



AN1337

APPLICATION NOTE

IP Address Filtering Using the M7010 and M7020 Network Search Engine

INTRODUCTION

This Application Note illustrates how IP address filtering can be effectively performed using ST Microelectronics M7010/20 Search Engines.

SUMMARY

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

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

For this reason, we need to take the protection of information 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 institution. Also, compliance to "policies" set for managing the access to IT resources is necessary.
- Ensuring that the integrity of information is maintained and that unauthorized changes are not made.

IP filtering, which is still the predominant technology, can be used to control access to Web resource. Each computer connected to the Internet has a unique address. These addresses are arranged in a hierarchy of domains, sub-domains and machine numbers. IP filtering works by restricting access to machines in a particular domain, sub-domain or even specific machine addresses.

With IP filtering, the licensee institution guarantees to 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 licensee institution's user community. The resource operator then simply checks the source IP address of each incoming request.

IP filtering complements and leverages the capability of VPN, Policy Management, CoS, QoS, VoIP and other applications.

IP ADDRESS BASICS

IP addressing scheme is an integral part in the process of routing IP data through the Internet. Each host on a 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. It is assigned by the Internet Network Information Center if the network is to be part of the Internet. The host number identifies a host and is assigned by the local network administrator.

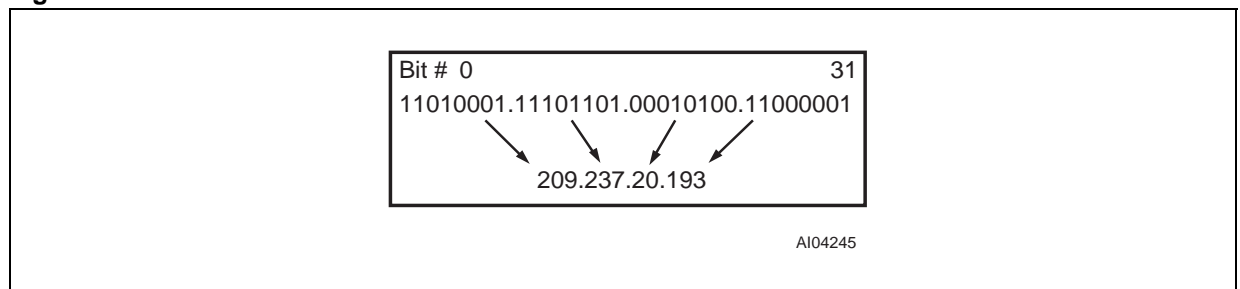
IP Address Format

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 8-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 8 bits at a time; each group of 8 bits being an octet. Each of the four octets is separated by a dot, and represented in a decimal format. This is known as dotted decimal notation. 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).

As an 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 (see Figure 1).

Figure 1. IP Address



IP Address Classes. IP addressing supports the following five address classes: Class A, Class B, Class C, Class D, and Class E.

In a class A address, the first octet is the network portion. Thus, the class A address of 209.237.20.193 has a major network address of 209. Octets 2, 3, and 4 represent the hosts. Class A addresses are used for networks that have more than 65,536 hosts. The Class A address of 127 is for a special function called a “loopback function.”

In a class B address, the first two octets identify the network portion. Therefore the class B address of 135.11.21.7 has a major network address of 135.11. Octets 3 and 4 (the next 16 bits) are for 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 of 205.45.9.37 has a major network address of 205.45.9. Octet 4 is for hosts. Class C addresses are used for networks with less than 254 hosts (see Table 1).

Class D addresses are known as multicast addresses and are used to send an IP datagram to a group of hosts on a network.

Class E addresses are reserved for experimental and future use.

Table 1. Address Classes

Address Class	First Octal (Decimal)	High-order Bits	IP Address Example
Class A	1 - 126	0	111.49.79.16
Class B	128 - 191	10	128.11.21.7
Class C	192 - 223	110	209.237.20.193
Class D	224 - 239	1110	230.100.80.0
Class E	240 - 247	11110	245.16.92.7

AN1337 - APPLICATION NOTE

Class Full Network Masks. Each of the address classes contains a set of class full network masks. The network mask defines which bits out of the 32 bit of the address are defined as the network portion and which is the host portion. The network mask is calculated by setting all bits to a value of 1 in the octets designated for the network portion and all bits to a value of 0 in the octets designated for the host portion. 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.

Table 2 summarizes the network and host portion of each address class:

Table 2. Address Class Network and Host Portions

Bit #	0	1		7	8					31		
	0	Network #				Host #						Class A
Bit #	0	2				15	16				31	
	10	Network #				Host #						Class B
Bit #	0	3						23	24		31	
	110	Network #				Host #						Class C

A subnet is an identifiably separate part of an organization's network. Typically, a subnet may represent all the machines at one geographic location, in one building, or on the same local area network. Having an organization's network divided into subnets allows it to be connected to the Internet with a single shared network address. Without subnets, 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.

Subnet Mask. 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 subnet number. The router knows which bits to consider by looking at a subnet mask. Using a mask saves the router having to handle the entire 32-bit address; it can simply look at the bits selected by the mask.

IP Subnet Addressing. All Classes of IP networks can be divided into smaller networks called subnetworks (or subnets). Dividing the major class network is called subnetting. Subnetting provides network administrators with several benefits. It provides extra flexibility, makes more efficient use of network address utilization, and contains broadcast traffic because a broadcast does not cross a router.

IP Subnet Mask. "Borrowing" bits from the host field and designating them as the subnet field creates a subnet address. The number of borrowed bits is variable and specified by the subnet mask.

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.


Subnetting addresses the problem of expanding the routing table. It also ensure that the subnet structure of a network is not visible outside of the organization's private network.

M7010/20 SEARCH ENGINES

The M7010 and M7020 are high-performance, pipelined, synchronous, Search Engines designed using the Associated Processing Technology (APT™). It is user configurable into tables as wide as 272-bits and cascaded depth of 992K 32-bit addresses. Its high speed of 83 million lookups per second and high capacity of incorporating nearly 1 million addresses can be employed in a variety of networking and communications applications requiring fast searches of various tables. It provides 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 orders-of-magnitude reduction in the search time.

The M7010 is organized as 8K x 136-bit, but can also be configured as 4K x 272-bit, 16K x 68-bit, or 32K x 34-bit (see Application Note, “AN1339 - 32-bit Applications Using the M7010, M7020 Network Search Engines”). M7020 is organized as 16K x 136-bit, but can also be configured as 8K x 272-bit, 32K x 68-bit, or 64K x 34-bit (see Application Note, “AN1339 - 32-bit Applications Using the M7010, M7020 Network Search Engines”). The M7010 can sustain 83 million searches per second on any sub-field of a 68-bit or 136-bit field, making it the fastest Search Engine 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 2. Variable Width Table Configuration



The M7010/20 contains a mask register for each data location. In addition, the device contains 16 68-bit global mask registers that can be dynamically selected in every SEARCH operation to select the search subfield. These mask registers provide an easy way for moving data to masks and enable selective lookups for subnets.

IP ADDRESS FILTERING

In Table 3, eight different networks have been arbitrarily selected.

- Some of them will be allowed access to the network, hence they will be entered in the data array in the *Include Table*.

Table 3. IP Address Filtering Example

	Network Address	Mask	# of Mask Bits	# of Hosts
I	162.11.35.64	FFFFFFE0	27	32
E	181.22.14.124	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	FFFFFFFC	30	4
E	175.43.31.70	FFFFFFFC	30	4

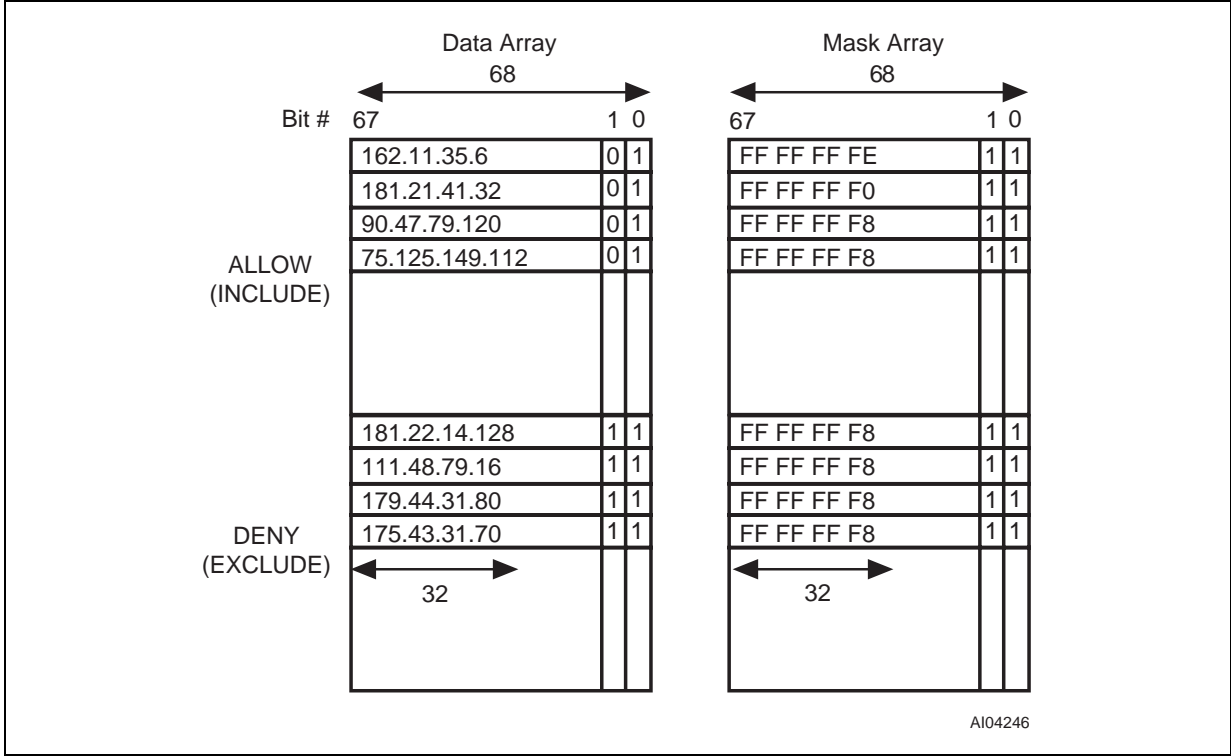
- “I” indicates “Include” and “E” indicates “Exclude.”
- Those IP addresses that are not permitted access to the network will be entered in the *Exclude Table* in the data array.

PROCEDURE FOR UPDATING TABLES

Two array segments have been created (see Figure 3). One part 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.

Upon initialization, the Search Engine 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; then, the table can be updated using the Learn command, if the entry is not found in the *Deny Table*.

Figure 3. Allow and Deny Tables for Updating Tables



Bit 1 is a table management bit. Bit 1 is "0" in the Include Table and "1" in the *Exclude Table*. Please note that in this example the table is 68-bit wide.

If any IP address is encountered that is not a part of either the Allow or Deny tables, then a new entry of this 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.

CONCLUSION

The need for filtering increases as we move higher up the network (OSI) layers. The proliferation of Internet services like: QoS, CoS, VPN are increasing demands, which 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 the M7010/20 devices, which offer IP filtering at 83 million times per second.

Performing IP filtering operation via the M7010/20 Search Engines provides higher speed and performance over legacy IP filtering approaches using algorithms. The M70X0 Search Engines offer a cost-effective alternative for improving the device performance.

AN1337 - APPLICATION NOTE

CONTACT INFORMATION

If you have any questions or suggestions concerning the matters raised in this document, please send them to the following electronic mail addresses:

apps.nvram@st.com (for application support)

ask.memory@st.com (for general inquiries)

Please remember to include your name, company, location, telephone number, and fax number.

Information furnished is believed to be accurate and reliable. However, STMicroelectronics assumes no responsibility for the consequences of use of such information nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of STMicroelectronics. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. STMicroelectronics products are not authorized for use as critical components in life support devices or systems without express written approval of STMicroelectronics.

The ST logo is registered trademark of STMicroelectronics
All other names are the property of their respective owners.

© 2001 STMicroelectronics - All Rights Reserved

STMicroelectronics GROUP OF COMPANIES
Australia - Brazil - China - Finland - France - Germany - Hong Kong - India - Italy - Japan - Malaysia - Malta - Morocco -
Singapore - Spain - Sweden - Switzerland - United Kingdom - U.S.A.

www.st.com