Principled Technologies®

TEST REPORT
JUNE 2009

Consolidating legacy Solaris BIND DNS servers onto Dell PowerEdge Solaris servers

Executive summary

Dell™ Inc. (Dell) commissioned Principled Technologies® (PT) to measure the DNS performance of three dual-socket blade servers running Sun™ Solaris™ 10 and to analyze the consolidation potential for customers staying with Solaris and BIND DNS server software, but migrating from older-generation servers to a current-generation x86 platform.

We tested the following single blades in their respective enclosures:

- Dell PowerEdge™ M610 running Solaris 10 5/09 and BIND 9.6.0-P1
- HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1
- IBM® BladeCenter® HS21 running Solaris 10 5/08 and BIND 9.3.4-P1

PT received the HP ProLiant BL460c blade directly from a third-party hardware reseller. Dell provided the other hardware for this test.

KEY FINDINGS

- The Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1, delivering 4.39 times the performance of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, could replace at least four of those servers. (See Figures 1 and 2.)
- The Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1, delivering 5.31 times the performance of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1, could replace at least five of those servers. (See Figures 1 and 2.)
- The idle power of the Dell PowerEdge M610 was on par with that of the HP ProLiant BL460c G1, and was less than that of the IBM BladeCenter HS21. (See Figure 3.)
- The Dell PowerEdge M610 achieved 3.96 times the performance per watt of the HP ProLiant BL460c G1and 5.81 times the performance per watt of the IBM BladeCenter HS21. (See Figure 3.)

Our goal was not to achieve maximum performance, but rather to focus on the performance typical of such a real-world scenario. We selected two common blade systems from approximately 3 years ago: the HP ProLiant BL460c G1 and the IBM BladeCenter HS21, each of which had two Intel Xeon E5345 processors, 4 GB of RAM, and two 73GB 10K SAS drives. We compared these to a new Dell PowerEdge M610 blade system with two Intel Xeon X5550 processors, 24 GB of RAM, and two 73GB 15K SAS drives.

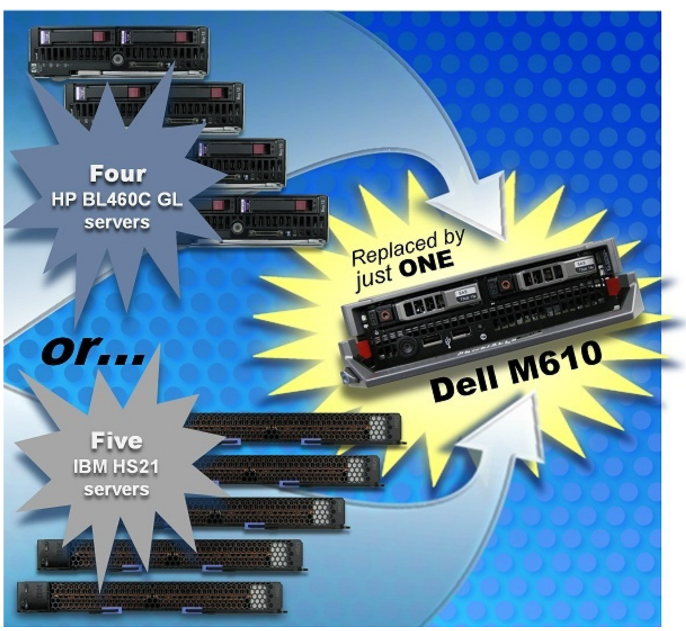

Figure 1: DNSPerf-based consolidation potential for the test configurations.

As Figure 1 depicts, our DNSPerf testing indicated that the Dell PowerEdge M610 can replace four or more HP ProLiant BL460c servers or five or more IBM BladeCenter HS21 servers.

We used the DNSPerf 1.0.1.0 (DNSPerf) test tool, available as a free download from http://www.nominum.com/services/measurement_tools.php. Nominum designed DNSPerf to test authoritative domain name servers. Nominum also provides a sample query file of 3 million records. Although Nominum provides this 3 million-record file for their RESPerf tool, we were able to adapt the file for this test. (We describe our modifications in Appendix B.) The workload's main reporting metric is queries per second (QPS). Unless we state otherwise, all results in this report are in QPS.

Figure 2 provides the performance scores of each blade server configuration when running our custom workload using DNSPerf. Each result is the median peak score of three benchmark runs. Higher scores are better. A higher number of QPS indicates the server can handle a greater load. The Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 achieved a score of 296,069, a 339.89 percent increase over the performance of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, which achieved a score of 67,305. The same Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 achieved a 431.05 percent increase over the performance of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1, which achieved a score of 55,751.

Server	QPS
Dell PowerEdge M610	296,069
HP ProLiant BL460c G1	67,305
IBM BladeCenter HS21	55,751

Figure 2: DNSPerf results for the test systems. Higher numbers are better.

Figure 3 details the power consumption, in watts, of the test servers while idle and during the median run of the DNSPerf benchmark. The idle power of the Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 was on par with that of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, and was less than that of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1. The average peak power of the Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 was higher than that of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, and was less than that of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1. The QPS per watt of the Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 was higher than that of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, and was higher than that of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1. The Dell PowerEdge M610 running Solaris 10 5/09 and BIND 9.6.0-P1 achieved 3.96 times the performance per watt of the HP ProLiant BL460c G1 running Solaris 10 5/08 and BIND 9.3.4-P1, and achieved 5.81 times the performance per watt of the IBM BladeCenter HS21 running Solaris 10 5/08 and BIND 9.3.4-P1.

Server	Idle power (watts)	Average peak power (watts)	QPS per watt	Relative performance per watt of the Dell PowerEdge M610
Dell PowerEdge M610	529.23	675.95	438	1.00
HP ProLiant BL460c G1	528.30	608.03	110	3.96
IBM BladeCenter HS21	677.99	739.88	75	5.81

Figure 3: Average power usage (in watts) of the test servers while idle and during the median DNSPerf run. Lower numbers are better.

Test methodology

Because our goal was to isolate the performance of the servers, we ran on an isolated network. During any test, the only systems on the network were the driver and the server under test. We explain our test system configuration information in Appendix A.

We created a custom workload that performed forward and reverse lookups against a single zone. The server under test was always fully authoritative for that zone. The final workload consisted of over 2.7 million queries. The zone we queried against comprised over 98,000 address records (A records) and over 72,000 pointer records (PTR records). We explain how we generated our test files in Appendix B.

For the Dell PowerEdge M610, we upgraded BIND to 9.6.0-P1, the latest version for which there was a Solaris package available. As we note in Appendix A, the motherboard and BIOS for the Dell PowerEdge M610 were pre-release versions. For the HP ProLiant BL460c G1 and the IBM BladeCenter HS21, we performed no special tuning.

DNSPerf execution

We used server-class systems, which we connected to our network via a gigabit switch, to generate the workload for our tests. For the HP ProLiant BL 460c G1 and the IBM BladeCenter HS21, we used two load-generating systems, one for each embedded NIC. In this configuration, the IBM BladeCenter HS21 had roughly 4 percent of the CPU available, while the HP ProLiant BL 460c G1 had roughly 5 percent available. However, adding more workload-generating systems did not increase the QPS, so we tested these blades with two workload-generating systems.

In the case of the Dell PowerEdge M610, we had more NICs available, but found that three workload-generating systems were sufficient to saturate the system. In this case, no CPU remained.

Each workload-generating system contained the DNSPerf driver application and executed a workload against a single zone, for which the server was authoritative.

By design, DNSPerf self-paces the workload so that the queue depth does not exceed the user-defined limit. The default queue depth is 20. We experimented using all three blades and found that we could raise the queue depth to 200 before queries started to time out. We used a queue depth of 200 in our testing.

DNSPerf does not provide a ramp-up period. We worked around this by using a script that ran DNSPerf two times. Because we considered the first run to be the ramp-up period, we used the result from the second run. Both periods were 5 minutes, and both runs used identical switches.

Finally, because DNSPerf writes its output to the screen, we redirected the DNSPerf output to a file.

We used the following switches to run DNSPerf:

```
dnsperf –s <server IP address> -d workload –l 300 –q 200 > file
```

Power measurement procedure

To record each server's power consumption during each test, we used three Extech Instruments (www.extech.com) 380803 Power Analyzer/Dataloggers. We connected the power cord from the server under test to each Power Analyzer's output load power outlet. We then plugged the power cord from the Power Analyzer's input voltage connection into a power outlet.

We used the Power Analyzer's Data Acquisition Software (version 2.11) to capture all recordings. We installed the software on a separate PC, which we connected to the Power Analyzer via an RS-232 cable. We captured power consumption at 1-second intervals.

To gauge the idle power usage, we recorded the power usage for 2 minutes while each server was running the operating system but otherwise idle. We summed the results from all three meters to get the idle power for the server.

We then recorded the power usage (in watts) for each server during the testing at 1-second intervals. To compute the average power usage, we averaged the power usage during the time the server was producing its peak performance results. We call this time the power measurement interval. We summed the results from all three meters to get the peak power for the server. Figure 3 shows the results of these measurements.

Appendix A – Test system configuration information

Figure 4 provides detailed configuration information about the test servers.

Servers	Dell PowerEdge M610	HP ProLiant BL460c G1	IBM BladeCenter HS21
Enclosure			
Model	Dell PowerEdge M1000e	HP BladeSystem c7000	BladeCenter H Type 8852
General dimension information			
Height (inches)	17.3	17.5	15.75
Width (inches)	17.6	17.5	19.0
Depth (inches)	29.7	32.0	28.0
U size in server rack (U)	10	10	9
Power supplies			
Total number	6	6	4
Wattage of each (W)	2,360	2,250	2880
Enclosure Power Management Policy	DCRedundant	Redundant	Redundant without performance impact
Cooling fans			
Total number	9	10	2 blowers
Dimensions (h x w) of each	3.1” x 3.5”	2.75” x 2.25”	4.5” x 11.5”
Voltage (V)	12	12	200-240
Amps (A)	7.0	16.5	5.5
General processor setup			
Number of processor packages	2	2	2
Number of cores per processor package	4	4	4
Number of hardware threads per core	2	1	1
System Power Management Policy	OS Control	HP Dynamic Power Savings mode	N/A
CPU			
Vendor	Intel	Intel	Intel
Name	Quad-Core Intel Xeon processor X5550	Quad-Core Intel Xeon processor E5345	Quad-Core Intel Xeon processor E5345
Stepping	D0	7	7
Socket type	LGA 1366	LGA 771	LGA 771
Core frequency (GHz)	2.66	2.33	2.33
Front-side bus frequency (MHz)	1,333	1,333	1,333
L1 cache	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)	32 KB + 32 KB (per core)
L2 cache	1 MB (4 x 256 KB)	2 x 4 MB (each 4 MB shared by two cores)	2 x 4 MB (each 4 MB shared by two cores)
L3 cache	1 x 8 MB	N/A	N/A
Thermal design power (TDP, in watts)	95	80	80
Platform			
Vendor and model no.	Dell PowerEdge M610	HP ProLiant BL460c G1	IBM BladeCenter HS21
Motherboard model no.	PWB5N793	SP438249	IBM 8853C2U

Servers	Dell PowerEdge M610	HP ProLiant BL460c G1	IBM BladeCenter HS21
Motherboard revision no.	X02 (pre-release version)	001	B1
BIOS name and version	Dell 0.2.15 (pre-release version)	HP BIOS I15 6/28/2007	IBM 1.10 1/31/08
BIOS settings	Enabled Logical Processor	Disabled Hardware Prefetcher and Adjacent Cache Line Prefetch	Disabled Hardware Prefetcher and Adjacent Cache Line Prefetch
Memory modules			
Total RAM in system (GB)	24	4	4
Number of types of memory modules	1	2	1
First type of memory module			
Vendor and model no.	Crucial CT51272BB1339	Hynix HYMP512F72CP8D2-Y5	Hynix HYMP512F72CP8D2-Y5
Type	PC3-10600	PC2-5300	PC2-5300
Speed (MHz)	1,333	667	667
Speed in the system currently running @ (MHz)	1,066	667	667
Timing/Latency (tCL-tRCD-iRP-tRASmin)	9-9-9-24	5-5-5-15	5-5-5-15
Size (GB)	24	2	4
Number of RAM modules	6 x 4 GB	2 x 1 GB	4 x 4 x 1 GB
Chip organization	Dual side	Dual side	Dual side
Second type of memory module			
Vendor and model no.	N/A	Samsung M395T2953EZ4-CE65	N/A
Type	N/A	PC2-5300	N/A
Speed (MHz)	N/A	667	N/A
Speed in the system currently running @ (MHz)	N/A	667	N/A
Timing/Latency (tCL-tRCD-iRP-tRASmin)	N/A	5-5-5-15	N/A
Size (GB)	N/A	2	N/A
Number of RAM modules	N/A	2 x 1 GB	N/A
Chip organization	N/A	Dual side	N/A
Hard disk			
Vendor and model no.	Seagate ST973451SS	HP DG072A8854	IBM 26K5777
Number of disks in system	2	2	2
Size (GB)	73	72	73.4
Buffer size (MB)	16	8	8
RPM	15,000	10,000	10,000
Type	SAS	SAS	SAS

Servers	Dell PowerEdge M610	HP ProLiant BL460c G1	IBM BladeCenter HS21
Operating system			
Name	Solaris 10 5/09 Operating System for x86-based systems (64-bit)	Solaris 10 5/08 Operating System for x86-based systems (64-bit)	Solaris 10 5/08 Operating System for x86-based systems (64-bit)
File system	UFS	UFS	UFS
Kernel	Solaris 10 5/08 s10x_u5wos_08 x86	Solaris 10 5/08 s10x_u5wos_10 x86	Solaris 10 10/08 s10x_u5wos_10 x86 with patch 113112-08
Language	English	English	English
Network card/subsystem			
Vendor and model no.	Broadcom BCM5709 NetXtreme II GigE	Broadcom BCM5708S NetXtreme II GigE	Broadcom BCM5708S NetXtreme II GigE
Type	Integrated	Integrated	Integrated
Blade switch modules	Two Cisco WS-CBS3130X-S Catalyst Blade Switch Modules	Two HP 1:10 Ethernet Blade Switch Modules	Internal Ethernet switch
Blade pass-through modules	Two Gigabit Ethernet pass-through modules	Two HP 4Gb Fibre Channel pass-through modules	Two IBM 39Y9323 Copper pass-through modules

Figure 4: Detailed configuration information about the test server systems.

Appendix B – How we created our workload

Because we wanted to test the performance of only the server, we limited this test to a single zone and performed only forward and reverse lookups. For simplicity, we also limited the test to IPv4 addresses. We did not update the zone at all during this test.

As we note above, we started with the 3 million-record query file that Nominum provides for download, queryfile-example-3million.gz. After unzipping the file, we translated underscores in names to dashes, to avoid potential problems with the names. We created our workload by removing the following types of records:

- SOA
- SRV
- AAAA
- AXFR
- TXT
- MX
- NS

We also removed the records with IPv6 addresses, those that queried the local zone. Finally, there were a few records with the string `dnsbugtest`. We removed those as well. The result was a workload with over 2.7 million queries in it. Finally, we appended `pti.` to the URLs, to ensure that the queries went to the correct domain.

Once we had the workload, we needed to create the data it would query. We ran our own custom program that scanned the workload and created two additional files, one consisting of A records and one consisting of PTR records. We sorted each of these new files and removed duplicates. To each of the more than 98,000 A records in the unique list of A records, we used our program to add a unique IP address. To each of the more than 72,000 records in the unique list of PTR records, we used our program to add a unique URL.

Once we had our source data, we manually edited the files to add the directives and SOA record to create our zone files.

Finally, we edited the named.conf file so that it used our zone files.

About Principled Technologies

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 you assess how it will fare against its competition, its performance, whether it's ready to go to market, 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.

Principled Technologies, Inc.
1007 Slater Road, Suite 250
Durham, NC 27703
www.principledtechnologies.com
info@principledtechnologies.com