

Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: 11 May 2022

T. Lemon  
Apple Inc.  
7 November 2021

Automatic Replication of DNS-SD Service Registration Protocol Zones  
draft-lemon-srp-replication-01

## Abstract

This document describes a protocol that can be used for ad-hoc replication of a DNS zone by multiple servers where a single primary DNS authoritative server is not available and the use of stable storage is not desirable.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 11 May 2022.

## Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Internet-Draft

DNS-SD SRP Replication

November 2021

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">Alternatives for maintaining SRP state</a>	<a href="#">3</a>
<a href="#">1.1.1.</a>	<a href="#">Primary authoritative DNS service</a>	<a href="#">3</a>
<a href="#">1.1.2.</a>	<a href="#">Multicast DNS Advertising Proxy</a>	<a href="#">4</a>
<a href="#">1.1.3.</a>	<a href="#">SRP Replication</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Implementation</a>	<a href="#">5</a>
<a href="#">2.1.</a>	<a href="#">Naming of a common service zone</a>	<a href="#">5</a>
<a href="#">2.1.1.</a>	<a href="#">Zone name based on network name</a>	<a href="#">5</a>
<a href="#">2.1.2.</a>	<a href="#">Zone name based on local configuration</a>	<a href="#">6</a>
<a href="#">2.1.3.</a>	<a href="#">Zone name based on DNS-SD discovery</a>	<a href="#">6</a>
<a href="#">2.2.</a>	<a href="#">Advertising one's own replication service</a>	<a href="#">7</a>
<a href="#">2.3.</a>	<a href="#">Discovering other replication services</a>	<a href="#">8</a>
<a href="#">2.3.1.</a>	<a href="#">Getting an SRP zone name from other SRP services</a>	<a href="#">8</a>
<a href="#">2.3.2.</a>	<a href="#">Establishing Replication with an existing Dataset</a>	<a href="#">9</a>
<a href="#">2.3.3.</a>	<a href="#">Establishing Replication When There Is No Existing Dataset</a>	<a href="#">11</a>
<a href="#">2.3.4.</a>	<a href="#">Conflicting precedence values</a>	<a href="#">11</a>
<a href="#">2.3.5.</a>	<a href="#">Handling primary transitions</a>	<a href="#">11</a>
<a href="#">2.4.</a>	<a href="#">Discovering the addresses of partners</a>	<a href="#">12</a>
<a href="#">2.5.</a>	<a href="#">Establishing Communication with a replication partner</a>	<a href="#">12</a>
<a href="#">2.6.</a>	<a href="#">Incoming connections</a>	<a href="#">13</a>
<a href="#">2.7.</a>	<a href="#">Eliminating extra connections</a>	<a href="#">13</a>
<a href="#">2.8.</a>	<a href="#">Initial synchronization</a>	<a href="#">13</a>
<a href="#">2.8.1.</a>	<a href="#">Sending candidates</a>	<a href="#">14</a>
<a href="#">2.9.</a>	<a href="#">Routine Operation</a>	<a href="#">15</a>
<a href="#">3.</a>	<a href="#">Protocol Details</a>	<a href="#">15</a>
<a href="#">3.1.</a>	<a href="#">DNS Stateful Operations considerations</a>	<a href="#">15</a>
<a href="#">3.1.1.</a>	<a href="#">DSO Session Establishment</a>	<a href="#">15</a>
<a href="#">3.1.2.</a>	<a href="#">DSO Session maintenance</a>	<a href="#">16</a>
<a href="#">3.2.</a>	<a href="#">DSO Primary TLVs</a>	<a href="#">16</a>
<a href="#">3.2.1.</a>	<a href="#">SRPL Session</a>	<a href="#">16</a>
<a href="#">3.2.2.</a>	<a href="#">SRPL Send Candidates</a>	<a href="#">17</a>
<a href="#">3.2.3.</a>	<a href="#">SRPL Candidate</a>	<a href="#">18</a>
<a href="#">3.2.4.</a>	<a href="#">SRPL Host</a>	<a href="#">19</a>
<a href="#">3.3.</a>	<a href="#">DSO Secondary TLVs</a>	<a href="#">20</a>
<a href="#">3.3.1.</a>	<a href="#">SRPL Candidate Yes</a>	<a href="#">20</a>
<a href="#">3.3.2.</a>	<a href="#">SRPL Candidate No</a>	<a href="#">21</a>
<a href="#">3.3.3.</a>	<a href="#">SRPL Conflict</a>	<a href="#">21</a>
<a href="#">3.3.4.</a>	<a href="#">SRPL Hostname</a>	<a href="#">21</a>
<a href="#">3.3.5.</a>	<a href="#">SRPL Host Message</a>	<a href="#">21</a>
<a href="#">3.3.6.</a>	<a href="#">SRPL Time Offset</a>	<a href="#">22</a>

<a href="#">3.3.7.</a>	SRPL Key ID . . . . .	<a href="#">22</a>
<a href="#">4.</a>	Security Considerations . . . . .	<a href="#">22</a>
<a href="#">5.</a>	Delegation of 'local.arpa.' . . . . .	<a href="#">23</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">23</a>
<a href="#">6.1.</a>	'srpl-tls' Service Name . . . . .	<a href="#">23</a>

<a href="#">6.2.</a>	DSO TLV type code . . . . .	<a href="#">23</a>
<a href="#">6.3.</a>	Registration and Delegation of 'local.arpa' as a Special-Use Domain Name . . . . .	<a href="#">24</a>
<a href="#">7.</a>	Informative References . . . . .	<a href="#">25</a>
<a href="#">8.</a>	Normative References . . . . .	<a href="#">25</a>
	Author's Address . . . . .	<a href="#">25</a>

## [1.](#) Introduction

The DNS-SD Service Registration Protocol provides a way for network services to update a DNS zone with DNS-SD information. SRP uses unicast DNS Updates, rather than multicast DNS, to advertise services. This has several advantages over multicast DNS:

- \* Reduces reliance on multicast
- \* Reduces traffic to devices providing services, which may be constrained devices operating on battery power
- \* Allows the advertisement of services on one network link to consumers of such services on a different network link

### [1.1.](#) Alternatives for maintaining SRP state

#### [1.1.1.](#) Primary authoritative DNS service

Ideally, SRP updates a primary authoritative DNS server for a particular zone. This DNS server acts as the sole source of truth for names within the DNS zone in which SRP services are published. Redundancy is provided by secondary DNS servers, if needed. However, this approach has some drawbacks.

First, it requires 100% availability on the part of a DNS primary authoritative server for the zone. If the primary server is not available for some period of time, new services appearing on the network cannot be registered until primary authoritative service is

restored.

The second drawback is that there is no automatic method for managing DNS authoritative service. This means that such a service requires an operator to set it up. What it means to set up such a service is that the following capabilities are provided:

- \* An host must be available to act as a primary authoritative DNS server
- \* The zone advertised by that server must be delegated, so that the local resolver can successfully answer queries in that zone

- \* The local resolver must be able to provide local browsing domain advertisements [RFC6763 [section 11](#)].

#### [1.1.2](#). Multicast DNS Advertising Proxy

An existing alternative to the use of DNS authoritative services for advertising SRP registrations the advertising proxy [[draft-tlsc-advertising-proxy](#)]. An advertising proxy advertises the contents of the SRP update zone using multicast DNS on links on which the need for such advertisements is anticipated. This works well for stub networks [[draft-lemon-stub-networks](#)], where services advertised on the stub network must be visible both on the stub network and on the adjacent infrastructure network, but do not generally need to be discoverable on other networks.

One drawback of the advertising proxy model, however, is that there is no shared database from which to advertise services registered by SRP. As a consequence, some of the guarantees provided by SRP, particularly first come, first served naming [[draft-ietf-dnssd-srp](#)]. Because advertising proxies are set up automatically on an ad-hoc basis, coordination between advertising proxies is not present, which means that if two devices claim the same name, but register with different SRP servers, the conflict is not detected until the service is advertised using mDNS. In practice, this results in frequent renaming of services, which means that consumers of services need to carefully follow each service that they use as the name changes over time.

An additional drawback is that, from the perspective of the SRP

client, SRP service is not unified: SRP servers tend to come and go, and whenever the SRP service with which a particular client has registered goes offline, the client has to notice that this has happened, discover a new SRP server, and re-register, or else it becomes unreachable.

### [1.1.3.](#) SRP Replication

This document describes a replication mechanism which eliminates the need for a single authoritative source of truth, as in the Primary Authoritative DNS model, while eliminating the drawbacks of the Advertising Proxy model. SRP Replication servers discover each other automatically. Each replication server maintains a copy of the SRP zone which is kept up to date on a best-effort basis.

SRP Replication has several benefits:

- \* As long as one SRP replication partner remains online at all times, SRP state is maintained when individual SRP replication partners go offline
- \* Name collisions when SRP clients change servers are avoided
- \* SRP service on a stub network can appear as an anycast service, so that SRP clients do not see an apparent change in servers and re-register when the server with which they most recently registered goes offline

## [2.](#) Implementation

SRP Replication relies on the fact that any given client is always registering with exactly one SRP server at any given time. This means that when an SRP server receives an SRP update from a client, it can be sure that no other SRP server has a more recent version of that SRP client's registration. Consequently, that SRP server can behave as if it is the source of truth for that client's registration, and other SRP servers can safely assume that any data they have about the client that is less recent can be replaced with the new registration data.

## [2.1.](#) Naming of a common service zone

In order for SRP replication partners to replicate a zone, they must agree upon a common name for the zone. We will describe two mechanisms for agreeing on a common zone here.

### [2.1.1.](#) Zone name based on network name

Network names aren't guaranteed to be unique, but tend to be unique for any given site. In the case of ad-hoc (permissionless) SRP-based service, such as an advertising proxy or an authoritative service using a locally-served zone [<https://www.iana.org/assignments/locally-served-dns-zones/locally-served-dns-zones.xhtml>], because the DNS zone name isn't required to be globally unique, a zone name based on the network name is an easy solution to generating a unique zone name.

When generating a zone name based on a network name, the zone name could be based on a locally configured global zone name, e.g. 'example.com'. It could be based on a locally-managed locally-served name, e.g. 'home.arpa'. Or it could be based on an unmanaged locally-served name, for which we propose to use the root name 'local.arpa.' For the rest of this section we will assume that the specific setting determines which of these domains will be used, and refer to whichever domain that is as DOMAIN.

For zone names based on the network name, the network type should be used as a differentiator, in case there are two different local network types with the same name. So, for example, 'WiFi.DOMAIN.'

#### [2.1.1.1.](#) Zone name based on WiFi SSID

If the zone being represented is a WiFi network, then the zone name for the network should be constructed using the WiFi SSID followed by 'WiFi.DOMAIN.'. For example, if the SSID is "Example Home" then the zone name would be 'Example Home.WiFi.DOMAIN.'. Note that spaces and special characters are allowed in domain names.

#### [2.1.1.2.](#) Zone name based on Thread network name

If the zone being represented is a Thread [Thread] network, then the

zone name for the network should be constructed using the Thread network name. For example, if the Thread network name is "openthread" then the zone name would be 'openthread.thread.DOMAIN.'

#### [2.1.2.](#) Zone name based on local configuration

The above examples assume that it makes sense for each separate subnet to be its own separate zone. However, since SRP guarantees name uniqueness using the first-come, first-served mechanism, it doesn't rely on mDNS's guarantee of per-link uniqueness. Consequently, it is not required that an SRP zone be constrained to the set of services advertised on a single link. For this reason, when it is possible to know that some set of links are all managed by the same set of SRP replication partners, and a name is known for that set of links, that name can be used. To avoid possible collisions, the subdomain 'srp' is used to indicate that this zone is an SRP zone. So in this case the link name would be the locally-known shared name, followed by 'srp.DOMAIN.'

An example of such a scenario would be Apple's HomeKit, in which all HomeKit accessories, regardless of which home network link they are attached to, all are shared in the same namespace. Suppose the HomeKit home's name is "Example Home". In such a situation, the domain name 'Example Home.srp.DOMAIN' could be used.

#### [2.1.3.](#) Zone name based on DNS-SD discovery

Another option for naming the local SRP Replication zone would be to use DNS-SD advertisements. This is particularly useful since each SRP replication partner advertises itself using DNS-SD, so there is a convenient place to put this information. To advertise a zone name based on DNS-SD discovery, the SRP replication partner should add two fields to the TXT record of the service instance. The first field is

the domain field: 'domain=name'. This indicates a proposed SRP replication zone name. The second is the join field. If 'join=yes' then other SRP replication servers are encouraged to use the domain name that appears in the domain field rather than creating a new domain.

#### [2.2.](#) Advertising one's own replication service

An SRP replication service only advertises its replication service if no other service for the domain it is replicating is already present, or else if it has successfully connected to and synchronized with all of the SRP services it sees advertised for the domain it is configured to replicate.

SRP replication service is advertised using DNS-SD [[RFC6763](#)]. The service name is '\_srpl-tls.\_tcp'. Each SRP replication partner should have its own hostname, which when combined with the service instance name and the local DNS-SD domain name will produce a service instance name, for example 'example-host.\_srpl-tls.\_tcp.local.' The domain under which the service instance name appears will be 'local' for mDNS, and will be whatever domain is used for service registration in the case of a non-mDNS local DNS-SD service.

SRP replication uses DNS port 853 [[RFC7858](#)] and is based on DNS Stateful Operations [[RFC8490](#)]. Therefore, the SRV record for the example we've given would be:

```
example-host._srpl-tls._tcp.local. IN SRV 0 0 853 example-  
host.local.
```

The TXT record for SRP replication advertises the following fields:

dataset a 64-bit number encoded as hexadecimal ASCII, produced with a high-quality random number generator [[RFC4086](#)]. The dataset ID is used by SRP servers to establish a common SRP dataset for a domain.

join 'yes' or 'no'. Indicates whether other SRP replication servers are invited to join in replicating the dataset.

prec a 32-bit number representing the precedence of the server advertising the TXT record, represented in decimal

domain the domain name that this dataset is intended to represent

This identifier need not be persistent across SRP replication partner restarts. So in our example the TXT record might look like this:



rec=1\008join=yes

(Note that each name/value pair in the TXT record is length-encoded, so the '#', the '\025', the '\006' and the '\008' are the lengths of the four name/value pairs.)

### [2.3.](#) Discovering other replication services

SRP Replication is a cooperative process. In order to ensure cooperation between SRP replication partners on a link, it is necessary that each replication partner be aware of other potential partners. This is accomplished by maintaining a continuous browse for services of the service type "\_srpl-tls.\_tcp".

An SRP Replication Partner MUST maintain an ongoing DNS-SD browse on the service name '\_srpl-tls.\_tcp' within the local browsing domain. The ongoing browse will produce two different types of events: 'add' events and 'remove' events. When the browse is started, it should produce an 'add' event for every SRP replication partner currently present on the network, including the partner that is doing the browsing. Whenever a partner goes offline, a 'remove' event should be produced. 'remove' events are not guaranteed, however.

When a new service is added, the SRP partner checks to see if it is in a compatible domain. If the SRP partner has a domain to advertise, it compares that domain to the domain advertised in the added service instance: if they are not the same, then this instance is not a candidate for connection, and should be ignored.

#### [2.3.1.](#) Getting an SRP zone name from other SRP services

If the SRP partner does not have a domain to advertise, then when it begins to browse for partners, it sets a timer for DOMAIN\_DISCOVERY\_TIMEOUT seconds.

If the SRP partner does not have a domain to advertise, and is therefore willing to join an existing domain, it checks to see if the TXT record for the service indicates that joining is permitted. If so, the SRP partner adopts the provided domain name. Once it has adopted such a domain name, it updates its own TXT record to indicate that domain name, and sets the 'join=yes' key/value pair in the TXT record. It also cancels the DOMAIN\_DISCOVERY\_TIMEOUT timer.

---

If the DOMAIN\_DISCOVERY\_TIMEOUT timer goes off, then the SRP partner MUST propose a zone name using one of the methods mentioned previously in [Section 2.1](#). It advertises that zone name in its TXT record, with 'join=yes'. It then sets a new timer for DOMAIN\_PROPOSE\_TIMEOUT seconds.

While waiting for the DOMAIN\_PROPOSE\_TIMEOUT timer to go off, any new 'add' events that arrive are examined to see if they are potential domains to join. If a potential domain to join is seen, and it is the same as the proposed domain, then the partner adopts that domain and treats it as its domain to advertise. It then cancels the DOMAIN\_DISCOVERY\_TIMEOUT timer.

At this point the SRP replication partner has a domain to advertise: either the one it produced, or one that it discovered.

### [2.3.2](#). Establishing Replication with an existing Dataset

Once an SRP replication partner has settled on a domain to advertise, it must either join other SRP replication partners in replicating that domain, or if it is the first, it must advertise its willingness to participate in replicating the domain. In order to do this, it must settle on a dataset ID.

The dataset ID is a random 64-bit number, generated by the first server to offer that dataset. There should always be exactly one dataset ID per domain, but the dataset ID has a separate purpose: it represents the set of data that is being replicated by a set of cooperating SRP replication partners. This data is then offered under the agreed-upon domain, but it's possible that there might be several sets of SRP replication partners that agree to replicate a particular domain, and then some event occurs which renders these partners visible to each other. When this happens, the independent sets of partners must converge on a single dataset. This is done using the dataset ID.

When more than one dataset ID is present for a particular domain, the dataset ID that is numerically lowest is preferred. This means that SRP replication partners that are currently replicating a dataset with a numerically higher dataset ID will have to abandon that dataset and join together in replicating the numerically lowest dataset. Servers that are not replicating the numerically lowest dataset will therefore stop advertising SRP replication service and begin attempting to join in replicating the preferred dataset.

When a set of servers are advertising a particular dataset ID, the

server with the lowest precedence is primary. The primary server is responsible for handing out precedence values to new partners as they

join in replicating the dataset. Precedence IDs are always allocated starting with the precedence that is one greater than the primary's precedence.

When an SRP replication partner has stopped advertising a particular dataset ID, or has just started and therefore hasn't started advertising a particular dataset ID, and there is a dataset ID present that it can join in replicating, it attempts to connect to the SRP replication partner that is primary for the dataset. If the startup handshake succeeds, the primary will assign a new precedence to the connecting partner as part of the handshake.

Once the synchronization phase has finished, the connecting partner will begin advertising the SRP service for the chosen domain using the new dataset ID and the precedence it received from the primary. The connecting server will then also attempt to connect to every SRP partner it sees advertising the same dataset ID and a lower precedence.

It is possible that an SRP partner will attempt to join in replicating a dataset, but the primary for that dataset may have discontinued service, but the advertisement for the primary is still in the cache. In this case, the SRP partner will attempt to reconfirm the primary's advertisement. In mDNS, this is done as described in [Section 10.4 of \[RFC6762\]](#). For DNS Push connections, this is done using the RECONFIRM message, described in [Section 6.5 of \[RFC8765\]](#). For regular (polled) DNS, the SRP partner must trigger a new DNS query. If the primary advertisement is successfully confirmed, this indicates that there is a problem connecting to the primary, in which case the connecting partner SHOULD discontinue attempting to connect for at least MIN\_RECONNECT\_AFTER\_FAILURE seconds.

Otherwise, the connecting partner will attempt to connect to the new primary if there is one. If there are no other servers advertising the dataset ID, then the connecting partner reverts to attempting to start its own replication of that dataset.

### [2.3.3.](#) Establishing Replication When There Is No Existing Dataset

When an SRP replication partner has attempted to discover partners with which to connect, and failed to do so, it then creates its own dataset ID and precedence and begins advertising that dataset. Both the dataset ID and precedence should be generated using a non-deterministic random number generator. The dataset ID should be a random number greater than or equal to zero and less than  $2^{64}$ . The precedence should be a random number greater than or equal to 0 and less than  $2^{15}$ . The reason for the upper limit is to allow for a large range of numbers toward which the precedence can increase.

The replication partner begins advertising this new dataset as soon as the dataset ID and precedence have been generated. As in the previous section, if a new dataset ID is seen shortly afterwards, this most likely indicates that two SRP replication instances came up at the same time; in this case as with the previous one, the lower dataset ID is preferred, and the partner advertising the higher dataset ID abandons that dataset ID to join the partner with the lower dataset ID.

The replication partner that first advertises the dataset is the primary replication partner for that dataset. It is responsible for assigning precedences to new partners.

### [2.3.4.](#) Conflicting precedence values

It is possible that two SRP replication partners that see different service advertisements could identify different SRP replication servers as primary and attempt to get their precedence values from those different servers. When this happens, it's possible that they might both get the same precedence value. When this occurs, as soon

each partner sees another partner advertising its precedence in an SRP replication advertisement, it must discontinue advertising and restart the dataset discovery process.

#### [2.3.5.](#) Handling primary transitions

An SRP partner either identifies itself as primary or not. When an SRP partner is primary, it never connects to other SRP servers--it only receives connections. When a non-primary partner connects to the primary partner, it knows it is connecting to the primary partner. If the connection with the primary drops, or if the primary's advertisement goes away, then the non-primary evaluates the set of advertisements that it sees. If its precedence is lowest, it identifies itself as primary.

Lemon

Expires 11 May 2022

[Page 11]

---

Internet-Draft

DNS-SD SRP Replication

November 2021

Non-primary servers receive updates from the primary whenever the maximum precedence value changes. Non-primary servers should track this precedence value. When a non-primary becomes primary, it should add ten to the most recently received precedence value, so as to skip any possible precedence assignments that haven't yet propagated.

#### [2.4.](#) Discovering the addresses of partners

When a partner is discovered, two new ongoing mDNS queries are started on the hostname indicated in the SRV record of the partner: one for A records, and one for AAAA records. Each time an address 'add' event is seen, either for an 'A' record or an 'AAAA' record, the partner adds the address to the list of addresses belonging to that partner.

#### [2.5.](#) Establishing Communication with a replication partner

When an address is added to a partner's address list, the partner first checks to see if the address is one of its own addresses. If so, then the partner is marked "me", and no connection is attempted to it. This is somewhat safer than comparing hostnames, since a hostname collision can result in renaming.

If the partner is not marked 'me', then the partner checks to see if it has an existing outgoing connection to that partner. If it does

not, then it checks to see whether it has disabled outgoing connections to that partner. If not, then it attempts to connect on the new address.

When a connection fails, it advances to the next address in the list, if there is one. If there are no remaining addresses, the partner sets a timer for RECONNECT\_INTERVAL seconds. When this timer expires, it starts again at the beginning of the list and attempts to connect to the first address, iterating again across the list until a connection succeeds or it runs out of addresses.

Additionally, when an address is added, it is checked against the list of unidentified incoming connections. If a match is found, and the partner is marked "me," then the unidentified connection is removed from the list and dropped. Otherwise, it is attributed to the matching partner, and the protocol is started at the point of receiving an incoming connection.

When an outgoing connection succeeds, the partner sends its server ID.

## [2.6.](#) Incoming connections

When an incoming connection is received, it is checked against the partner list based on the source address of the incoming connection. If the address appears on the list of addresses for a partner, then the connection is attributed to that partner. If no matching partner is found, a timer of UNIDENTIFIED\_PARTNER\_TIMEOUT seconds is set, and the incoming connection is added to the list of "unidentified" connections.

If a matching partner is found, then the partner waits for an incoming partner ID. When such an ID is received, it is compared to the partner's server-id. If the incoming server ID is the same as or greater than the partner's server ID, the connection is dropped. Otherwise, the connection proceeds to the "initial synchronization" state.

## [2.7.](#) Eliminating extra connections

When an outgoing connection succeeds, the partner sends its server ID to the partner. When an incoming connection succeeds, the partner waits for a server ID. Because both connections are partner connections, and we only need one connection, the partner with the higher server ID acts as the client and the partner with the lower server ID acts as the server. If the server IDs are equal, then the connecting server generates a new server ID, updates its TXT record, and re-does the comparison.

## [2.8.](#) Initial synchronization

The connecting partner begins the session by sending its server ID. The receiving partner waits for a server ID, and when it receives one, does the server ID comparison mentioned earlier. If the connection survives the comparison, then the server sends a response to the session message and waits for the client to request a list of update candidates.

The connecting partner waits for a response to the initial session message, and when it is received, requests that the server send candidates.

### [2.8.1.](#) Sending candidates

When a partner receives a "send candidates" message that it is expecting to receive, it generates a candidate list from the list of known SRP clients. This list includes SRP clients that have registered directly with the partner, and SRP clients that have been received through SRP replication updates. Each candidate contains a hostname, a time offset, and a key identifier.

The key identifier is computed as follows:

```

uint32_t key_id(uint8_t *key_data, int key_len) {
    uint32_t key_id = 0;
    for (int i = 0; i < key_data_len; i += 4) {
        key_id += ((key_data[i] << 24) | (key_data[i + 1] << 16) |
                    (key_data[i + 2] << 8) | (key_data[i + 3]));
    }
    return key_id;
}

```

When a partner receives a candidate message during the synchronization process, it first searches for an SRP registration with a hostname that matches the hostname in the candidate message. It then compares the key ID to the key ID in the candidate message. If the key ID doesn't match, it sends back a candidate response status of "conflict". If the key ID does match, it compares the time provided to the time the existing host entry was received. If the time of the update is later, it sends a "send host" response. If it is earlier or the same, it sends a "continue" response. If there is no matching host entry for the candidate message, the partner sends a "send host" response.

When a partner receives a candidate response with a status of "send host", it generates a host message, which contains the hostname, the time offset, and the SRP message that was received from the host. The partner then applies the SRP update message as if it had been received directly from the SRP client. The host update time sent by the partner is remembered as the time when the update was received from the client, for the purposes of future synchronization.

When a partner is finished iterating across its list of candidates, it sends a "send candidates" response.

When a partner receives a "send candidates" response, if it is the server, it sends its own "send candidates" message, and processes any proposed candidates.

When a partner that is a server receives a "send candidates" response, it goes into the "routine operation" state. When a partner that is a client sends its "send candidates" response, it goes into the "routine operation" state.



## [2.9.](#) Routine Operation

During routine operation, whenever an update is successfully processed from an SRP client, the partner that received that update queues that update to be sent to each partner to which it has a connection, whether server or client. If there are no updates pending to a particular client, the update is sent immediately. Otherwise, it's sent when the outstanding update is acknowledged.

When during routine operation a partner receives a host update from its partner, it immediately applies that update to its local SRP zone. This is based on the assumption that a new update is always more current than a copy of the host information in its database.

## [3.](#) Protocol Details

The DNS-SD SRP Replication Protocol (henceforth SRPL) is based on DNS Stateful Operations [[RFC8490](#)]. Each SRP replication partner creates a listener on port 853, the DNS-over-TLS [[RFC7858](#)] reserved port. This listener can be used for other DNS requests as well.

Participants in the protocol are partners. To distinguish between partners, the terms "partner" and "partner" are used. "Partner" refers to the partner that is communicating or receiving communication. "Partner" refers to the other partner. Partners can be clients or servers: a partner that has established a connection to a partner is a client; a partner that has received a connection from a partner is a server.

### [3.1.](#) DNS Stateful Operations considerations

DNS Stateful Operations is a DNS per-connection session management protocol. DNS Push session management includes session establishment as well as session maintenance.

#### [3.1.1.](#) DSO Session Establishment

An DSO session for an SRPL connection can be established either by simply sending the first SRPL message, or by sending a DSO Keepalive message. [Section 5.1 of \[RFC8490\]](#).

### [3.1.2.](#) DSO Session maintenance

DSO sessions can be active or idle. As long as the SRPL protocol is active on a connection, the DSO state of the connection is active. DSO sessions require occasional keepalive messages. The default of fifteen seconds is adequate for SRPL.

An idle DSO session must persist for long enough that there is a chance for the browse that identifies it to succeed. Therefore, the minimum DSO session inactivity timeout is  $2 \times \text{UNIDENTIFIED\_PARTNER\_TIMEOUT}$  seconds.

### [3.2.](#) DSO Primary TLVs

Each DSO message begins with a primary TLV, and contains secondary TLVs with additional information. The primary TLVs used in the SRPL protocol are as follows:

#### [3.2.1.](#) SRPL Session

DSO-TYPE code: SRPLSession. Introduces the SRPL session. The SRPL session TLV contains no data, just the type and length. SRPL Session primary TLV requests do not include any secondary TLVs. SRPL Session requests are DSO requests: the recipient is expected to send a response TLV. Both request and response TLVs have the same format. An SRPL Session response may include an SRPL Precedence secondary TLV.

##### [3.2.1.1.](#) SRPL client behavior

The SRPL Session request is sent by a partner acting as a client to its partner once the TLS connection to the partner, acting as a server, has succeeded. The SRPL session message establishes the DSO connection as an SRP protocol connection. If it is the first DSO message sent by the partner acting as a client, then it also establishes the DSO session.

When the SRPL partner acting as a client receives a response to its SRPL session message, it sends an SRPL Send Candidates message.

When an SRPL partner receives an SRPL Precedence secondary TLV in an SRPL Session response, if it thinks it is connected to the primary partner, it sets its precedence to the assigned value. If it thinks it is connecting to a non-primary, then it disconnects and waits `NON_PRIMARY_RESETTLE_TIME` seconds before reconnecting. It also attempts to reconfirm the service advertisement for the partner it thinks is primary.

#### [3.2.1.2.](#) SRPL server behavior

An SRPL partner acting as a server that receives an SRPL Session request checks to see if the connection on which it was received is already established. If so, this is a protocol error, and the SRPL partner MUST drop the connection.

When an SRPL partner that is primary receives an SRPL Session request, the SRPL Session response MUST include an SRPL Precedence secondary TLV, which assigns a precedence to the connecting SRPL partner.

#### [3.2.2.](#) SRPL Send Candidates

DSO-TYPE code: SRPLSendCandidates. Requests the partner to send its candidates list. The SRPL Send Candidates message contains no additional data. The SRPL Send Candidates primary TLV does not include any secondary TLVs. SRPL Send Candidates messages are DSO requests: the recipient is expected to send a response TLV. Both request and response TLVs have the same format.

##### [3.2.2.1.](#) SRPL client behavior

An SRPL partner acting as a client MUST send an SRPL Send Candidates request after it has received an SRPL Session response. It MUST NOT send this request at any other time.

An SRPL partner acting as a client expects to receive an SRPL Send Candidates message after it has received an SRPL Send Candidates response. If it receives an SRPL Send Candidates message at any other time, this is a protocol error, and the SRPL partner should drop its connection to the server.

##### [3.2.2.2.](#) SRPL server behavior

An SRPL partner acting as a server expects to receive an SRPL Send Candidates request after it has sent an SRPL Session response. If it receives an SRPL Candidates request at any other time, this is a protocol error, and it MUST drop the connection.

An SRPL partner acting as a server MUST send an SRPL Send Candidates request after it has sent an SRPL Send Candidates response.

An SRPL partner acting as a server MUST enter the "normal operations" state after receiving an SRPL Send Candidates response from its partner.

### [3.2.3.](#) SRPL Candidate

DSO-TYPE code: SRPLCandidate. Announces the availability of a specific candidate SRP client registration. The SRPL Candidate message contains no additional data. SRPL Candidate messages are DSO requests: the recipient is expected to send a response TLV. Both request and response TLVs have the same format.

#### [3.2.3.1.](#) Required secondary TLVs

The SRPL Candidate request MUST include the following secondary TLVs: SRPL Hostname, SRPL Time Offset, and SRPL Key ID. If an SRPL partner receives an SRPL Candidate request that doesn't contain all of these secondary TLVs, this is a protocol error, and the partner MUST drop the connection.

The SRPL Candidate response MUST include one of the following status TLVs: SRPL Candidate Yes, SRPL Candidate No, or SRPL Conflict. If an SRPL partner receives an SRPL Candidate response which does not contain exactly one of these TLVs, this is a protocol error, and the partner MUST drop the connection.

#### [3.2.3.2.](#) SRPL partner common behavior

SRPL partners expect to receive SRPL Candidate messages between the time that they have sent an SRPL Send Candidates message and the time that they have received an SRPL Send Candidates response. If an SRPL Candidate message is received at any other time, this is a protocol error, and the partner MUST drop the connection.

Partners MUST NOT send SRPL Candidate requests if they have sent any SRPL Candidate or SRPL host requests that have not yet received responses. Partners receiving SRPL Candidate requests when they have

not yet responded to an outstanding SRPL Candidate request or SRPL Host request MUST treat this as a protocol failure and drop the connection.

When a partner receives a valid SRPL Candidate message, it checks its SRP registration database for a host that matches both the SRPL Hostname and SRPL Key ID TLVs. If such a match is not found, the partner sends an SRPL Candidate response that includes the SRPL Candidate Yes secondary TLV.

If a match is found for the hostname, but the Key ID doesn't match, this is a conflict, and the partner sends an SRPL Candidate response with the SRPL Conflict secondary TLV.

If a match is found for the hostname, and the key ID matches, then the partner computes the update time of the candidate by subtracting the value of the SRPL Time Offset TLV from the current time in seconds. This computation should be done when the SRPL Candidate message is received to avoid clock skew. If 'candidate update time' - 'local update time' is greater than SRPL\_UPDATE\_SKEW\_WINDOW, then the candidate update is more recent than the current SRP registration. In this case, the partner sends an SRPL Candidate response and includes the SRPL Candidate Yes secondary TLV. The reason for adding in some skew is to account for network transmission delays.

#### [3.2.4.](#) SRPL Host

DSO-TYPE code: SRPLHost. Provides the content of a particular SRP client registration. The SRPL Host message contains no additional data. SRPL Host messages are DSO requests: the recipient is expected to send a response TLV. Both request and response TLVs have the same format.

##### [3.2.4.1.](#) Required secondary TLVs

The SRPL Host request MUST include the following secondary TLVs: SRPL Hostname, SRPL Key ID, and one or more SRPL Host Message TLVs. If an SRPL partner receives an SRPL Candidate request that doesn't contain all of these secondary TLVs, this is a protocol error, and the

partner MUST drop the connection.

The SRPL Host message always includes at least one SRPL Host Message TLV, which contains the most recent update the SRP server has received for that host. However, in some cases an update for a host may update some, but not all, service instances that reference that host; in this case, the SRPL Host request MUST include all of the previously received SRP updates that would be required to reconstruct the current state of the host registration on the server sending the SRPL Host request.

#### [3.2.4.2.](#) SRPL partner common behavior during synchronization

SRPL partners expect to receive either zero or one SRPL Host requests after sending an SRPL Candidate response with a SRPL Candidate Yes secondary TLV. If an SRPL Host request is received at any other time during synchronization, this is a protocol error, and the partner MUST drop the connection. The only time that an SRPL Host request would not follow a positive SRPL Candidate response would be when the candidate host entry's lease expired after the SRPL Candidate request was sent but before the SRPL Candidate response was received.

SRPL partners send SRPL Host requests during synchronization when a valid SRPL Candidate response has been received that includes an SRPL Candidate Yes secondary TLV. The host request is generated based on the current candidate (the one for which the SRPL Candidate request being responded to was sent).

#### [3.2.4.3.](#) SRPL partner common behavior during normal operations

When an SRPL partner during normal operations receives and has successfully validated an SRP update from an SRP client, it MUST send that update to each of its connected partners as an SRPL Host request. If the connection to a particular partner is not busy, and there are no updates already queued to be sent, it MUST send the SRPL Host message immediately. Otherwise, it MUST queue the update to send when possible. The queue MUST be first-in, first-out.

After an SRPL partner has sent an SRPL Host request to a partner, and before it receives a corresponding SRPL Host response, it MUST NOT send any more SRPL Host messages to that partner.

When an SRPL partner receives an SRPL Host request during normal operations, it MUST apply it immediately. While it is being applied, it MUST NOT send any other SRPL Host requests to that partner.

When an SRPL Host request has been successfully applied by an SRPL partner, the partner MUST send an SRPL Host response.

If an SRPL partner receives an SRPL Host request while another SRPL Host request is being processed, this is a protocol error, and the partner MUST drop the connection to its partner.

### [3.3.](#) DSO Secondary TLVs

In addition to the Primary TLVs used to send requests between SRPL partners, we define secondary TLVs to carry formatter information needed for various SRPL requests.

#### [3.3.1.](#) SRPL Candidate Yes

DSO-TYPE code: SRPLCandidateYes. In an SRPL Candidate response, indicates to the partner that an SRPL Host message for the candidate is wanted and should be sent.

Appears as a secondary TLV in SRPL Candidate responses. MUST NOT appear in any other SRPL request or response. MUST NOT appear in addition to either SRPL Conflict or SRPL Candidate No secondary TLVs.

#### [3.3.2.](#) SRPL Candidate No

DSO-TYPE code: SRPLCandidateNo. In an SRPL Candidate response, indicates to the partner that an SRPL Host message for the candidate is not wanted and should not be sent.

Appears as a secondary TLV in SRPL Candidate responses. MUST NOT appear in any other SRPL request or response. MUST NOT appear in addition to either SRPL Conflict or SRPL Candidate Yes secondary TLVs.

#### [3.3.3.](#) SRPL Conflict

DSO-TYPE code: SRPLConflict. In an SRPL Candidate response, indicates to the partner that an SRPL Host message for the candidate is not wanted and should not be sent. Additionally indicates that the proposed host conflicts with local data. This indication is informative and has no effect on processing.

Appears as a secondary TLV in SRPL Candidate responses. MUST NOT appear in any other SRPL request or response. MUST NOT appear in addition to either SRPL Candidate Yes or SRPL Candidate No secondary TLVs.

#### [3.3.4.](#) SRPL Hostname

DSO-TYPE code: SRPLHostname. In an SRPL Candidate or SRPL Host request, indicates to the partner the hostname of an SRP registration. The hostname is represented in DNS wire format [Section 3.1 of \[RFC1035\]](#).

Required as a secondary TLV in SRPL Candidate and SRPL Host requests. MUST NOT appear in any other SRPL request or response.

#### [3.3.5.](#) SRPL Host Message

DSO-TYPE code: SRPLHostMessage. In an SRPL Host request, conveys four data objects in order:

- \* the lease time and key lease time returned to the client, represented as two unsigned 32-bit numbers in units of seconds.
- \* the time offset at which the message was received, represented as a 32-bit unsigned number of seconds. The time offset is computed as the difference between the time when the SRPL Host Message TLV is being constructed for transmission, and the time when the SRP update contained in the SRPL Host Message was received.

- \* the SRP Update message received from the SRP client. This contains the contents of the message, but not any IP, UDP, TCP or TLS headers that may have encapsulated it.

The content of the SRPL Host Message is used to update the host on



the partner receiving the request. Note that the SRP message being sent can't be modified by the SRPL partner sending it, so in order to validate the message (assuming that the signature includes a nonzero time), the validation process should adjust the current time by the time offset included in the SRPL Time Offset TLV when comparing against the signature time when checking for replay attacks. The computation of the current time of signing should be done when the message is received to avoid clock skew that might result from processing delays.

Required as a secondary TLV in SRPL Host requests. MUST NOT appear in any other SRPL request or response.

#### [3.3.6.](#) SRPL Time Offset

DSO-TYPE code: SRPLTimeOffset. In an SRPL Candidate, conveys the difference between the time the SRP update was received from the SRP client and the current time on the partner generating the request, in seconds.

Required as a secondary TLV in SRPL Candidate and SRPL Host requests. MUST NOT appear in any other SRPL request or response.

#### [3.3.7.](#) SRPL Key ID

DSO-TYPE code: SRPLKeyID. In an SRPL Candidate, conveys the key ID of the SRP client.

Required as a secondary TLV in SRPL Candidate requests. MUST NOT appear in any other SRPL request or response.

### [4.](#) Security Considerations

SRP replication basically relies on the trustworthiness of hosts on the local network. Since SRP itself relies on the same level of trust, SRP replication doesn't make things worse. However, when the option to have a central SRP server is available, that is likely to be more trustworthy.

## [5.](#) Delegation of 'local.arpa.'

In order to be fully functional, the owner of the 'arpa.' zone must add a delegation of 'local.arpa.' in the '.arpa.' zone [[RFC3172](#)]. This delegation should be set up as was done for 'home.arpa', as a result of the specification in [Section 7 of \[RFC8375\]](#).

## [6.](#) IANA Considerations

### [6.1.](#) 'srpl-tls' Service Name

IANA is requested to add a new entry to the Service Names and Port Numbers registry for srpl-tls with a transport type of tcp. No port number is to be assigned. The reference should be to this document, and the Assignee and Contact information should reference the authors of this document. The Description should be as follows:

Availability of DNS-SD SRP Replication Service for a given domain is advertised using the "\_srpl-tls.\_tcp.<domain>." SRV record gives the target host and port where DNS-SD SRP Replication Service is provided for the named domain.

### [6.2.](#) DSO TLV type code

The IANA is requested to add the following entries to the 16-bit DSO Type Code Registry. Each type mnemonic should be replaced with an allocated type code, both in this table and elsewhere in the document. RFC-TBD should be replaced with the name of this document once it becomes an RFC.

Internet-Draft

DNS-SD SRP Replication

November 2021

Type	Name	Early Data	Status	Reference
SRPLSession	SRPL Session	No	STD	RFC-TBD
SRPLSendCandidates	SRPL Send Candidates	No	STD	RFC-TBD
SRPLCandidate	SRPL Candidate	No	STD	RFC-TBD
SRPLHost	SRPL Host	No	STD	RFC-TBD
SRPLCandidateYes	SRPL Candidate Yes	No	STD	RFC-TBD
SRPLCandidateNo	SRPL Candidate No	No	STD	RFC-TBD
SRPLConflict	SRPL Conflict	No	STD	RFC-TBD
SRPLHostname	SRPL Hostname	No	STD	RFC-TBD
SRPLHostMessage	SRPL Host Message	No	STD	RFC-TBD
SRPLTimeOffset	SRPL Time Offset	No	STD	RFC-TBD
SRPLKeyID	SRPL Key ID	No	STD	RFC-TBD

Table 1

### 6.3. Registration and Delegation of 'local.arpa' as a Special-Use Domain Name

IANA is requested to record the domain name local.arpa.' in the Special-Use Domain Names registry [[SUDN](#)]. IANA is requested, with the approval of IAB, to implement the delegation requested in [Section 5](#).

IANA is further requested to add a new entry to the "Transport-Independent Locally-Served Zones" subregistry of the the "Locally-Served DNS Zones" registry [[LSDZ](#)]. The entry will be for the domain 'local.arpa.' with the description "Ad-hoc DNS-SD Special-Use Domain", listing this document as the reference.

## [7.](#) Informative References

## [8.](#) Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC3172] Huston, G., Ed., "Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa")", [BCP 52](#), [RFC 3172](#), DOI 10.17487/RFC3172, September 2001, <<https://www.rfc-editor.org/info/rfc3172>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", [RFC 6762](#), DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", [RFC 7858](#), DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8375] Pfister, P. and T. Lemon, "Special-Use Domain 'home.arpa.'", [RFC 8375](#), DOI 10.17487/RFC8375, May 2018, <<https://www.rfc-editor.org/info/rfc8375>>.
- [RFC8490] Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", [RFC 8490](#), DOI 10.17487/RFC8490, March 2019, <<https://www.rfc-editor.org/info/rfc8490>>.
- [RFC8765] Pusateri, T. and S. Cheshire, "DNS Push Notifications", [RFC 8765](#), DOI 10.17487/RFC8765, June 2020, <<https://www.rfc-editor.org/info/rfc8765>>.

[SUDN] "Special-Use Domain Names Registry", July 2012,  
<<https://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>>.

[LSDZ] "Locally-Served DNS Zones Registry", July 2011,  
<<https://www.iana.org/assignments/locally-served-dns-zones/locally-served-dns-zones.xhtml>>.

#### Author's Address

Ted Lemon  
Apple Inc.  
One Apple Park Way

Lemon

Expires 11 May 2022

[Page 25]

---

Internet-Draft

DNS-SD SRP Replication

November 2021

Cupertino, California 95014  
United States of America

Email: [mellon@fugue.com](mailto:mellon@fugue.com)

Lemon

Expires 11 May 2022

[Page 26]