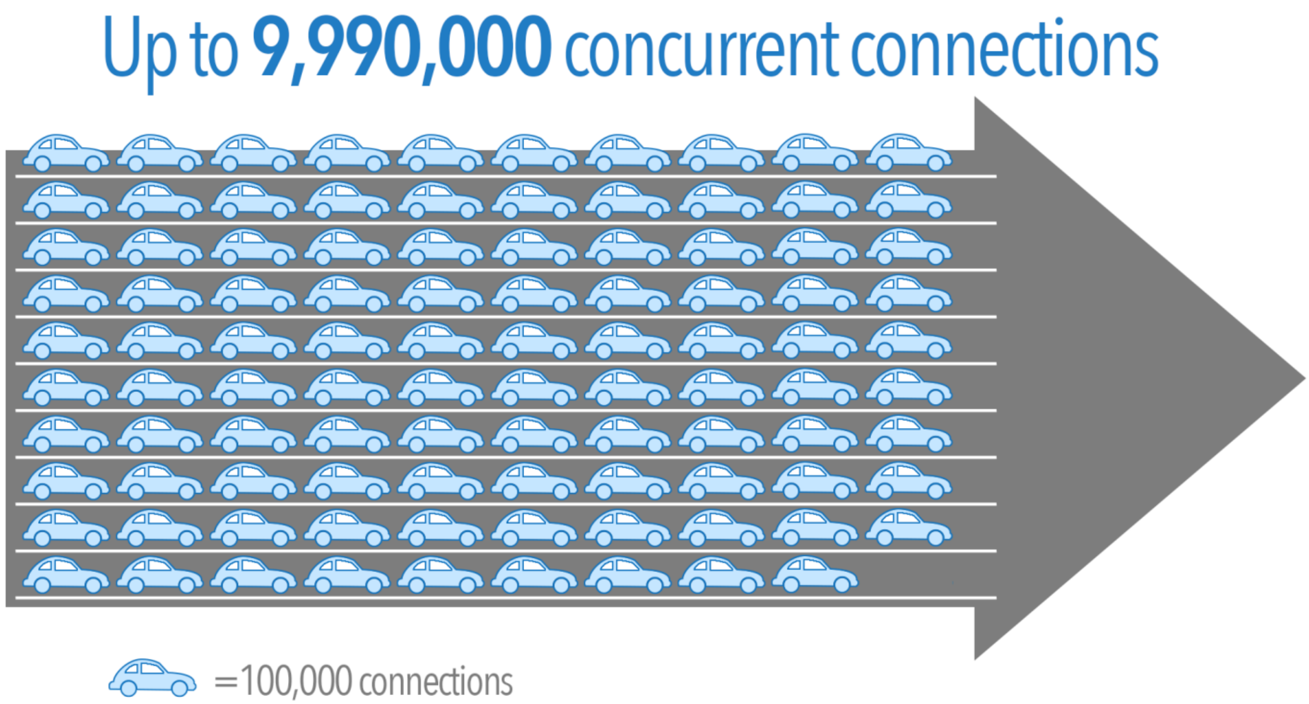

Figure 2: The Barracuda FDC T740 handled nearly 10 million simultaneous users.

Maximize the connections per second your servers can handle

Of course, real users don't always stay on the page or click through on a Web page. Sometimes they cancel their requests, which requires build up and break down of Web requests—more for the load balancer and servers to handle. When using these types of requests, the Barracuda FDC handled 1.33 million connections per second, again supporting an extremely high number of requests when you consider connections per minute or even hour. (See Figure 3.)

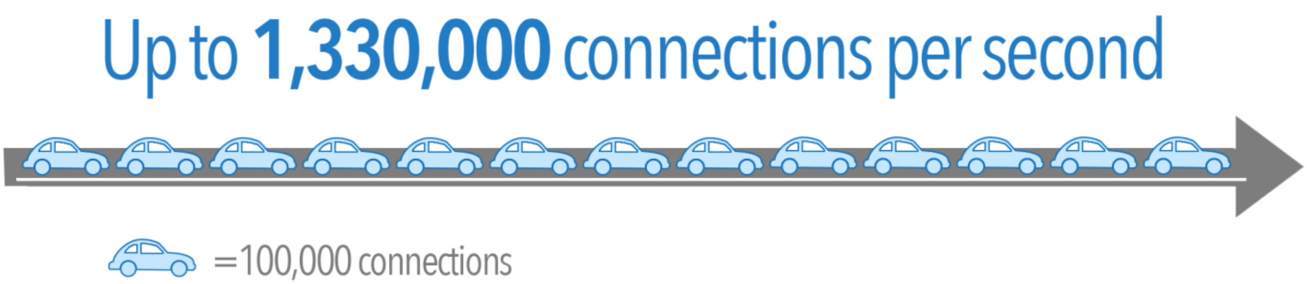

Figure 3: The Barracuda FDC T740 handled 1.3 million connections per second.

CONCLUSION

Traffic jams aren't pretty. They slow everyone down and cause a great deal of frustration. Frustrated Web users, waiting to access a page, have one major difference from people stuck in gridlock—they can choose to leave, and they will. Don't let your customers disappear due to slow Web performance.

Built on DPDK for fast packet-processing performance, the Barracuda FDC T740 has the power to direct a high volume of requests to your servers. In our tests, the Barracuda FDC handled up to 60.71 Gb/s of throughput, a figure made even more impressive when you consider the appliance's list price. If your infrastructure can handle a large number of users making varied requests, the less likely you are to have problems where users time out or get frustrated. We found that the Barracuda FDC handled 9.99 million simultaneous Web connections and 1.33 million connections per second.

These numbers show that the DPDK-powered Barracuda Load Balancer FDC T740 can balance the Web traffic load that the modern data center requires, and do so at an affordable cost.

APPENDIX A – HOW WE TESTED

To test the Barracuda FDC T740, we used Ixia devices to create the test load and measure performance. Ixia provides test systems for network infrastructure to verify that a network and its components function reliably under load.

Our testbed infrastructure comprised two similar Ixia appliances (see Figure 4), one simulating a variable number of Web clients, and the second simulating 16 Web servers, and two Dell Networking 8024f switches in a stacked configuration (see Figure 5). Each Ixia appliance connected to the switch stack via eight 10GbE cables. We connected the Barracuda FDC T740 to the switch stack via eight 10GbE cables as one LACP LAG. The MTU size on each link was 1,500 bytes as we choose not to use jumbo packets on the back-end network.

Ixia appliance	Network ports	Testbed function
PerfectStorm ONE 1GE/10GE Fusion	16 10GbE (fibre optic)	Web servers
Xcellon-Ultra XT80-V2	16 10GbE (fibre optic)	Web clients

Figure 4: Details on the Ixia appliances used to generate Web load and measure performance.

Stack member	Port range and type	VLANs	Link characteristic	Access
1	1-16 (fibre optic)	702	unlagged	Ixia server appliance
1	21-28 (fibre optic)	N/A	Switch stack	Stack member #2
1	21-24 (copper)	702, 712	LAG 1	Barracuda FDC
2	1-16 (fibre optic)	712	unlagged	Ixia client appliance
2	21-28 (fibre optic)	N/A	Switch stack	Stack member #1
2	21-24 (copper)	702, 712	LAG 1	Barracuda FDC

Figure 5: Network port assignments on the Dell Networking 8024f switch stack.

We used the Ixia IxLoad application (version 6.60.0.406) on these appliances to specify and generate the loads, consume the loads, and measure performance stats, etc. The two appliances together could generate just under 80Gb/s of throughout. We tested our infrastructure's ability to support such throughputs by connecting the two appliances back-to-back through the switch stack (omitting for the moment the Barracuda FDC) and found it could pass 92 percent of that capacity.

Configuring the Barracuda Load Balancer FDC T740

As previously mentioned, we configured the eight 10GbE ports in an active LACP LAG. We created and assigned VLANs to this LAG. VLAN 702 was the back-end server network (172.47.0.0/16) and VLAN 712 was the client-side network (172.46.0.0/16). We assigned IP address 172.47.0.82 to the Barracuda FDC as the gateway from the server-side to client-side. Conversely, the front-end clients sent their requests to the Barracuda FDC's virtual interface with IP address 172.46.0.82.

We configured the Barracuda FDC to distribute the layer 4 requests to 16 servers (simulated on the Ixia PerfectStorm One appliance) with IP addresses 172.47.0.100 to 172.47.0.115. We configured the Barracuda FDC to impersonate the clients to the back-end servers, and the backend servers sent their return traffic to the clients through the load balancer (i.e., we did not make use of the direct server return feature). Logging and statistic generation on the Barracuda FDC were set to their default values.

Configuring the Ixia IxLoad application for each test

For each test, we modelled the device in Ixia as a load balancer and disabled the direct server return option. The IxLoad application will dynamically alter the number of users, connections, and connection attempts, in order to reach the configured target metric or the device's maximum.

Measuring throughput

We chose to use only HTTP GET requests, rather than HTTP PUTs or a mixture of the two, to follow typical Web use patterns. We looked at three sizes for amount of application data returned by the server; namely, page sizes of 1KB, 5KB and 1MB. The Web clients and servers used the HTTP 1.1 protocol, with no limit on the duration of the HTTP connection. Each client could initiate up to 100 simultaneous connections. We gradually increased the number of clients over 30 seconds and ran each test for 10 minutes.

For the 1KB page-size tests, we used an IxLoad target of 40 Gb/s; for the 5KB page-size tests, we used a target of 50 Gb/s. For the 1MB page-size test, we choose a target of 80 Gb/s.

Measuring concurrent connections

We used HTTP GET requests. We looked at three sizes for amount of application data returned by the server; namely, page sizes of 1B, 1KB, and 5KB, though we found little variation in the results. The results presented here are for 1KB page sizes. The Web clients and servers used the HTTP 1.0 protocol so that the connection was dropped and restarted after each request. Each client could initiate up to 100 simultaneous connections. We gradually increased the number of clients over 30 seconds, and ran each test for 10 minutes.

We used an IxLoad target of 10,000,000 concurrent connections for these tests.

Measuring new connection creation/destruction rate

We used HTTP GET requests and page sizes of 1B. The Web clients and servers used the HTTP 1.0 protocol so that the connection was dropped and restarted after each request. Each client could initiate up to 100 simultaneous connections. We gradually increased the number of clients over 30 seconds, and ran each test for 10 minutes.

We used an IxLoad target of 3,000,000 connections/s.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.