TEST REPORT
JANUARY 2009

BIND DNS performance on Dell and Sun servers running Solaris

Executive summary

Dell Inc. (Dell) commissioned Principled Technologies (PT) to measure the DNS performance of the following two servers running the Sun Solaris operating system:

- Dell PowerEdge R900 running BIND 9.3.5-P1 on Solaris 10
- Sun Fire V440 running BIND 8.3.3 on Solaris 9

Our goal was to compare a typical existing installation with a typical new purchase configuration, so we could analyze what customers might experience if they were staying with Solaris and its bundled BIND DNS server software, but were migrating from RISC to x86. Our goal was not to achieve maximum performance, but rather to focus on the performance typical of such a real-world scenario. We selected a common system configuration for an existing Sun environment—a 3-year-old Sun Fire v440 with 32 GB of RAM and all of its disk drive bays full (four 73GB 10K drives)—and compared it to a moderate configuration of a new six-core Intel Xeon 7400 series-based Dell PowerEdge R900 server with 64 GB of RAM.

KEY FINDING

- Using only 1 of the 24 cores in a Solaris Container, the Dell PowerEdge R900 running BIND 9.3.5-P1 on Solaris 10 achieved 45.5 percent higher performance than the Sun Fire V440 running BIND 8.3.3 on Solaris 9 without virtualization. (See Figure 1.)

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.

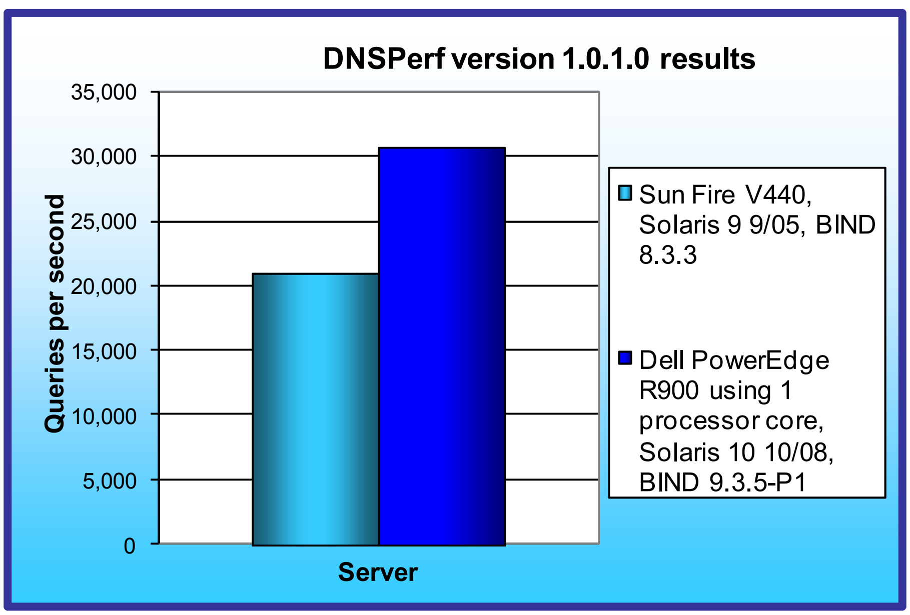

Figure 1: DNSPerf results for the test configurations. Higher numbers are better.

Figure 1 details the DNSPerf performance of each configuration when running our custom workload. Each result is the median peak score of three benchmark runs. A higher number of QPS indicates the server can handle a greater load. When using 1 of its 24 cores in a Solaris Container, the Dell PowerEdge R900 running BIND 9.3.5-P1 on Solaris 10 achieved a score of 30,666, a 45.5 percent increase over the performance of the Sun Fire V440 running BIND 8.8.8 on Solaris 9 without virtualization, which achieved a score of 20,903.

Figure 2 details the results, each of which is the median peak score of three benchmark runs. Higher scores are better.

Server	QPS
Dell PowerEdge R900 running BIND 9.3.5-P1 on Solaris 10 using a single zone with one dedicated processor core	30,666
Sun Fire V440 running BIND 8.3.3 on Solaris 9	20,903

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

We also wanted to explore how a customer might consolidate multiple DNS servers onto a single server. Because we believe that a Dell PowerEdge R900 can handle the DNS requirements for many installations with plenty of power left over, we decided to devote one-fourth of the PowerEdge R900 to running six DNS servers. Note: During these tests, mpstat data from the PowerEdge R900 showed that approximately 70 percent of the total processor resources remained available.

For each of the six instances, we created a corresponding Solaris Container that contained a standalone DNS server. We configured each of the six Solaris Containers with a single processor and a dedicated IP address. The six Solaris Containers simultaneously running DNSperf on the Dell PowerEdge R900 were up to 41.3 percent higher than the performance of DNSperf on the Sun Fire V440. The average performance of those six Containers on the Dell PowerEdge R900 was 26.7 percent higher than the performance of the Sun Fire V440.

Figure 3 details the results, each of which is the median peak score of three benchmark runs. Higher scores are better.

Server	QPS
Average score for a single Zone on a Dell PowerEdge R900 running BIND 9.3.5-P1 on Solaris 10 using six Zones, each with one dedicated processor core	26,483
Sun Fire V440 running BIND 8.3.3 on Solaris 9	20,903

Figure 3: DNSPerf results for the test systems when we used six Solaris Zones on the Dell PowerEdge R900. Higher 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 Sun Fire V440, we performed no special tuning.

For the Dell PowerEdge R900, we created a single Solaris Container using the methodology we detailed in *Migrating from Solaris 9 Enterprise Edition on Sun Fire V440 with BIND 8.3.3 to Solaris 10 Enterprise Edition on Dell PowerEdge R900 with BIND 9.3.5-P1* (http://www.principledtechnologies.com/Clients/Reports/Dell/DNS_migration_guide.pdf). The container had one dedicated network interface and the processor set in its resource pool contained only one processor. We used the svccfg tool to set the number of BIND 9.3-P1 threads to three. Since BIND 8.3.3 is single threaded, we did not have the opportunity to set the number of threads for the Sun Fire V440.

DNSPerf execution

We used a server-class system, which we connected to our network via a gigabit switch, to generate the workload for our tests. The 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 the Sun Fire V440 and found that we could raise the queue depth to 100 before queries started to time out. We used a queue depth of 100 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 100 > file
```

Appendix A – Test system configuration information

Figure 4 provides detailed configuration information about the test servers.

Server	Dell PowerEdge R900	Sun Fire V440
General processor setup		
Number of processor packages	4	4
Number of cores per processor package	6	1
Number of hardware threads per core	1	1
CPU		
Vendor	Intel	Sun
Name	Xeon 7450	UltraSPARC IIIi
Socket type	PGA604	959
Core frequency (GHz)	2.4	1.59
Front-side bus frequency (MHz)	1,066	JBus running at up to 200
L1 cache	32 KB instruction 32 KB data	64 KB 4-way data cache 32 KB 4-way instruction cache
L2 cache	9 MB shared	1 MB on-chip, 4-way set associative
L3 cache	12 MB shared	N/A
Platform		
Vendor and model number	Dell PowerEdge R900	Sun Fire V440
Motherboard model number	CN-0947H-13740-87N-000P	540-5418-06
Motherboard revision number	X00	50
BIOS name and version	Dell BIOS 1.1.9	OBP 4.22.33 2007/06/10 12:42 Sun Fire V440v Netra 440
BIOS settings	Default	Default
Memory module(s)		
Vendor and model number	Samsung M395T5750EZ4-CE65	Micron MT36VDDF25672G-335D2
Type	PC2-5300F	PC2700R DDR 333
Speed (MHz)	667	333
Speed in the system currently running @ (MHz)	667	333
Size (GB)	2	2
Number of RAM modules	32	16
Chip organization	Double-sided	Double-sided
Hard disk		
Vendor and model number	7 x Seagate ST973451SS, 1 x Fujitsu MBC2073RC	4 x Fujitsu MAP3735NC
Number of disks in system	8	4
Size (GB)	73	73.5
Buffer size (MB)	16	8
RPM	15,000	10,000
Type	SAS/Serial Attached SCSI	Ultra320 SCSI
Controller	PERC 6/i	LSI 1030
Driver	mega_sas (00.00.01.23)	SCSI Disk Driver 1.435

Server	Dell PowerEdge R900	Sun Fire V440
Operating system		
Name	Solaris 10 10/08 Operating System for x86-based systems (64-bit)	Solaris 9 9/05 Operating System for SPARC-based systems (64-bit)
Build number	Solaris 10 10/08 s10x_u6wos_07b x86	Solaris 9 9/05 HW s9s_u9wos_06b SPARC
File system	UFS	UFS
Kernel	SunOS 5.10 Generic_137138-09	SunOS 5.9 Generic_118558-34
Language	English	English
DNS software		
Name	BIND 9.3.5-P1 (part of the Solaris 10 distribution, 10/08)	BIND 8.3.3 (part of the Solaris 9 distribution, 9/05)
Build number	10.2.0.1.0	9.2.0.1.0
Graphics		
Vendor and model number	ATI ES 1000	None
Resolution	1,280 x 1,024	N/A (Used command line interface)
Network card/subsystem		
First network card/subsystem		
Vendor and model number	Broadcom BCM5708	Broadcom BCM5421
Type	4 x integrated	2 x integrated
Driver	Broadcom 3.5.6	CE Ethernet Driver v1.158
Second network card/subsystem		
Vendor and model number	Intel PRO/1000 PT server adaptor	N/A
Type	2 x dual-port, 5 x single-port, 4 x PCI card	N/A
Driver	Intel version 5.2.8	N/A
Optical drive		
Vendor and model number	Toshiba Samsung TS-L333A	Toshiba Samsung SD-C2612
USB ports		
Number	4	4
Type	USB 2.0	USB 1.0
Power supplies		
Total number	2	2
Wattage of each	1,570	680
Cooling fans		
First type of cooling fan		
Number	4	2
Vendor and model number	Nidec UT094	NMB BG0903-B044-00S
Dimensions	88 mm(h) x 80 mm(w)	95 mm diameter
Voltage	12	12
Amps	1.76	1.34

Server	Dell PowerEdge R900	Sun Fire V440
Second type of cooling fan		
Number	4	1
Vendor and model number	Nidec V34809-35DELF	NMB 4715kl-04W-B59
Dimensions	120 mm(h) x 140 mm(w)	125 mm (h) x 120 mm (w)
Voltage	12	12
Amps	3.3	1.3

Figure 4: Detailed system configuration information for the two test configurations.

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.

Once we had the files for our DNS server, we configured DNS in a Solaris Container using the methodology we detailed in *Migrating from Solaris 9 Enterprise Edition on Sun Fire V440 with BIND 8.3.3 to Solaris 10 Enterprise Edition on Dell PowerEdge R900 with BIND 9.3.5-P1* (http://www.principledtechnologies.com/Clients/Reports/Dell/DNS_migration_guide.pdf).

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