

Cisco Nexus 3000 Series Switches: Understand, Configure and Tune the Forwarding Table

What You Will Learn

This document explains the structure of the Cisco Nexus® 3000 Series Switches forwarding table, the logic of the forwarding table, and how to configure and tune the forwarding table for optimal results.

Cisco Nexus 3000 Series Switches Forwarding Table Structure

Cisco Nexus 3000 Series Switches are equipped with a switch-on-a-chip (SoC) forwarding application-specific integrated circuit (ASIC). This ASIC contains two hardware tables for the forwarding logic:

- Longest prefix match (LPM) table, also called the routing table
- Host table

Cisco Nexus 3000 Series Forwarding Table Logic

The routing entries are defined from /0 to /31 for IPv4 and from /0 to /64 for IPv6 entries and hosted in the LPM table. The IPv6 routing entries from /65 to /127 are in a dedicated fixed-size routing table called LPM_128. The host entries are defined by /32 for IPv4 and /128 for IPv6.

The LPM table can contain routing entries as well as host entries. The entries can be used by IPv4 and IPv6 packets. The table is shared between routes (LPM) and hosts. The host entries get access to this table only if:

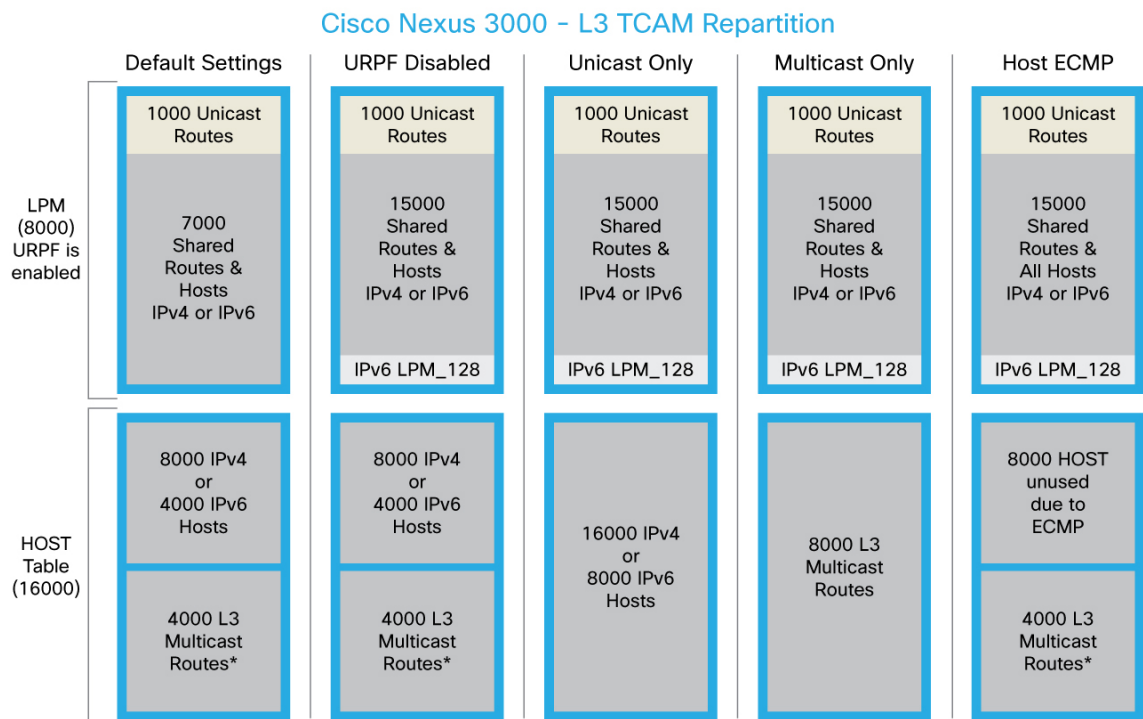
- The host table is already filled by host entries or multicast entries
- Equal-cost multipath (ECMP) for hosts (/32 or /128) is enabled

The host table can contain host entries as well as multicast routes (mroutes). The multicast route table size is a user-configurable value. The remaining entries in this table can be used by the host entries.

An mroute takes two entries in the host table: one for (S, G), and one for (*,G). When there are IPv4 and IPv6 host entries, space is allocated on a first come, first served basis, unless the IPv6 host table is configured manually as discussed in the next section (see implementation 4). By default, the multicast table consumes half the host table size.

Figure 1 shows the forwarding table.

Figure 1. Cisco Nexus 3000 Series Forwarding Table



Cisco Nexus 3000 Series Forwarding Table Configuration

There are four configuration options for the forwarding tables:

- Option 1: Disable Unicast Reverse Path Forwarding (URPF) check for the LPM table routing entries. This is a global command that doubles the size of the LPM table when the feature is disabled. By default, URPF is enabled.
- Option 2: Enable ECMP for (/32 or /128) routes. This configuration will force all /32 (for IPv4) or /128 (for IPv6) routes to be programmed exclusively in the LPM table so that they can be subject to ECMP.
- Option 3: Configure the size of the mroute table between 0 and 8000. Note that an mroute consumes two entries in the host table: one for (S,G) and one for (*,G). Therefore, 8000 mroutes will occupy the entire host table of 16,000 entries. Any host entry will then occupy the shared LPM table space.
- Option 4: Configure the size of the IPv6 host limit (/128) space dedicated to IPv6 in the host table. This manual configuration can help ensure space for a specific number of IPv6 entries in the table if required.

The forwarding tables can be respectively configured as follows:

- Option 1: Disable **urpf checks**: “**system urpf disable**”.
- Option 2: Enable ECMP for host entries with “**hardware profile unicast enable-host-ecmp**”.
- Option 3: Configure the size of the mroute table with “**hardware profile multicast max-limit <0-8000>**”.
- Option 4: Configure the size of the IPv6 host table with “**hardware profile ucast6 max-limit <0-8000>**”.

Note: Any configuration change is disruptive and a switch reload is needed. Remember to save the running configuration prior to the reload operation.

What Are the Sizes of the LPM and Host Tables?

- The LPM table size by default is 8000 (8191) entries. The size can be doubled, to 16,000 (16,384) entries by disabling URPF.
- 2. The host table size is fixed. Depending on the model of hardware, it will be 8000 entries or 16,000 entries. The Cisco Nexus 3064-PQ (first generation) has 8000 entries; all other Cisco Nexus 3064 Switches (3064-E and 3064-X) and Cisco Nexus 3000 Series Switches (3048 and 3016) have 16,000 entries. For the purpose of this document, the 16,000-entry value will be used.
- By default, multicast consumes half the host table. The other half is used by host entries.

What Is the Default Forwarding Table Configuration?

The LPM table has the URPF check enabled, and therefore the LPM table size is 8000 entries, not 16,000 entries. The IPv6 entries are in a different table, called LPM_128. LPM_128 can contain 128 IPv6 routes and accepts 256 entries, taken from a dedicated table space in addition to the 16,000 LPM table size.

The first 1024 entries in the LPM table are reserved for routes only. The remaining 15,000 entries can be shared between routes and hosts. This reserved 1024 value is not configurable.

The host table by default has 4096 mroute entries defined, using 8192 entries. The remaining 8192 are available for host entries: IPv4 or IPv6. Changing the size of the multicast table conditions the remaining host table size.

How Many Entries Are Consumed in the Forwarding Table?

- An IPv4 route consumes one entry in the LPM table.
- For each routing entry, the network and broadcast IP addresses are automatically programmed in the host table and consume one entry each.
- A routing IPv6 /0 to /64 entry consumes two entries in the LPM table. An IPv6 /65 to /127 route consumes two entries in the LPM_128 table, which is distinct from the LPM table.
- A host IPv4 entry consumes one entry in the host table.
- A host IPv6 entry can consume up to two entries in the host table: unique-local and global. One link-local entry is added as a generic entry in the LPM table, with FE80::/8 as the route.
- An IPv4 mroute consumes two entries in the host table: one for (*,G) and one for (S,G).

Verifying the Table Use

You can verify table use as shown here. Note that the output may vary from one version of software to another. The display shown here uses software release version 5.0(3)U3.

```
N3K# show hardware profile status
  Total LPM Entries = 8191.      -> This is the LPM table size, when URPF is
  enabled (default)
  Total Host Entries = 16384.    -> This is the size of the host table size on
  Trident+
  Reserved LPM Entries = 1024.  -> This is a reserved 1024 entries from the LPM
  table for routing
```

```
Reserved Host Entries = 4000.    -> This value can be ignored; removed in
5.0(3)U(4)

Max Host4/Host6 Limit Entries (shared)= 12384/6192*    -> Shared table for host
IPv4 and IPv6 with respective capacity

Max Host6 Limit Entries = 1000.

Max Mcast Limit Entries = 2000. -> Default 2000 mroutes value (using 4000
entries)

Used LPM Entries (Total) = 6.
Used IPv4 LPM Entries = 4.
Used IPv6 LPM Entries = 2.
Used IPv6 LPM_128 Entries = 1.
Used Host Entries in LPM (Total) = 0.
Used Host4 Entries in LPM = 0.
Used Host6 Entries in LPM = 0.
Used Mcast Entries = 0.
Used Mcast OIFL Entries = 0.
Used Host Entries in Host (Total) = 7.
Used Host4 Entries in Host = 7.
Used Host6 Entries in Host = 0.
```

Conclusion

The Cisco Nexus 3000 Series forwarding architecture provides a customizable capability for IPv4 and IPv6, unicast and multicast, routing table and host table entries. This document explains the table structure, the configuration possibilities, and the architecture.

For More Information

<http://www.cisco.com/go/nexus3000>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)