



Technical Report

DNS Load Balancing in ONTAP

Configuration and Best Practices

Justin Parisi, NetApp
October 2016 | TR-4253

Abstract

This document explains how to configure NetApp® storage systems with NetApp ONTAP® management software for use with DNS load balancing methodologies. In particular, this document covers the on-box DNS feature available in ONTAP, various configuration methods, and best practices.

Document Classification

Public

TABLE OF CONTENTS

1	Version History	3
2	Domain Name Systems (DNS) in ONTAP	3
2.1	What Is DNS?	3
3	DNS Load Balancing	7
3.1	Round-Robin DNS	7
3.2	On-Box DNS Load Balancing.....	7
3.3	Deciding How to Configure the On-Box DNS Zone.....	10
4	Configuring On-Box DNS Load Balancing.....	15
4.1	Configuring On-Box DNS on the Storage Virtual Machine	15
4.2	Configuring Windows DNS Server to Work with On-Box DNS.....	15
4.3	Configuring BIND-Style DNS Servers to Work with On-Box DNS.....	30
4.4	Configuring Clients to Use ONTAP Data LIFs as DNS Servers	33
5	Conclusion	36

LIST OF TABLES

Table 1)	On-box DNS algorithm calculations.....	9
Table 2)	Data LIF options for on-box DNS load balancing in ONTAP.	13

LIST OF BEST PRACTICES

Best Practice 1:	ONTAP Version Recommendation: On-Box DNS	8
Best Practice 2:	Geometric Mean Configuration.....	10
Best Practice 3:	Windows DNS Configuration Recommendations	11
Best Practice 4:	BIND DNS Configuration Recommendations	12
Best Practice 5:	Recommendations for Data LIFs Acting as DNS Servers	13

LIST OF FIGURES

Figure 1)	Example of off-box DNS round-robin method using A records.....	7
Figure 2)	On-box DNS load balance example.	8
Figure 3)	Factors to consider in setting up on-box DNS load balancing on Windows DNS servers.	11
Figure 4)	On-box DNS with multiple subnets in same SVM.	12

1 Version History

Version	Date	Document Version History
Version 1.0	July 2016	Initial release
Version 2.0	October 2016	Updated for ONTAP 9.1

2 Domain Name Systems (DNS) in ONTAP

ONTAP enables storage administrators to present multiple logical interfaces (data LIFs) per storage virtual machine (SVM) across multiple nodes to clients for NAS access. In NAS environments, clusters can have up to 24 nodes, so the number of potential data LIFs in a cluster is large. This potentially can create confusion about access for clients if they rely on mounting through IP addresses. Clients can overload a node with requests if they continuously mount the same data LIFs, and attempting to remember specific IP addresses can be challenging.

Management of these IP addresses can also be challenging. When an IP address needs to change, more points must be considered if clients are accessing a known IP by the address. If adding new data LIFs or removing data LIFs, administrators must make clients explicitly aware of these changes.

To simplify client access to these data LIFs as well as the management of the NAS networking components from the storage side, [Domain Name System \(DNS\)](#) is often implemented to obfuscate multiple data LIFs behind a single host name.

For general DNS best practices in ONTAP, refer to [TR-4379: Name Services Best Practice Guide](#).

The following Requests for Comments (RFCs) cover DNS standards and provide general information about DNS:

- [RFC 1035 – Domain Names](#)
- [RFC 1123 – Requirements for Internet Hosts](#)
- [RFC 2181 – Clarifications to the DNS Specification](#)

2.1 What Is DNS?

DNS is a hierarchical naming system for devices on a network that provides a way to associate human-readable names to less readily memorized items, such as IP addresses, service records, and so on. DNS relegates the issuance of these records to one or more servers that act as authoritative sources on the network.

DNS Terminology

The following section covers different types of DNS terminology used with on-box DNS.

A/AAAA Records

A/AAAA records ([RFC-1101](#)) map host names to IP addresses. An A record maps a host name to an IPv4 address. An AAAA record maps host names to IPv6 addresses. These maps are used for forward DNS lookups.

Canonical Name (CNAME)

This is an alias of a host name.

Service (SRV) Records

SRV records ([RFC-2782](#)) define a DNS record for a specific domain service, including LDAP, CIFS, NFS, Exchange, and so on. These records can point to multiple A/AAAA records to provide round-robin load balance functionality and high availability.

Pointer (PTR) Records

PTR records map IP addresses to canonical names. This mapping is used for reverse DNS lookups.

Name Server (NS) Records

NS records are used to delegate a subdomain to a set of name servers. These records can be authoritative or nonauthoritative records.

Start of Authority (SOA) Records

This type of record defines which name server is the authoritative answer for a DNS request. If a name server that does not have an SOA record issues a response to a DNS request, the response returns to the client as a “nonauthoritative” response.

SOA records contain the following information:

- Primary name server from the DNS domain
- Time stamp of updates
- Zone refresh time
- Failed refresh retry times
- SOA record timeout
- Negative time to live (TTL) (how long failed resolvers live in failure cache)

DNS Forwarder

A DNS forwarder is a DNS server on a network that forwards DNS queries for external DNS names to DNS servers outside that network. You can also forward queries according to specific domain names by using conditional forwarders, which override regular DNS forwarders.

Conditional Forwarder

A conditional forwarder is a DNS server on a network that forwards DNS queries according to the DNS domain name in the query. For example, a DNS server can be configured to forward all of the queries it receives for names ending with `example.newname.com` to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers. A conditional forwarder is used when a DNS server's domain differs from the desired DNS domain name.

For example:

`example.newname.com → netapp.com`

A conditional forwarder requires the data LIFs to be added to DNS as name servers and to have a Start of Authority (SOA) record. It also requires having a forward lookup zone and reverse lookup entries created. Windows 2008 and later might require SOA records. Windows 2003 DNS does not require SOA records.

Stub Zones

From the [Microsoft TechNet article on stub zones](#):

A stub zone is a copy of a zone that contains only those resource records necessary to identify the authoritative Domain Name System (DNS) servers for that zone. A stub zone is used to resolve names between separate DNS namespaces. This type of resolution might be necessary when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

A stub zone consists of:

- The start of authority (SOA) resource record, name server (NS) resource records, and the glue A resource records for the delegated zone.
- The IP address of one or more master servers that can be used to update the stub zone.

The master servers for a stub zone are one or more DNS servers authoritative for the child zone, usually the DNS server hosting the primary zone for the delegated domain name.

A stub zone is required if conditional forwarding does not work because the name servers are not Start of Authority (SOA) servers and the DNS zone created is not a stub zone.

For a comparison of stub zones and conditional forwarders, see:

[Contrasting stub zones and conditional forwarders](#)

Primary Zones

A primary zone is a DNS zone that is the primary source of information for a zone and that stores a master copy of zone data in a local file or in the database. Unlike stub zones, primary zones allow creation of records (A, AAAA, SRV, and so on).

DNS Delegations

A DNS delegation delegates requests in the same domain to the DNS servers specified in the delegation zone. For example, for `cdot.netapp.com` in the DNS domain of `netapp.com`, use a delegation.

For more information on zone delegations, see the Microsoft TechNet article on [delegating zones and understanding zone delegation](#).

Subdomains

A subdomain is a DNS domain that is part of the primary DNS domain. For example, `dns.domain.com` is a subdomain of `domain.com`.

DNS options in ONTAP

NetApp ONTAP offers a variety of options for controlling DNS configurations, including:

- Dynamic DNS (IPv4 and IPv6)
- On-box DNS load balancing
- Ability to use data LIFs as name servers and/or name records

The following DNS configuration options are available at **advanced privilege** in ONTAP 9.1 and later:

PARAMETERS

`-vserver <vserver name>` - Vserver

Use this parameter to specify the Vserver whose DNS mapping is modified.

`[-domains <text>, ...]` - Domains

Use this parameter to specify a domain for the Vserver.

`[-name-servers <IP Address>, ...]` - Name Servers

Use this parameter to specify the IP addresses of the DNS name servers for this Vserver.

`[-state {enabled|disabled}]` - (DEPRECATED)-Enable/Disable DNS

Note: This parameter has been deprecated and might be removed in a future version of Data ONTAP.

Use this parameter with the value `enabled` to specify that the DNS server mapping is active. Use this parameter with the value `disabled` to specify that the DNS server mapping is not active.

`[-timeout <integer>]` - Timeout (secs)

Use this parameter to specify a timeout value (in seconds) for queries to the DNS servers.

`[-attempts <integer>]` - Maximum Attempts

Use this parameter to specify the number of times to attempt queries to the DNS servers.

`[-is-tld-query-enabled {true|false}]` - Is TLD Query Enabled? (privilege: advanced)

Use this parameter to enable or disable top-level domain (TLD) queries. If the parameter is set to `false`, the resolver will not attempt to resolve a name that has no "." characters in it. The default value for this parameter is `true`.

`[-require-source-address-match {true|false}]` - Require Source and Reply IPs to Match (privilege: advanced)

Use this parameter to allow dns responses sourced from an IP that does not match where the vserver sent the request. If the parameter is set to `false`, the resolver will allow response from an IP other than the one to which the request was sent.

`[-require-packet-query-match {true|false}]` - Require Packet Queries to Match (privilege: advanced)

Use this parameter to check if the query section of the reply packet is equal to that of the query packet. If the parameter is set to `false`, the resolver will not check if the query section of the reply packet is equal to that of the query packet.

Note: DNS configuration should be done at the storage virtual machine (SVM) level in ONTAP 8.3 and later.

3 DNS Load Balancing

An added benefit of using DNS host names to point to multiple IP addresses is having the ability to leverage various load balancing mechanisms with DNS servers. DNS load balancing is a way to distribute client requests for host names across multiple IP addresses without needing client interaction. Generally, DNS load balancing is done [round-robin](#). Load balancing can also be done through third-party load balancers or through the ONTAP feature known as on-box DNS load balancing.

3.1 Round-Robin DNS

Round-robin DNS is the most common form of DNS load balancing. It is offered by default in DNS servers and is a simple way to offer IP addresses to clients requesting them.

To create a round-robin A/AAAA record, simply create another A/AAAA record with the same name as the original record.

Figure 1) Example of off-box DNS round-robin method using A records.

cluster	Host (A)	10.10.10.10
cluster	Host (A)	10.10.10.11
cluster	Host (A)	10.10.10.12

- For more information on round-robin DNS in Windows, see the following:

[Configuring Round-Robin DNS in Windows](#)

- For more information on round-robin DNS in BIND, see the following:

[Round-Robin Load Distribution](#)

- For information about ONTAP networking best practices, see:

[TR-4182: Best Practices for Clustered Data ONTAP Network Configurations.](#)

Round-Robin DNS Limitations

Round-robin balancing does not take into account such things as server load, network connectivity, and so on. If a server experiences issues in a round-robin configuration, the DNS server might still issue the IP address for the problematic server, which can create issues for clients. Because of this possibility, round-robin DNS might not be an ideal method for enterprise NAS environments because these environments might require a more discerning load balancing methodology. Fortunately, ONTAP offers an integrated, simple, intelligent load balancing solution for DNS, free of charge—no license is required.

3.2 On-Box DNS Load Balancing

ONTAP offers the ability to leverage the DNS service on each node to service DNS requests from clients. ONTAP also can issue data LIF IP addresses based on an algorithm that take into account CPU load and port throughput on the node. This process provides the least-used data LIF to make sure of proper load balancing across the cluster for mount requests. After a mount/map is successful, the client continues to use that connection until remount.

This approach differs from round-robin DNS, because the external DNS server services all requests and has no insight into how busy a node in the cluster is.

On-Box DNS Considerations

Use of DNS load balancing is not necessary when using NFSv4.x referrals, because the connection is made to the local node regardless of which IP address is returned from DNS.

Additionally, round-robin DNS issues IP addresses with a time to live (TTL). The TTL caches the DNS request in Windows for 24 hours by default. On-box DNS issues a TTL of 0, which means that DNS is never cached on the client and a new IP is always issued based on load.

DNS Time to Live (TTL) during LIF migrations

In rare cases, on-box DNS may respond with a TTL of 24 hours, particularly during LIF migrations (manual or automatic during storage failover events). In those cases, the DNS server will store the record in negative cache (as in, don't use this LIF) until the DNS server cache is purged or the expiration time is reached. This issue is covered in [bug 1027140](#) and is resolved in ONTAP versions listed in the public bug report. For information on flushing DNS server caches, please review the associated server's documentation.

On-Box DNS interaction with pNFS

On-box DNS does not apply to pNFS data traffic, which redirects traffic for I/O consistently during mounts. However, on-box DNS can assist in load balancing connections to the metadata servers (MDS) in the cluster. For more information about pNFS, see [TR-4067: NFS Best Practices and Implementation Guide](#) and [TR-4063: Parallel Network File System Configuration and Best Practices for Clustered Data ONTAP](#).

Best Practice 1: ONTAP Version Recommendation: On-Box DNS

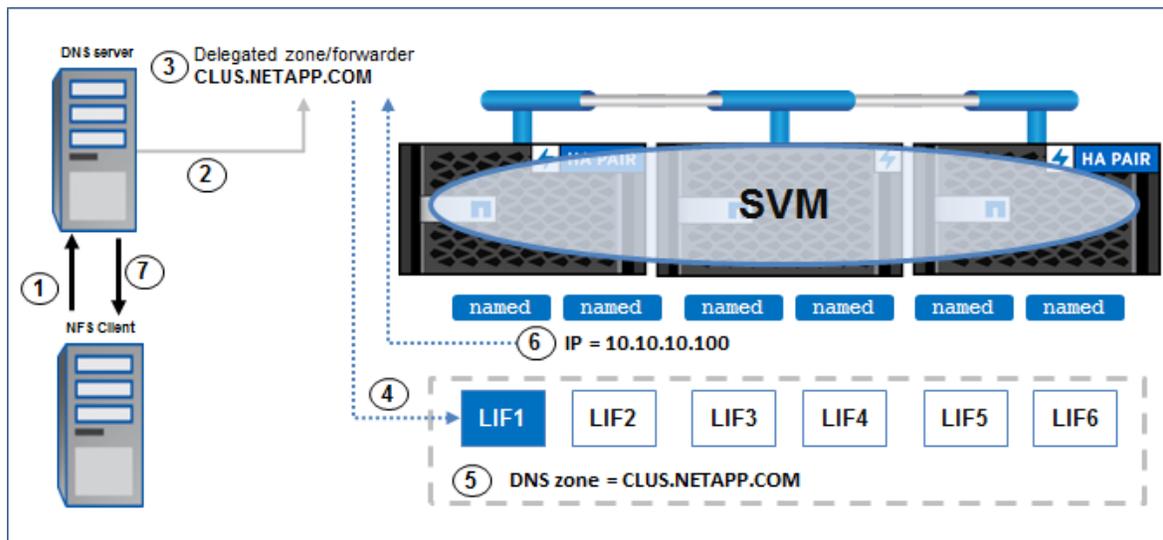
When using on-box DNS in ONTAP, make sure that the cluster runs one of the following versions for best results:

- 8.1.3P1 or later
- 8.2.4 or later
- 8.3.2 or later
- 9.0 or later

How On-Box DNS Load Balancing Works

Each node in the cluster has a service running ([named](#)) that handles incoming DNS requests from clients. The node also issues IP addresses based on a calculated weight that is determined by using an [algorithm](#) based on CPU utilization and node throughput.

Figure 2) On-box DNS load balance example.



When a client attempts to access the cluster by DNS host name, the following process takes place:

1. The client issues a DNS request and uses the DNS server specified in its configuration.
2. The DNS server looks for the host name in the request.
3. When using on-box DNS, the host name is either a DNS delegation or a conditional forwarder. The record contains a list of data LIF IP addresses to use for DNS requests.
4. The request is forwarded or delegated to one of the data LIF IP addresses on a round-robin basis.
5. The data LIF receives the request if the LIF has the DNS zone configured and is set to listen for DNS queries (which opens port 53 on the LIF).
6. The node receiving the request checks the DNS weights for each node and issues an IP address based on the calculated load.
7. The IP address is returned to the DNS server, which then returns the IP address to the client.

Note: In ONTAP versions earlier than 8.2, on-box DNS load balancing did not work with ifgrps or VLANs; with implementations that have those configurations, use external round-robin DNS. ONTAP versions 8.2 and later allow on-box DNS load balancing on ifgrps and VLANs.

The On-Box DNS Algorithm

The ONTAP on-box DNS algorithm is covered in [patent number US8271652](#). You can find complete details at the patent location. An abstract from that patent follows:

DNS name resolution is integrated into each node in a network storage cluster, to allow load balancing of network addresses, using a weighted random distribution to resolve DNS requests. A node in the cluster gathers statistics on utilization of resources, such as CPU utilization and throughput, on nodes in the cluster and distributes those statistics to all other nodes. Each node uses the same algorithm to generate weights for the various IP addresses of the cluster, based on the statistics distributed to it. The weights are used to generate a weighted list of available network addresses. In response to a DNS request, a DNS in a given node randomly indexes into the weighted address list to resolve requests to a network address. The weights are chosen so that the DNS is likely to pick an IP address which has a low load, to balance port and node usage over time.

The algorithm incorporates a series of weights assigned to data LIFs participating in the DNS load balancing group. These weights are refreshed every minute and use CPU weight and throughput weight to calculate a final weight.

Table 1) On-box DNS algorithm calculations.

CPU weight

$\text{cpu_weight} = 100.0 - (\% \text{ of CPU being used}) / \text{Number of IP addresses on node where IP address resides}$

Throughput weight

$\text{thpt_weight} = 100.0 - (\% \text{ of port throughput being used}) / \text{Number of IP addresses on port where IP address resides}$

Final weight

$\text{final_weight} = (\text{thpt_weight} + \text{cpu_weight}) / 2$

Geometric Mean Versus Arithmetic Mean

In ONTAP versions before the fix for [bug 619247](#), the DNS load balance algorithm used an [arithmetic mean](#) rather than a [geometric mean](#). The arithmetic mean was known to return IP addresses for nodes with low throughput and 100% CPU utilization, so it was changed. Current versions of ONTAP use the geometric mean by default. NetApp does not recommend that you change this option.

Best Practice 2: Geometric Mean Configuration

Do not modify the geometric mean for load balancing unless directed by NetApp Technical Support.

This behavior is controlled through a CLI option in **advanced privilege** mode:

```
cluster::*> network options load-balancing show
Geometric Mean Algorithm for load balancing: true
```

3.3 Deciding How to Configure the On-Box DNS Zone

This section covers how to decide which DNS zone methodology to use to configure on-box DNS load balancing.

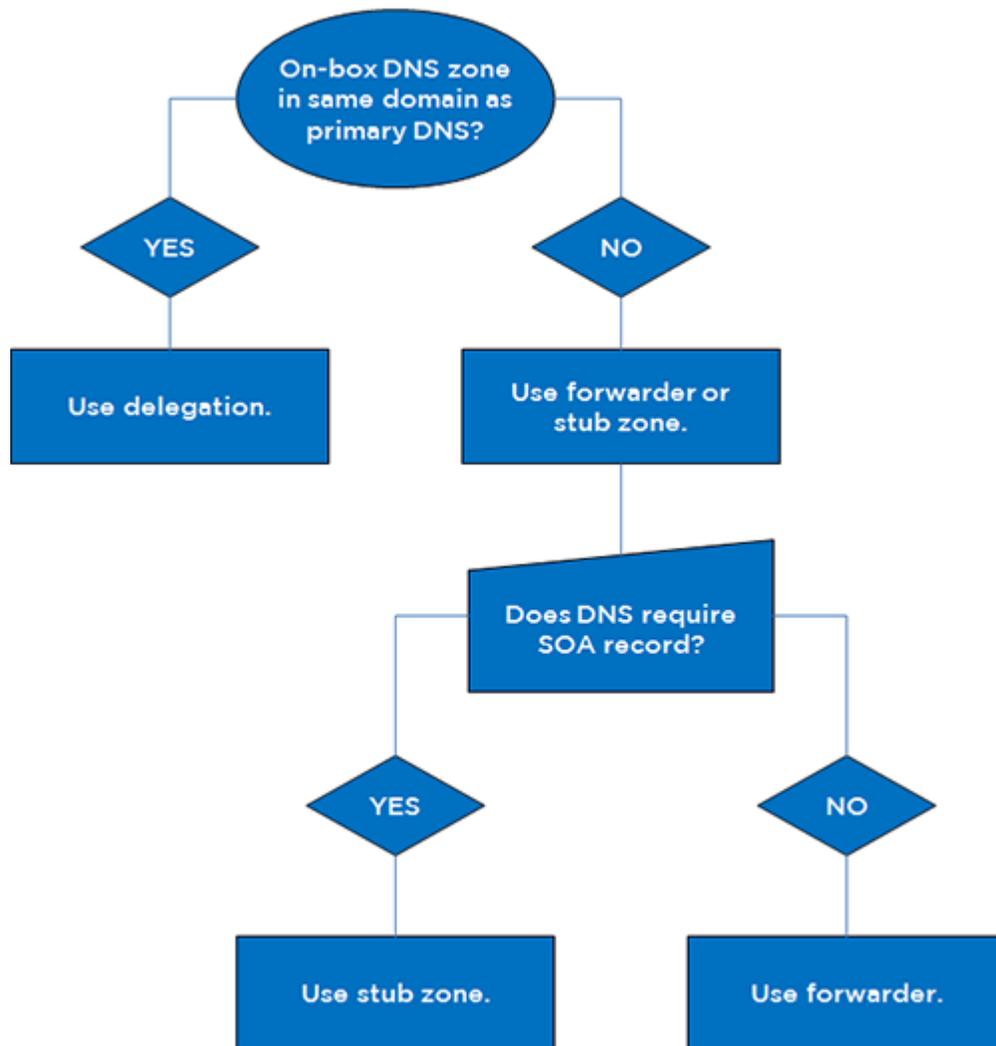
Note: The same concepts apply to non-Windows DNS servers (such as BIND). DNS is an Internet standard, as covered in [RFC 1035](#).

Deciding Which Configuration to Use in Windows DNS

When configuring on-box DNS load balancing, a design decision needs to be made about whether to use conditional forwarding, a stub zone, or a DNS zone delegation. [This blog](#) covers use-case scenarios for when to use which type of forwarding zone.

The design decision is based on a variety of factors, as shown in Figure 3.

Figure 3) Factors to consider in setting up on-box DNS load balancing on Windows DNS servers.



In some cases, it might make sense to configure clients to reference the data LIFs acting as DNS listeners directly as name servers. For guidance on doing so, see the section “Configuring Clients to Use ONTAP Data LIFs as DNS .”

Best Practice 3: Windows DNS Configuration Recommendations

Use the following guidance to decide which type of DNS zone to use with Windows DNS servers.

- For data LIFs named with a DNS zone in the same domain as the primary DNS server, use DNS delegations.
- For data LIFs named with a DNS zone in a different DNS domain than the primary DNS server, use a stub zone unless SOA records are not required. In those cases, use forwarders.

Deciding Which Configuration to Use with BIND DNS

When configuring on-box DNS load balancing, a design decision needs to be made about whether to use forwarding, a subdomain zone, or a DNS zone delegation.

Best Practice 4: BIND DNS Configuration Recommendations

Use the following guidance to decide which type of DNS zone to use with BIND.

- Use forwarders if you do not use caching name servers and allow recursive requests.
- Ideally, use a zone delegation if the DNS domain is not a child domain. Delegations allow you to specify SOA and NS records; forwarders do not. Additionally, delegations can be replicated to slave DNS servers automatically with BIND zone files, while forwarders are manually added to `named.conf`.
- If the DNS domain is a child domain, use subdomains.

Note: If using BIND9 DNS servers with on-box DNS, be sure you run ONTAP 8.2.3 or later because of [bug 892388](#).

Using On-Box DNS with Data LIFs in Different Subnets and Networks

In ONTAP, it is possible to have a configuration in which DNS servers live in a different physical or virtually segmented network or IP space than the data LIFs to which clients connect and still use on-box DNS to serve the desired data LIFs to clients.

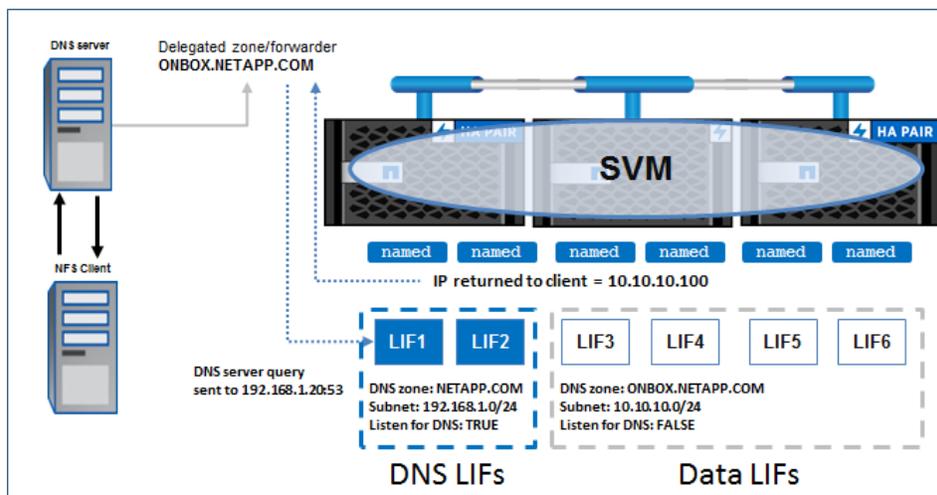
To do so, configure the LIFs that can communicate with the DNS servers to listen for DNS queries. The data LIFs that participate in the DNS zone should be configured to use the desired DNS zone and *not* listen for DNS queries (`-listen-for-dns-query false`).

Doing so enables the DNS server to communicate to the SVM using the DNS LIFs. It also enables the server to return a list of IP addresses to clients that might not be able to communicate with it.

Note: A data LIF that has `-listen-for-dns-query` set to “true” must also have a `-dns-zone` specified; otherwise, the cluster does not allow that LIF to listen for DNS queries.

The following diagram illustrates a similar configuration.

Figure 4) On-box DNS with multiple subnets in same SVM.



The data LIF configuration looks like the following example (the data LIF called “data1” can communicate with the DNS servers; the data LIF called “dns-zone” cannot):

```
cluster::*> net int show -vserver SVM -fields dns-zone,listen-for-dns-query,address
(network interface show)
vserver lif      address      dns-zone      listen-for-dns-query
-----
SVM      data1      10.63.57.237 domain.netapp.com      true
SVM      dns-zone  10.10.10.200 onbox.domain.netapp.com false
```

Enabling On-Box DNS on Data LIFs in ONTAP

For a data LIF to serve DNS queries, the `-listen-for-dns-query` option must be set to “true.” For the SVM to return data LIFs in DNS queries, the desired data LIFs participating in the DNS zone must be assigned the DNS zone with `-dns-zone`. Any data LIF that acts as an SOA for DNS queries must have network connectivity to the DNS servers to which the clients point. This can be done nondisruptively.

Best Practice 5: Recommendations for Data LIFs Acting as DNS Servers

As a best practice, configure multiple data LIFs as DNS servers if at all possible to ensure resiliency and load balancing of DNS requests. It also makes sense to set the `lb-weight` for LIFs serving DNS requests to 0 so that they don’t get used in the DNS zone for data traffic.

The data LIFs that participate in on-box DNS load balancing depend on the configuration of the following network interface options.

Table 2) Data LIF options for on-box DNS load balancing in ONTAP.

Network Interface Option	What It Does	Privilege Level
<code>-dns-zone</code>	Specifies the DNS zone of the data LIF participating in the on-box DNS load balance operation. Multiple DNS zones can be specified in an SVM.	Admin
<code>-listen-for-dns-query</code>	Specifies that the data LIF will listen for DNS queries on port 53 and act as an SOA.	Admin
<code>-lb-weight</code>	Use this parameter to modify the load balancing weight of the data LIF. A valid load balancing weight is any integer between 1 and 100 or the word “load.” If you specify the same load balancing weight for all data LIFs in a DNS zone, client requests are uniformly distributed, similar to round-robin DNS. A data LIF with a low load balancing weight is made available for client requests less frequently than one that has a high load balancing weight.	Advanced

It is possible to designate only specific data LIFs in a DNS zone to participate as the name servers through the `listen-for-dns-query` option while leaving other data LIFs to be used only for data traffic in the DNS zone. It is also possible to have data LIFs in the same SVM that do not participate in the on-box DNS load balancing zone but still can serve data traffic.

Manually Modifying the lb-weight of Data LIFs Participating in On-Box DNS

If multiple data LIFs are used in on-box DNS load balancing, it is possible to modify the lb-weight of specific data LIFs to be featured sooner in the load balancing algorithm. One use case for this is to favor nodes in a cluster using SSDs or All Flash FAS (AFF) systems in the weighting of data LIFs rather than nodes using spinning disks or favoring nodes with more RAM/CPU.

For example, if a 4-node cluster has an HA pair with FAS8xxxs operating as AFF personalities and 2 nodes are FAS3xxx nodes with SAS shelves, it might make sense to configure the data LIFs owned by the AFF nodes to have higher weights than the nodes with SAS shelves. Doing so would take advantage of the enhanced performance capabilities of the AFF systems.

Some guidelines:

- Setting a LIF to a weight of 100 means that the data LIF will almost always be used in DNS requests.
- Setting a LIF to a weight of 1 means that a data LIF will virtually never be used in DNS requests.
- If all lb-weights are the same, round-robin DNS is used.
- Keep in mind [how the on-box DNS load balancing algorithm works](#) when deciding whether to manually configure the lb-weights of data LIFs.

Configuring ONTAP to Enable/Disable Sending of SOA Records

In some cases, such as with non-Windows DNS servers, it might be necessary to disable the sending of SOA records from the cluster to get on-box DNS zones working with multiple subnets. You can disable these records with this **advanced privilege** command:

```
cluster::> set advanced
cluster::*> network options send-soa modify -enable true
```

Note: If you use multiprotocol NAS (CIFS/SMB and NFS) on the same cluster and choose to disable `send-soa`, be sure that both environments function properly with sending of SOA records disabled.

Disabling the sending of SOA records renders the on-box DNS zone as a nonauthoritative responder to DNS requests.

Using Data LIFs as Authoritative Name Servers for Clients

Because data LIFs can be configured to listen on port 53 for DNS requests and act as SOA servers, they can also be used as name servers on clients and act as independent DNS servers. This configuration can be useful in environments in which DNS servers might not be able to be modified or when clients do not have access to DNS servers in the domain.

To use data LIFs as name servers, simply configure the client's DNS configuration (`resolv.conf` on Linux clients, DNS property boxes on Windows clients). For details and examples of this use, see the section in this document on [configuring clients to use ONTAP Data LIFs as DNS servers](#).

4 Configuring On-Box DNS Load Balancing

This section covers configuration of on-box DNS load balancing in ONTAP.

4.1 Configuring On-Box DNS on the Storage Virtual Machine

To configure on-box DNS for the cluster, select the appropriate data LIFs to participate in the load balance. Be sure to designate data LIFs to act as DNS servers that listen for DNS queries.

Configuration Steps 1) Setting up on-box DNS load balancing on the cluster.

1. Enable DNS zones on the data LIF.

```
::> net int modify -vserver [SVM] -lif [LIF] -dns-zone [cdot.domain.com]
```
2. Set the desired LIF to listen for DNS queries (8.2 and later only).

```
::> net int modify -vserver [SVM] -lif [LIF] -listen-for-dns-query true
```
3. Configure the lb-weight in **advanced privilege** mode on the data LIF to “load” or the desired lb-weight.

```
::*> net int modify -vserver [SVM] -lif [LIF] -lb-weight load
```

4.2 Configuring Windows DNS Server to Work with On-Box DNS

The following configuration steps can be used to configure on-box DNS on Windows DNS servers. The following scenarios are covered in this section:

- Delegations
- Stub zones
- Conditional forwarders

Setting Up DNS Delegations in Windows DNS

The following steps show how to set up DNS delegations in Windows DNS servers. The server version used in the example is Windows 2008R2, but the same steps apply for other Windows servers. For official steps, refer to the [Microsoft TechNet documentation](#).

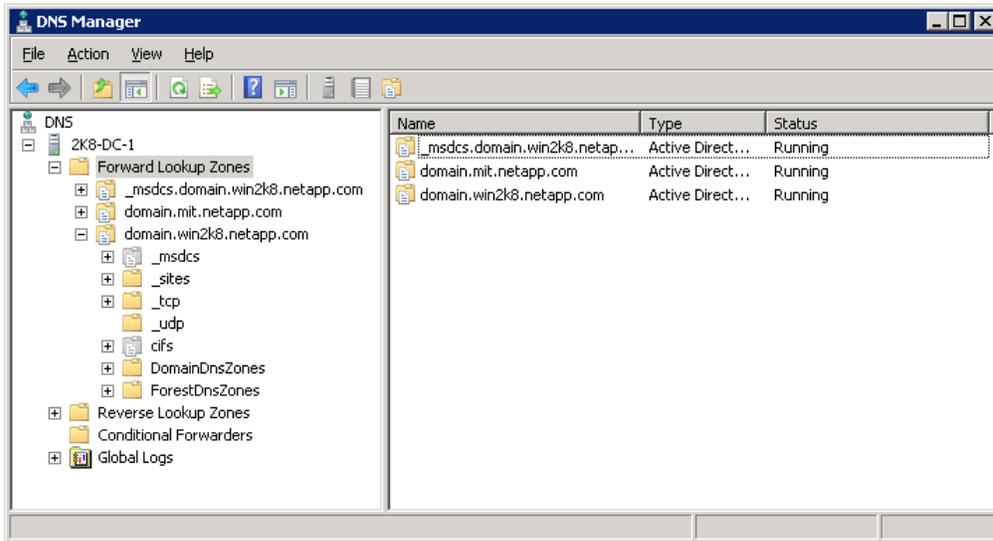
[DNS delegations](#) are used for:

- Delegating management of a DNS namespace to another location in your organization
- Dividing large zones into smaller zones to distribute load among multiple servers or create better fault tolerance
- Extending the namespace to add additional subdomains

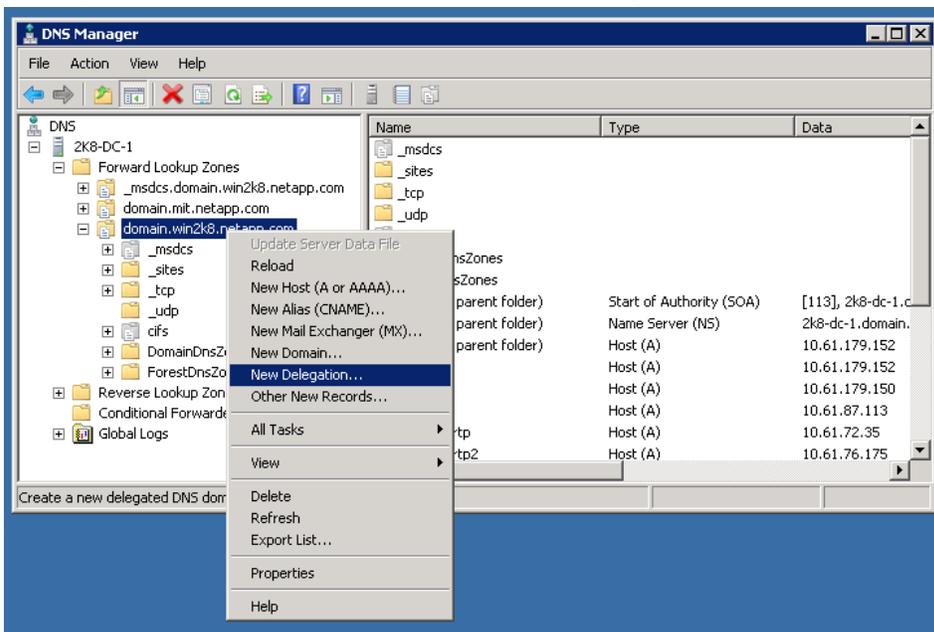
In the case of on-box DNS, delegations can be used to redirect DNS zone traffic to data LIFs on an SVM. Generally speaking, a delegation would be used if the data LIF DNS zones will be in the same DNS domain as the DNS servers. One example is if the data LIFs use *cluster.domain.com* and the DNS servers' domain is *domain.com*.

Configuration Steps 2) Setting up DNS delegations.

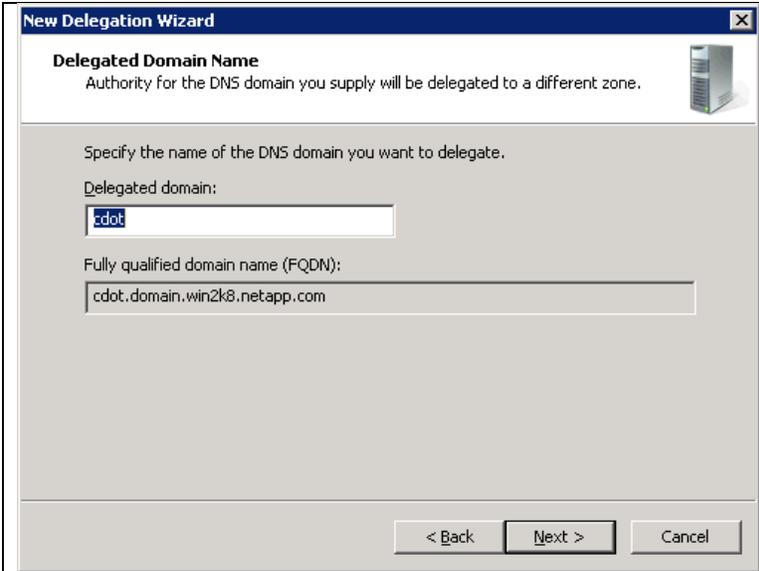
1. Open the DNS Manager console.



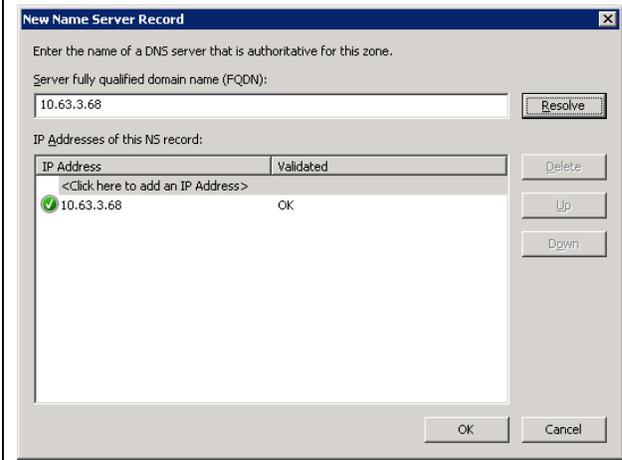
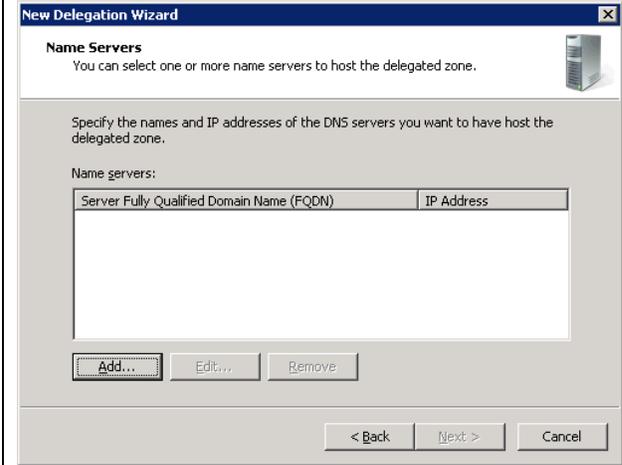
2. Right-click the DNS domain and select New Delegation.

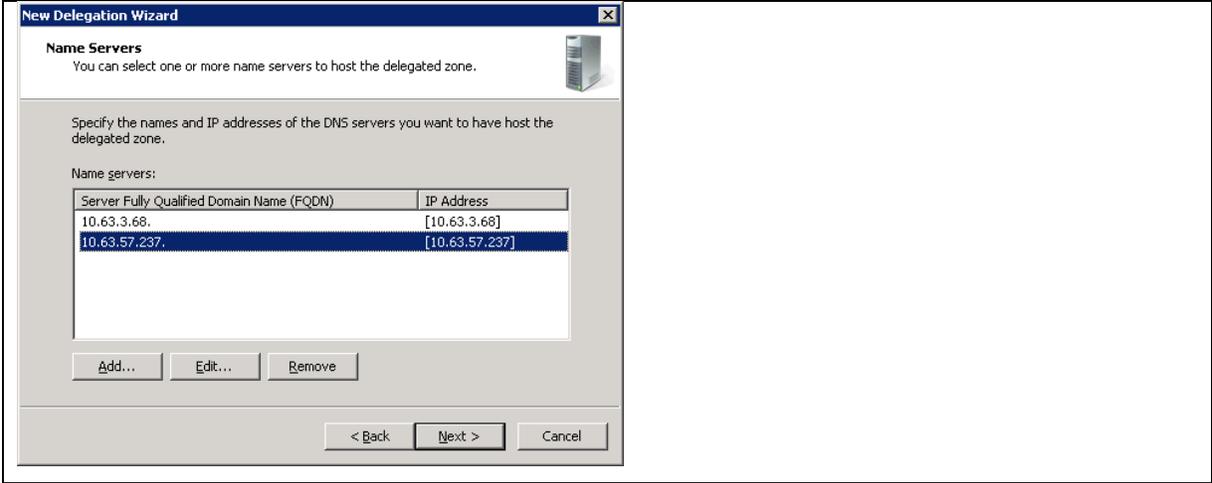


3. Enter the name of the delegated domain.



4. Add the data LIFs as name servers (one at a time).



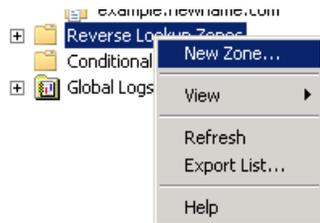


Configuration Steps 3) Setting up reverse lookup zones and PTR records.

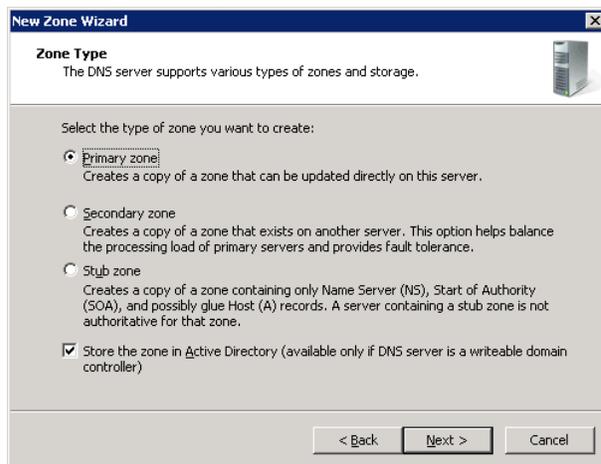
1. Create the reverse lookup zone and PTR records.

On-box DNS does not support reverse lookups for IPv4 earlier than ONTAP 8.2. IPv6 support was added in ONTAP 8.3. If you want to force clients to use the host name only for Kerberos, do not create PTR records. Doing so prevents direct IP mounts and makes sure that load balancing is enforced. However, in some cases, PTR records are required for Kerberized NFS to work.

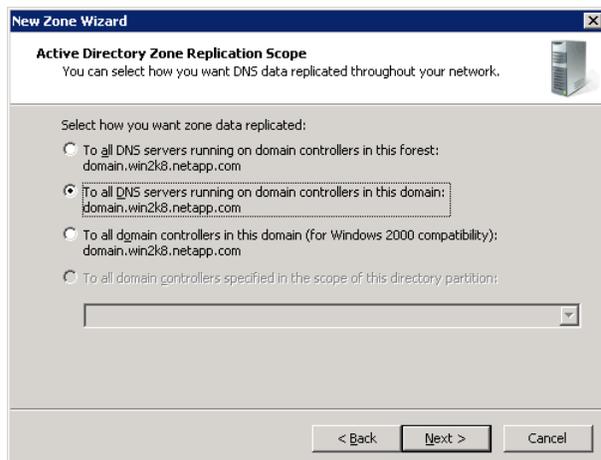
2. Create the reverse lookup zones for the data LIFs.



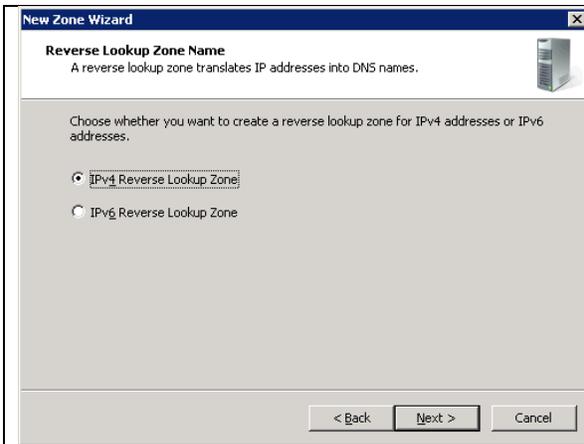
3. Select Primary Zone, because DNS in ONTAP cannot service reverse lookups.



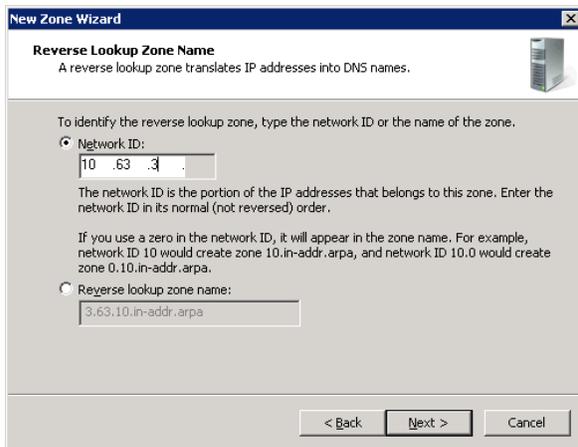
4. Select a zone replication policy to use.



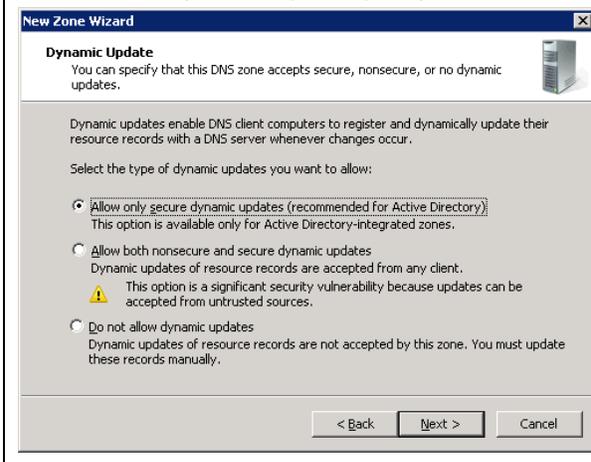
5. Select IPv4 or IPv6 for the lookup zone, depending on what the ONTAP version supports and what the data LIFs use.

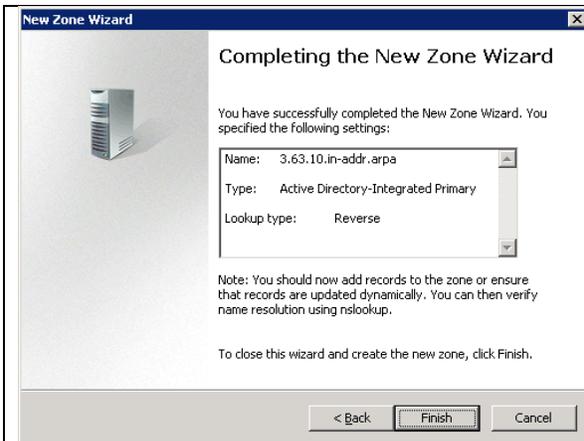


6. Enter the network ID/subnet (the first three octets of the IP address).



7. Select a dynamic update policy.





8. Repeat steps 6 through 11 for other subnets.
9. Test DNS lookups for the new zone by using `nslookup` or `dig`.

```
C:\>nslookup cdot
Server: UnKnown
Address: ::1
```

```
Non-authoritative answer:
Name: cdot.domain.win2k8.netapp.com
Address: 10.63.57.237
```

```
C:\>nslookup cdot
Server: UnKnown
Address: ::1
```

```
Non-authoritative answer:
Name: cdot.domain.win2k8.netapp.com
Address: 10.63.3.68
```

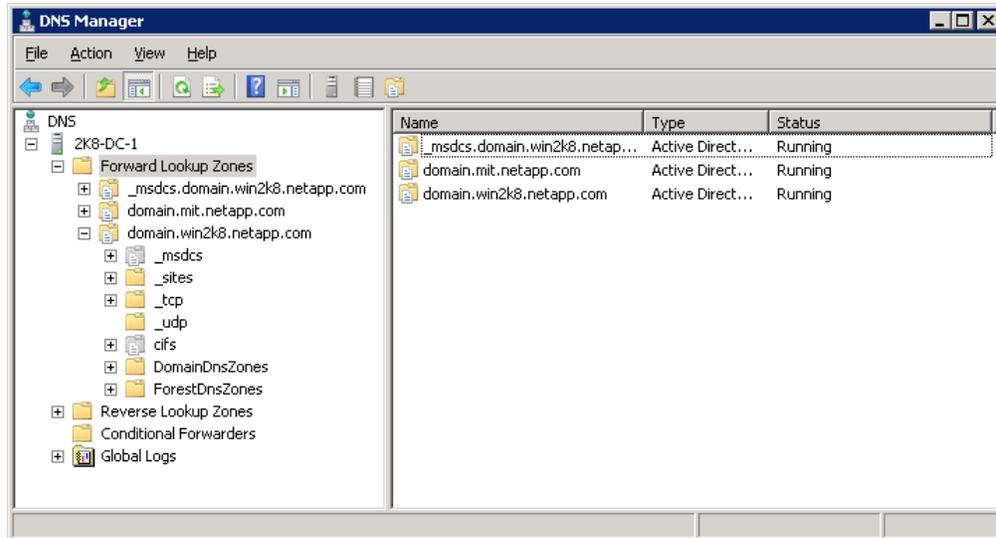
Setting Up DNS Stub Zones in Windows DNS

The following steps show how to set up DNS stub zones in Windows DNS servers. The server version used in the example is Windows 2008R2, but the same steps apply for other Windows servers. For official steps, refer to the [Microsoft TechNet documentation](#).

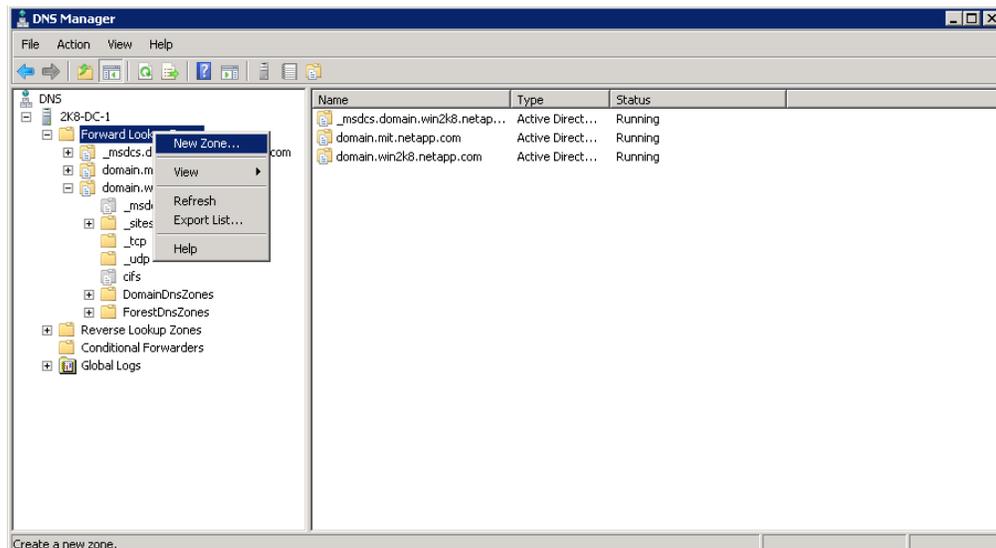
Stub zones are used when a DNS zone needs to be integrated with Active Directory and/or when the zone requires SOA records. With on-box DNS, this is an ideal setup, because data LIFs that listen as DNS servers can be listed as SOA records in stub zones.

Configuration Steps 4) Setting up stub zones.

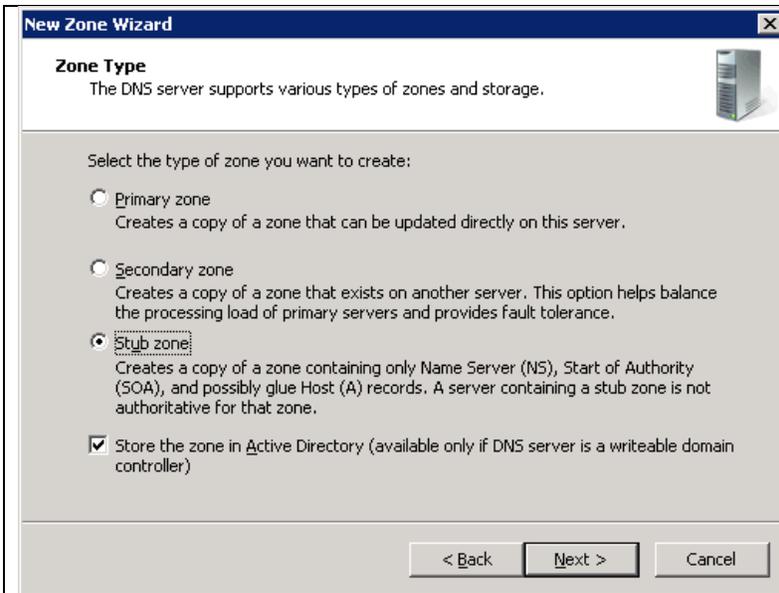
1. Open the DNS Manager console.



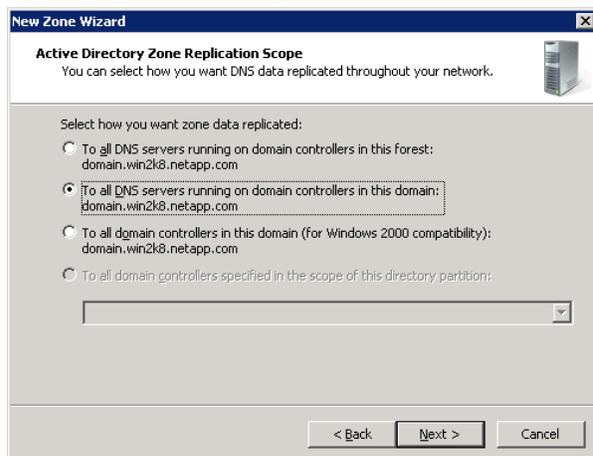
2. Right-click Forward Lookup Zones and select New Zone.



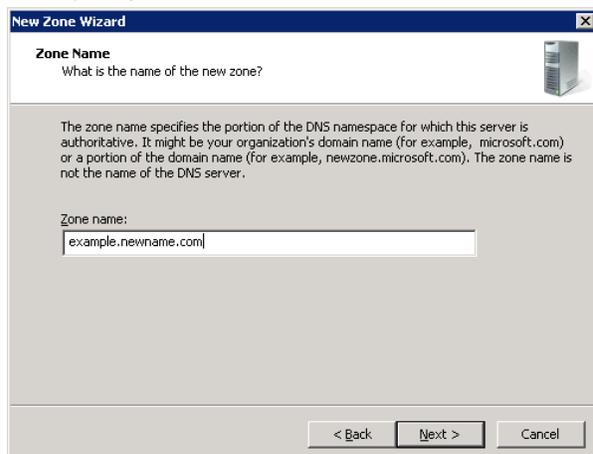
3. Select Stub Zone as the zone.



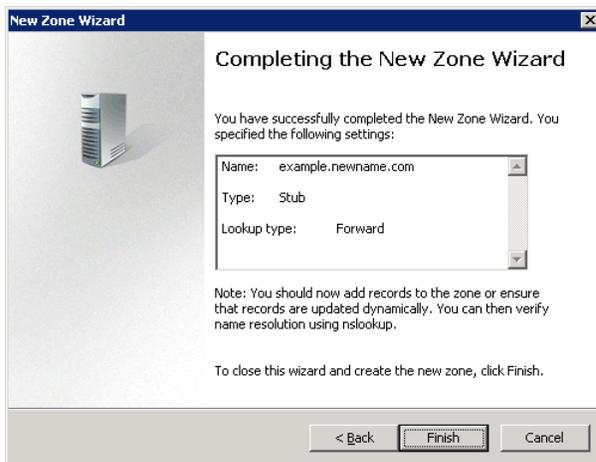
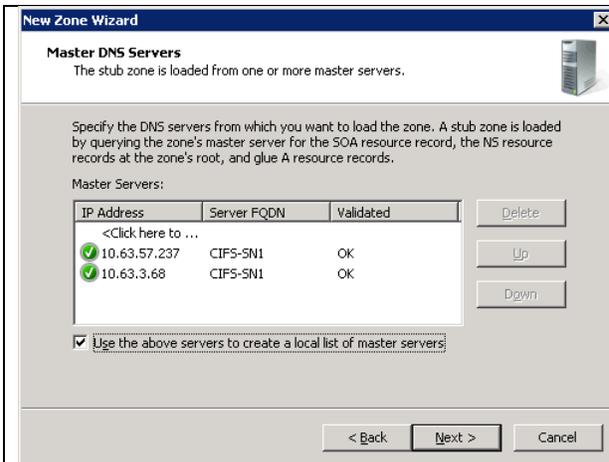
4. Select how zone replication should function.



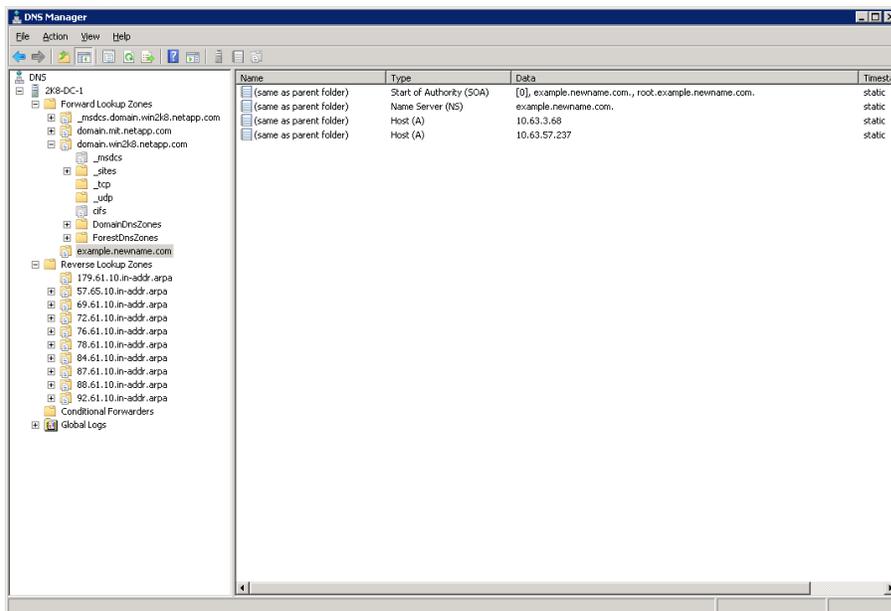
5. Specify the zone name.



6. Add all data LIFs that are configured for on-box DNS to the master DNS server list. Select the Use the Above Servers to Create a Local List of Master Servers check box.



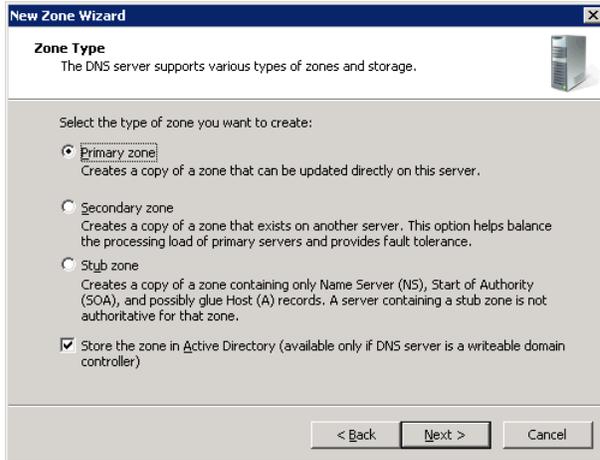
7. Verify that the stub zone has the SOA and NS records.



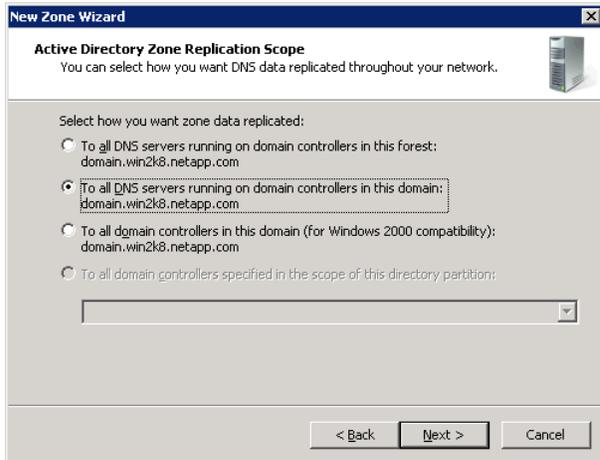
8. Create the reverse lookup zones for the data LIFs.



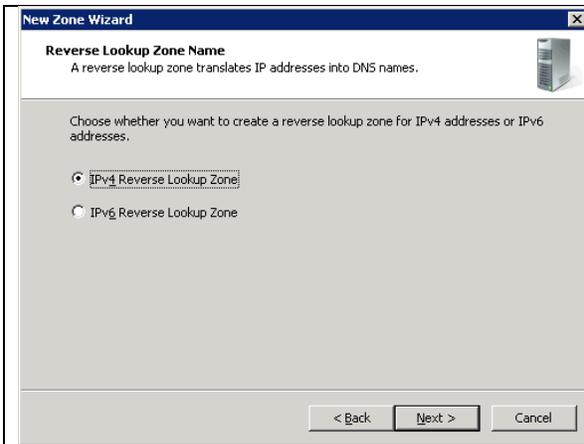
9. Select Primary Zone, because DNS in ONTAP cannot service reverse lookups.



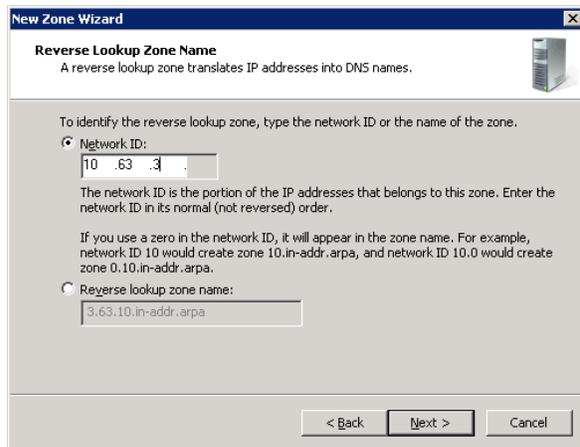
10. Select a zone replication policy to use.



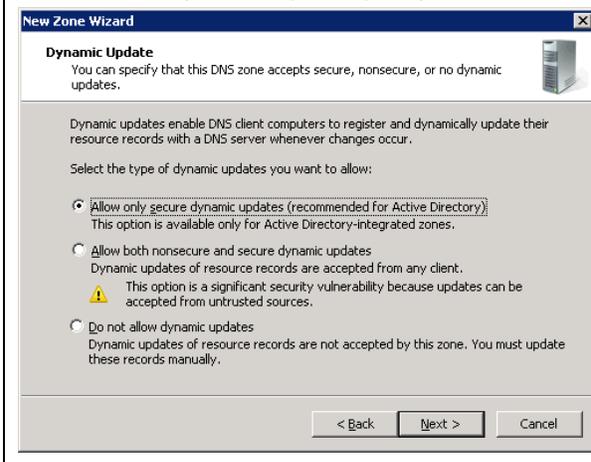
11. Select IPv4 or IPv6 for the lookup zone, depending on what the ONTAP version supports and what the data LIFs use.

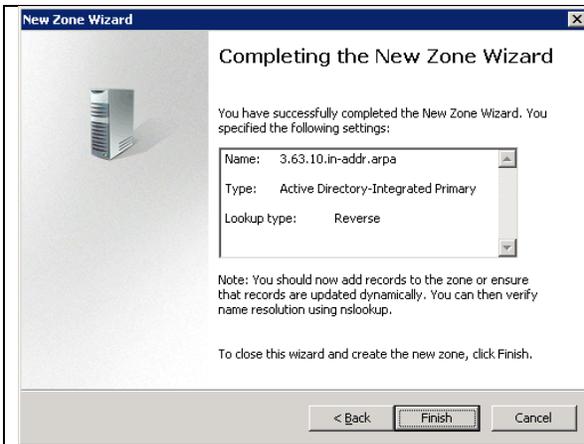


12. Enter the network ID/subnet (the first three octets of the IP address).



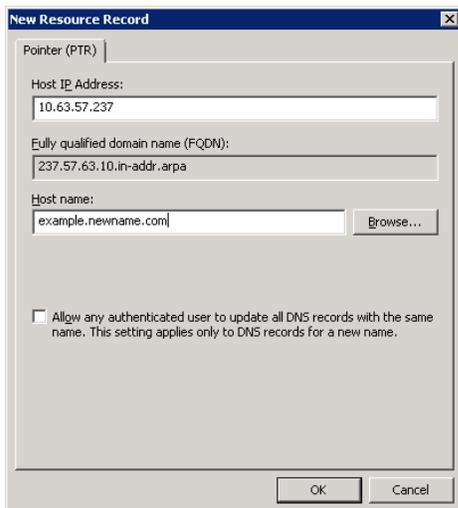
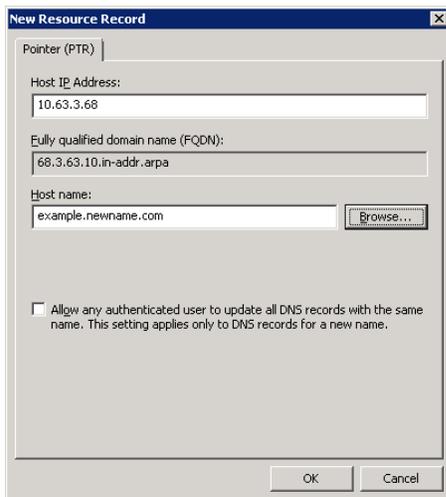
13. Select a dynamic update policy.





14. Repeat steps 8 through 13 for other subnets.

15. Add the PTR records for the data LIFs, because ONTAP does not support reverse name lookups.



16. Use nslookup to test the forward and reverse lookups in DNS.

```
C:\>nslookup example.newname.com
Server: localhost
```

```

Address:  ::1

Name:     example.newname.com
Addresses: 10.63.57.237
          10.63.3.68

C:\>nslookup 10.63.57.237
Server:   localhost
Address:  ::1

Name:     example.newname.com
Address:  10.63.57.237

C:\>nslookup 10.63.3.68
Server:   localhost
Address:  ::1

```

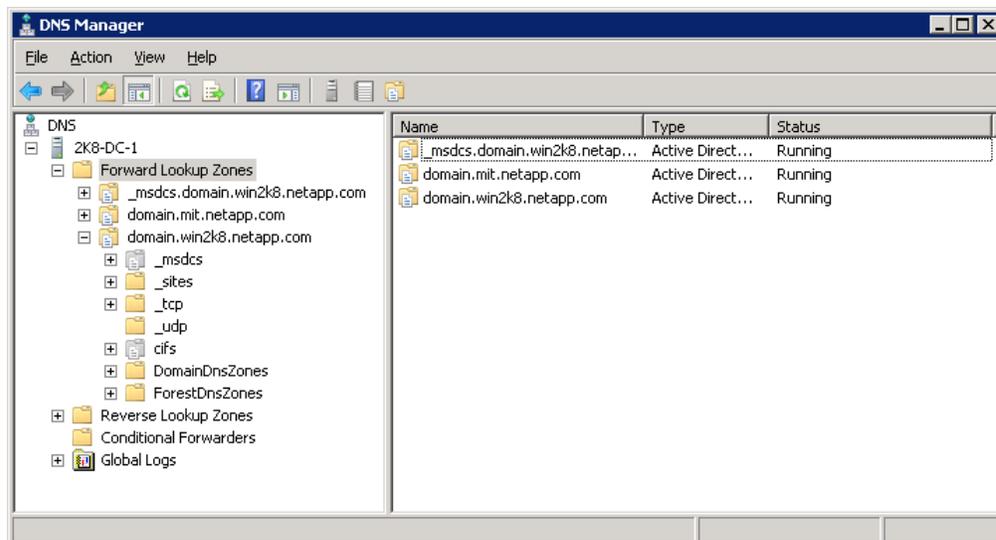
Setting Up Conditional Forwarders in Windows DNS

The following steps show how to set up DNS conditional forwarders in Windows DNS servers. The server version used in the example is Windows 2008R2, but the same steps apply for other Windows servers. For official steps, refer to the [Microsoft TechNet documentation](#).

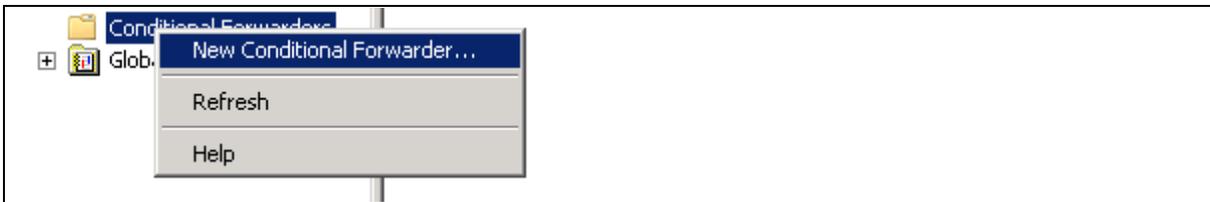
[Conditional forwarders](#) are used to forward DNS queries according to the DNS domain name in the query to a DNS server in the DNS domain. In most cases, conditional forwarders are appropriate to use with on-box DNS when the data LIF DNS domain name to be forwarded is in a different domain than the DNS domain of the main DNS servers. An example is when queries to *example.different.com* are forwarded with a conditional forwarder configured in the DNS domain *domain.com*.

Configuration Steps 5) Setting up conditional forwarders: Windows 2008.

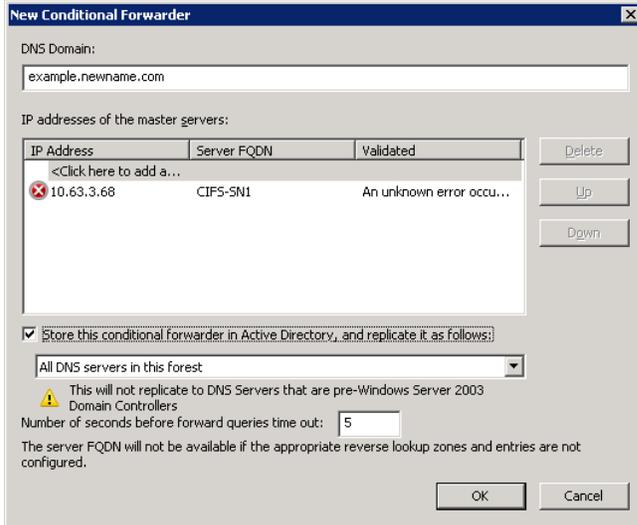
1. Open the DNS Manager console.



2. Right-click Conditional Forwarders and select New Conditional Forwarder.



3. Enter the DNS domain and data LIFs. If an error occurs, the [server might not be sending SOA record requests](#). Either correct that issue or use a [stub zone](#) instead.



4. Click OK and use nslookup to test the forwarded zone.

```
C:\>nslookup example.newname.com
Server: localhost
Address: ::1
```

```
Name: example.newname.com
Addresses: 10.63.57.237
          10.63.3.68
```

```
C:\>nslookup 10.63.57.237
Server: localhost
Address: ::1
```

```
Name: example.newname.com
Address: 10.63.57.237
```

```
C:\>nslookup 10.63.3.68
Server: localhost
Address: ::1
```

4.3 Configuring BIND-Style DNS Servers to Work with On-Box DNS

In many cases, Windows servers are used for DNS resolution, particularly when Active Directory is present in an environment. This is because Active Directory requires DNS for functionality as well as the simple integration and GUI provided by Windows.

However, some environments leverage Linux-based DNS servers, such as BIND or BIND9. In those configurations, the same concepts apply when considering the design of on-box DNS, as covered in the section "Deciding How to Configure the On-Box DNS."

In the following example, the DNS server used is a CentOS/RHEL 7 box using BIND as its DNS server.

The following configurations are covered:

- Data LIFs with DNS zones in the same domain as the primary DNS server
- Data LIFs with DNS zones in a different domain than the primary DNS server

On-Box DNS Configuration: Data LIFs in Same Domain as BIND Server

To use data LIFs in the same domain as the parent domain of the BIND server, use a **subdomain entry** in the zone file. [Subdomains](#) allow the DNS server to pass the requests for a specific zone on to the appropriate servers through zone transfers, providing fault tolerance. If subdomains are not used, the DNS server might think the request is an A/AAAA record request and the lookup will fail with NXDOMAIN (domain does not exist).

In BIND servers, adding zones is as simple as modifying configuration files. To add a subdomain, the following needs to be done:

- Add a zone configuration for the on-box DNS subdomain to the master zone file.
- Add NS and A (glue) records for the data LIFs that will be listening for DNS queries.
- Add an NS record for the parent DNS server.

The following shows how to set up a subdomain for a zone in the same DNS domain as the parent DNS server. This is the SVM's on-box DNS configuration:

```
ontap-tme-prod::> net int show-zones -vserver parisi
(network interface show-zones)

Vserver      Interface Name  DNS Zone                Listen For
-----      -
parisi
              data            onbox.bind.parisi.com   true
              data2          onbox.bind.parisi.com   false

2 entries were displayed.
```

This is the DNS server's domain/host name:

```
# hostname
dns.bind.parisi.com
```

Configuration Steps 6) Sample subdomain zone to add to master zone file.

```
$ORIGIN onbox.bind.parisi.com.
@                IN      NS      onbox.bind.parisi.com.
                 IN      NS      dns.bind.parisi.com.
onbox.bind.parisi.com.  IN     A       10.193.67.226
```

Once these steps are taken, on-box DNS requests are returned for that zone from the cluster:

```
[root@centos7 ~]# nslookup onbox
Server:      10.193.67.227
Address:    10.193.67.227#53

Non-authoritative answer:
Name:   onbox.bind.parisi.com
Address: 10.193.67.226

[root@centos7 ~]# nslookup onbox
Server:      10.193.67.227
Address:    10.193.67.227#53

Non-authoritative answer:
Name:   onbox.bind.parisi.com
Address: 10.193.67.229
```

Adding PTR Records to BIND DNS Servers

In some cases, it might be necessary to add PTR records to BIND DNS servers so that reverse lookups work for the SVM data LIFs participating in the DNS zone. Adding PTR records particularly comes into play when Kerberos is involved.

Adding PTR records is done just the way any other PTR record addition is made. Add the necessary entries to the zone file for the desired reverse lookup zone.

Example:

```
[root@dns named]# cat 67.193.10.in-addr.arpa.zone
$TTL 86400
@ IN SOA      bind.parisi.com.      root.parisi.bind.com. (
    2013042202 ;Serial
    3600       ;Refresh
    1800       ;Retry
    604800    ;Expire
    86400     ;Minimum TTL
)

67.193.10.in-addr.arpa.      IN      NS      dns.bind.parisi.com.

225 IN      PTR      centos7.bind.parisi.com
227 IN      PTR      dns.bind.parisi.com
226 IN      PTR      onbox.cluster.com
229 IN      PTR      onbox.cluster.com
```

Example of working reverse lookups:

```
[root@centos7 ~]# dig PTR 10.193.67.226

; <<> DiG 9.9.4-RedHat-9.9.4-29.el7_2.3 <<> PTR 10.193.67.226
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 44516
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;10.193.67.226.                IN      PTR

;; AUTHORITY SECTION:
.                10793  IN      SOA     a.root-servers.net. nstld.verisign-grs.com.
2016062700 1800 900 604800 86400
```

BIND9 Configurations and Other Third-Party DNS Servers

BIND9 DNS servers use the same general configuration, but the location of the files is different. For example, `named.conf` for BIND9 is stored in `/etc/bind` rather than `/etc/named`. Be sure to check your DNS application's product documentation and man pages for details.

For other third-party DNS servers, such as those that implement GUIs, the concepts behind the design are the same.

- Use subdomains for on-box DNS configurations in the same DNS domain as the parent.
- Use forwarders for on-box DNS configurations in different DNS domains.

For additional information, contact the provider of the third-party GUI.

On-Box DNS Configuration: Data LIFs in Different Domain Than BIND Server

To use data LIFs in a different domain than the parent domain of the BIND server, add a **forwarding zone** entry into `named.conf`. Keep in mind that the forwarding zone might not replicate to other DNS servers, so plan accordingly.

The forwarding entry needs:

- Name of the DNS zone.
- Type of "forward."
- Forwarder entries to IP addresses of the data LIFs to be used as DNS servers.
- If using multiple DNS servers, add the zone to those as well, because `named.conf` might not be configured to replicate to those other servers.

This is the SVM's on-box DNS configuration:

```
ontap-tme-prod::> net int show-zones -vserver parisi
(network interface show-zones)

```

Vserver	Interface Name	DNS Zone	Listen For DNS Query
parisi	data	onbox.cluster.com	true
	data2	onbox.cluster.com	false

```
2 entries were displayed.
```

This is the DNS server's domain/host name:

```
# hostname  
dns.bind.parisi.com
```

Configuration Steps 7) Sample configuration for forwarding zone in named.conf for BIND.

```
zone "onbox.cluster.com" IN {  
    type forward;  
    forwarders {10.193.67.226;};  
};
```

Once these steps are taken, the following are the results of nslookup for that zone:

```
[root@centos7 ~]# nslookup onbox.cluster.com  
Server:      10.193.67.227  
Address:     10.193.67.227#53  
  
Non-authoritative answer:  
Name:   onbox.cluster.com  
Address: 10.193.67.229  
  
[root@centos7 ~]# nslookup onbox.cluster.com  
Server:      10.193.67.227  
Address:     10.193.67.227#53  
  
Non-authoritative answer:  
Name:   onbox.cluster.com  
Address: 10.193.67.226
```

4.4 Configuring Clients to Use ONTAP Data LIFs as DNS Servers

In some cases, clients might need to be configured to use the cluster data LIFs as DNS servers. Instances when this might be necessary include:

- Clients do not have network access to primary DNS servers.
- Primary DNS servers cannot be modified to use zones, delegations, or forwarders.
- General preference.

Clients can use multiple name servers and zones when resolving host names, so clients can use both primary DNS domains as well as the data LIF domains configured on the cluster. It is also possible to use data LIFs as local DNS servers as well as use on-box DNS with general DNS zone configuration on the same SVM.

The first step to configure on-box DNS as the client's DNS name server is to configure on-box DNS on the SVM, enabling at least one data LIF to listen for DNS queries. It is also necessary to confirm that the cluster is sending SOA records.

```
cluster::> net int modify -vserver SVM -lif data -dns-zone cluster.local -listen-for-dns-query  
true  
  
cluster::> net int show -vserver SVM1 -lif data -fields dns-zone,listen-for-dns-query,address  
(network interface show)  
vserver lif address dns-zone listen-for-dns-query  
-----  
SVM1 data 10.193.67.220 cluster.local true  
  
cluster::> set advanced  
cluster:*> network options send-soa show  
Enable sending SOA: true
```

Next, configure the client to use the data LIF as a DNS name server and add the search domain configured on the data LIF.

Configuring Linux Clients with resolv.conf

In Linux clients, such configuring would be done with `resolv.conf` files. The following client shows that, before configuring `resolv.conf`, the DNS domain `cluster.local` could not be resolved.

```
# dig cluster.local

; <<>> DiG 9.9.4-RedHat-9.9.4-29.el7_2.3 <<>> cluster.local
;; global options: +cmd
;; connection timed out; no servers could be reached
```

When the client is configured to use the data LIF with that DNS zone, it can resolve properly.

Configuration Steps 8) On-box data LIFs as DNS servers—Linux clients.

```
# cat /etc/resolv.conf
# Generated by NetworkManager
search cluster.local
nameserver 10.193.67.220

# dig cluster.local

; <<>> DiG 9.9.4-RedHat-9.9.4-29.el7_2.3 <<>> cluster.local
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15220
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 1
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags;; udp: 4096
;; QUESTION SECTION:
;cluster.local.                IN      A

;; ANSWER SECTION:
cluster.local.                0       IN      A      10.193.67.220

;; AUTHORITY SECTION:
cluster.local.                86400   IN      NS     cluster.local.

;; Query time: 24 msec
;; SERVER: 10.193.67.220#53(10.193.67.220)
;; WHEN: Tue Jun 21 13:02:44 EDT 2016
;; MSG SIZE rcvd: 72
```

Reverse lookup works as well:

```
# dig 10.193.67.220

; <<>> DiG 9.9.4-RedHat-9.9.4-29.el7_2.3 <<>> 10.193.67.220
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: REFUSED, id: 60475
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 1
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags;; udp: 4096
;; QUESTION SECTION:
;10.193.67.220.                IN      A

;; Query time: 12 msec
;; SERVER: 10.193.67.220#53(10.193.67.220)
;; WHEN: Tue Jun 21 13:03:18 EDT 2016
;; MSG SIZE rcvd: 42
```

Other DNS servers can be added to the configuration and resolve names properly. For example, if we add a Google DNS server, we can resolve `google.com`:

```
# cat /etc/resolv.conf
# Generated by NetworkManager
search cluster.local
nameserver 10.193.67.220
nameserver 8.8.8.8

# nslookup google.com
;; Got recursion not available from 10.193.67.220, trying next server
Server:      8.8.8.8
Address:     8.8.8.8#53

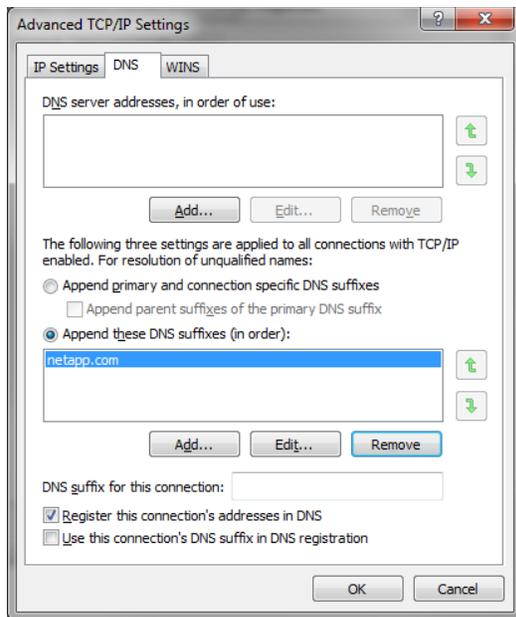
Non-authoritative answer:
Name:   google.com
Address: 216.58.219.206
```

Configuring Windows Clients to Use On-Box DNS as Independent DNS Servers

Windows clients can also use ONTAP data LIFs as DNS servers for data access on an SVM. Windows configurations generally use a GUI, but CLI utilities such as PowerShell can also be used. This example covers the GUI configuration and leverages the use of the data LIFs as DNS servers in addition to an existing DNS configuration.

Configuration Steps 9) On-box data LIFs as DNS servers—Windows clients.

This is the existing DNS configuration for the Windows client:



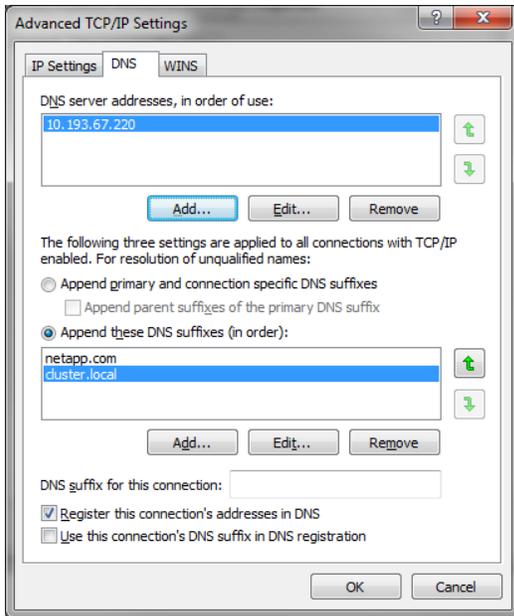
The DNS servers are being pulled through DHCP. The DNS suffixes have been manually configured.

As it currently stands, nslookup requests for the data LIF's zone (cluster.local) fail:

```
C:\>nslookup cluster.local
Server:  dns.netapp.com
Address: 10.193.67.200

*** dns.netapp.com can't find cluster.local: Non-existent domain
```

To leverage the cluster's data LIFs as DNS servers to return cluster data LIFs when cluster.local is queried, the configuration should look like this:



Here, we added only the data LIF participating in on-box load balancing as a DNS server. Other DNS servers would need to be added in addition.

Once the new server is added, flush DNS cache (Windows caches DNS for 24 hours by default) and try the nslookup for the cluster zone:

```
C:\> nslookup cluster.local
Server: cluster.local
Address: 10.193.67.220

Name: cluster.local
Address: 10.193.67.220
```

5 Conclusion

On-box DNS load balancing is a viable alternative to using external solutions, such as round-robin DNS load balancing. Being able to load balance DNS requests based on load helps alleviate the overall impact to a scale-out cluster and provides an intelligent method to serve NAS connectivity in enterprise environments.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.