CommunityDNS

Performance testing of BIND, NSD and CDNS platforms on identical hardware.

May 2010

CommunityDNS
Bath University Innovations Centre
Broad Quay
Bath
BA1 1UD
UK
Feedback@CommunityDNS.net

The Purpose of the Test

As more and more zones are being signed with DNSSEC and more DNS queries requesting a signed answer, CommunityDNS (CDNS) ran tests to asses the readiness of the two main DNS server packages, BIND and NSD, for the added work load this will impose on standard server hardware and what other issue or limitations this might impose.

The same tests were applied on CDNS' platform using the same hardware and tests run on BIND and NSD to understand the working implications on all three platforms in both DNSSEC and non-DNSSEC environments.

Executive Summary

Whilst the results of the test were enlightening, the journey for conducting the tests proved just as informative.

When using standard hardware platforms initial tests illustrated that whilst BIND and NSD were reaching their maximum capacities for each of the test levels, it was unclear as to why query processing reached a plateau for CommunityDNS' platform when the platform suggested its maximums had not yet been reached. Through analysis network hardware and cabling were replaced, yielding noticeable increase in the amount of queries processed. What was discovered was the necessity to change from a 100MB switch to a 100GB switch. The 100MB router approached maximum processing at approximately 13,000 queries per second. Improvements were also noticed when replacing CAT-5 cabling with CAT-6. This is of particular importance since packet sizes will increase from 512 bytes to 4,000 bytes. All tests were rerun under 100GB port speeds and CAT-6 cabling.

Testing revealed CommunityDNS' platform capacity under-utilized when using 100MB switch, router and CAT-5 cabling.

Network and bandwidth impact due to increase in packet size from 512 bytes size to 4,000 byte size.

The test itself tested BIND, NSD and CommunityDNS platforms with four different zone files, each either unsigned or with a realistic sample of signed records. Using the following zone sizes, BIND, NSD and CommunityDNS (or CDNS) platforms peaked at the following queries per second (q/sec) levels:

7,691 records	Unsigned	Signed
BIND	53,600	39,000
NSD	94,400	77,100
CDNS	124,000	103,000

BIND: CDNS processes 131% more q/sec for unsigned and 164% for signed. ***NSD:*** CDNS processes 31% more q/sec for unsigned and 34% for signed.

240,419 records	Unsigned	Signed
BIND	37,500	28,000
NSD	79,000	61,700
CDNS	122,300	107,000

BIND: CDNS processes 226% more q/sec for unsigned and 282% for signed. ***NSD:*** CDNS processes 55% more q/sec for unsigned and 73% for signed.

19,405,229 records	Unsigned	Signed
BIND	57,500	25,500
NSD	83.000	53,300
CDNS	120,500	89,300

BIND: CDNS processes 110% more q/sec for unsigned and 250% for signed. ***NSD:*** CDNS processes 45% more q/sec for unsigned and 68% for signed.

57,873,014 records	Unsigned	Signed
BIND	0	0
NSD	0	0
CDNS	120,500	89,300

BIND: Was unable to load the file of this size.
NSD: Was unable to load the file of this size.

Capacity and Scaling. One other aspect to glean from this study deals with overall capacity and scaling. Tests of the CommunityDNS platform in both unsigned and signed environments illustrate consistently high performance in the handling of zone files from small to the very large, whether unsigned or signed. Such scalability and capacity have a direct bearing on the ability to continue answering legitimate queries while not being impacted by a D/DoS attack.

Degradation. In an *unsigned* environment, there is only a 2.8% loss in performance on the CommunityDNS platform when handling zone file sizes from 7,691 names to 57,873,014 names. For file sizes of 7,691 to 19,405,229 names there is a 2.4% loss for CommunityDNS. In comparison NSD's performance incurs a 12.1% degradation in performance. While BIND actually gains in overall percentage of -7.2% degradation it is consistently the lowest performer prior to maxing out. Please note, during the test neither BIND nor NSD could handle a zone file of 57,873,014 names.

Unsigned zone files from 7,691 names to 19,405,229 names:

CDNS: 2.4% degradation
BIND: -7.2% degradation
NSD: 12.1% degradation

In a *signed* environment, there is a 13.3% loss in performance on the CommunityDNS platform when handling zone file sizes from 7,691 names to 57,873,014 names. For file sizes of 7,691 to 19,405,229 names there is a 3.6% loss for CommunityDNS. In comparison BIND looses approximately 34.6% in performance whereas NSD's performance incurs a 30.9% degradation in performance across the respective zone file sizes. Please note, during the test neither BIND nor NSD could handle a zone file of 57,873,014 names. ***What this means for TLD operators who use BIND, when moving from an unsigned zone to a signed zone there will be an immediate loss of 34.6% in capacity.***

Signed zone files from 7,691 names to 19,405,229 names:

CDNS: 3.6% degradation
BIND: 34.6% degradation
NSD: 30.9% degradation

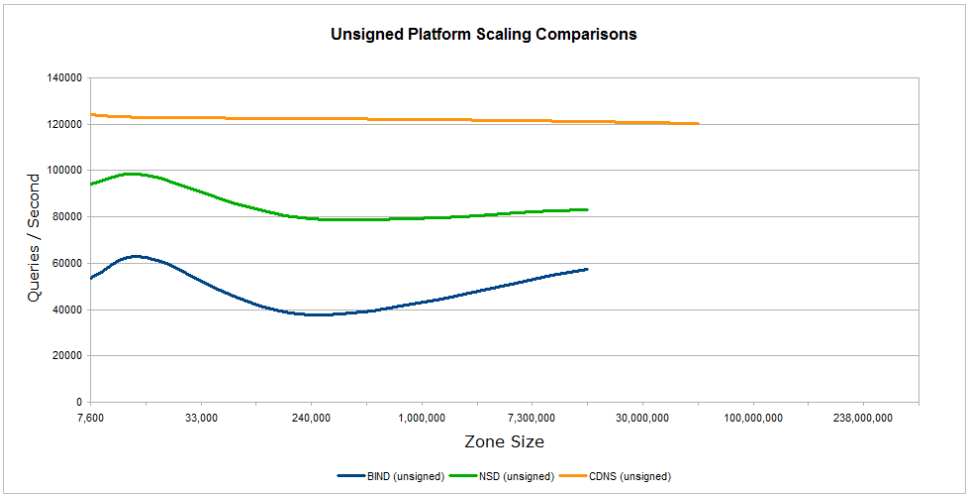

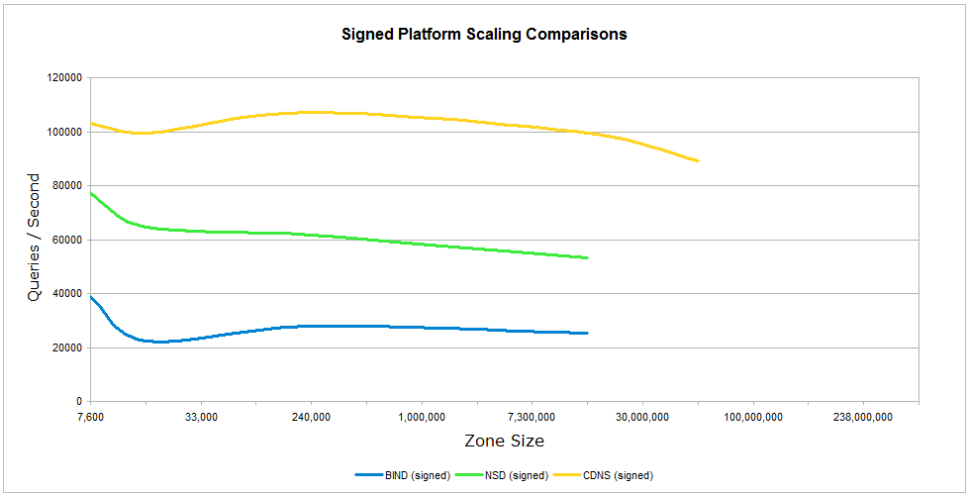

Processing Peaks:
Queries/Second

7,691 names Unsigned:
BIND: 53,600
NSD: 94,400
CDNS: 124,000

7,691 names Signed:
BIND: 39,000
NSD: 71,000
CDNS: 103,000

240,419 names Unsigned:
BIND: 37,500
NSD: 79,000
CDNS: 122,300

240,419 names Signed:
BIND: 28,000
NSD: 61,700
CDNS: 107,000

19,405,229 names Unsigned:
BIND: 57,500
NSD: 83,000
CDNS: 121,000

19,405,229 names Signed:
BIND: 25,500
NSD: 53,300
CDNS: 99,300

Note: BIND and NSD were unable to load zone size of 57,873,014 signed and unsigned names. CDNS processing peaks: 120,500 unsigned, 89,300 signed.

The Test Environment

Hardware

The choice of test equipment was to try and reflect the typical equipment that might already be in current use.

Server 1 – Sun X2100 (Client)

- 2 x Dual-Core AMD Opteron Processor 2218
- 8Gb RAM
- Dual Broadcom Corporation NetXtreme BCM5715 Gigabit Ethernet

Server 2 – Dell (R200)

- Single Dual core Intel Xeon E3120 @ 3.16GHz
- 8Gb RAM (Max capacity of server)
- Dual Broadcom Corporation NetXtreme BCM5721 Gigabit Ethernet

Server 3 – Sun X2100 (used for .COM evaluation)

- 2 x Dual-Core AMD Opteron Processor 2218
- 16Gb RAM
- Dual Broadcom Corporation NetXtreme BCM5715 Gigabit Ethernet

Software

Server software
- NSD v3.2.5
- Bind v9.7.0-P2
- CommunityDNS v2.3.2

Data

The source data used for the tests was real ccTLD and gTLD zone data.

For the unsigned queries the DO bit was not set (active), for signed DNSSEC queries the DO bit and EDNS0 was set. The test data comprised of 35% NX domains (names known not to be registered in the zone), 50% names registered in the zone for which there were no DS records, and 15% names registered in the zone where DS records were known to exist. Where a zone had a limited number of DS Records, we undertook signing of names within a zone to satisfy the 15% threshold. The created DS records were identical except for the last byte of both DS Records, which were both randomised.

The ZSK used was 1024 bits and the KSK was 2048 bits.

The list of query names were pre-generated and chosen at random for each zone under evaluation and sent from the Client server to the relevant server under evaluation (listed above).

Client and Server relationship.
Should "invalid" (incorrectly formed) queries be sent by the Client with the DO bit and EDNS0 set, the default for the CDNS' software is that these queries be ignored (not answered) by the Server thereby enabling the Server to concentrate on allocating resources to only valid DNSSEC queries. To maximise performance and to ensure the same conditions for all platforms under evaluation, the CDNS default option was set and no incorrectly formed queries were used in the list of query names in the BIND, NSD, and CDNS evaluation.

The CDNS option to record the source IP address of such "invalid" DO bit and EDNS0 set queries, (which may be the source of a Denial of Service attack by consuming the available return bandwidth) – this function was not activated nor needed as only correctly formed queries were sent to the Server/platforms under evaluation.

Tests Conducted

For this study the following, identical tests were performed on BIND, NSD and CDNS platforms.

Throughput: CPU utilisation / Queries per second
Capacity: CPU utilisation / Number of Resource Records in Zone
Resilience: Ability to withstand a DNSSEC Denial of Service attack.
Efficiency: Software efficiency in resource management

Summary test environment considerations

Initially the unsigned test were conducted using a 100Mbps switch, which whilst coping well with unsigned responses, when DNSSEC answers were required the throughput topped out at between 13,000 and 18,000 signed answers/second depending on the number of Resource Records being returned. This was the throughput ceiling of the 100Mbps switch handling the (return) DNSSEC answers. As a result, we then had to re-run both the signed and unsigned tests using Gigabit switches and Cat 6 cabling and in so doing removed the previous DNSSEC throughput ceiling.

It has been noted that when fully stressed, BIND and in some instances NSD consumes 100% of the CPU. Whilst CDNS was able to achieve higher levels of throughput and response, at best it was only able to consume 89% of available resources. Further evaluation will be conducted as to where that performance bottleneck may be arising. Ideas/suggestions welcome.

Zones containing a varying number of Resource Records were chosen; specifically 7,691; 240,419; 19,405,229 and 57,873,014. The zone sizes were chosen to reflect:

- Small TLD or Large Corporate zone file: 7,691 records
- Small TLD zone file: 240,419 records
- Midsize TLD zone file: 19,405,299 records
- Large TLD zone file: 57,873,014 records

The Client used in all tests was a SUN X2100. The Server used in all tests except the 57,873,014 (.COM zone) was the Dell R200 (as specified as Server 2 above). For the 57,873,014 resource record test, the server used was a Sun x2100 (as specified as Server 3 above).

Whilst the CDNS platform was able to load the large zone (.COM zone data), we were unable to get BIND and NSD to load the .COM zone on the hardware under evaluation.

The results were as follows:

Small TLD or Large Corporate Zone
7,691 records

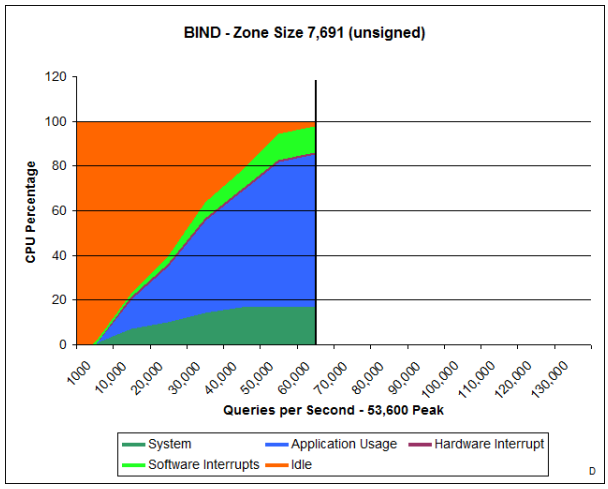

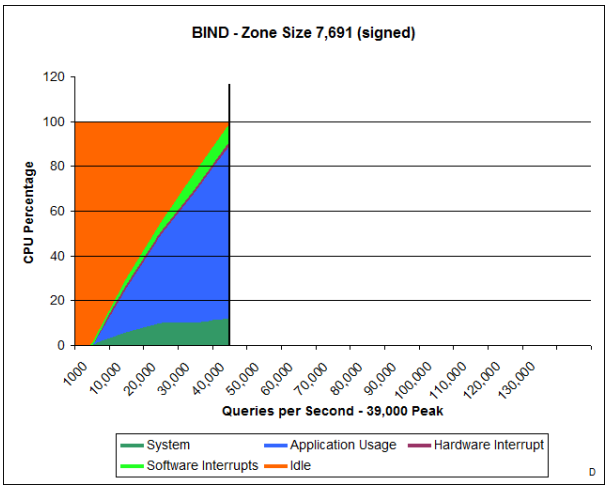

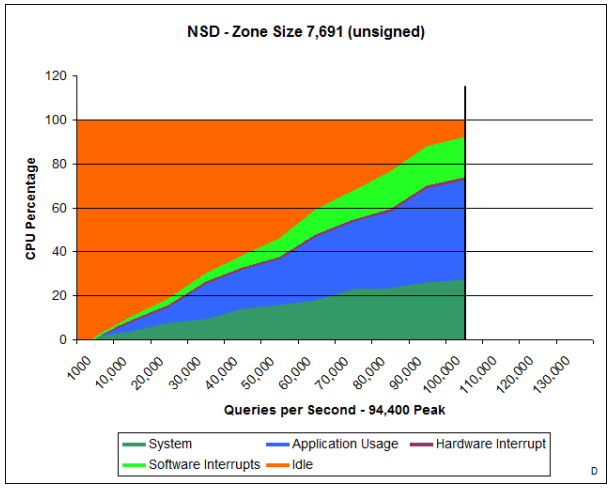

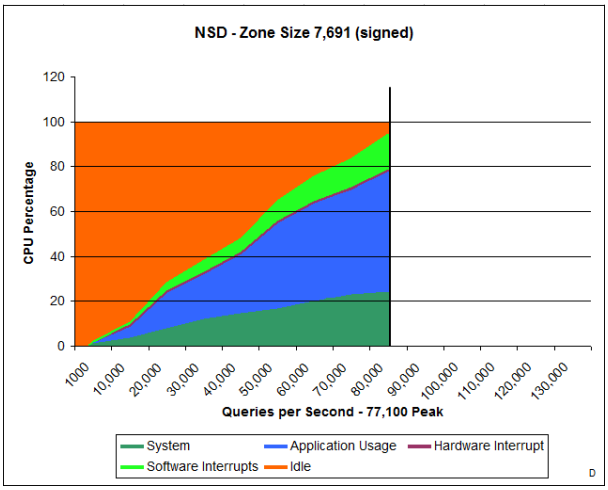

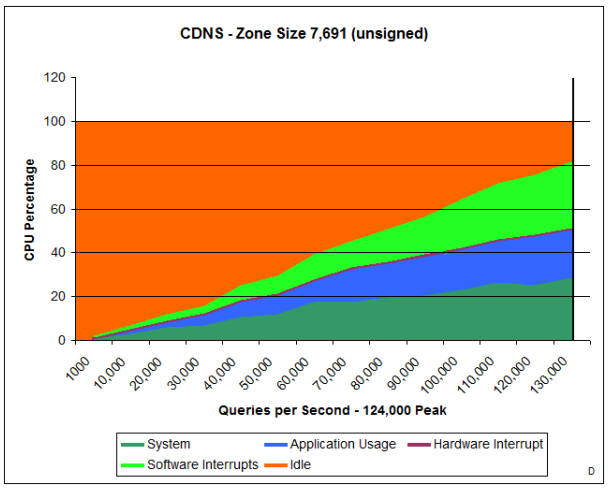

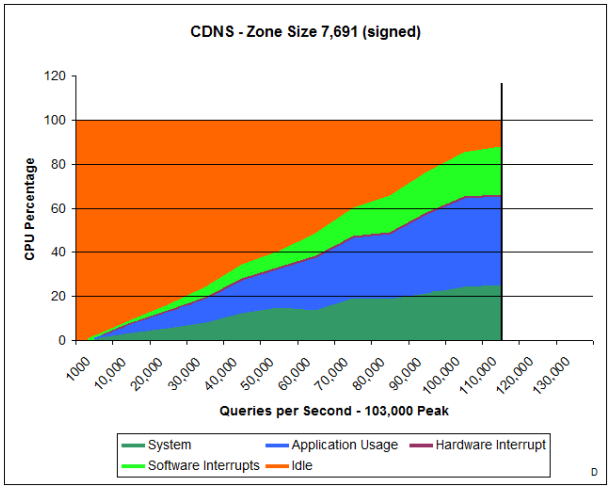

Summary - Small TLD or Large Corporate Zone
7,691 records

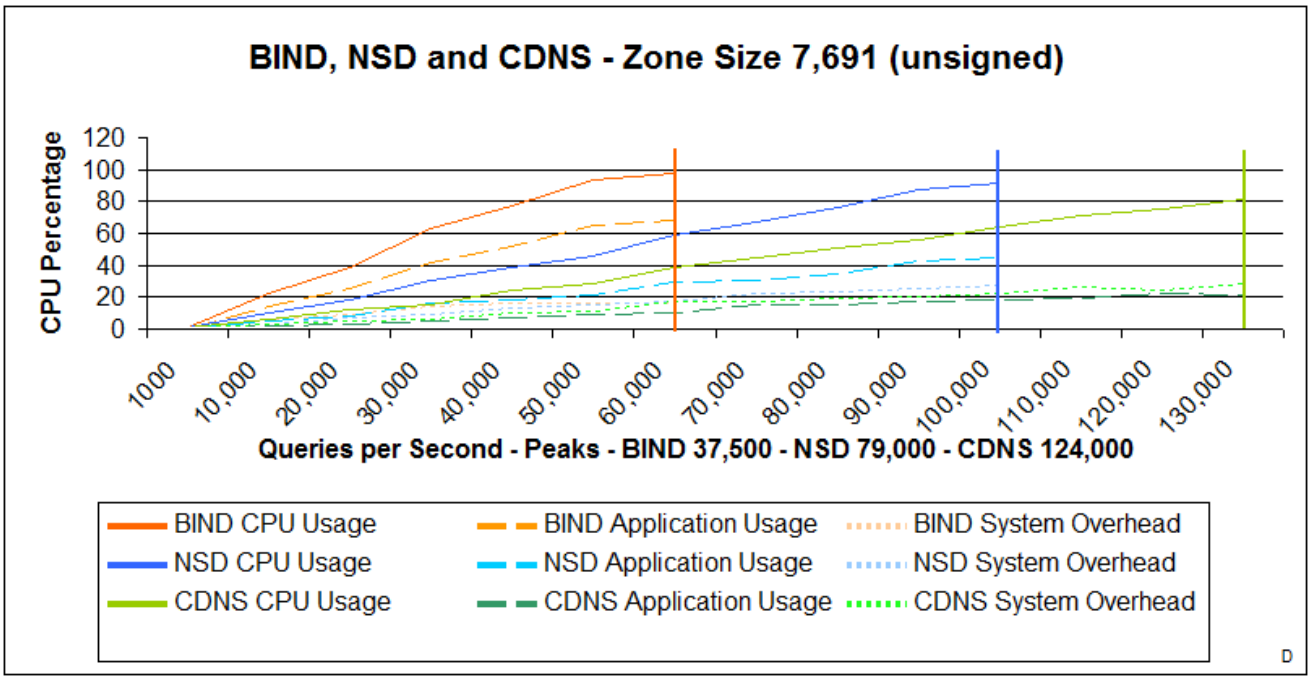

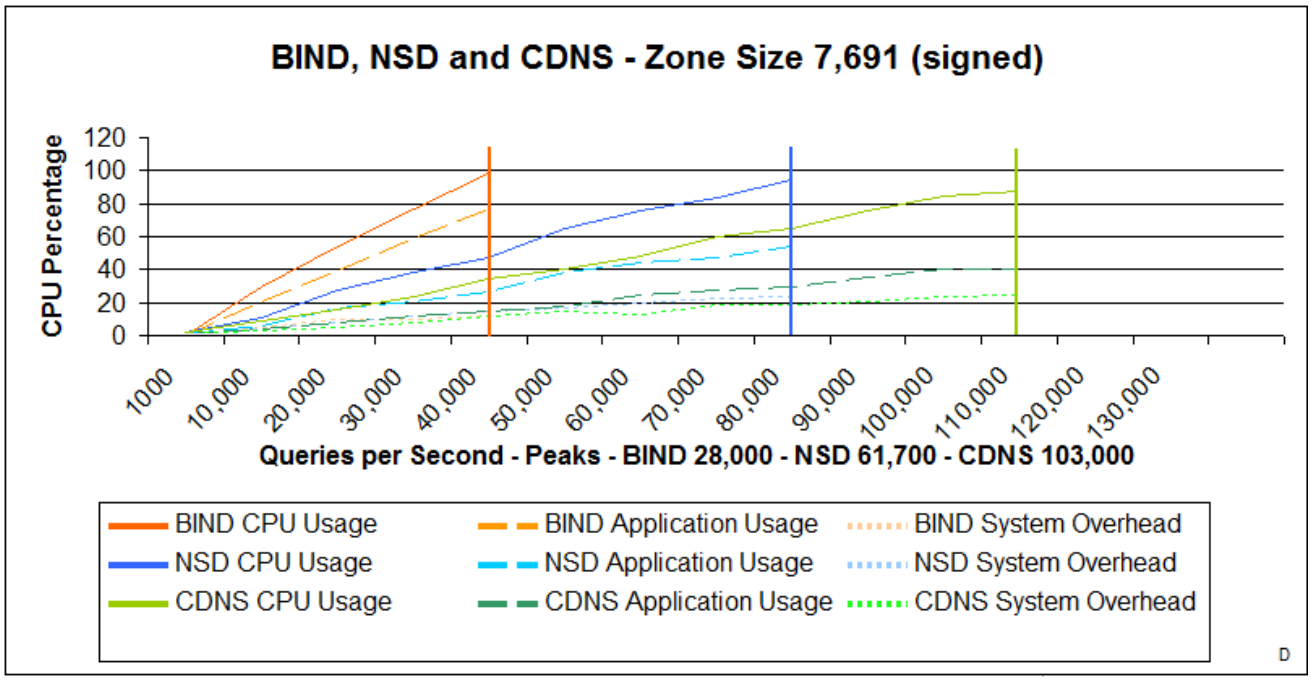

Summary - Small TLD or Large Corporate Zone
7,691 records

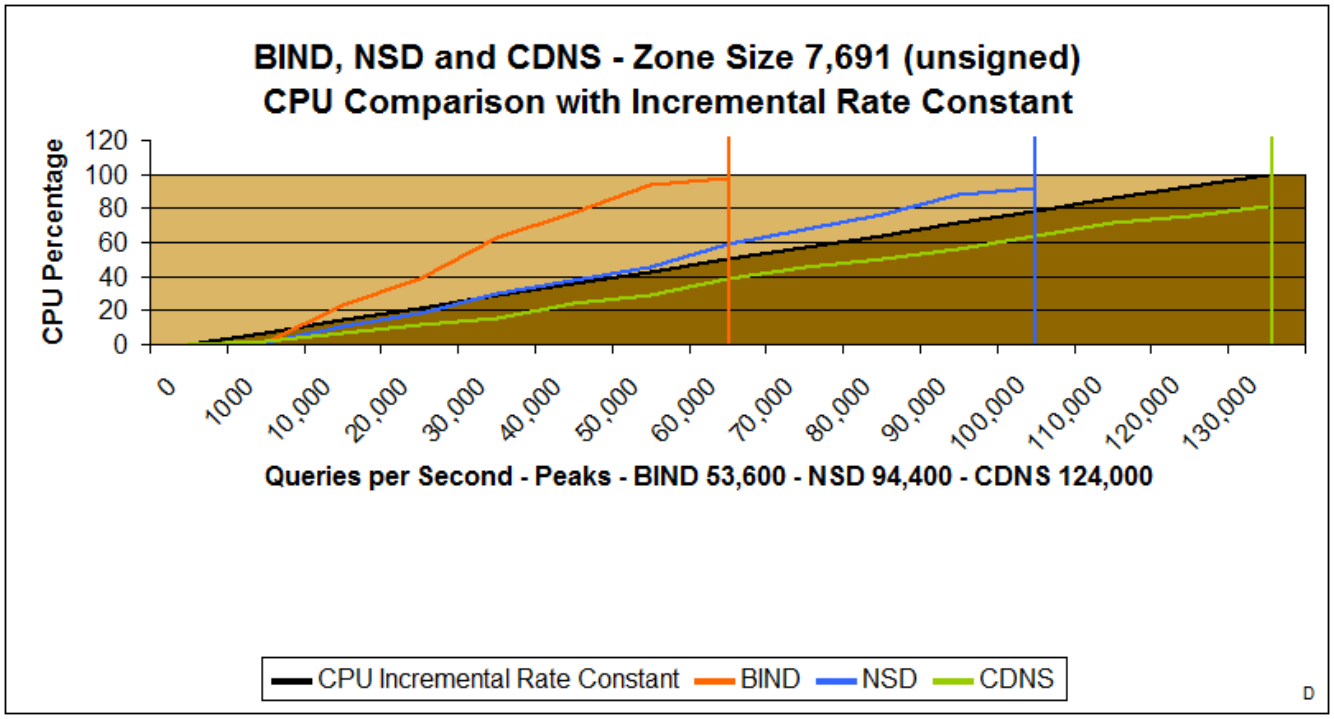

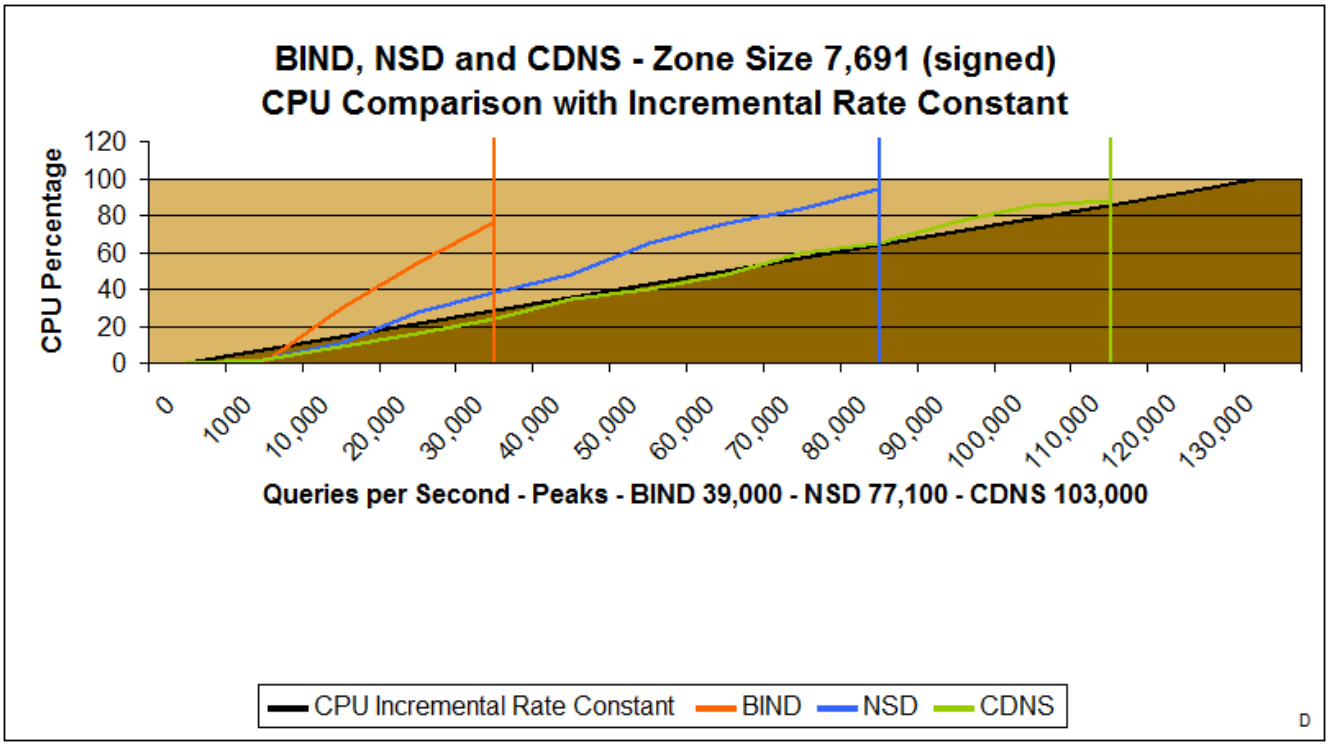

Small TLD Zone performance evaluation
240,419 records

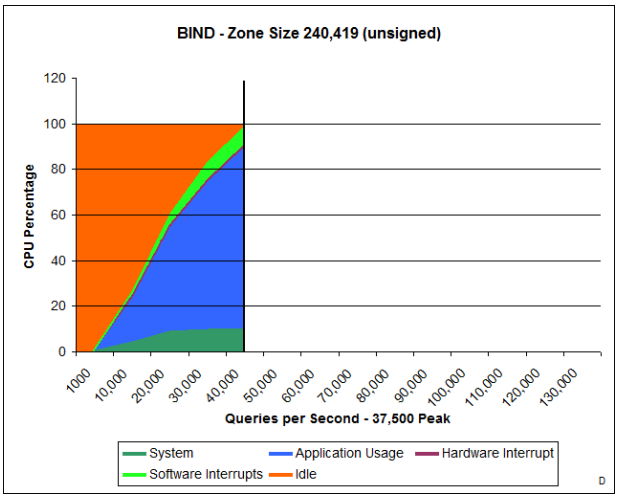

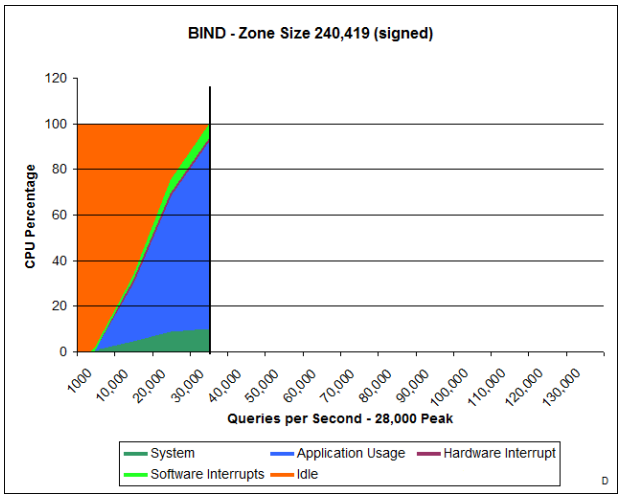

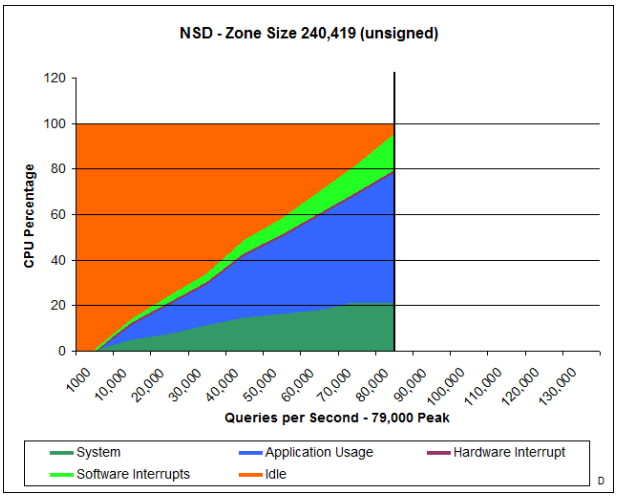

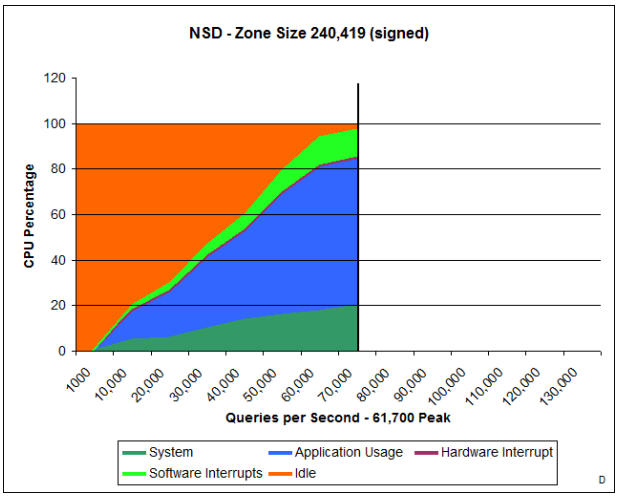

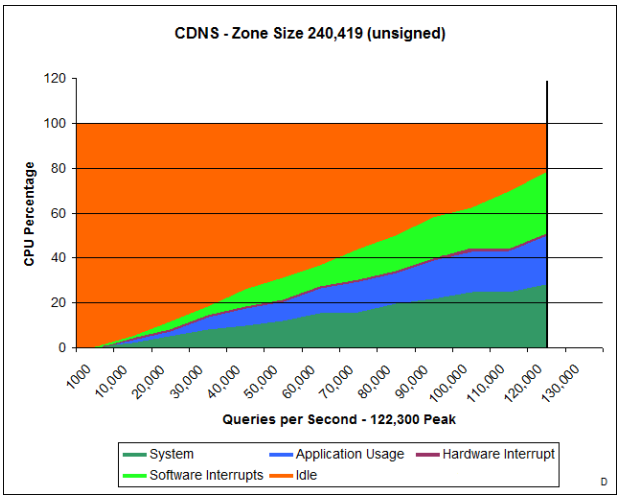

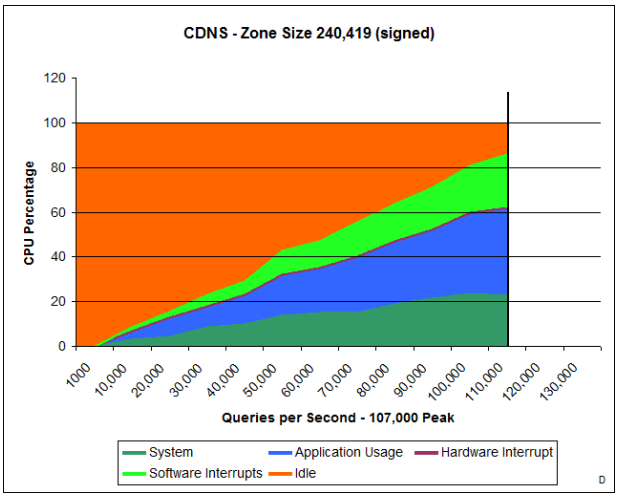

Small TLD Zone performance evaluation
240,419 records

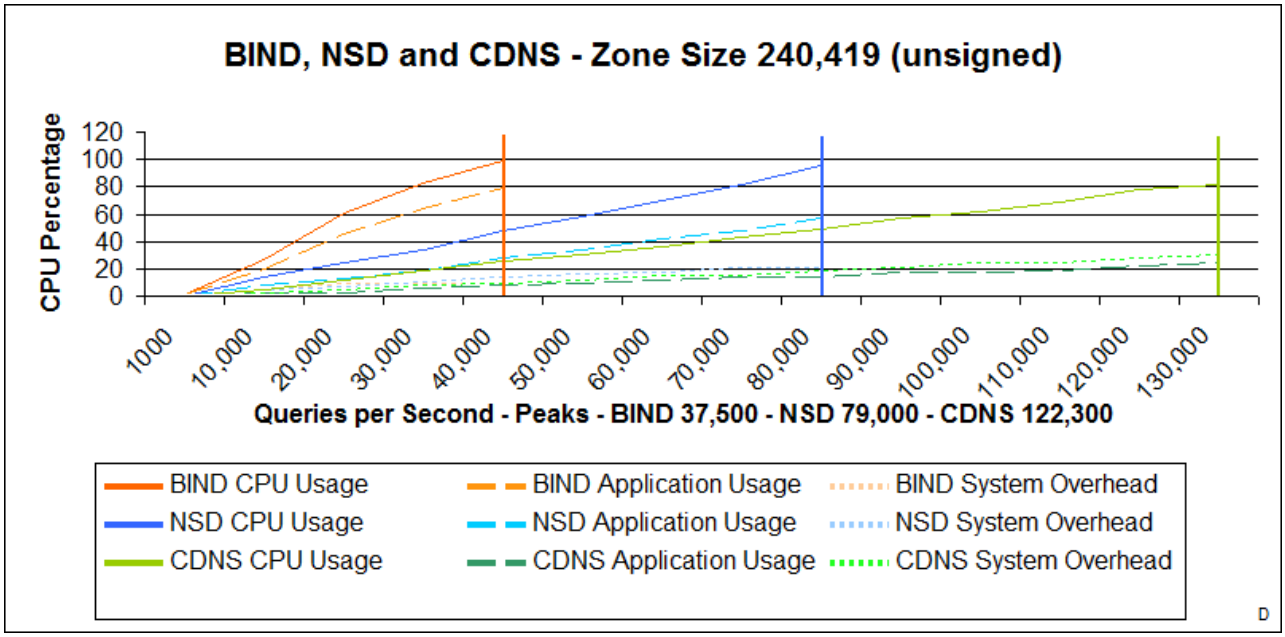

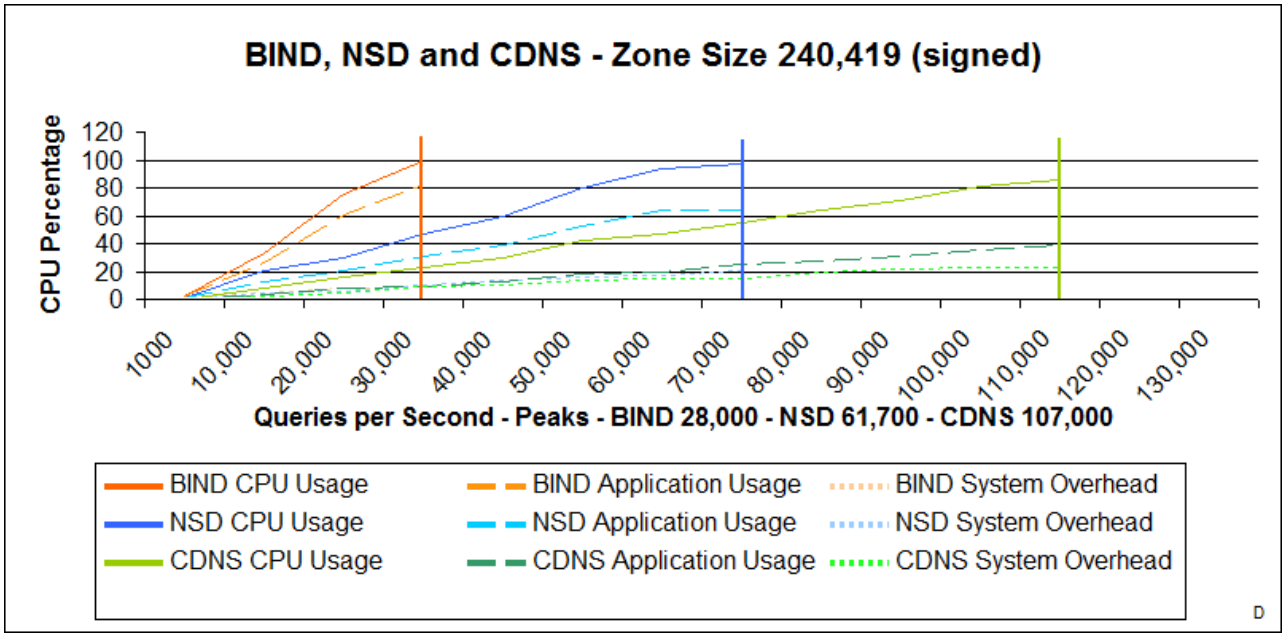

Summary - Small Zone software CPU/Data Set efficiency
240,419 records

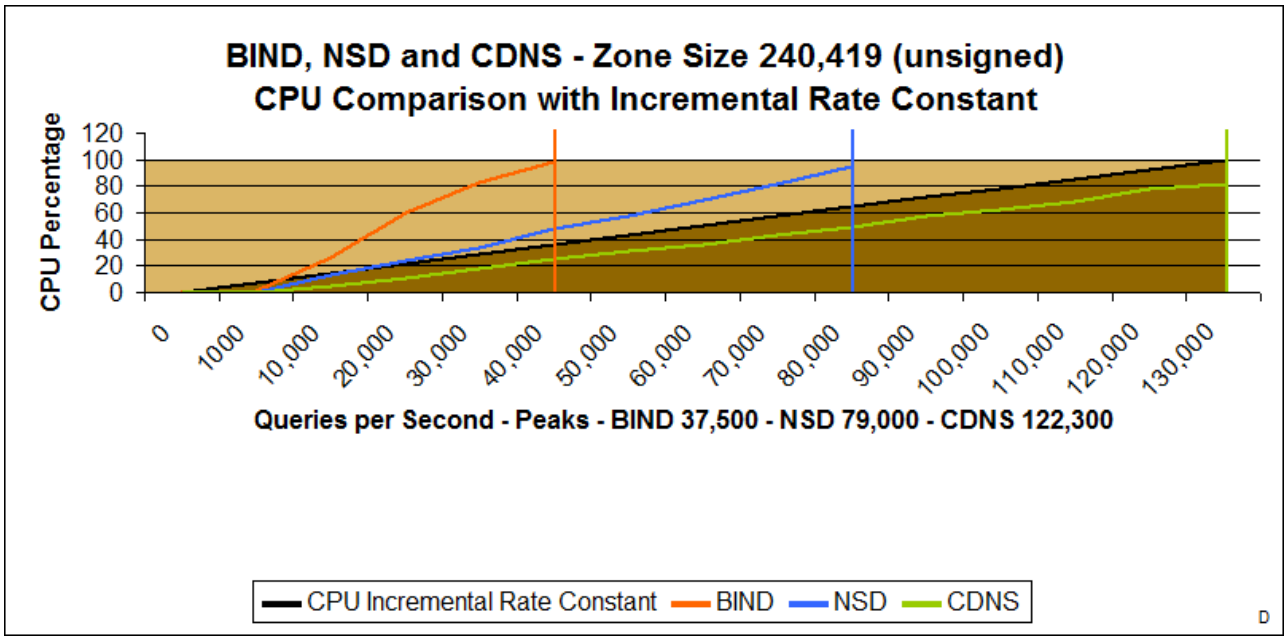

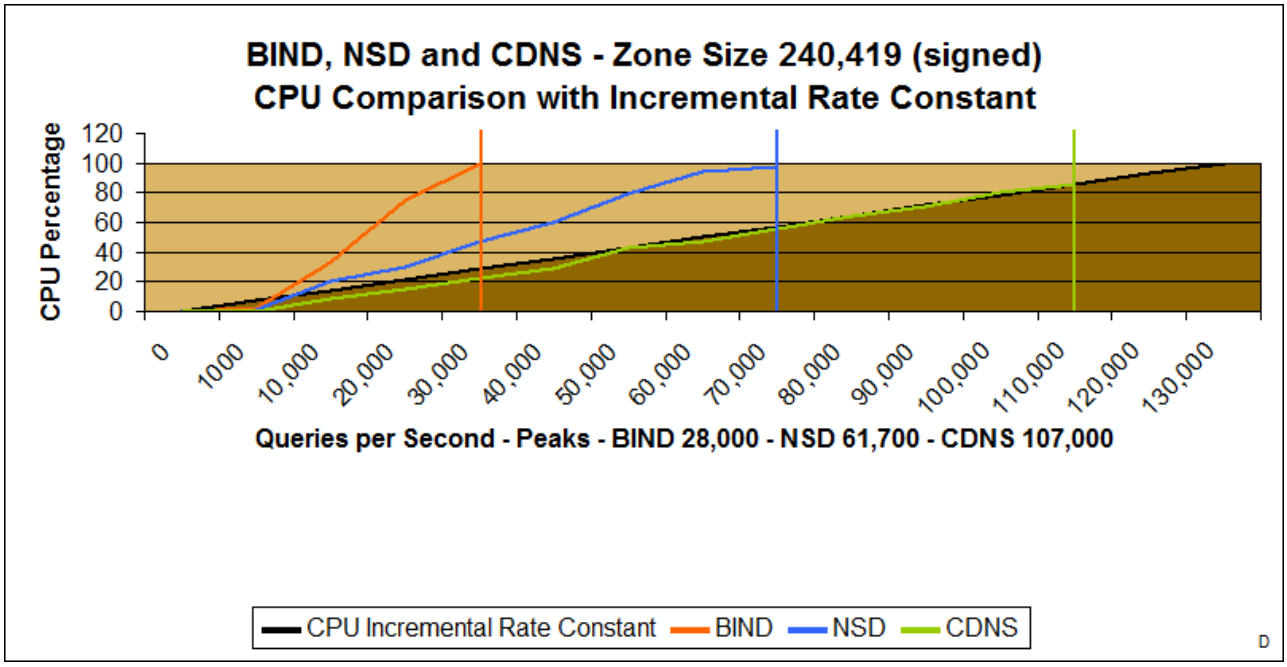

Mid-Sized TLD Zone performance evaluation
19,405,229 records

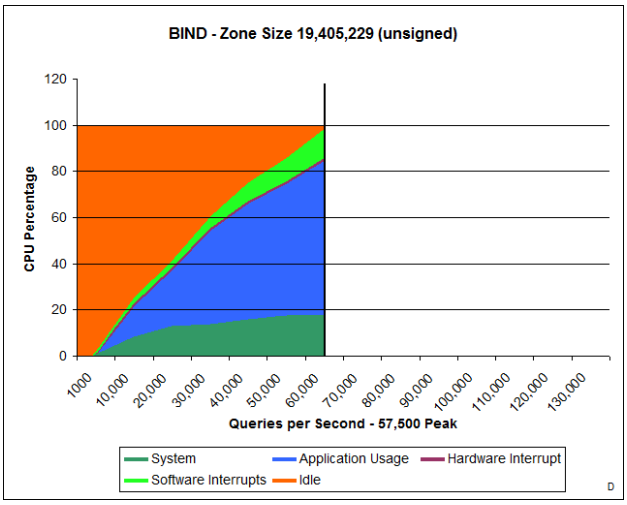

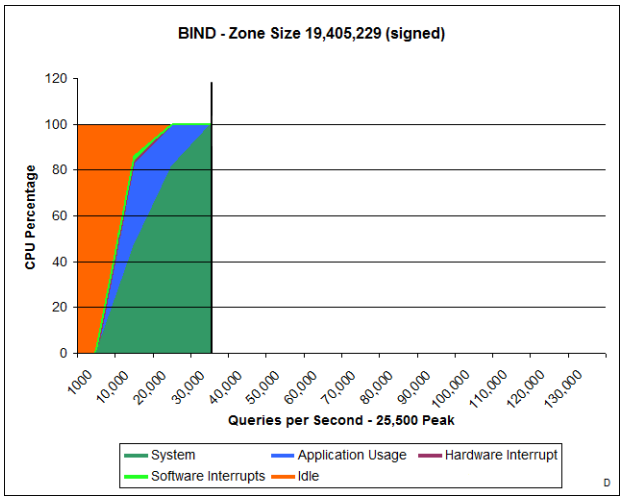

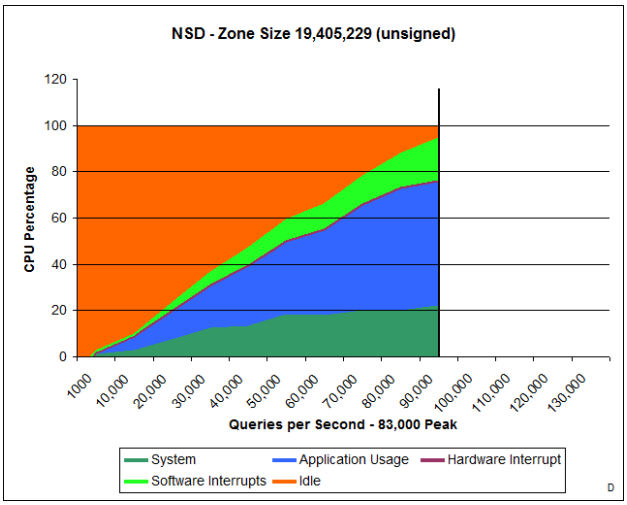

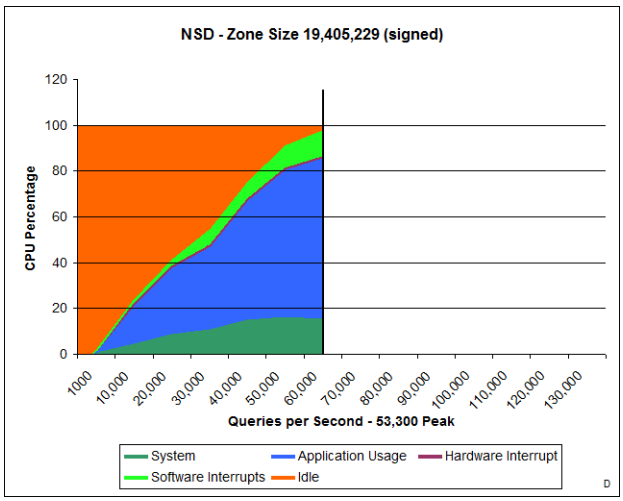

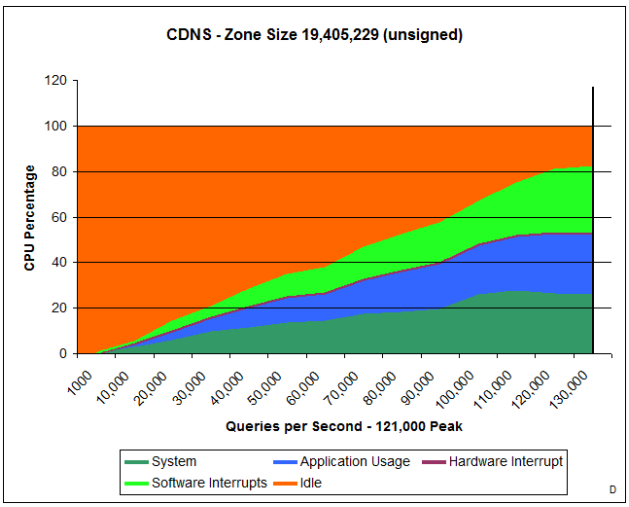

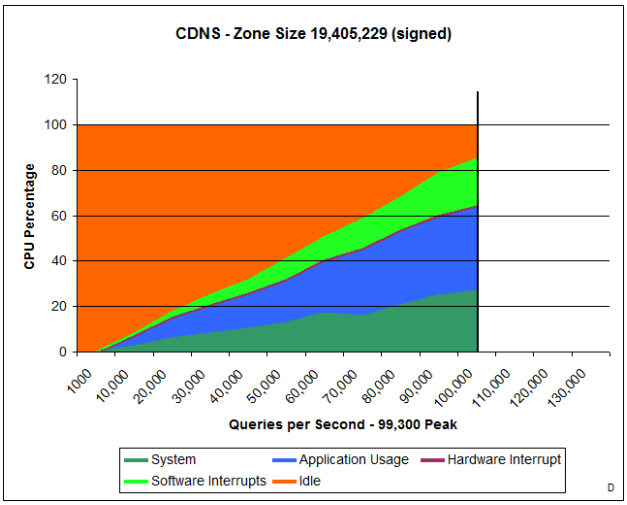

Mid-Sized TLD Zone performance evaluation
19,405,229 records

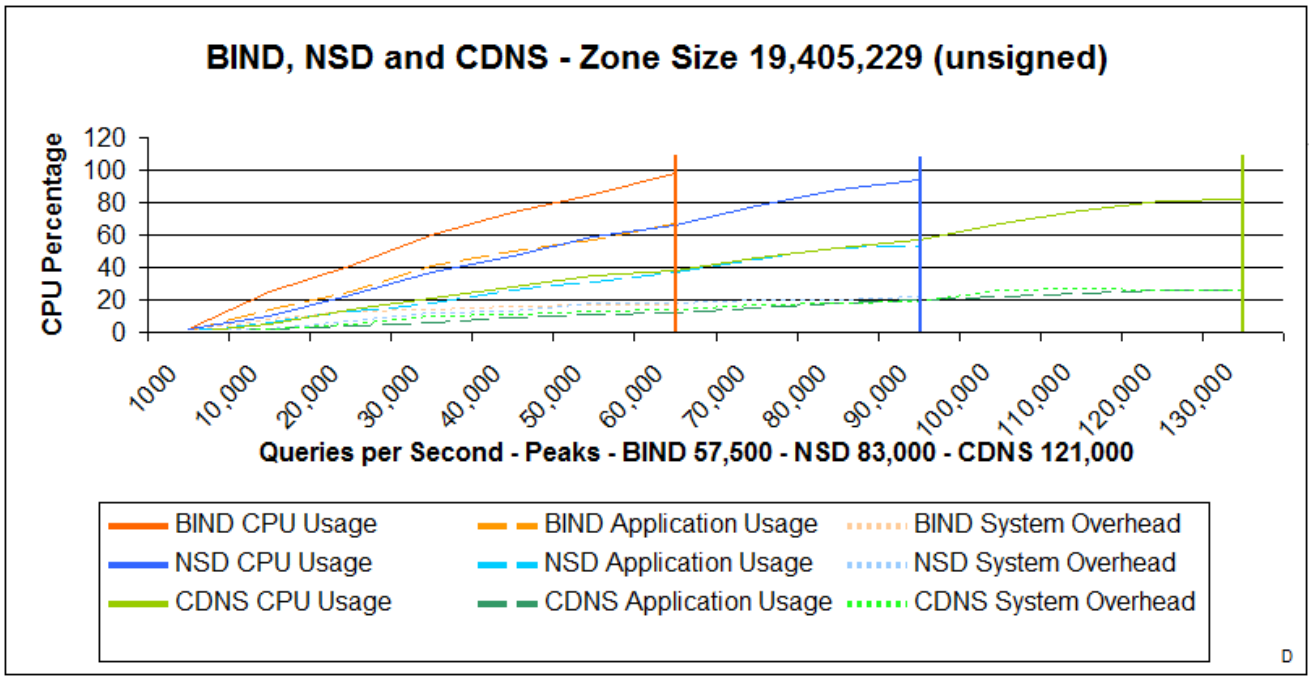

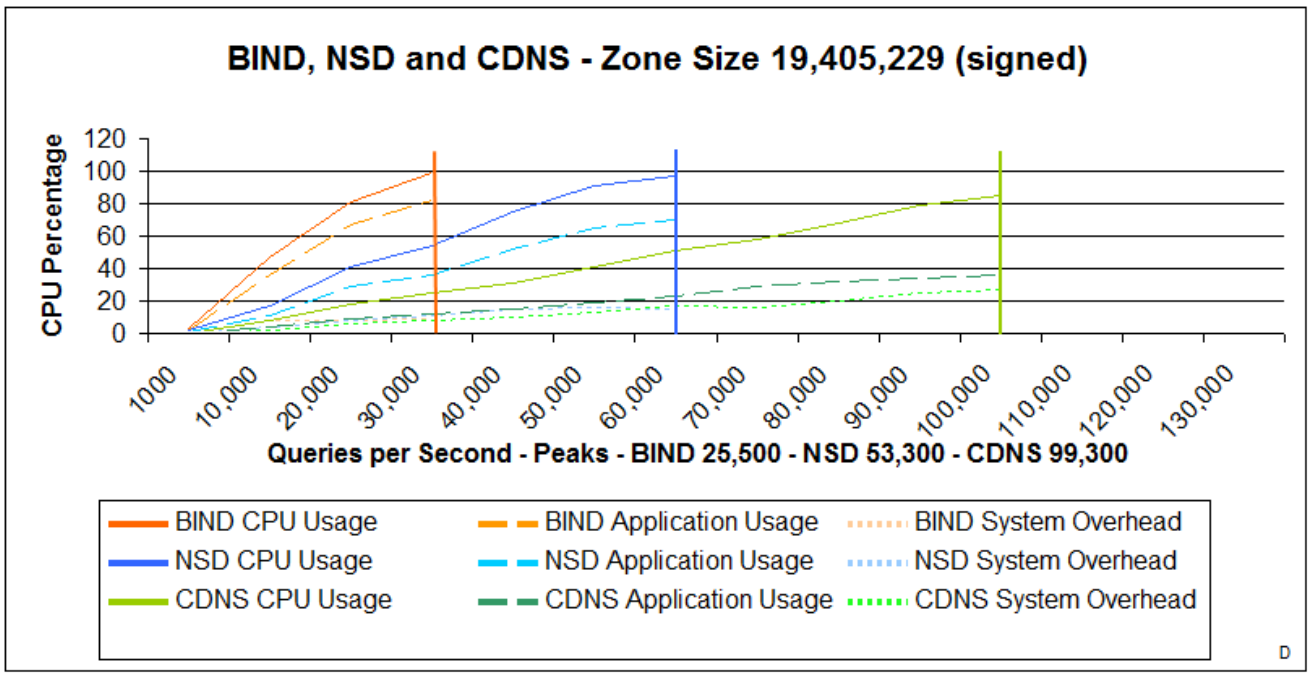

Summary – Mid-Sized Zone software CPU/Data Set efficiency
19,405,229 records

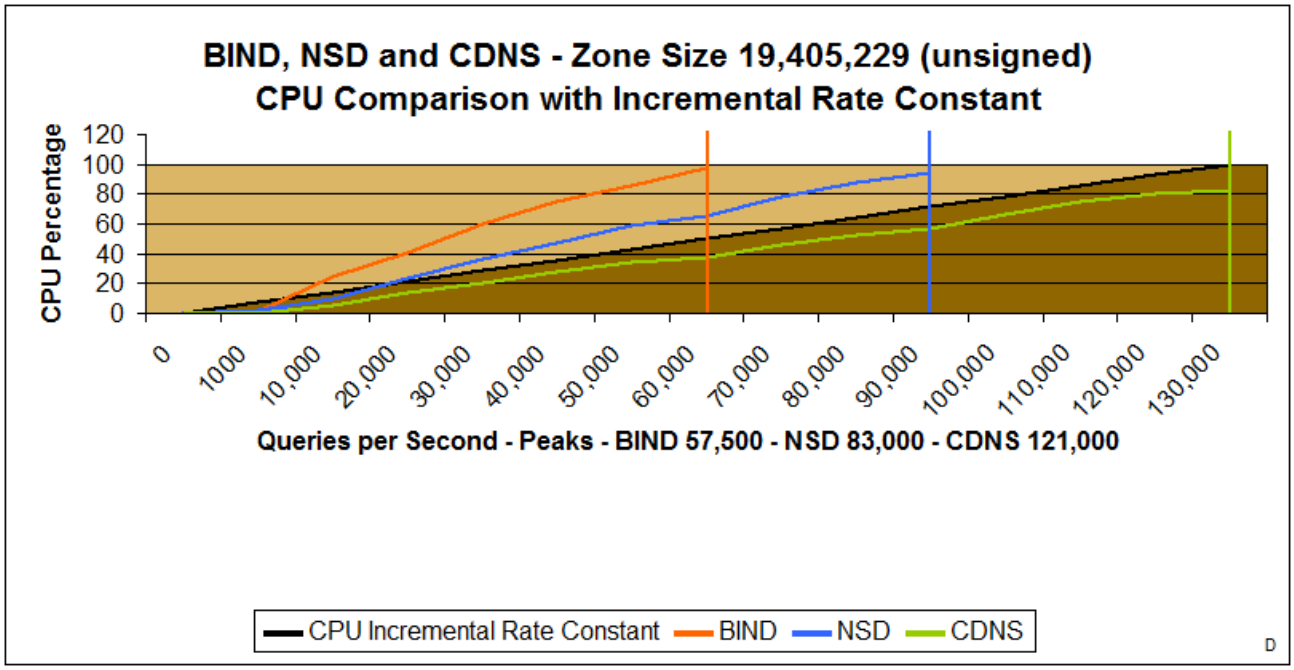

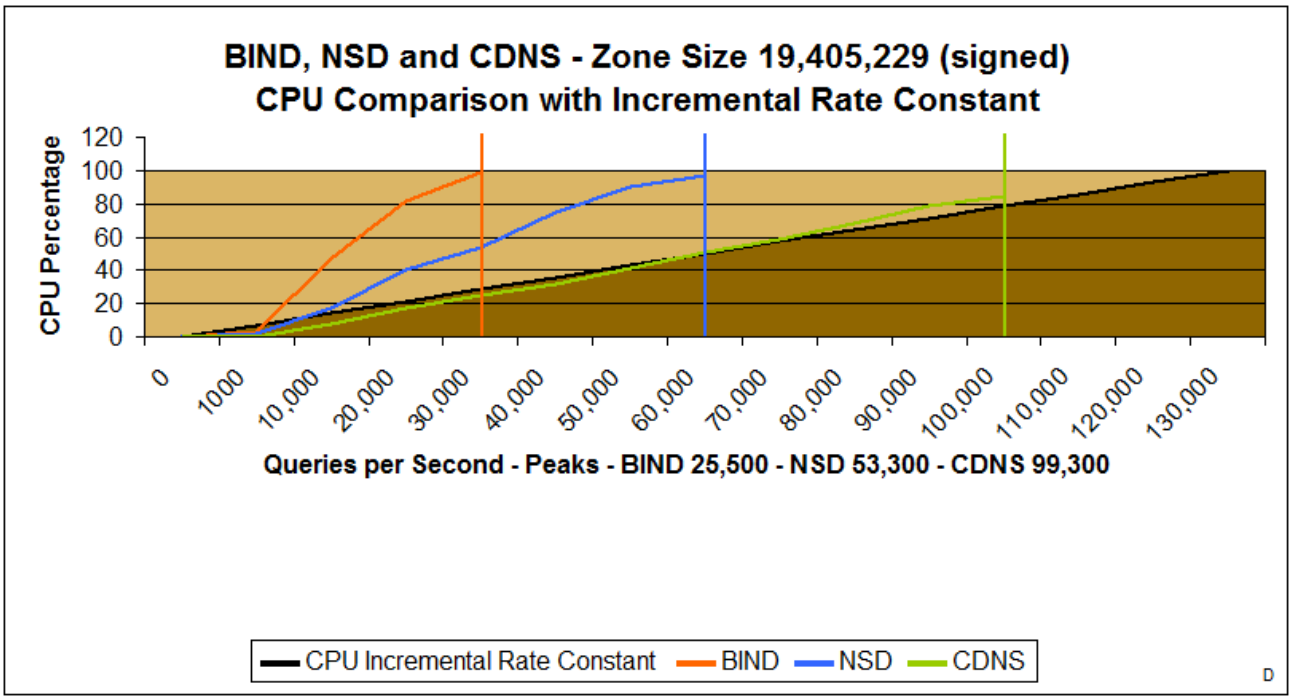

Summary – Large-Sized Zone software CPU/Data Set efficiency
57,873,014 records*

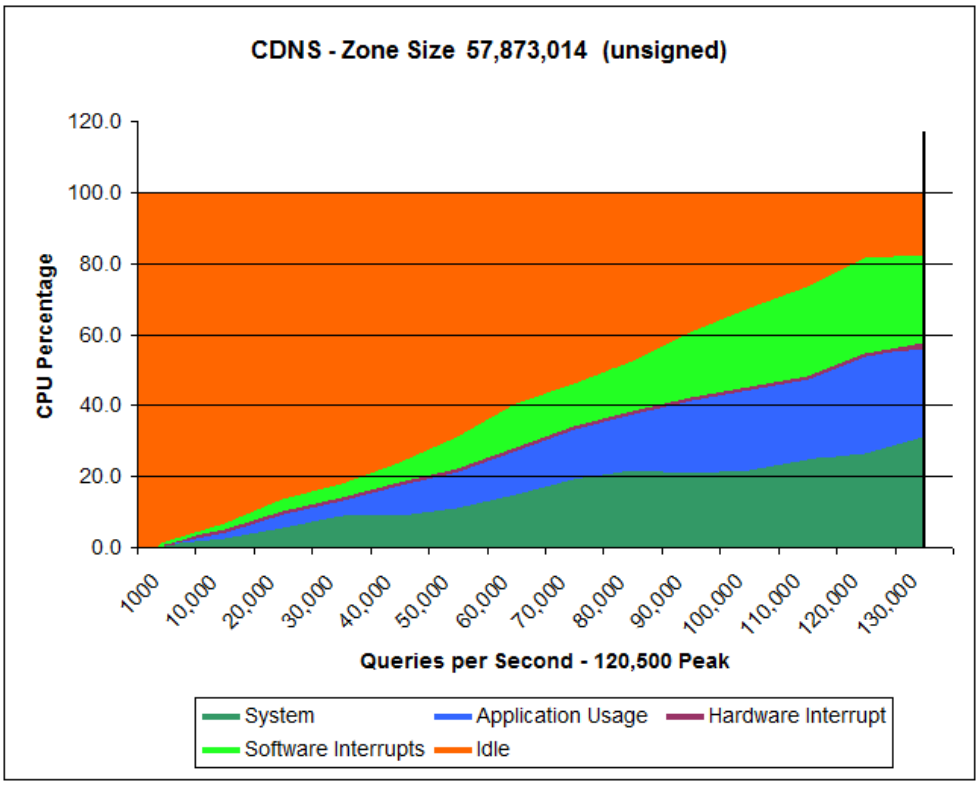

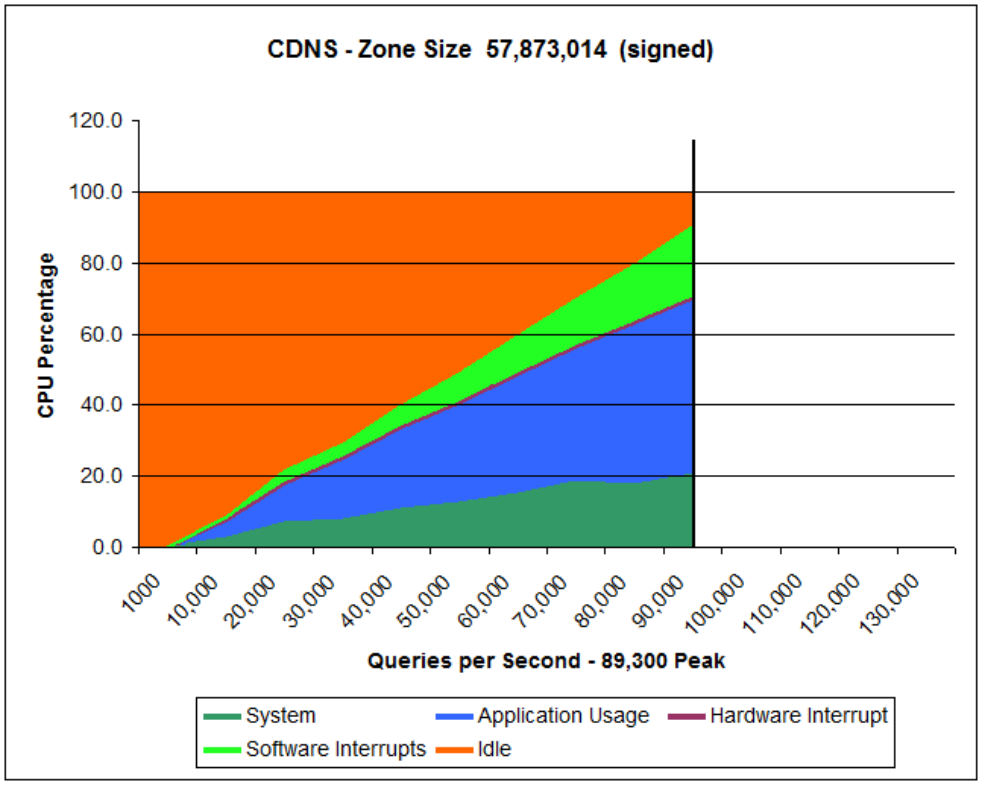

*** BIND and NDS were unable to load the zone file of this size**

Summary – Large-Sized Zone software CPU/Data Set efficiency
57,873,014 records*

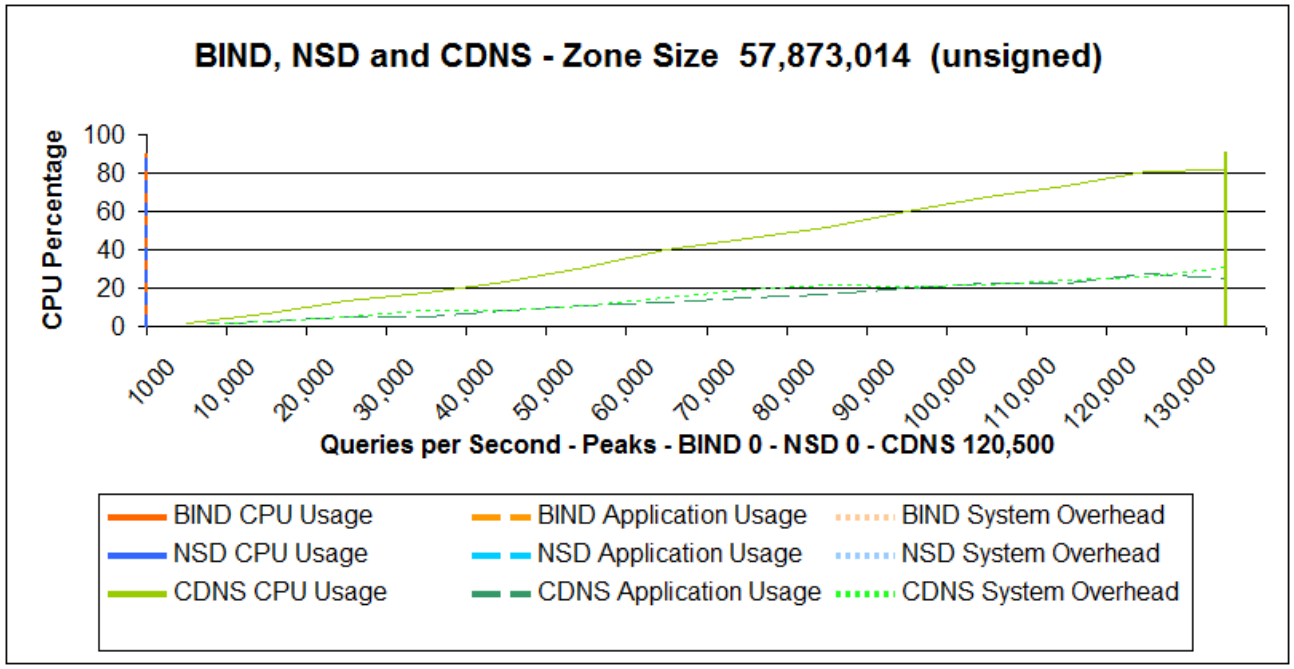

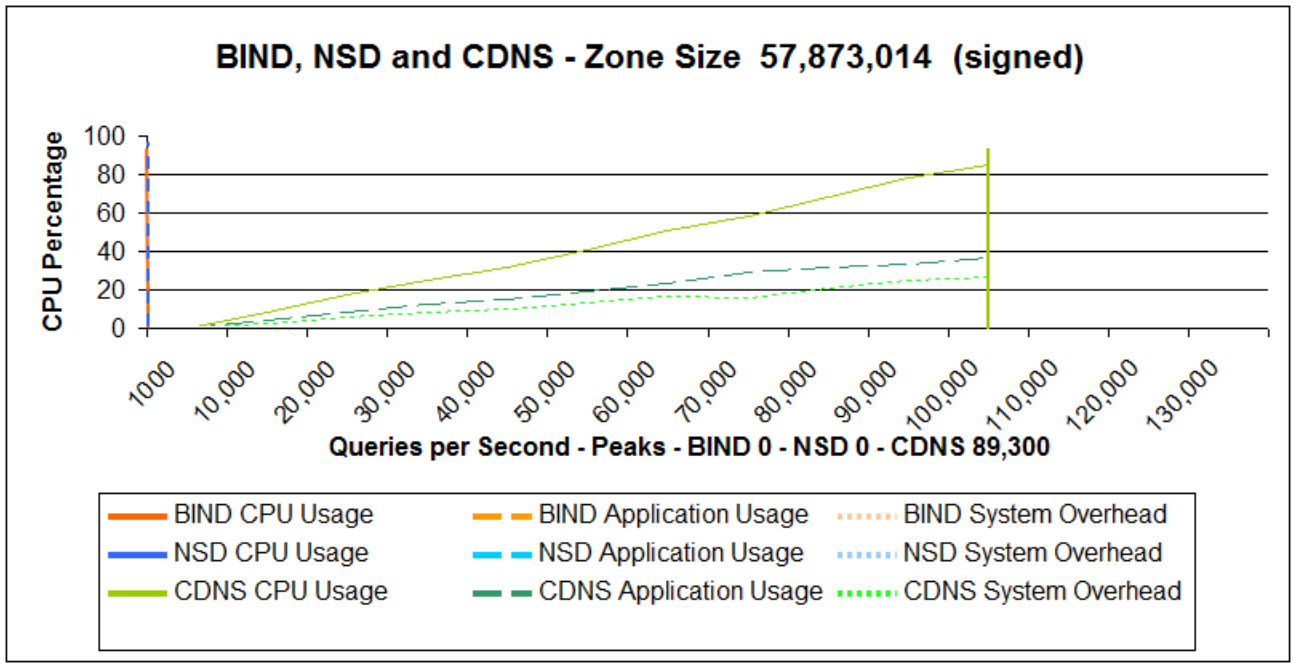

*** BIND and NDS were unable to load the zone file of this size**

Summary – Large-Sized Zone software CPU/Data Set efficiency
57,873,014 records*

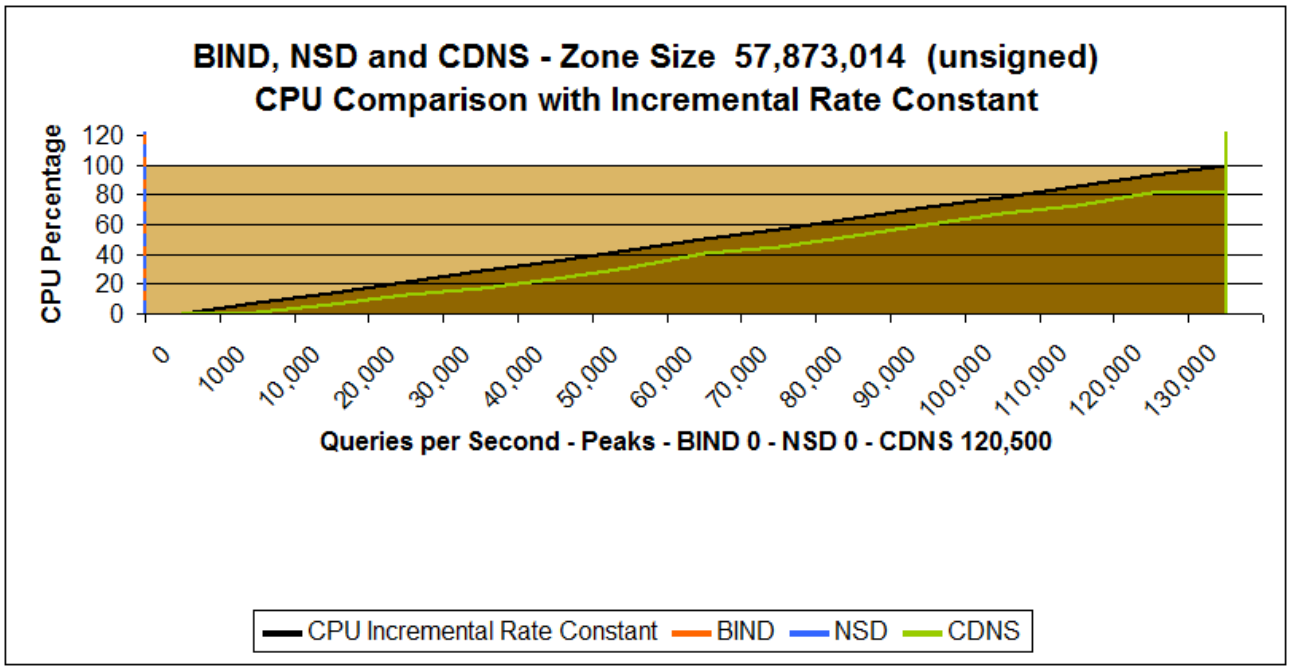

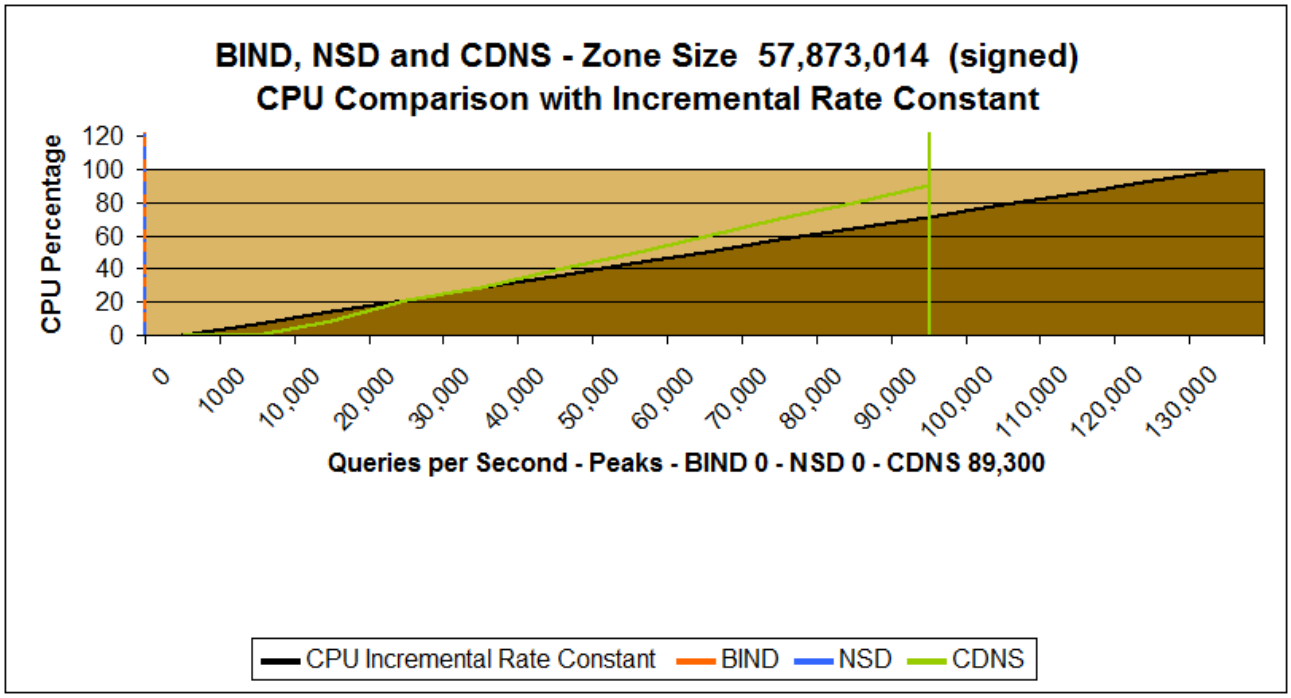

*** BIND and NDS were unable to load the zone file of this size**

About CommunityDNS

The global Anycast DNS provider with the network engineered for security, optimized for speed and designed for resilience, CommunityDNS provides:

Security:
- An extensive global DNS Anycast network with platforms hardened to military specifications.
- Extensive encryption throughout CommunityDNS' complete Anycast infrastructure.

Speed:
- 8 times faster than BIND
- 10 times faster than NSD
- Over 11 times faster than Oracle's fastest database on reads.
- Over 4 times faster than Oracle's fastest database on writes.
- Never impacted by volumes associated with DoS/DDoS attacks, always delivering legitimate queries.
- The sheer speed of CommunityDNS' secure technology provides the speed and capacity essential for mitigating attacks and increased infrastructure demands associated with DNSSEC implementation.

Resilience:
- Maintained 100% uptime since first providing DNS services in 1996.
- Provides 99.999% SLA support.
- Stability of platform performance and capacity, currently supporting 70% of today's Internet.
- Successfully tested to authoritatively support 500,000,000 domains.
- Successfully tested to support 585,000,000,000 queries per day.

Serious about security, resilience and vision CommunityDNS fully supports:
- IPv6
- DNSSEC (NSEC, NSEC3, NSEC3+OptOut)
- IDNs (Internationalized Domain Names)
- Onsite authoritative resolving of 125,000,000+ domains
- Reverse lookup
- Helping ISPs identify client virus/bot infection
- Name server aliasing
- Complete white-label services
- Updates through standard AXFR, IXFR (with or without TSIG) and FTP over VPN.
- Real time monitoring information
- Automated process controls preventing loading of flawed zone files.

CommunityDNS – When speed, scale, capacity and always-on matter!