

SIP.edu Cookbook



Contents

[Introduction](#)

[Getting Started](#)

[DNS](#)

[Proxies](#)

[Gateways](#)

[User Agents](#)

[Directory Considerations](#)

[Security Considerations](#)

[Deployments](#)

[Glossary](#)

[Contacts](#)

[Related Links](#)

DNS Configuration

Jeremy George <jeremy.george@yale.edu> (May 12, 2003)

Locating SIP Services

How users know a server - its name and how they locate that server on the net - its address have long been separated to provide a lasting identifier without hindering the flexibility required by local network administrators. Similarly, services need to be separated from the machines that provide the services, and for the same reason.

In addition it's convenient to be able to list one's username@domain as identical to one's email address just by changing the application prefix, without regard to what set of machines are actually providing the services. For example, email:alice.smith@bigu.edu and sip:alice.smith@bigu.edu.

RFC3263 specifies DNS as the preferred mechanism for determining the IP address, port and transport of the host to which a SIP request is sent. Transport must be determined because SIP requests can be sent via UDP, TCP, SCTP or TLS over TCP for secure, encrypted sessions, unlike many more limited protocols.

DNS provides two record types relevant to SIP requests: SRV and NAPTR. Some implementations will use SRV records only.

SRV Records

SRV records (specified in RFC2782) are in the form of

```
"_Service._Proto.Name TTL Class SRV Priority Weight Port Target"
```

For example,

```
"_sip._udp.bigu.edu 43200 IN SRV 10 10 5060 sipserver.bigu.edu."
```

The service is SIP.

The transport is UDP. Other values could be TCP, SCTP or TLS.

The cache lifetime is 12 hours (43,200 seconds.) This could be any positive signed 32 bit integer.

The class is IN (this is always true.)

The record type is SRV.

The priority is 10. With multiple SRV records the priority determines the proxy query order. Lower values are queried first.

The weight is 10. With multiple SRV records of similar priority, the weight determines proportionally how often a proxy is queried. Higher values are queried more often. So, a weight of 20 would be queried twice as often as one of 10. A weight of 30 would be queried three times as often as one of 10.

The port is 5060.

The proxy server FQDN is sipserver.bigu.edu and as is required in DNS the FQDN is terminated with a dot.

An example DNS implementation with redundant proxy servers might look like this:

```
bigu.edu IN SOA ns.bigu.edu. root.bigu.edu. (
    2003032001
    10800
    3600
    604800
    86400 )

bigu.edu.          43200 IN NS      ns.bigu.edu.
;
ns.bigu.edu.      43200 IN A      10.0.0.20
sipserver1.bigu.edu. 43200 IN A      10.0.0.21
sipserver2.bigu.edu. 43200 IN A      10.0.0.22
;
_sip_udp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver1.bigu.edu.
_sip_udp.bigu.edu. 43200 IN SRV 1 0 5060 sipserver2.bigu.edu.
_sip_tcp.bigu.edu. 43200 IN SRV 0 4 5060 sipserver1.bigu.edu.
_sip_tcp.bigu.edu. 43200 IN SRV 0 2 5060 sipserver2.bigu.edu.
_sips_tcp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver1.bigu.edu.
_sips_tcp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver2.bigu.edu.
```

In this configuration UDP SIP requests will always be sent to sipserver1 because the priority value is lower than sipserver2. Should that request fail, sipserver2 would be queried. In effect, sipserver2 is an emergency backup to sipserver1.

Two TCP SIP requests will be sent to sipserver1 for each one sent to sipserver2 because the weight for sipserver1 is twice that for sipserver2. In this case sipserver1 and sipserver2 are both being queried and load balanced. A weight such as this might be used where one machine is significantly more powerful than another.

Secure SIP requests will be equally load balanced between the two servers.

A SIP UA wanting to initiate a call to sip:bigcheese@bigu.edu will first send a DNS SRV lookup to bigu.edu. With a successful return, the SIP URI may be re-written to sip:bigcheese@sipserver1.bigu.edu. Absent information on bigu's preference for transport, the calling UA will choose the transport it prefers among the responses it gets.

In this example the calling UA could not use SCTP as a transport. The available choices are UDP, TCP, TLS or allow the call to fail.

In some circumstances allowing the call to fail may be the best choice. For example, the Health Insurance Portability and Accountability Act of 1996's (HIPAA) final security rule reported in the Federal Register February 20, 2003 says "we decided to make use of

encryption in the transmission process an addressable implementation specification. Covered entities are encouraged, however, to consider use of encryption technology for transmitting electronic protected health information, particularly over the Internet." The choice of which protocols to prefer or even accept is implementation-specific and should be considered based on the sensitivity of the transmitted information and the likelihood of hostile interception.

NAPTR Records

According to RFC3263 "NAPTR records provide a mapping from a domain to the SRV record for contacting a server with the specific transport protocol in the NAPTR service field." In other words NAPTR records provide a mechanism for the called domain to specify which protocols it prefers a SIP request to use.

NAPTR records (specified in RFC3403) are in the form of

```
"domain-name TTL Class NAPTR order preference flags service regexp target"
```

For example,

```
"bigu.edu. IN NAPTR 60 50 "s" "SIP+D2U" "" "_sip._udp.bigu.edu."
```

The domain name being queried. This is the portion to the right of the at sign in sip:alice.smith@bigu.edu.

The cache lifetime is 12 hours. This could be any positive signed 32 bit integer.

The class is IN (this is always true.)

The record type is NAPTR.

The order is 60. The preference is 50. The meaning of order and preference in NAPTR records is different from preference and weight in SRV records. As with SRV records however, lower values have higher precedence. The order specifies the order in which records are read. If a capability match is found between calling and called parties, that protocol must be used and other records discarded. If the order fields are all the same, the preference field is examined. Lower values have higher precedence but calling parties may override the called domain preference and select a higher preference value transport.

The flag is s. Flags are application specific. In this case the flag specifies that the lookup is terminal and all the information is present to find the appropriate SRV record in the regexp or target field.

The service is SIP+D2U. This specifies the SIP protocol over UDP. Other possible values are SIP+D2T for SIP over TCP, SIP+D2S for SIP over SCTP and SIPS+D2T for secure SIP over TLS over TCP. TLS over UDP is not defined. All valid service field values are registered with IANA.

The regexp is blank. The target is _sip._udp.bigu.edu. The regexp (regular expression) and target fields in combination make up the substitution field. One, and only one, must be used; the other must be blank. Regular expressions are powerful substitution mechanisms. However, they also can be complex and hence error prone. Unless you have a commanding need for the flexibility of a regular expression, we recommend you specify a static target.

The addition of NAPTR records to the previous example might look like this.

```

bigu.edu IN SOA ns.bigu.edu. root.bigu.edu. (
    2003032001
    10800
    3600
    604800
    86400 )

bigu.edu.          43200 IN NS      ns.bigu.edu.
;
ns.bigu.edu.       43200 IN A      10.0.0.20
sipserver1.bigu.edu. 43200 IN A      10.0.0.21
sipserver2.bigu.edu. 43200 IN A      10.0.0.22
;
_sip._udp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver1.bigu.edu.
_sip._udp.bigu.edu. 43200 IN SRV 1 0 5060 sipserver2.bigu.edu.
_sip._tcp.bigu.edu. 43200 IN SRV 0 4 5060 sipserver1.bigu.edu.
_sip._tcp.bigu.edu. 43200 IN SRV 0 2 5060 sipserver2.bigu.edu.
_sips._tcp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver1.bigu.edu.
_sips._tcp.bigu.edu. 43200 IN SRV 0 0 5060 sipserver2.bigu.edu.
;
bigu.edu. IN NAPTR 0 0 "s" "SIPS+D2T" "" _sips._tcp.bigu.edu.
bigu.edu. IN NAPTR 1 0 "s" "SIP+D2T" "" _sip._tcp.bigu.edu.
bigu.edu. IN NAPTR 2 0 "s" "SIP+D2U" "" _sip._udp.bigu.edu.

```

NAPTR records are not necessary but if they are present RFC3263 mandates at least three records. It further states that they should be listed in this precedence: 1) SIPS+D2T, 2) SIP+D2T and SIP+D2U. In common English this means that TLS over TCP should be used if the calling party has the capability. Failing that TCP should be used and UDP is permitted only as a last resort to keep the call from failing.

In practice BIND implementations are often smart enough to return the SRV records in the target field and the A records to which they point among the glue data so that additional DNS lookups are not required.

In terms of call flow then a SIP User Agent Client first sends a DNS NAPTR request to the domain specified in the Request-URI. If valid records are returned an appropriate transport is identified. Depending on the richness of the glue data in the first request, a second request is sent to the value in the substitution field.

If no NAPTR records are returned, a DNS SRV request is sent based on the transport preferred by the UAC. If valid records are returned, the request is sent to the preferred proxy.

As a last resort if no SRV records are found a DNS A record request is sent for the domain in the Request-URI. In the case of `sip:alice.smith@bigu.edu` the request would be for the IP address of `bigu.edu`. If a valid IP address is returned, the request is sent to that address using UDP.

The best practice for sites wanting just to get started may be to implement SRV records but not NAPTR records. Those wishing complete detailed information on service location are referred to RFCs 2782, 3263 and 3403.