One cluster == One DNS record

Andrew Jackson
AndrewJackson947@gmail.com

Andrey Borodin
amborodin@acm.org

Problem: Currently there is no built-in option to use a single DNS name to define a cluster of postgres instances for users to connect to. Currently libpq support target_session_attrs and looking up multiple IP's from an A DNS record but if that record/port combination houses standbys and primaries, it will fail if it finds a primary when you specify a standby (for example). Because of this additional infrastructure is required to implement the client side of HA (VIPs, CNAME's, LDAP connection parameter lookup, reverse proxy, etc).

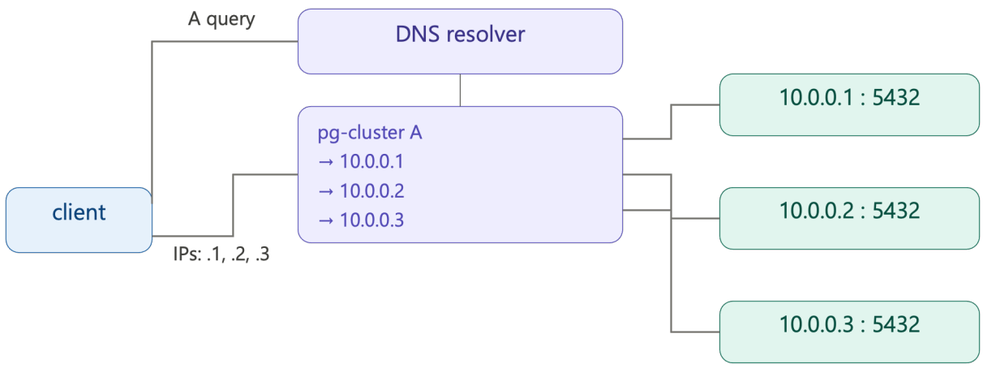

Patch idea: Implement try_all_addrs connection parameter that enables HA postgres clusters to sit behind a single DNS name/port.

```
$ cat /etc/hosts | grep pg-load
127.0.0.1 pg-loadbalancetest
127.0.0.2 pg-loadbalancetest
127.0.0.3 pg-loadbalancetest

psql “host=pg-loadbalancetest
try_all_addrs=1 target_session_attrs=read-
write”
```

The idea here is that it may be a reasonable assumption, in many environments, that every service sitting behind a given A record/port combination is homogeneous (all primaries, all standbys, multiple IPs pointing to the same hardware, etc), but for admins that want to set their HA cluster behind a single A record this opt in connection parameter will still allow it.

Pros:
-Very small change
-Backwards Compatible: functionality is entirely opt-in, preserving the “fail fast” default functionality.

Open Questions:
- Makes changes to PqConnectPoll. This function is incredibly complex, is synchronous and is probably due for a general refactor that this patch may make more difficult
- Performance: Hosts are connected in sequential order, would be more optimal if multiple hosts could be tried at once.

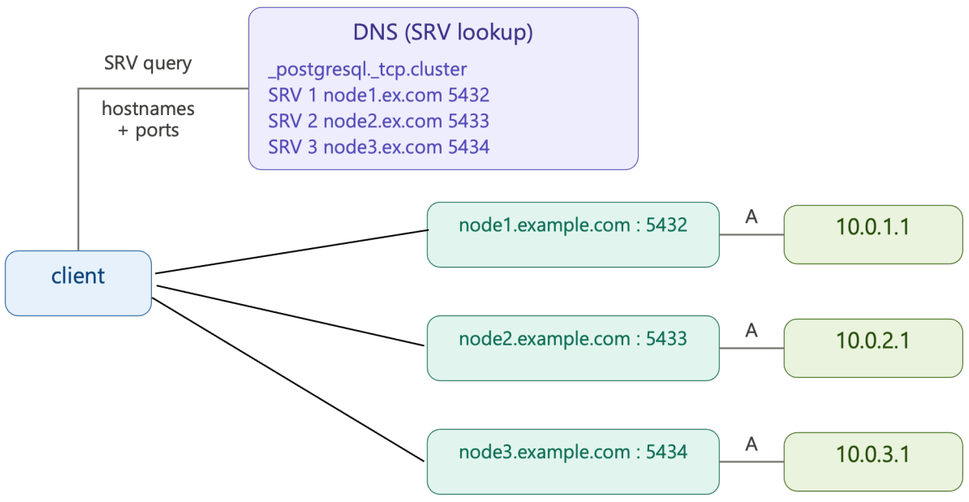

Patch idea: implement DNS SRV record support in libpq, enabling a single SRV name to resolve to multiple hostname/port combinations for HA cluster discovery.

```
$ dig _postgresql._tcp.mydb.prod.example.com SRV
;; ANSWER SECTION:
_postgresql._tcp.mydb...  SRV 1 1 5432
node1.example.com.
_postgresql._tcp.mydb...  SRV 2 1 5433
node2.example.com.
_postgresql._tcp.mydb...  SRV 3 1 5434
node3.example.com.

psql ”postgres+srv=mydb.prod.example.com
target_session_attrs=read-write"
```

SRV records return several hostnames. libpq resolves each SRV target to an IP via A/AAAA lookup and then tries host/port pairs one by one following RFC 2782 priority and weight ordering until a suitable server is found.

Pros:
-Port info in DNS: each SRV record includes hostnames and port
-Priority and weight support: RFC 2782 priority/weight fields.

Open Questions:
- New URL schema: a lot of infrastructural components might be surprised
- Platform DNS API: SRV lookups require res_query() / DnsQuery() which are not uniformly available. The patch must handle portability across Linux, macOS, and Windows cleanly.

← Commitfest items →

PGConf.dev 2026