NeuStar UltraDNS Infrastructure Summary

UltraDNS Whitepaper

ULTRADNS®
A NeuStar® Service

NEUSTAR®
Trusted to bring networks together

Table of Contents

I. NeuStar UltraDNS system and network architecture..... ..3
II. Anycast implementation..... 4
III. DDoS mitigation platform.....6
IV. Global presence 9

I. The NeuStar UltraDNS system and network architecture

A. UltraDNS architecture is comprised of three different levels:

- System level - multiple separate, yet related meshes of servers that have a primary-and-secondary or master-and-slave relationship.
- Mesh level - multiple nodes that have identical data sets, which are synchronized via replication over the wide area network.
- Node level - system components, co-located at the same network point of presence and functioning together to provide the DNS protocol service.

B. UltraDNS data model

The UltraDNS node is designed around a data model maintained within a commercial database. The data model contains information about principal objects managed by the system (e.g., users, DNS zones, and resource records) and additional information required to control processes operating on the data (e.g., service configuration parameters and ACL info). Various functionalities of the UltraDNS system are provided by numerous disparate processes, which primarily serve as a conduit between the database and end-user requests.

C. UltraDNS name servers

UltraDNS name servers answer Internet protocol DNS queries based on authoritative DNS data maintained in the database. One major NeuStar UltraDNS innovation was the ability to make an authoritative DNS server that was capable of answering DNS from a database-reliant system at a speed comparable to that of a memory-resident system, such as BIND. NeuStar UltraDNS uses network deployment and routing control to allow the scalability of such a system by linear addition of hardware to meet load requirements along with DNS-specific caching algorithms and associated cache invalidation mechanisms. With this configuration, the UltraDNS system has confirmed load capacity one order of magnitude above the combined load of all existing TLDs.

D. UltraDNS nodes

Each node is designed to provide both security and scalability for the UltraDNS network. By using dedicated hardware, UltraDNS partitions each major part of

the network to function independently, thereby ensuring access control to each point, as well as growth capability. Hardware can be transparently added to an existing node without affecting service to that node. In this case, once a new device is added, it immediately begins announcing Anycast addresses and is included in the pool of servers available to answer queries within that node. If a server were to fail, it would immediately stop announcing Anycast addresses and queries would be answered by the next functioning server in that node.

E. Software

UltraDNS has developed a non-BIND proprietary code built from the ground up. In 2004, the code base underwent an extensive third-party security audit which found no vulnerabilities that could be abused remotely to acquire restricted privileges or cause failure of directory-resolution capabilities on the UltraDNS system. In addition to supporting standard DNS specifications and RFCs, there are numerous features and enhancements that have been incorporated into the UltraDNS system to ensure robustness, security, and redundancy well above what is capable with legacy DNS server implementations.

F. Border Gateway Protocol (BGP)

UltraDNS has incorporated BGP announcement generating code directly into the UltraDNS DNS resolver. This allows for name servers or complete nodes to be removed from the pool of active systems upon the detection of anomalous data, or the failure of any key element of that name server or node. Should the entire node fail, BGP announcements for that node are withdrawn and queries automatically routed only to operational nodes. The code is fully compliant with the following RFCs: 2453, 2080, 2328, 2460, 2373, 2463, 2464, 2236, 1812, 1771.

II. Anycast Implementation

A. DNS Infrastructure

The primary mechanism for addressing and route announcements pioneered and used by UltraDNS is known as Anycast. This technique involves the

announcement of the same IP addresses by multiple nodes at the same time. The result is a DNS infrastructure that has the highest level of performance and lowest level of latency and packet loss possible, while providing the ability to increase the number of available name servers globally during times of need (DDoS attacks, etc) without modifications by users or networks external to UltraDNS.

By injecting a BGP route from each node, the system leverages IP routing to deliver user queries to a topologically nearby node. This results in:

- A reduction of network latency for DNS transactions, compared to "standard" DNS services deployments.
- A reduction in the number of queries that are routed to distant servers, thereby reducing the likelihood of encountering congested routers.
- A resulting reduction in the number of dropped query packets that cause DNS timeouts/retries.
- Improved end-user performance and reliability.

The UltraDNS mechanism has been adopted by most major root/TLD operators as a "Best Current Practice."

Diverse network connectivity is deployed within the UltraDNS network. Primary connectivity is provided by international network providers: Verizon, NTT/Verio, Global Crossing and Level 3. Each node is multi-homed with 100 Mbps (Fast Ethernet) connections from each provider. For robustness and redundancy, a carefully architected matrix of network announcements is used to ensure that minor and catastrophic failures of any elements within the UltraDNS network does not result in resolution failures for end users. NeuStar has also implemented additional connections at most nodes to local public-switched peering fabrics. The company employs a liberal peering policy, and over 50 networks are peered directly with UltraDNS at one or more locations, using these public peering facilities.

B. Global Anycast network

Added reliability is achieved by announcing up to six global IP addresses from each device in the UltraDNS TLD server network infrastructure. This provides additional redundancy in the face of network routing problems that can be caused by third parties. In the unlikely event that one or more IP addresses become unreachable, queries from users are able to failover to an alternate global IP address. The fundamental design creates thirty-two unique combinations of IP addresses, network providers, and physical node locations. Even in the event of the catastrophic failure of a provider's entire network backbone, physical location or region, queries from any location in the world continue to have a functional set of address/route/location options. In any Internet region, no more than two IP addresses of the six that are announced result in packets reaching the same node over the same path. And no more than three IP addresses of the six that are announced result in packets reaching the same physical location irrespective of what networks were traversed.

III. DDoS mitigation platform — The DNS Shield

NeuStar has devoted significant resources in research and development to thwart Distributed Denial of Service (DDoS) attacks and improve both the security and reliability of the Internet's critical DNS infrastructure.

The success of UltraDNS in defending itself from DDoS attacks in Q4 of 2002 prompted an evaluation of preparedness for more complex DDoS attacks in the future. As the results of forensic diagnostics were applied to advanced simulation models for projecting evolution of attack scenarios, the concept of an out-of-band signaling network for DNS was identified as the most secure defense against the next progression of DDoS attacks.

In late 2002, following a 1-hour DDoS attack against the root servers, an attack was launched against the UltraDNS TLD server network. The attack was of such intensity that UltraDNS upstream providers were themselves forced to apply mitigation steps to stabilize their own networks. During this 48-hour attack, UltraDNS was able to maintain operations without end-user impact. This was

largely due to the relatively unsophisticated attack tools used.

Studies have shown that a carefully crafted DDoS attack with sophisticated attack tools could create a situation where normal mitigation efforts could fall short. As a result, a process of laboratory simulations was begun to test the effectiveness of a probable solution. The testing indicated that in the most advanced implementations, at up to approximately 1 Gb/s of sustained attack traffic, the solution would be effective. However, above that level, the solution would begin to fail because of the physical limitations of router and server interfaces, which would be overrun. In late 2003, the evolution of Trojan armies, or "botnets" reached the point that attacks greater than 1Gb+ became relatively simple to mount, and likely to be used. However, until mid-2004, the "perfect DNS DDoS" attack was still theoretical.

In June of 2004, a four-hour attack against the DNS systems of Akamai effectively shut down its service despite a significant deployment of DNS, security, and network assets well in excess of the 1Gb threshold. Forensic analysis indicated that the attack closely matched the profile of the UltraDNS identified “perfect DNS DDoS” and, as predicted by company engineers, it was unstoppable. Since it became clear that even some of the world's largest DNS networks can be taken down by an array of botnets, a stronger, more resilient solution had to be created - the "DNS Shield."

As a result of the laboratory simulations, an UltraDNS extension to the initial mitigation solution had already developed, moving to the development of a private-network alternative to the original approach. Since there is currently no commercially available or practical method of filtering or processing the “Perfect DNS DDoS” packet, the UltraDNS solution sought to provide a mechanism that met the primary objective of an authoritative DNS system while under a severe attack – that of enabling recursive servers to continue to resolve queries normally for the bulk of Internet users.

NeuStar UltraDNS has achieved this objective by identifying the largest sources

of legitimate DNS queries for the zones that UltraDNS is authoritative for, and then deploying complete and self-contained authoritative UltraDNS Nodes onto local segments that contain these “Trusted Recursive Servers” via point-to-point Ethernet circuits. Queries from these Trusted Recursive Servers for the UltraDNS zones are then asked and answered in a fully isolated and protected environment. This topology provides for unprecedented sub 5-millisecond query response times within the networks where Local Nodes are installed, while ensuring that only queries from the Trusted Recursive Servers are able to reach the nodes.

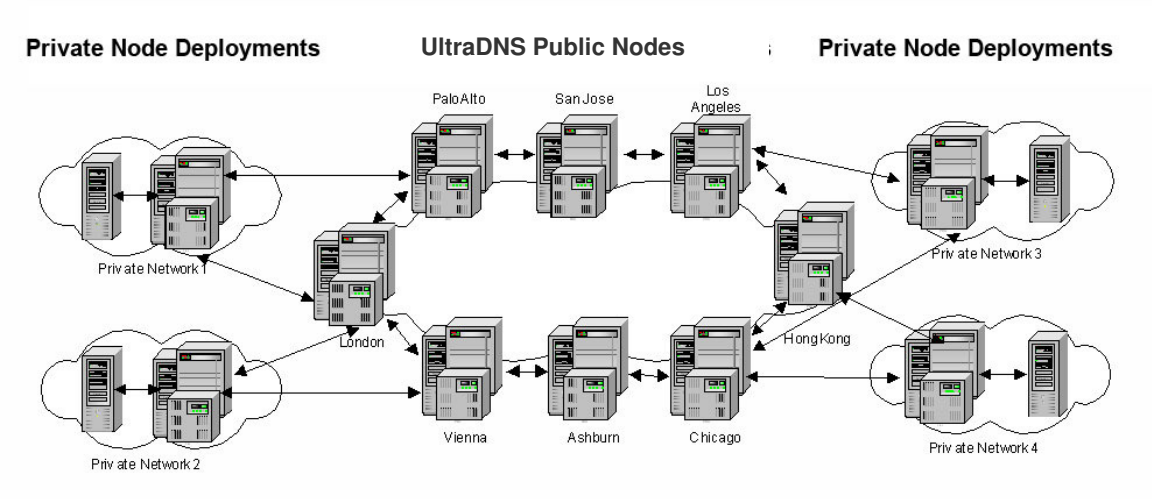

Figure 1 The DNS Shield

Local nodes are functionally identical to UltraDNS public nodes and include the use of the announced Anycast IP addressing scheme via BGP as well as protected connectivity to the UltraDNS Replication System to assure data consistency. Local nodes employ the same operational standards as public nodes so anomalies identified the same way are handled accordingly. Should local nodes within a Host ISP’s network fail, local routes are withdrawn and the trusted Recursive Servers automatically follow normal external announcements and paths to UltraDNS public nodes.

To avoid significant customer service calls and issues, Host ISPs (the ISPs

controlling trusted Recursive Servers that are sources of the queries) protect access to “their” UltraDNS local nodes. When appropriate, they are also encouraged to permit their customers to access local nodes via their own recursive servers, which are configured to forward queries for UltraDNS zones to Host ISP trusted Recursive Servers. However, the Host ISP is fully responsible for making this decision and managing it. Host ISPs must confirm their understanding that if they have not maintained the integrity of the local isolated network, they will not experience the benefits of this system during a DDoS attack.

By inviting the largest service providers to connect directly and privately to the UltraDNS directory infrastructure via the DNS Shield authoritative DNS information stored by UltraDNS can be assured as being valid and always available to participating ISP end users. The UltraDNS local node DNS Shield provides DDoS-resistant DNS resolution to almost 100-million Internet users. This number of protected users is soon expected to be in excess of 200 million globally. In addition, NeuStar has embarked on a project to deploy local nodes at appropriate locations in emerging countries, with local announcements only. The first of these was formally agreed to at the ICANN meeting in Cape Town, and is now deployed at the JINX exchange in Johannesburg, South Africa.

Connectivity to the DNS Shield is by invitation only to the largest ISPs and recursive DNS server operators. Smaller enterprises and lower usage recursive DNS server operators will not have direct access to the DNS Shield. However, the UltraDNS DNS Shield concept can be extended and such smaller networks can utilize their upstream service providers in a similar fashion. This forced hierarchy preserves the decentralized and public nature of the DNS, but introduces a now required level of security, authentication and responsibility into the entire model to maintain DNS stability.

IV. Global Presence

NeuStar UltraDNS maintains a mesh of fourteen globally synchronized DNS server systems.

These systems are multi-homed at all locations and are publicly and privately peered. Advanced replication ensures that data is replicated between and within all nodes. UltraDNS Servers are located in:

- Switch and Data, CA, USA
- Switch and Data, NY, USA
- Equinix Inc, CA, USA
- Equinix Inc, TX, USA
- Equinix Inc, Chicago, USA
- Equinix Inc, VA, USA
- Equinix Inc, Hong Kong
- Terremark, FL, USA
- Metromedia Fiber Network Inc (AboveNet): UK
- CNNIC, China
- JINX, Johannesburg, South Africa
- Restena, Luxembourg
- NIXI, India
- Equinix, Sydney, Australia

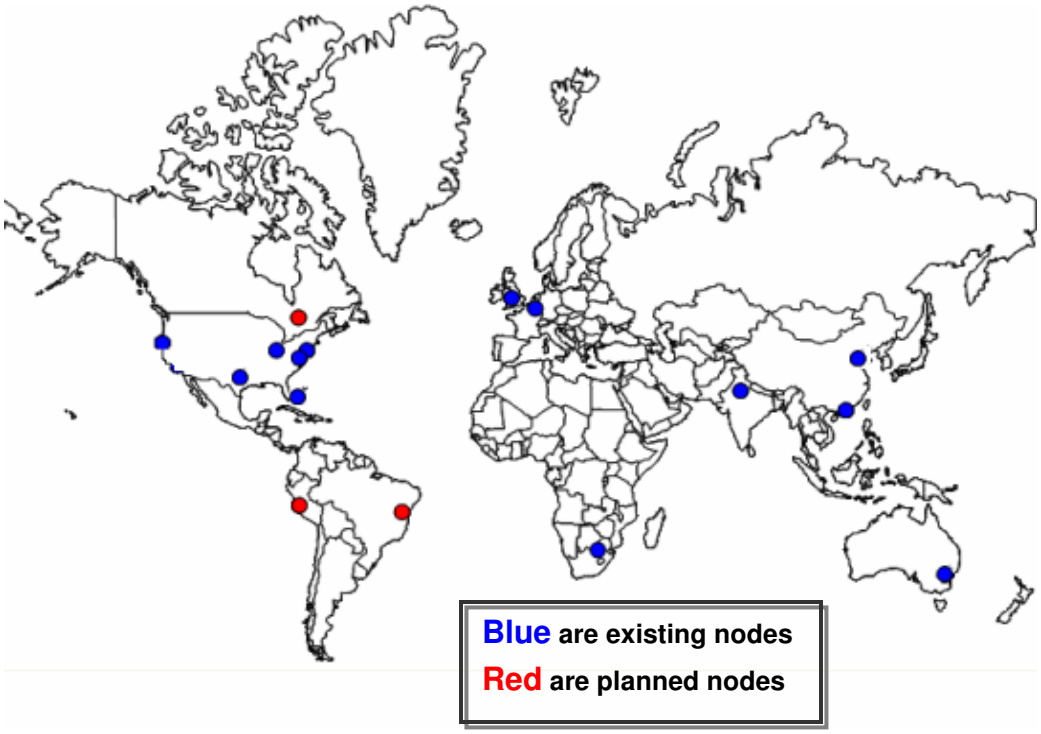

About NeuStar and NeuStar UltraDNS Services

NeuStar, Inc. (NYSE: NSR) is a provider of clearinghouse and directory services to the global communications and Internet industry. Through its UltraDNS Services, NeuStar provides solutions to organizations that rely on the Internet Protocol (IP) for their critical business processes, applications, and communication services. NeuStar offers these solutions via one of the world's largest proprietary non-BIND Directory Services Platforms.

Providing integral Authoritative DNS services at the root level for top-level domain (TLD) registries and second-level domain (SLD) registrants, we have thousands of enterprise, service provider, and core infrastructure customers. NeuStar UltraDNS Services power the resolution of over 20 million domains around the world, and NeuStar is the only company of its kind that is actively involved at every level of the DNS tree.

NeuStar Corporate Headquarters
NeuStar, Inc.
46000 Center Oak Plaza
Sterling, VA 20166
Press Contact
Marc Abshire
Director of Global Media Relations
Phone: (571) 434-5151
Email: **marc.abshire@NeuStar.biz**
UltraDNS
East Coast Office
21631 Ridgetop Circle 2nd floor
Sterling, VA 20166
West Coast Office
1000 Marina Blvd.
(Suite 400)
Brisbane, CA 94005
Midwest Office
150 N. Michigan Avenue
(28th Floor)
Chicago, IL 60602
European Office
London, England
4 Lombard Street
London EC3V 9HD
+44 (0) 207 933 8632
Customer Support
Email: Support@UltraDNS.com
Phone: (888) 367-4820
Sales & Information
Call (888) 367-4812 or **click here** to send a sales or information request.