

NAME

systemd-resolved.service, systemd-resolved - Network Name Resolution manager

SYNOPSIS

systemd-resolved.service

/lib/systemd/systemd-resolved

DESCRIPTION

systemd-resolved is a system service that provides network name resolution to local applications. It implements a caching and validating DNS/DNSSEC stub resolver, as well as an LLMNR resolver and responder. Local applications may submit network name resolution requests via three interfaces:

- The native, fully-featured API **systemd-resolved** exposes on the bus. See the [API Documentation](#)^[1] for details. Usage of this API is generally recommended to clients as it is asynchronous and fully featured (for example, properly returns DNSSEC validation status and interface scope for addresses as necessary for supporting link-local networking).
- The glibc [getaddrinfo\(3\)](#) API as defined by [RFC3493](#)^[2] and its related resolver functions, including [gethostbyname\(3\)](#). This API is widely supported, including beyond the Linux platform. In its current form it does not expose DNSSEC validation status information however, and is synchronous only. This API is backed by the glibc Name Service Switch ([nss\(5\)](#)). Usage of the glibc NSS module [nss-resolve\(8\)](#) is required in order to allow glibc's NSS resolver functions to resolve host names via **systemd-resolved**.
- Additionally, **systemd-resolved** provides a local DNS stub listener on IP address 127.0.0.53 on the local loopback interface. Programs issuing DNS requests directly, bypassing any local API may be directed to this stub, in order to connect them to **systemd-resolved**. Note however that it is strongly recommended that local programs use the glibc NSS or bus APIs instead (as described above), as various network resolution concepts (such as link-local addressing, or LLMNR Unicode domains) cannot be mapped to the unicast DNS protocol.

The DNS servers contacted are determined from the global settings in `/etc/systemd/resolved.conf`, the per-link static settings in `/etc/systemd/network/*.network` files, the per-link dynamic settings received over DHCP and any DNS server information made available by other system services. See [resolved.conf\(5\)](#) and [systemd.network\(5\)](#) for details about systemd's own configuration files for DNS servers. To improve compatibility, `/etc/resolv.conf` is read in order to discover configured system DNS servers, but only if it is not a symlink to `/run/systemd/resolve/resolv.conf` (see below).

systemd-resolved synthesizes DNS resource records (RRs) for the following cases:

- The local, configured hostname is resolved to all locally configured IP addresses ordered by their scope, or — if none are configured — the IPv4 address 127.0.0.2 (which is on the local loopback) and the IPv6 address `::1` (which is the local host).
- The hostnames "localhost" and "localhost.localdomain" (as well as any hostname ending in ".localhost" or ".localhost.localdomain") are resolved to the IP addresses 127.0.0.1 and `::1`.
- The hostname "gateway" is resolved to all current default routing gateway addresses, ordered by their metric. This assigns a stable hostname to the current gateway, useful for referencing it independently of the current network configuration state.
- The mappings defined in `/etc/hosts` are resolved to their configured addresses and back.

Lookup requests are routed to the available DNS servers and LLMNR interfaces according to the following rules:

- Lookups for the special hostname "localhost" are never routed to the network. (A few other, special domains are handled the same way.)
- Single-label names are routed to all local interfaces capable of IP multicasting, using the LLMNR protocol. Lookups for IPv4 addresses are only sent via LLMNR on IPv4, and lookups for IPv6 addresses are only sent via LLMNR on IPv6. Lookups for the locally configured host name and the "gateway" host name are never routed to LLMNR.
- Multi-label names are routed to all local interfaces that have a DNS sever configured, plus the globally configured DNS server if there is one. Address lookups from the link-local address range are never routed to DNS.

If lookups are routed to multiple interfaces, the first successful response is returned (thus effectively

merging the lookup zones on all matching interfaces). If the lookup failed on all interfaces, the last failing response is returned.

Routing of lookups may be influenced by configuring per-interface domain names. See [systemd.network\(5\)](#) for details. Lookups for a hostname ending in one of the per-interface domains are exclusively routed to the matching interfaces.

See the [resolved D-Bus API Documentation](#)^[1] for information about the APIs systemd-resolved provides.

/ETC/RESOLV.CONF

Three modes of handling /etc/resolv.conf (see [resolv.conf\(5\)](#)) are supported:

- A static file /usr/lib/systemd/resolv.conf is provided that lists the 127.0.0.53 DNS stub (see above) as only DNS server. This file may be symlinked from /etc/resolv.conf in order to connect all local clients that bypass local DNS APIs to **systemd-resolved**. This mode of operation is recommended.
- **systemd-resolved** maintains the /run/systemd/resolve/resolv.conf file for compatibility with traditional Linux programs. This file may be symlinked from /etc/resolv.conf and is always kept up-to-date, containing information about all known DNS servers. Note the file format's limitations: it does not know a concept of per-interface DNS servers and hence only contains system-wide DNS server definitions. Note that /run/systemd/resolve/resolv.conf should not be used directly by applications, but only through a symlink from /etc/resolv.conf. If this mode of operation is used local clients that bypass any local DNS API will also bypass **systemd-resolved** and will talk directly to the known DNS servers.
- Alternatively, /etc/resolv.conf may be managed by other packages, in which case **systemd-resolved** will read it for DNS configuration data. In this mode of operation **systemd-resolved** is consumer rather than provider of this configuration file.

Note that the selected mode of operation for this file is detected fully automatically, depending on whether /etc/resolv.conf is a symlink to /run/systemd/resolve/resolv.conf or lists 127.0.0.53 as DNS server.

SIGNALS

SIGUSR1

Upon reception of the SIGUSR1 process signal **systemd-resolved** will dump the contents of all DNS resource record caches it maintains into the system logs.

SIGUSR2

Upon reception of the SIGUSR2 process signal **systemd-resolved** will flush all caches it maintains. Note that it should normally not be necessary to request this explicitly – except for debugging purposes – as **systemd-resolved** flushes the caches automatically anyway any time the host's network configuration changes.

SEE ALSO

[systemd\(1\)](#), [resolved.conf\(5\)](#), [dnsssec-trust-anchors.d\(5\)](#), [nss-resolve\(8\)](#), [systemd-resolve\(1\)](#), [resolv.conf\(5\)](#), [hosts\(5\)](#), [systemd.network\(5\)](#), [systemd-networkd.service\(8\)](#)

NOTES

1. API Documentation
<http://www.freedesktop.org/wiki/Software/systemd/resolved>
2. RFC3493
<https://tools.ietf.org/html/rfc3493>