

# DELEG records: Omnibus vs. Discrete

---

Paul Hoffman  
IETF 121, Dublin  
November 2024

# Two proposals for DELEG records

---

- A single **omnibus RRtype** with sub-fields for all delegation features, the same as or modelled closely after SVCB records
- **Individual RRtypes** for each new delegation feature

# Reusing SVCB

---

- Authoritative and recursive software already knows how to parse and display the records
  - Will need to be updated for each new DELEG feature, but the parsing will be relatively minor
- However, reusing or defining new SVCB parameter keys might limit the features that can be added to DELEG, even somewhat

# Individual RRtypes for each feature

---

- Easy to do: it's what we already do for everything else in the DNS
- However, we have to reinvent new parsing and display rules for every new RRtype

# Cross-feature requirements and exclusions

---

- Assume a new DELEG feature is being added that requires that an earlier feature also appears in the DELEG omnibus record or the record set
- Assume a new DELEG feature is being added that requires that an earlier feature **MUST NOT** appear in the DELEG omnibus record or the record set
- Can these be more easily and understandably be expressed with the omnibus record or with individual RRtypes

# Interactions with provisioning

---

- An omnibus record might be easier to add to EPP/RDAP/etc than a bunch of individual records because it is a single record to update
- However, after there are many DELEG features, updating one master record requires updating all the features in one transaction instead of just updating what changes

# Can convert between the two

---

- The current SVCB-based proposals can be trivially converted to individual RRtypes
- If the WG wants to model both methods, that's easy
  - The hard part is saying how the resolver will respond when seeing the information, not the record format

# I have opinions

---

- ... but I'm not a DNS software developer
- ... and many of you are
- Mic line!