# On the DNS Deployment of Modern Web Services

Shuai Hao*†, Haining Wang*, Angelos Stavrou‡, and Evgenia Smirni†
*University of Delaware, Newark, DE, USA
†College of William and Mary, Williamsburg, VA, USA
‡George Mason University, Fairfax, VA, USA
Email: {haos,hnw}@udel.edu, astavrou@gmu.edu, esmirni@cs.wm.edu

***Abstract*—Accessing Internet services relies on the Domain Name System (DNS) for translating human-readable names to routable network addresses. At the bottom level of the DNS hierarchy, the authoritative DNS (ADNS) servers maintain the actual mapping records and answer the DNS queries. Today, the increasing use of upstream ADNS services (i.e., third-party ADNS-hosting services) and Infrastructure-as-a-Service (IaaS) clouds facilitates the establishment of web services, and has been fostering the evolution of the deployment of ADNS servers. To shed light on this trend, in this paper we present a large-scale measurement to study the ADNS deployment patterns of modern web services and examine the characteristics of different deployment styles, such as performance, life-cycle of servers, and availability. Furthermore, we focus specifically on the DNS deployment for subdomains hosted in IaaS clouds.**

## I. Introduction

As a hierarchical distributed database system, the Domain Name System (DNS) is one of the most important components of Internet infrastructure, providing the mapping between the domain names and network-level addresses to direct clients to specific Internet services. In DNS hierarchy, the Root and Top-Level-Domain nameservers are mainly used as the querying referrals, while the authoritative DNS (ADNS) servers, administered by the service providers, are responsible for storing the name-to-address records and returning answers to the clients.

Deploying authoritative nameservers requires extra hardware resources and additional maintenance support. Also, the critical roles of DNS service in web infrastructure make it an attractive target to attackers. Thus, web service providers are increasingly adopting the upstream authoritative DNS servers, including the top sites (e.g., Amazon and Twitter) that have the ability to maintain their own ADNS infrastructures. In addition, to save a large amount of investment for infrastructure, many of today's popular web services are directly built upon Infrastructure-as-a-Service (IaaS) clouds such as Amazon EC2 and Windows Azure. The traditional web service providers are also migrating extended services into clouds to use the "illusively-infinite" computing and storage resources. The IaaS infrastructure greatly facilitates the establishment of modern web services and also promotes the process of delegating the authoritative name resolution to third-party ADNS service providers. Besides traditional web-hosting providers such as Dyn [6] and Ultradns [14], the Content Delivery Networks (CDN) and cloud service providers also offer the ADNS services that integrate the name resolution into their CDNs or cloud infrastructures [1], [4].

Existing DNS measurements studied the characteristics of DNS activities and operations [16], [17], [21], [24], [26], the root or top-level-domain servers [20], [22], [29], [30], [36], or the DNS resolvers [15], [18], [35]. Some works involving the characteristics of ADNSes mainly focused on the comparison with local DNS (LDNS) servers, but none of them explored various ADNS deployments for web services. Complementary to these prior works, we present a large-scale measurement study in attempt to answer the following questions: (1) how do modern web services deploy their ADNS servers? (2) what are the characteristics of different ADNS deployment patterns? and (3) in particular, how do the cloud-hosting subdomains administer their ADNS servers?

We first collect the authoritative DNS server information for top-ranking websites on Alexa's list [2] and eliminate the redundant domain records. This constructs our dataset with about 2.3 million nameservers for about 0.94 million websites. We then develop a systematic method to explore ADNS server deployment patterns and perform the geo-distributed probing experiments. In particular, by directly issuing DNS queries to each ADNS server, we examine their deployment details and characteristics. Next, we focus on the DNS deployment of web services whose subdomains are hosted in cloud infrastructure. We extract the subdomain list from an existing dataset [5], reproduce the ADNS servers of subdomains for comparing with the original results, and examine their deployment. We summarize our major findings and contributions as follows:

- We use a simple heuristic method to determine the ADNS deployment patterns. In fact, it is fairly easy to recognize the pattern for an individual website from its NS records, but it is much more difficult when looking for millions of websites in such a large-scale study.
- We validate the use of ADNS proxy infrastructure by examining the transition delay and the TTL aging.
- We first quantify the usage and profile the characteristics of ADNS servers in terms of the deployment patterns.
- We find that most top-ranked websites deploy their own DNS servers but emerging popular social sites tend to use the upstream DNS-hosting services. We also observe few servers being used in private deployment.
- We find that the ADNS deployment patterns remain stable. The change of private servers is more frequent than that of upstream servers. The websites using upstream services change frequently their hosting domains but have