

SVCB (and HTTPSSVC)

Service binding and parameter specification via the DNS

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

DNSOP virtual interim - April 2020

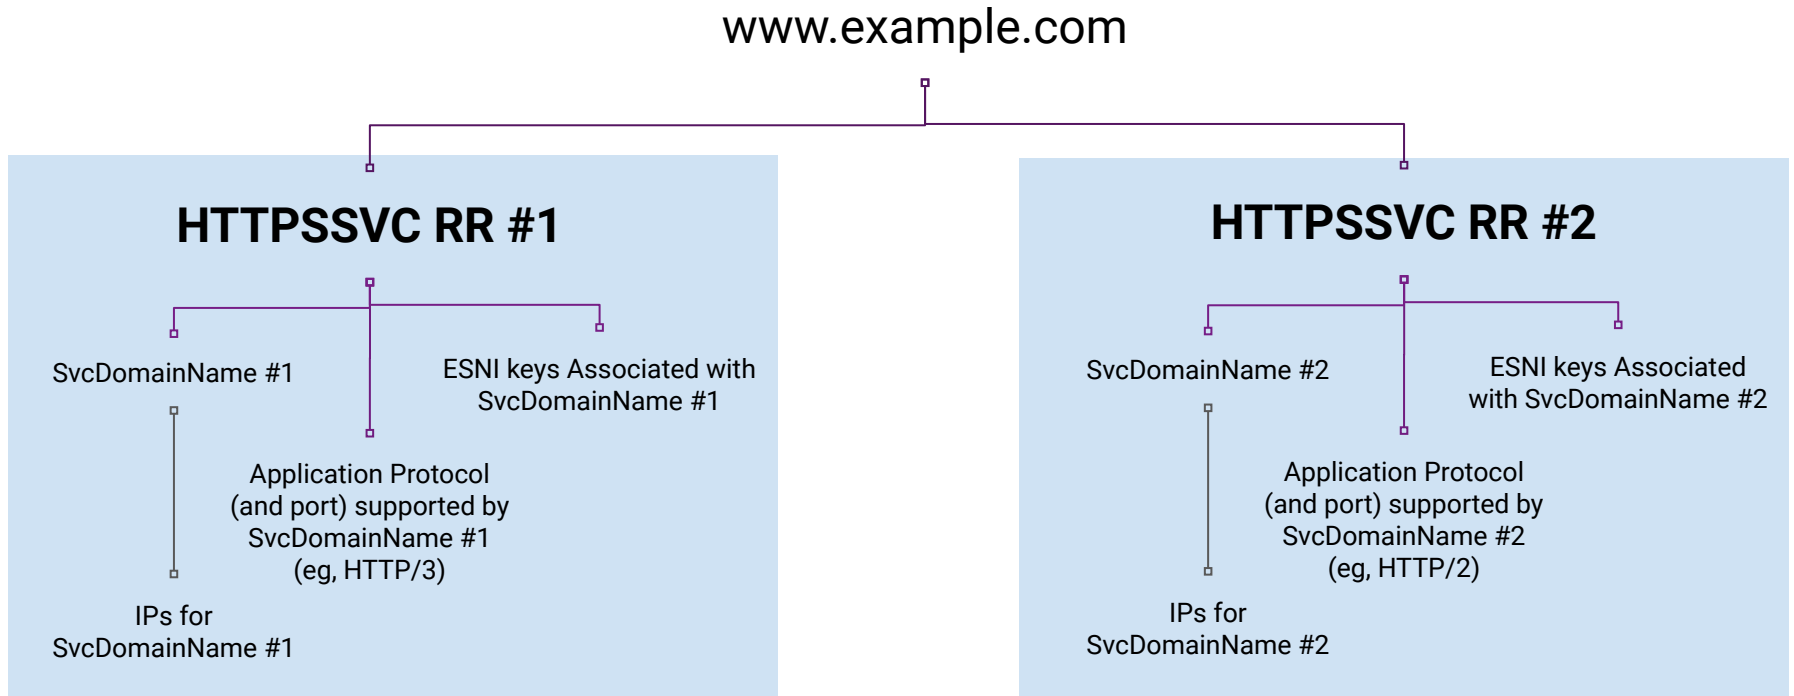
<https://tools.ietf.org/html/draft-ietf-dnsop-svcb-httpssvc-02>

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 -01

- Decoupled from Alt-Svc. Any changes to Alt-Svc will happen later.
- Improved specificity
 - Priority zero (“0”) defines AliasForm, instead of being reserved for AliasForm.
 - Expanded description of recursive resolver behavior
 - Much more precise description of the intended ALPN behavior
- Security adjustments
 - Match the HSTS specification's language on HTTPS enforcement
 - New text regarding resolution timeouts (clients must fail hard to avoid a downgrade attack)
 - Removed the 'empty esniconfig' fallback mechanism and simplified ESNI connection logic
- SvcFieldValue tweaks
 - Repeated SvcParamKeys are no longer allowed.
 - In the wire format, SvcParamKeys must be in sorted order.
 - The "=" sign may be omitted in a key=value pair if the value is also empty.

Topic of interest: [ALPN](#)

- Goals: Retain TLS security guarantees, make QUIC-only services expressible but hard to configure by accident, avoid duplicating ESNI keys, don't require holistic RRSet validation at the authoritative, minimize fragility.
- Unified backend:
 - HTTPSSVC 1 backend.example. alpn=h3 esniconfig=... port=8443
 - “I support HTTP/3 on UDP 8443 **and HTTP/2 on TCP 8443**, both with this ESNI config”
 - Clients attempting HTTP/2 will normally **also offer HTTP/1.1** in their ClientHello.
- Split backend:
 - HTTPSSVC 1 backend-h3.example. alpn=h3 **no-default-alpn** esniconfig=...
 - HTTPSSVC 2 backend-h2.example. esniconfig=...
 - “I support QUIC on backend-h3 (UDP 443) and HTTP/2 on backend-h2 (TCP 443)”
- ALPN MUST indicate the whole stack (e.g. HTTP/3 over QUIC over UDP)
- **Seeking final comments on this topic**

Topic of interest: Ports

- `_1234._https.asdf.example. HTTPSSVC 1 backend.example. port=5678`
 - “When loading `https://asdf.example:1234`, try HTTP/2 on `backend.example:5678`”
- SVCB (deliberately) conflates TCP and UDP port numbers
 - Enables sharing ESNIConfig (which can be large) between HTTP/3 and HTTP/2
 - Can use separate RRs if the configurations are different
- Should we remove “port=...”?
 - If so, should we preserve the URI’s port, or send everything to the scheme’s default port?
 - Arguments for removal: compatibility with port-restricting gateways, simplicity.
 - Arguments against removal: would prevent running HTTP/3 on a different port from HTTP/2; multiplexing customers by port on a shared IP; non-Web API use-cases.
 - **Seeking WG input.** Current plan: no change.

Topic of interest: Naming of records

- Poll of working groups yielded that all options are hated by some people, so proposing “least bad” options
- Proposed: Leave “SVCB” as-is
 - (refer to as “a SVCB record”)
- Proposed: Rename “HTTPSSVC” to “HTTPS”
 - (but refer to as “a SVCB-form HTTPS record”)

- Updates will be made shortly (after substantive changes stabilize)

Next steps...

- Continue work on clarity and remove TODOs
- Finalize RR names
- Aiming to start WGLC and early codepoint allocation before IETF 108
- Q: Interest in interop testing around or during IETF 108?

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!