

SVCB (and HTTPS)

Service binding and parameter specification via the DNS

Ben Schwartz <bemasc@google.com>
Mike Bishop <mbishop@evequefou.be>
Erik Nygren <erik+ietf@nygren.org>

IETF 106 - November 2019

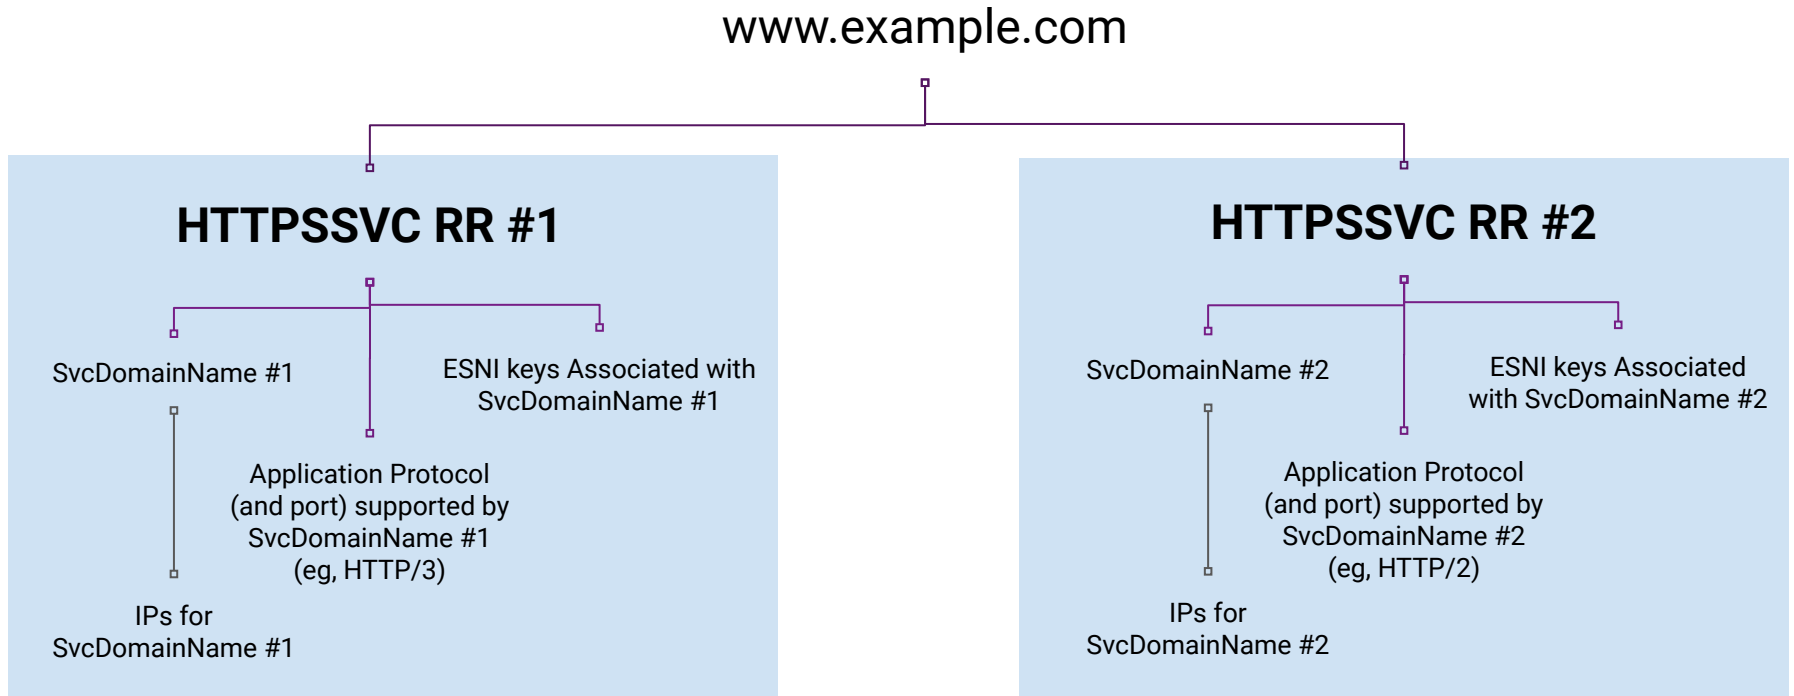
<https://tools.ietf.org/html/draft-ietf-dnsop-svcb-httpsvc-01>

SVCB Overview

- Goal: bootstrap optimal connections from a single DNS query
- In “AliasForm”, it acts like CNAME but can be at the apex
- In “ServiceForm” it is an extensible service description, currently supporting:
 - TLS ALPN
 - Port
 - Encrypted SNI configuration
 - IP hints
- HTTPSSVC is an SVCB-compatible RR type specialized for HTTPS
 - Indicates origin defaults to HTTPS
 - Avoids underscore prefixes
 - Improves compatibility with wildcard domains
 - Compatible with existing CNAME delegations

Example: HTTPSSVC and Multi-CDN hosting

Clients may end up on one or more service endpoints (i.e. sets of servers) which may have different capabilities and keys, such as on different CDNs. HTTPSSVC provides a way to tie these together.




AliasForm (SvcFieldPriority=0)

- Covers many “SRV” and “ANAME” use-cases



Service Form (SvcFieldPriority>0)

- Covers ESNI use case and other protocol improvements

 Lower SvcFieldPriority
means preferred

svc.example.net. 7200 IN HTTPSSVC **2** svc3.example.net. alpn=h3 port=8003 \
esniconfig=...

SvcFieldValue encodes protocol, port,
ESNI keys, and other params



svc.example.net. 7200 IN HTTPSSVC **3** svc2.example.net. alpn=h2 port=8003 \
esniconfig=...

“Please use QUIC to UDP svc3.example.net:8003 with this ESNI configuration, or use HTTP/2 to TCP svc2.example.net:8002 with this other ESNI configuration.”

Changes since IETF 105

- Support non-HTTP protocols by generalizing from HTTPSSVC to SVCB
 - SVCB uses `_port._scheme` prefixes to support arbitrary protocols
- Adopted by the working group
- Made AliasForm vs. ServiceForm implicit based on SvcFieldPriority
- Many minor changes
 - Relaxed IP hint handling requirements
 - Added and removed descriptions of various optional optimizations
 - Reduced emphasis on conversion to and from Alt-Svc
 - Terminology updates and other clarifications

Major remaining design questions

- How to balance ESNI strictness against reliability and misconfiguration ([#73](#))
 - Current requirements prevent fallback from ESNI to non-ESNI unless the server specifically indicates that this is allowed, potentially creating a “footgun” for server admins who don’t realize that not all networks allow QUIC.
- How should we limit the alias chain length? ([#57](#))
 - Current text has a rough suggestion of “8”, but there’s (almost?) no need for more than 1.
- Should we remove the “0” in AliasForm? ([#63](#))
 - Would match CNAME presentation form
 - Might break the pattern of SRV, URI, MX, NAPTR
- What to name the RRTYPE
- ...



Next steps...

- Continue refining requirements with potential implementers
- Work on clarity and remove TODOs
- Hoping to be ready for WGLC before IETF 107

Current workspace:

<https://github.com/MikeBishop/dns-alt-svc>

Editor's draft:

<https://mikebishop.github.io/dns-alt-svc/draft-ietf-dnsop-svcb-httpssvc.html>

Feedback on mailing list(s) and to authors most welcome!

FAQs

- Why are there IP hints?
 - The IP hints are a performance optimization that avoids one DNS roundtrip when
 - SvcDomainName != “.”, i.e. there is a layer of indirection, AND
 - the recursive resolver is not SVCB-aware
 - This is only useful in the multi-CDN and “CNAME at apex” cases.
 - Avoiding a DNS roundtrip in this case has been a strong requirement from ESNI
- Why not have two RR types for AliasForm and ServiceForm?
 - AliasForm and ServiceForm could be separate RR types, but
 - This would significantly increase load on client, recursive, and authoritative servers
 - Clients would be incentivized to only implement the RR corresponding to ServiceForm

Comparison between SVCB & ANAME

(for the “zone apex CNAME” issue)

SVCB

Pros:

- Doesn't require any changes to DNS servers

Cons:

- Only respected by compliant clients
- Adds a roundtrip if the recursive is not cooperating

ANAME

Pros:

- Doesn't require any changes to clients

Cons:

- Requires complex changes to participating authoritative servers, especially when DNSSEC or ECS is also in use

Neither may fully replace the need or use-cases for the other.