WHITEPAPER

Infoblox 4030 DNS Caching Appliance

Assessing Performance Metrics for a New Generation of Recursive DNS Servers

Assessing Performance Metrics for a New Generation of Recursive DNS Servers

ISPs are facing an ever-increasing volume of DNS queries from customers. This is due to a "perfect storm" of factors: the proliferation of mobile devices, including smartphones and tablets, and the widespread adoption of virtualization and cloud computing. Luckily, a new generation of recursive DNS servers has been developed to address the need for higher query performance.

However, current methods for measuring the performance of recursive DNS servers can cause misleading assessments of this next-generation technology. A new interpretation of the current means of measurement must be adopted to account for and acknowledge the dynamic, high performance capabilities of these new recursive DNS servers. This revised approach will allow evaluators to fully appreciate these servers' ability to continue serving queries from cache even when their capacity to serve queries that require recursion is reached.

The Essential Role of Recursive DNS Servers

The Domain Name System, or DNS, is the Internet's hierarchical, distributed naming system. DNS maps "domain names" to a variety of types of data, including IP addresses. Most resources on the Internet are identified by domain names. DNS distributes the responsibility for administering domain names, allowing organizations around the Internet to manage a set of related domain names locally while making information about them available throughout the Internet. These related domain names are grouped into administrative units called zones.

The Internet's DNS zones are organized hierarchically, with a single "root" zone at the very top, top-level zones one step below it, second-level zones one step below them, and so on. Each zone contains information delegating responsibility for the zones below it, called subzones, to a set of DNS servers. Thanks to this system, a DNS server can begin the search for information about any domain name at the root zone, following a succession of delegations to identify the DNS servers with information about the zone that contains that domain name. One of these DNS servers responds with the information about the domain name being sought — for example, the IP address of a given web site.

The Internet's DNS servers generally perform one of two functions: they have complete, "authoritative" data about one or more zones and answer queries from other DNS servers about those zones, or they query other DNS servers to find answers on behalf of clients and cache the results. The latter variety of DNS servers are referred to as "recursive" DNS servers, after the method they use to answer clients' queries, recursion.

When a user types a domain name into a web browser or a mobile user runs an app, a query is sent to a recursive DNS server, often run by the user's ISP. The name server uses a process called recursive name resolution, or just "recursion," to find the answer to the query. The name server first looks in its own cache to see if it already knows the answer to the query. If it does, the name server simply returns it. If it doesn't, the name server checks its cache for cached data "close to" the answer. For example, if the query received requested the address of www.google.com, the names of the name servers for google.com would be close to the answer.

If nothing helpful is found in the cache, the name server starts looking for the authoritative server that has the answer, starting by asking one of the root servers. The root server holds the locations of the authoritative servers for various top-level zones, for example, .com, .net, .org, etc. The root returns a list of the authoritative name servers for the top-level zone, and the recursive server then sends the query to one of those authoritative servers. This process continues until the recursive name server finds the set of name servers authoritative for the zone that contains the answer. These return either the answer or an indication that the answer isn't available, such as an error indicating that no such domain name exists. Not surprisingly, the number of queries the recursive name server sends can add up pretty quickly, especially for domain names not frequently looked up.

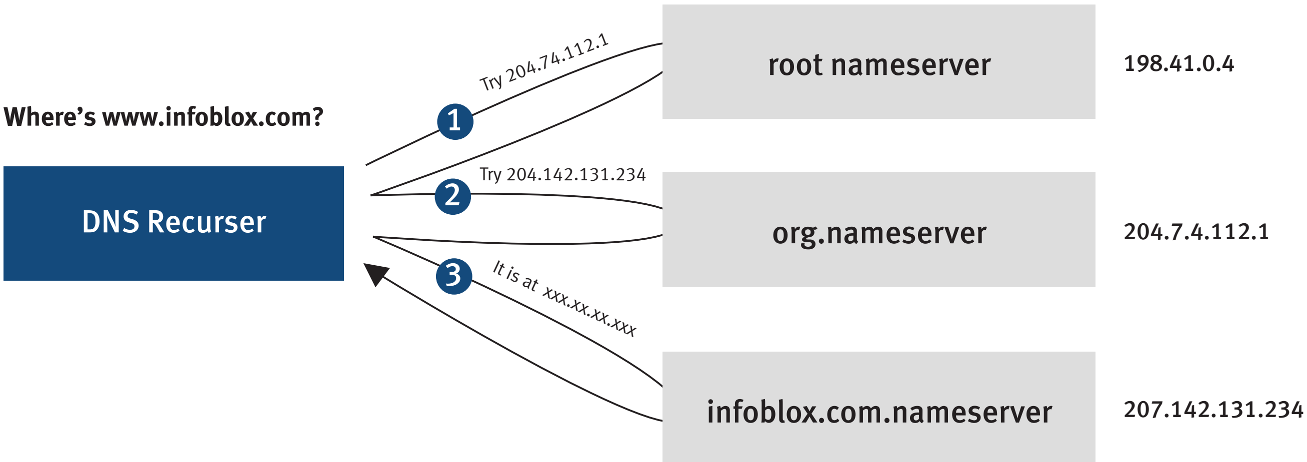

Figure 1: In this example, a recursive name server, labeled "DNS Recurser," consults three authoritative name servers successively to resolve the address of www.wikipedia.org.

Written at the University of California at Berkeley in 1984, the BIND (Berkeley Internet Name Domain) name server was one of the first DNS implementations, and it's still used extensively today on the Internet. BIND name servers can process a limited number of recursive queries in a given time interval, as well as maintain a limited cache of the answers to a number of frequently asked queries.

All Responses Are Not Created Equal

A cached response is one whose answer is stored in the server's cache, while a non-cached response is one that requires tracking down an authoritative name server for the requested domain name. Predictably, the response time for a cached response is generally much faster than for that of a non-cached response. Since a BIND recursive name server can process only a given number of queries at one time, the number of queries it can process at a given time will be much greater if most of them are served from cache.

The processing capacity of a BIND recursive name server dictates how many queries it can process at a given time. When the number of queries begins to exceed the processing capacity, the server will first queue all incoming queries until a limit is reached, and then it will drop all future queries until queries currently being processed are completed and capacity is available. The effect is that BIND recursive name servers sometimes — and, in some cases, depending on capacity, often — temporarily stop responding.

For an ISP, enterprise, or any organization dependent on its name servers for business functions and customer service, these temporary shutdowns can result in mission-critical content being unavailable. Depending on the timing of this unavailability, the results can vary from the merely annoying to the catastrophic.

Clearly, a new generation of recursive name servers is needed — especially in light of the unprecedented increase in the number of devices connected to the Internet and the subsequent increase in the load on recursive name servers. For example, many popular mobile apps send a dozen or more DNS queries when they start. Traditional BIND name server technology simply cannot keep up in this new environment.

Variables Affect Query Response Time

The challenges to a new generation of recursive name servers go beyond the mere volume of queries. Query response time is affected by a number of variables. Some of these are limited to the server itself, while others lie outside the server's control.

- One of the latter is *Internet latency*, which is caused by network congestion anywhere in the path between the recursive name server and any of the authoritative name servers that are involved in answering a query. As Internet latency increases, the capacity of a name server to answer non-cached responses decreases. An increase in congestion near the recursive DNS server can also lead to queries being dropped.
- *Time to Live (TTL)* is how long a particular set of resource records may be cached on a recursive DNS server. TTLs can be of any duration, but generally these are short-lived (from a few minutes to a few hours). This enables administrators to change the records on the authoritative name servers and propagate those changes relatively quickly.
- *The Cache Hit Ratio (CHR)* measures the ratio of queries that are answered by cached responses to the total number of queries, and is generally considered during testing to be a static figure.

All these factors affect the performance of recursive name servers. But in order to evaluate a new generation of recursive servers — along with the way cached responses are handled, or not handled, once processing capacity is reached — CHR is the single variable whose interpretation most needs to change.

Flaws Hinder Determination of the True Cache Hit Ratio

Current methods of determining the CHR for BIND recursive name servers provide misleading results. In the lab or at a product presentation to an ISP, a server manufacturer names a fixed CHR, typically 85 to 90%, and then proceeds to test the maximum query rate the server can process before it stops responding.

This methodology misrepresents the performance of the name server being tested. Besides not being evaluated in a real-world network, testing against a “fixed” CHR does not account for the way recursive name servers are being used in today’s deployments.

For example, while the number of Internet users has increased dramatically in recent years, most traffic connects to a small number of sites, such as social media sites and search engines. The resulting CHR is dynamic and changes from day to day and even hour to hour, depending on both time and current events. For example, the Super Bowl halftime show this year produced an unprecedented surge of DNS requests as fans looked online for merchandise and news updates using their mobile devices.

Clearly, in this new epoch in the evolution of IT, a new approach is needed to measure the growing number of queries in light of the rapidly fluctuating query patterns brought on by changing consumer behavior and the proliferation of Internet devices.

The Next Generation of Recursive Name Servers Is Dynamic

Drawing upon its expertise in DNS technology, Infoblox is creating the next generation of recursive name servers. Created specifically for ISPs, large enterprises and telcos, these carrier-grade recursive name servers feature high performance, robust, and cost-effective DNS caching. The single most advanced feature of the Infoblox 4030, the first in a line of new recursive name servers from Infoblox, is its ability to continue to respond to DNS queries from cache even when the number of outstanding recursive queries has reached its limit. This unique split architecture allows the cache to continue to respond at full speed independently of recursive load and Internet latency.

The three dozen separate DNS queries sent by a single smartphone as it starts up will no longer contribute to name server overload. Optimized for DNS caching, the Infoblox 4030 can achieve performance of up to a million queries per second. Protected against security threats (including DoS and DDoS attacks), this new generation of recursive name server supports DNS blacklisting, DNS anycast, NXDOMAIN redirection, and DNSSEC.

Of particular advantage to ISPs looking to accommodate increased customer DNS traffic, the Infoblox 4030 is designed to reduce capital investment in general, and integration and support costs in particular, and to offer high scalability to handle future growth. Administration via the Infoblox Grid™ translates to higher efficiency, and the Infoblox 4030’s hardened form factor with no root access assures high security.

The Infoblox 4030 Requires a New Way of Interpreting Metrics

Because the Infoblox 4030 does not stop answering queries from cache when capacity limits are reached for cache misses, these name servers need a new way of interpreting metrics in order to be properly evaluated for performance.

With the Infoblox 4030, the CHR can now approach 99% as the number of users increases, most of whom are accessing popular sites, because queries continue to be answered from cache. Popular URLs, such as social media sites and search engines, remain reachable for customers no matter the volume of queries being received.

So looking at a static CHR to evaluate the Infoblox 4030 no longer makes sense. Instead, a new approach to evaluation is needed for this next generation of carrier-grade recursive name servers.

Two test scenarios examined the differences between the Infoblox 4030 and traditional BIND recursive name servers. Three graphs of the results show how a new approach to interpreting metrics can clearly demonstrate the improved performance of the Infoblox 4030 over a BIND 9.8 recursive name server.

The graph in Figure 2 shows that the response latency in a BIND 9.8 recursive name server increases substantially as the query rate increases. Note how the average latency increases dramatically after ten seconds have elapsed, indicating the limit for outstanding recursive queries has been reached.

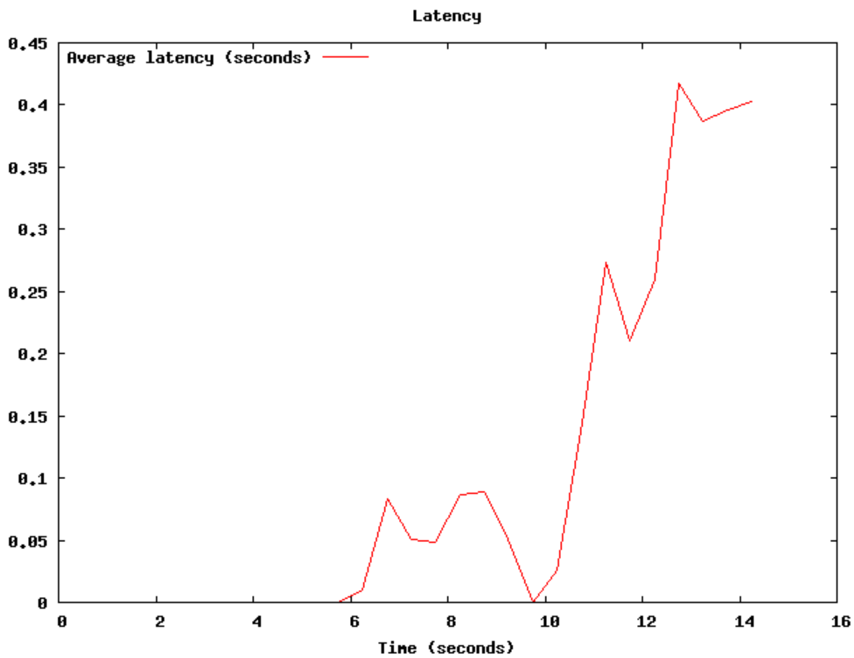

Figure 2: The average query response latency of a BIND 9.8 recursive DNS server changes over time.

The graph in Figure 3 shows the same response latency results under similar conditions for the Infoblox 4030. Note how the response latency of the Infoblox 4030 remains constant over time, unlike the BIND 9.8 name server in Figure 2 where the response latency increases dramatically after ten seconds when the maximum outstanding recursive requests limit is reached.

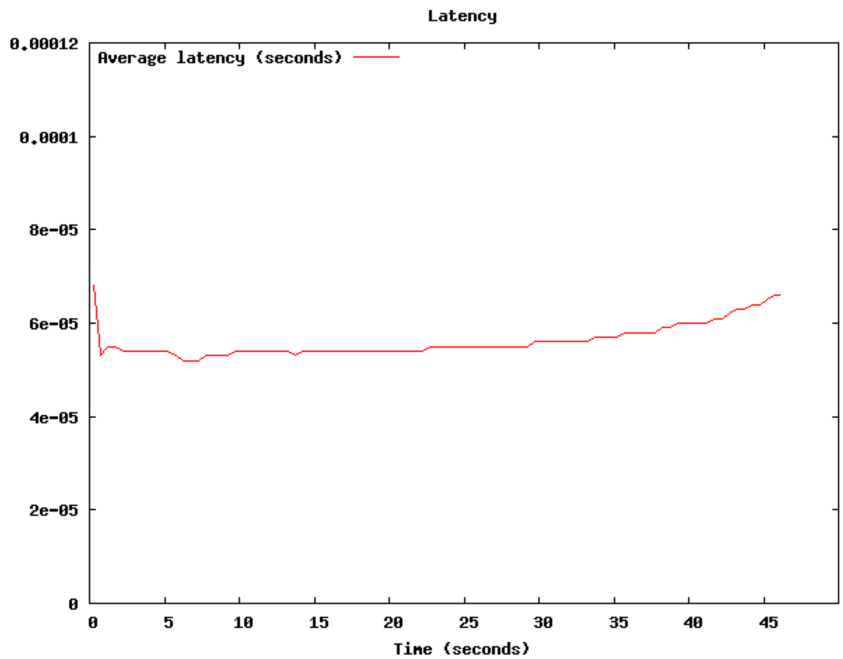

Figure 3: The average query response latency of an Infoblox 4030 remains constant over time.

The constant response latency of the Infoblox 4030 is a clear metric assessment of the Infoblox 4030's ability to continue to process queries from cache without interruption, even after the processing limit for recursive queries has been reached.

The graph in Figure 4 shows how BIND 9.8's response rate changes over time as the query rate increases. Note how the response rate drops off at 35k queries per second. This is a result of the total number of outstanding recursive requests hitting the processing limit.

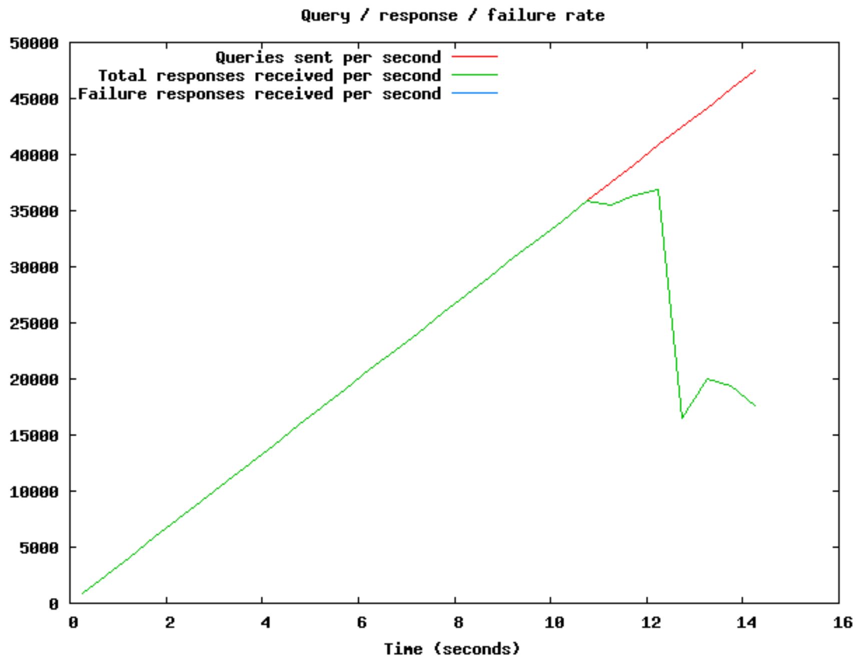

Figure 4: The query response rate of a BIND 9.8 recursive name server changes as the query rate increases over time.

Effect of New Technology on ISPs, Large Enterprises and Telcos

ISPs especially, but also telcos and large enterprises can benefit from this new generation of recursive name server technology from Infoblox in a number of ways. The Infoblox 4030 can:

- Improve query response rate for better end-user experience,
- Ensure the highest customer satisfaction,
- Add performance to generate profitable service streams,
- Accommodate a larger customer base and increased traffic,
- Ease the process of migrating to IPv6,
- Reduce the number of recursive name servers needed,
- Lower TCO for managing DNS infrastructure.

However, adopting and applying a new approach to assessing performance metrics is essential in order to understand the capabilities of this new recursive name server technology. The interpretation of these performance metrics must account for the increased capacity of the new Infoblox 4030 — and those that will follow in this line of new generation recursive name servers from Infoblox — including unimpeded and non-stop processing of queries from cache.

ISPs and others can assess the performance capabilities of this next generation of recursive name servers with this new, more accurate interpretation of metrics. Adding this new generation of recursive name servers to their infrastructure can lead to higher performance, improved customer satisfaction, better end-user experience, added revenue stream opportunities, and lower costs for infrastructure management.

Corporate Headquarters:
+1.408.625.4200
+1.866.463.6256
(toll-free, U.S. and Canada)
info@infoblox.com
www.infoblox.com