



Authoritative DNS Capacity and Performance Guidelines

This chapter provides information on Authoritative DNS capacity and performance guidelines to help with system sizing for 64-bit Cisco Prime Network Registrar 8.3.5.4 and later.

- [DNS System Deployment Limits, on page 1](#)
- [DNS Database Architecture , on page 2](#)
- [DNS System Sizing , on page 3](#)

DNS System Deployment Limits

Cisco Prime Network Registrar makes the following recommendations on maximum Authoritative DNS System configuration sizes. The following recommendations are as per Cisco Prime Network Registrar Authoritative DNS server which can be a primary, primary HA, or secondary server. A redundant DNS architecture will contain multiple of these types of servers all servicing the same data. Therefore, the capacity can be expanded horizontally by introducing a new set of servers. These recommendations are guidelines to ensure a properly functioning DNS deployment.



Note DNSSEC enabled zones (Cisco Prime Network Registrar 9.1 and later versions) will include auto-generated RRs that significantly increase the number of RRs in the zone.

- Maximum of 25 million RRs per Authoritative DNS server (primary, HA pair, or secondary server), ideally not to exceed 2 million RRs per zone. Multiple DNS primary servers can be used for deployments requiring more RRs.
- Maximum of 10000 zones per Authoritative DNS server (primary, HA pair, or secondary server). Multiple DNS primary servers can be used for deployments requiring more zones.
- Maximum of 4 secondary servers per primary or HA pair.
- Maximum of 2 tiers of secondary servers (first tier secondaries and second tier secondaries).
- Maximum of 2 second tier secondary servers per first tier secondary server.

DNS Database Architecture

The Authoritative DNS servers utilize a combination of in-memory cache and on-disk databases to store and maintain authoritative RR data. For sizing purpose, assume an each RR requires 300 bytes of memory for the RR cache and 300 bytes of disk space for the RR DB. The CSET DB has a higher disk space requirement for each RR since it records changes to the RR set, but those changes are capped to the number of history changes kept per zone.

RR DB

- Database that stores all RRs (protected and unprotected) for the zones configured on a DNS server.
- On primary DNS servers, RR data edits are written to the RR DB either through administrative actions (that is, RR adds), or DNS updates and zone scavenging. On secondaries, the RR DB is written through zone transfers.
- The RR DB is required for all ADNS servers (primary/secondary).

RR Cache

- Increases query performance by storing a subset of the RR DB data (stores entire name sets).
- Most active RR data is stored to RR cache dynamically as part of RR DB lookups generated by DNS query processing.
- The memory foot print of the RR Cache is capped by a configurable DNS server attribute (*mem-cache-size*). When the maximum cache size has been reached, the DNS server will remove older entries from the cache to make room for newer entries. Each RR requires approximately 300 bytes of memory.
- DNS server reload/restart causes the RR cache to be deleted. When the server starts up again, it is rebuilt based on query traffic.
- The RR cache is required for all ADNS servers (primary/secondary).

CSET DB

- Database that stores RR changes (adds, deletes, protection changes, and refreshes) needed to respond to the incremental zone transfer requests (IXFRs).
- RR changes are first stored in the RR DB and then persisted to the CSET DB.
- For DNS servers that do not need to service incremental zone transfers (that is, secondaries that do not send outbound IXFRs), server performance can be increased by disabling persisted change sets (*csetdb-persist-csets*). By default, changes are automatically persisted to the CSET DB.
- DNS maintains only a limited configurable number of changes (*csetdb-htrim-max-cset-kept*) and automatically trims entries when the maximum has been reached. Trimming helps limit the database size. For deployments with DNS updates, it is recommended that the number of changes kept is increased to avoid full zone transfers.
- If the CSET DB is deleted, the DNS server will create an empty database and respond with full zone transfers (AXFRs) until new zone history data is populated into the database.

HA DB

- Database that stores state information about the DNS HA pair as well as data about RR changes during a communications interrupted or partner down event.
- Only applicable on primary HA DNS servers (main and backup).
- If the HA DB is deleted, HA synchronization causes all zone data to be pushed from the HA main to the HA backup.

DNS System Sizing

A Cisco Prime Network Registrar DNS deployment can be categorized as small, medium, or large depending on the number of RRs/zones, DNS update activity, and recovery time during an outage or update. The number of zones can have an impact on the size of the deployment, primarily the number of RRs is the deciding factor. Also, if the DNS deployment requires a large number of RRs/zones, it is recommend that multiple DNS deployments be used - ideally segregating the data appropriately so that related zones/RRs are configured together.



Note To ensure a properly functioning Authoritative DNS system, it is important to monitor system disk space and memory. If the Authoritative DNS server runs out of memory, it will crash. If it runs out of disk space, it will no longer be able to service requests and the databases may become corrupt and unusable.

Regional Management of DNS Deployments

The regional server provides license management of all Cisco Prime Network Registrar local clusters, and allows for central management and replication of Cisco Prime Network Registrar DNS deployments. Follow the below recommendations for system sizing and configuration adjustments to be made when using regional DNS cluster management:

- A minimum of 4 CPUs
- A minimum of 8 GB of RAM
- Disk space should be at minimum an aggregate of the disk size of all the managed DNS (main) primary clusters.
- On large DNS deployments, replication of unprotected RRs should be disabled (*poll-replica-rrs*).

Small Deployment

- 1-1000 RRs and 1-100 zones
- Mainly static data; zone edits are primarily done by administrators.
- Typically consists of one primary and a secondary server.
- DNS Caching server is not required or can be handled by hybrid mode.
- DNS can be recovered from a shadow backup within a matter of minutes with little to no impact on production.

- A minimum of 2 CPUs
- A minimum of 4 GB of RAM
- A minimum of 10 GB of disk space

Medium Deployment

- 1000-100,000 RRs and 100-1000 zones
- A pretty even mix of static and dynamic data; 100 updates per second or less.
- Typically consists of one primary and two to four secondaries.
- Typically consists of two to four DNS Caching Servers. DNS Caching Servers must be deployed on separate machines or VMs.
- DNS can be recovered from a shadow backup within an hour with minimal impact to production.
- A minimum of 4 CPUs
- A minimum of 8 GB of RAM
- A minimum of 25 GB of disk space. On the primaries, the number of change sets kept (*csetdb-htrim-max-cset-kept*) should be increased. The value will depend on how many DNS updates are handled by the system, but should be between 1000 and 5000.

Large Deployment

- 100,000-25,000,000 RRs and 1000-10,000 zones
- Dynamic data makes up a larger percentage of the data; thousands of updates per second.
- Typically consists of two primaries (DNS HA pair) and four secondaries.
- Typically consists of four or more DNS Caching servers.
- DNS recovery is complex and must be done during a maintenance window; DNS servers can take an hour or more to recover from a shadow backup.
- A minimum of 8 CPUs
- A minimum of 16 GB of RAM. The DNS RR cache memory size (*mem-cache-size*) should be increased (approximately 300 bytes per RR, but not to exceed 2,000,000 KB).
- A minimum of 100 GB of disk space. On the primaries, the number of change sets kept (*csetdb-htrim-max-cset-kept*) should be increased. The value will depend on how many DNS updates are handled by the system, but should be between 5000 and 10,000.