Application Note 005: Internet Protocol Address Filtering

Cypress Semiconductor Corporation © 2001 does not assume responsibility for use of any circuitry or application specifications described herein. No circuit patent licenses are implied, and Cypress Semiconductor Corporation reserves the right to change said circuitry or application specifications at any time and without notice. Cypress Semiconductor Corporation assumes no liability arising out of any customer application or assistance arising from the information on this page and any pages linked hereto. No license, either express or implied, is granted to intellectual property including patent rights of Cypress Semiconductor Corporation or any third party. Cypress Semiconductor Corporation products are not intended, authorized, or warranted to be suitable for life-support applications or any other critical applications that could involve the potential risks of death, personal injury, or severe property or environmental damage. All product and company names mentioned in this document are the trademarks of their respective holders.

1 INTRODUCTION

This application note illustrates how Internet protocol (IP) address filtering can be effectively performed using Cypress Semiconductor Corporation's CYNSE70032/LNI7010 and/or CYNSE70064/LNI7020 Network Search Engines (NSEs).

1.1 Summary

There is a direct relationship between the value of the Internet and the number of sites that are connected to it. As the Internet grows, the value of each site's connection to it grows because it provides the organization that owns it with access to an ever-expanding user/customer population.

In many situations, organizations agree to share limited-access resources with other organizations as part of a consortia, financial associations, or other resource-sharing collaborations. In such an agreement, an enterprise will define its user community as one that has access to some network resource. This community is typically large, numbering perhaps in the tens of thousands, and membership may be volatile over time, reflecting for example the characteristics of a student body at a large university. The operator of the network resource, which may a web site or a resource reached by other protocols such as Telnet terminal emulation or other information retrieval protocol, must decide whether users seeking access to the resource are actually members of the user community that the licensee organization originally defined as part of their license agreement.

For this reason, the protection of information must be taken seriously. There are two major aspects to this.

- Access to information is controlled in an appropriate way, with people able to see the information appropriate to their roles inside and outside of the organization. Also, compliance to policies set for managing access to IT resources is necessary.
- Steps must be taken to ensure that information integrity is maintained, and that unauthorized changes cannot be made.

IP filtering, still the predominant technology, can be used to control access to world-wide web (Web) resources. Each computer connected to the Internet has a unique address. These addresses are arranged in a hierarchy of domains, subdomains, and machine numbers. IP filtering works by restricting access to machines in a particular domain, subdomain, or even to specific machine addresses.

With IP filtering, the licensed organization guarantees the resource operator that all traffic coming from a given set of IP addresses (perhaps all IP addresses on one or more networks) represent legitimate traffic on behalf of the license organization's user community. The resource operator then simply checks the source IP address of each incoming request.

IP filtering compliments and leverages the capability of virtual private network (VPN), Policy Management, class of service (CoS), quality of service (QoS), voice over IP (VoIP), and other applications.

1.2 IP Address basics

The IP addressing scheme is an integral part in the process of routing IP data through the Internet. Each host on a transmission-control protocol/Internet protocol (TCP/IP) network is assigned a unique 32-bit logical address. The IP address is divided into two main parts: the network number and the host number.

The network number identifies the network. If the network is to be a part of the Internet, the network number is assigned by the Internet Network Information Center. The host number identifies a host and is assigned by the local network administrator.

1.3 IP Address Format

Dotted-Decimal Notation. In order to make Internet addresses more user-friendly, IP addresses are often expressed as four-decimal numbers. Each number is separated by a dot. This format is known as "dotted-decimal notation." Dotted-decimal notation divides the 32-bit Internet address into four eight-bit fields and specifies the value of each field independently as a decimal number (with the fields separated by dots).

The 32-bit IP address is grouped eight bits at a time; each group of eight bits being an octet. Each of the four octets is separated by a dot, and represented in a decimal format. Every bit in an octet has a binary weight (128, 64, 32, 16, 8, 4, 2, and 1). The minimum value for an octet is 0 (all bits set to 0), and the maximum value for an octet is 255 (all bits set to 1). For example, the IP address can be represented as 209.237.20.193. Each of the decimal digits represents a string of four binary digits. Thus, the IP address is a string of 0s and 1s, as shown in Figure 1.

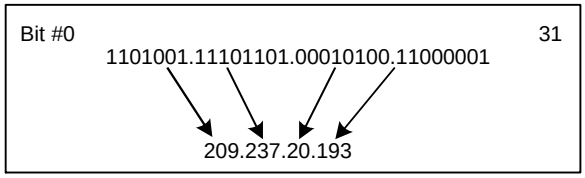

FIGURE 1. IP ADDRESSES

1.4 IP Address Classes

IP addressing supports the following three address classes: Class A, Class B, and Class C. In a class A address, the first octet is the network portion. Thus, the class A address 209.237.20.193 has a major network address of 209. Octets 2, 3, and 4 represent hosts. Class A addresses are used for networks that have more than 65,536 hosts.

In a class B address, the first two octets identify the subnetwork. Thus, the class B address of four (the next 16 bits) are for the hosts. Class B addresses are used for networks that have between 256 and 65,536 hosts.

In a class C address, the first three octets represent the network portion. The class C address 295.45.9.37 has a major network address of 295.45.9. Octet 4 is for hosts. Class C addresses are used for networks with less than 254 hosts. See Table 1.

	FIRST OCTAL (DECIMAL)	HIGH-ORDER BITS	IP ADDRESS EXAMPLE
Class A	1–126	0	111.49.79.16
Class B	128–191	0	128.11.21.7
Class C	192–223	110	209.237.20.193

TABLE 1. IP ADDRESS CLASSES

1.5 Class Full Network Masks

Each of the address classes contains a set of class full network masks. The network mask defines which bit(s) of the 32-bit address are defined as the network portion, and which are the host portion. The network mask is calculated by setting all bits in the octets designated for the network portion to a value of 1, and all bits in the octets designated for the host portion to a value of 0. For example, the Class A network mask is defined as 255.0.0.0. Similarly the Class B network mask is 255.255.0.0. And the Class C network mask is 255.255.255.0. Figure 2 below summarizes the network and host portion of each address class.

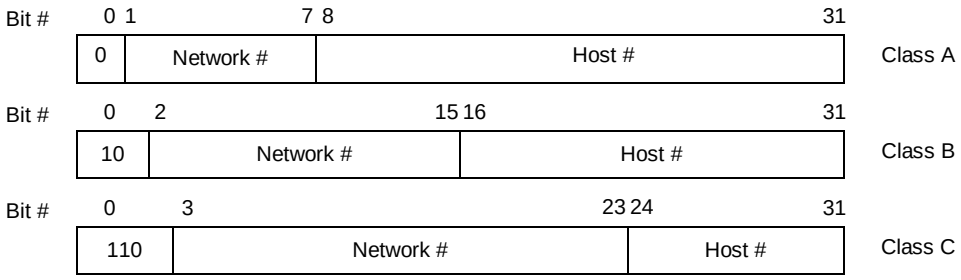

FIGURE 2. NETWORK AND HOST PORTION OF EACH ADDRESS CLASS

2 SUBNETWORK MASK

A subnetwork is an identifiably separate part of an organization's network. Typically, a subnetwork may represent all the machines at one geographic location, in one building, or on the same local area network (LAN). Having an organization's network divided into subnetworks allows it to be connected to the Internet with a single shared network address. Without subnetworks, an organization could get multiple connections to the Internet, one for each of its physically separate subnetworks. However, this requires an unnecessary use of a limited number of network numbers.

Once a packet has arrived at an organization's gateway or connection point with its unique network number, it can be routed within the organization's internal gateways using the subnetwork number. The router knows which bits to consider by looking at a subnetwork mask. Using a mask saves the router having to handle the entire 32-bit address because it can simply look at the bits selected by the mask.

2.1 IP Subnetwork Addressing

All classes of IP networks can be divided into smaller networks called subnetworks. Dividing the major class network is called subnetworking. Subnetworking provides network administrators with several benefits, such as extra flexibility, more efficient use of network address utilization, and broadcast traffic delivery (because a broadcast will not cross a router).

2.2 IP Subnetwork Mask

“Borrowing” bits from the host field and designating them as the subnetwork field creates a subnetwork address. The number of borrowed bits is variable and is specified by the subnetwork mask.

2.3 How a Router Routes a Packet

When a router receives a packet, it makes a routing decision based on the destination address portion of the packet. It then looks up the destination address in its routing table. If the destination address is within a known network/subnetwork, the router forwards the packet to the next-hop gateway for that destination network/subnetwork. Once the packet leaves the router, it is the responsibility of the next-hop gateway to forward the packet to its final destination. If the router does not have the destination network in its routing table, it may forward the packet to a predetermined default gateway (if configured) and let the default gateway handle getting the packet to the destination network. Otherwise, it will drop the packet and inform the sending host that the network is not reachable.

Subnetworking addresses the problem of expanding the routing table. It also ensure that the subnetwork structure of a network is not visible outside of the organization’s private network.

3 LNI7010/20 SEARCH ENGINES

The LNI7010/LNI7020 devices are high-performance, pipelined, synchronous NSEs designed using the Associated Processing Technology™ (APT). They are user-configurable into tables as wide as 272 bits with cascaded depths of up to 992K 32-bit addresses. Their high speed of 83 million lookups per second and high capacity of incorporating nearly a million addresses can be employed in a variety of networking and communications applications that require fast searches of various tables. They provide a performance advantage over other memory-based search algorithms such as binary or tree-based searches by comparing the desired information against the entire list of pre-stored entries in a single cycle, thereby giving an order-of-magnitude reduction in search time.

The LNI7010 NSE is organized as 8K x 136 bits, but it can also be configured as 4K x 272 bits or 16K x 68 bits. The LNI7020 NSE is organized as 16K x 136 bits, but can also be configured as 8K x 272 bits or 64K x 72 bits. The LNI7010 NSE can sustain 83 million searches per second on any subfield of a 68- or 136-bit field, and the 4-meg LNI7040 NSE can sustain 100 million searches per second on any subfield of a 72- or 144-bit field. This makes them the fastest search engines in the market. These high-speed, high-capacity chips can be employed in a variety of networking and communications applications that require fast searches of various tables. Figure 3 shows various configurations of these devices.

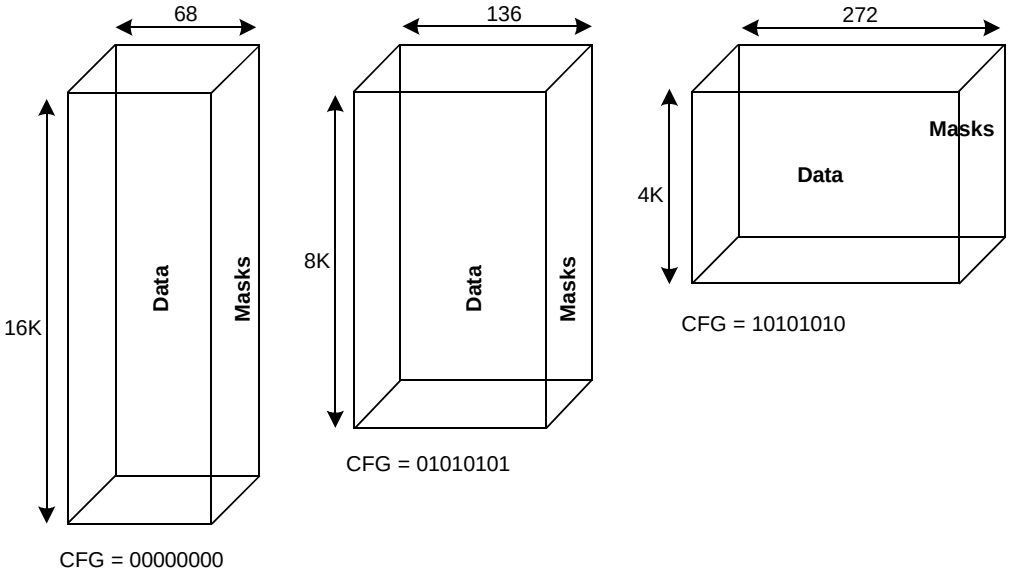

FIGURE 3. VARIABLE-WIDTH TABLE CONFIGURATIONS

The LNI7010/LNI7020 NSEs contain mask registers for each data location. In addition, the devices contain 16 68-bit global mask registers (GMRs) that can be dynamically selected in every search operation to select the search subfield. These mask registers provide an easy way to move data to masks and to enable selective lookups for subnetworks.

4 IP ADDRESS FILTERING

In Table 2 below, eight different networks have been arbitrarily selected. Some of them will be allowed access to the network, so they will be entered in the data array in *Include* Table. Those IP addresses that are not permitted access to the network will be entered in the *Exclude* Table.

TABLE 2. ARBITRARILY SELECTED IP ADDRESSES

	NETWORK ADDRESS	MASK	# MASKS	# HOSTS
I	162.11.35.64	FFFFFFE0	27	32
E	181.22.14.128	FFFFFFE0	27	32
I	181.21.41.32	FFFFFFF0	28	16
E	111.49.79.16	FFFFFFF8	29	8
I	90.47.79.120	FFFFFFF8	29	8
E	179.44.31.80	FFFFFFF8	29	8
I	75.125.159.112	FFFFFFF4	30	4
E	175.43.31.70	FFFFFFF0	30	4

5 PROCEDURE FOR UPDATING TABLES

Upon initialization, the NSE data array should be written with all bits set to 0 and the mask array should be written with all bits set to FF. Upon receiving the new IP address over the databus, it is necessary to search for the entry in the *Allow* table . Initially, the entry will not be found. The table can then be updated using the LEARN command if the entry is not found in the *Deny* table.

Two segments have been created, as illustrated in Figure 4. One segment is for the authorized IP addresses to access the Internet, and the other is to prevent access for unauthorized IP addresses. The corresponding mask is set in the mask array. Bit 1 is a table management bit, and is set to 0 in the *Include* table and to 1 in the *Exclude* table. **Note**. In this example, the table is 68 bits wide. If any IP address is encountered that is not a part of either *Allow* or *Deny* tables, a new entry of the IP address can be added in the last location of the *Deny* table. The corresponding mask for this new IP address should be entered as FFFFFFFF in the mask array.

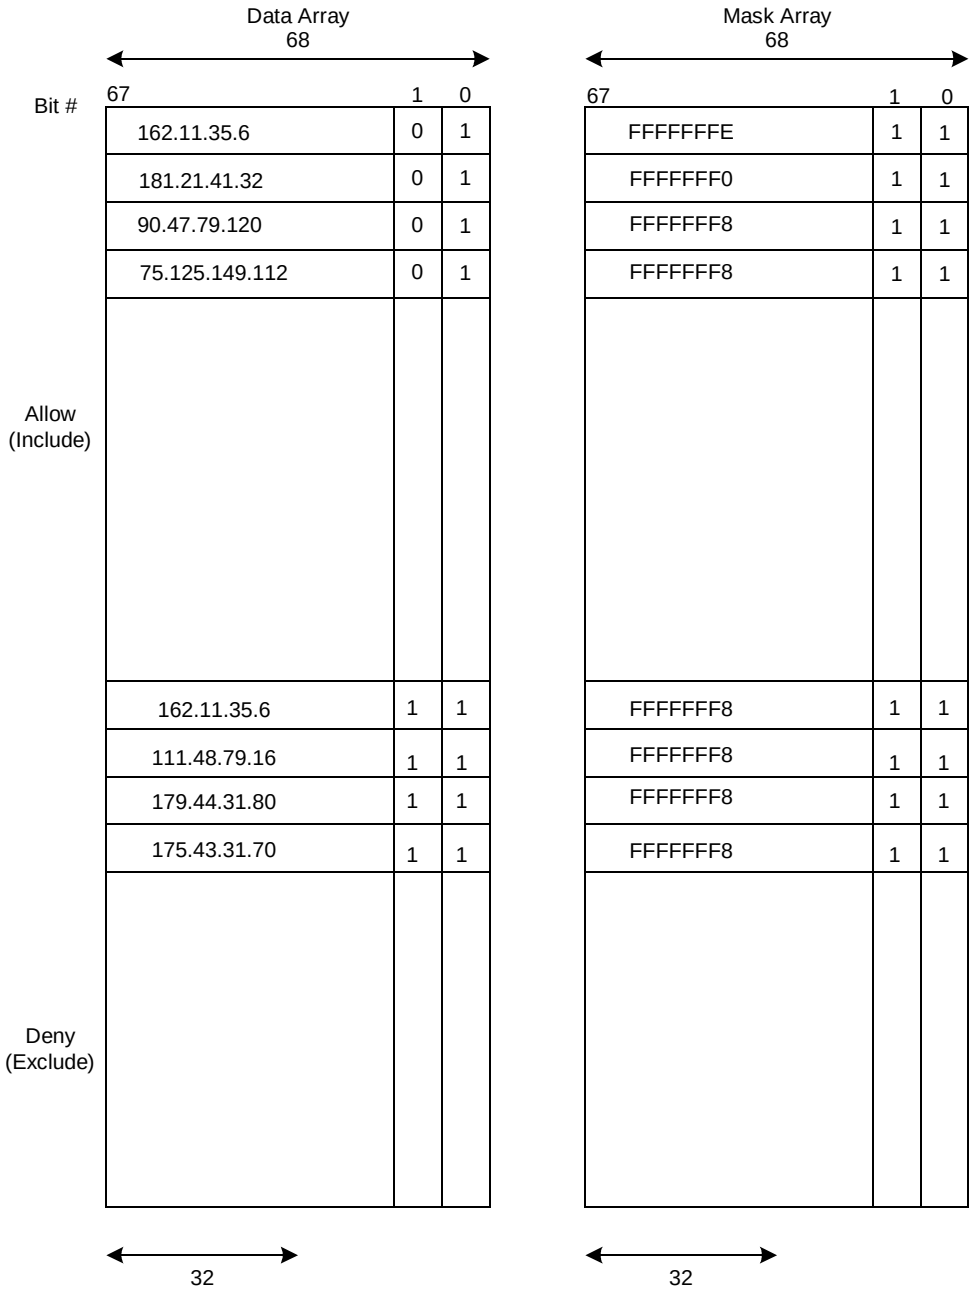

FIGURE 4. ALLOW AND DENY TABLES (AN EXAMPLE)

The need for filtering increases as we move higher up the network (OSI) layers. The proliferation of Internet services like QoS, CoS, and VPN are increasing demands that cannot be met by the traditional software-based algorithms alone. Applying various algorithms for IP address filtering requires a considerable amount of processing time. The algorithms used currently add higher processing costs and have increased latency, unlike Lara's LNI7010/LNI7020 devices, which offer IP filtering at 83 million times per second.

Performing IP filtering operations via the LNI7010/LNI7020 NSEs provides higher speed and performance over legacy IP filtering approaches using algorithms. The LNI70X0 NSEs also offer a cost-effective alternative for improving the device performance.