# SIPS, IPPS, or Oops!
# An Analysis of the Security and Privacy of DNS Service Discovery (abstract #29)

Joseph Call
*William & Mary*

Mostafa Ahmed
*William & Mary*

Stephen Herwig
*William & Mary*

The Domain Name System (DNS) is colloquially known as the "telephone book of the Internet," as it allows a client to look up the IP address for the domain name it is trying to access. However, the DNS also provides several methods that allow clients to first discover the domain name for a service by knowing just its generic service name (such as `sip`, `ipp`, or `ssh`). This extra level of indirection provides tremendous flexibility for domain administrators to publish and reconfigure service deployments, and for client software to autodiscover a service instance. **In this preliminary work, we investigate the security and privacy of the DNS service discovery ecosystem by asking the following questions: (1)** ***What services do domains publish?*** **(2)** ***Do domains publish these services securely?*** **and (3)** ***Do these services themselves provide strong authentication and privacy for their clients?***

**Overview of Publishing Services on the DNS.** In our preliminary work, we investigate the use of the following DNS records for service discovery: `SRV` [5], `NAPTR` [5], and `PTR` [1, 7]. Briefly, for a `SRV` query, a client queries a name of the form `_<Service>._<Proto>.<Apex-Domain>` (such as `_sip._udp.example.com` for voice-over-IP services) and the response contains the domain names and port numbers of all service replicas in that DNS zone. A client may first issue a `NAPTR` query to determine the valid names for such a `SRV` query. Whereas `SRV` and `NAPTR` allow a client to discover *identical replicas*, the DNS-SD specification [1] allows a client to enumerate *distinct service instances* using a series of `PTR` and `SRV` queries. Similar to `NAPTR`, DNS-SD optionally supports enumerating a zone's service types.

**Methodology.** To automate the discovery of services, we develop a DNS scanner using the popular `miekg/dns` Go package [4], and implement `SRV`, `NAPTR`, and `DNS-SD` probes. We use our scanner to scan the Tranco Top 1 Million domains [6]. Since performing an exhaustive scan of all services is impractical (there are over 12,000 registered `_<Service>._<Proto>` label pairs [3]), our scanner first uses a probablistic tuning phase to estimate service popularity, and then performs a complete scan with this smaller estimated list.

**Preliminary Analysis.** Figure 1 shows the number of instances our scanner found for the top-22 most popular services, broken down by the service's underlying transport protocol as determined by its `_<Service>._<Proto>` label pair. We observe that the services surprisingly comprise conferencing, directory (contacts and calendars), mail, NAT traversal,

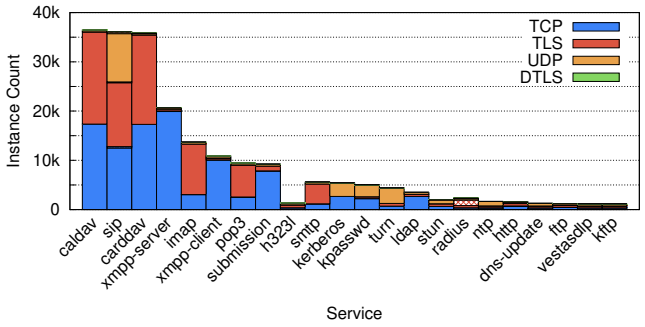


Figure 1: Service instance advertisements by service type.

and authentication protocols. Many of these instances are not using TLS. Of those that do, Figure 2 reveals that a significant portion have invalid TLS certificates.

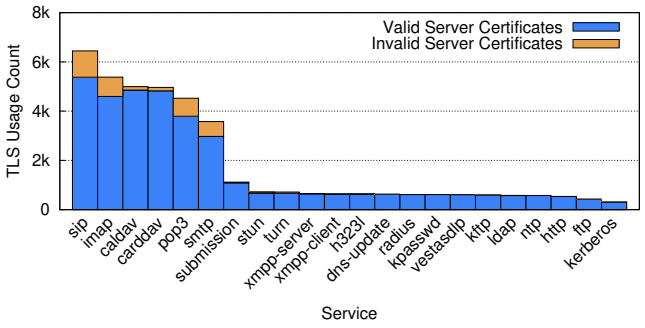


Figure 2: Count of service instances with valid and invalid TLS certificates.

Similar to prior work [2], we find that 91.6% of domains that advertise a service, and 89.9% of domains that host a service, do so without proper use of DNNSEC. The most common reason for for failure to DNSSEC validate a response is lack of a signature (a missing `RRSIG` record).

**Future Work.** Several of the most popular services (such as `caldav` and `cardav`) are HTTP-based, and for future work, we will investigate the HTTP authentication methods that these instances use. Additionally, as many services use UDP, we also plan to assess the prevalence amplification attack vectors. Finally, we will incorporate additional DNS service discovery methods into our scanner, including the recently proposed `SVCB` record [8].

## References


[1] Stuart Cheshire and Marc Krochmal. DNS-Based Service Discovery. RFC 6763, February 2013.

[2] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. A longitudinal, end-to-end view of the DNSSEC ecosystem. In *USENIX Security Symposium*, 2017.

[3] Michelle Cotton, Lars Eggert, Dr. Joseph D. Touch, Magnus Westerlund, and Stuart Cheshire. Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry. RFC 6335, August 2011.

[4] Miek Gieben. Alternative (mor granular) approach to a DNS library. https://github.com/miekg/dns.

[5] Arnt Gulbrandsen and Dr. Levon Esibov. A DNS RR for specifying the location of services (DNS SRV). RFC 2782, February 2000.

[6] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, and Wouter Joosen. TRANCO: A research-oriented top sites ranking hardened against manipulation. In *Network and Distributed System Security Symposium (NDSS)*, 2019.

[7] Paul Mockapetris. Domain names - implementation and specification. RFC 1035, November 1987.

[8] Benjamin M. Schwartz, Mike Bishop, and Erik Nygren. Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records). RFC 9460, November 2023.