Application Note 008: Implementing an Access List Coprocessor with NDSE Technology for Multigigabit Network Devices

Marketing

Cypress Semiconductor Corporation

San Jose, CA 95134

July 10, 2001

1. ACCESS LIST COPROCESSOR

Simple router and switch designs use Layer 2 and Layer 3 information to redirect network packets between the physical ports. A necessary addition to this function is to control the flow of information between networks attached to the router / switch. In particular it is often desirable to restrict certain protocols to only be transmitted by some ports, or limited firewall functions to be provided where a port is connected to a public network. In router / switch designs, this is implemented by creating an access list which defines the protocols, addresses, and services that are allowed or disallowed. These access lists are then applied to a specific port or set of ports. Every packet transmitted through a port then needs to be interrogated to determine if the packet should be allowed or blocked. As networks evolve, access lists can become complex with the rules processing creating a significant overhead to the process of packet delivery. This is particularly important when considering network devices which support multi-gigabit ports, where latency and throughput can be compromised by any non-deterministic delays in packet processing. The function of resolving the “Access List” can be off-loaded to an “Access List Coprocessor” built using “Network Search Engine Technology”, thus enabling wire speed routing and switching even where complex access lists are deployed..

The applications of the “Access List Coprocessor” are not limited to switch and router designs. Any networking device that requires multi-protocol access control (such as a firewall) could use the processor in a similar design to that described here.

Current access list designs process the table sequentially, that is, they compare the parameters with the first record in the list and then the next until a match is found. The allow or deny bit of the match is the queried to see if the packet is to be routed or not. As all matching may also be masked, there may be more than one record in the list that will be a satisfactory match, however it is the result of the first match which is used. All access list have an implicit ‘deny all’ entry on the end of the list which means that if a match is not found, then the packet will be denied.

2. IMPLEMENTATION

Architecture. To implement an ‘Access List Coprocessor’, Network Search Engine technology (such as the LNI7020 product from Lara Networks, Inc.) may be deployed, either in conjunction with a Network Processor, or with a standard RISC microprocessor. A typical system architecture is shown in Figure 1.

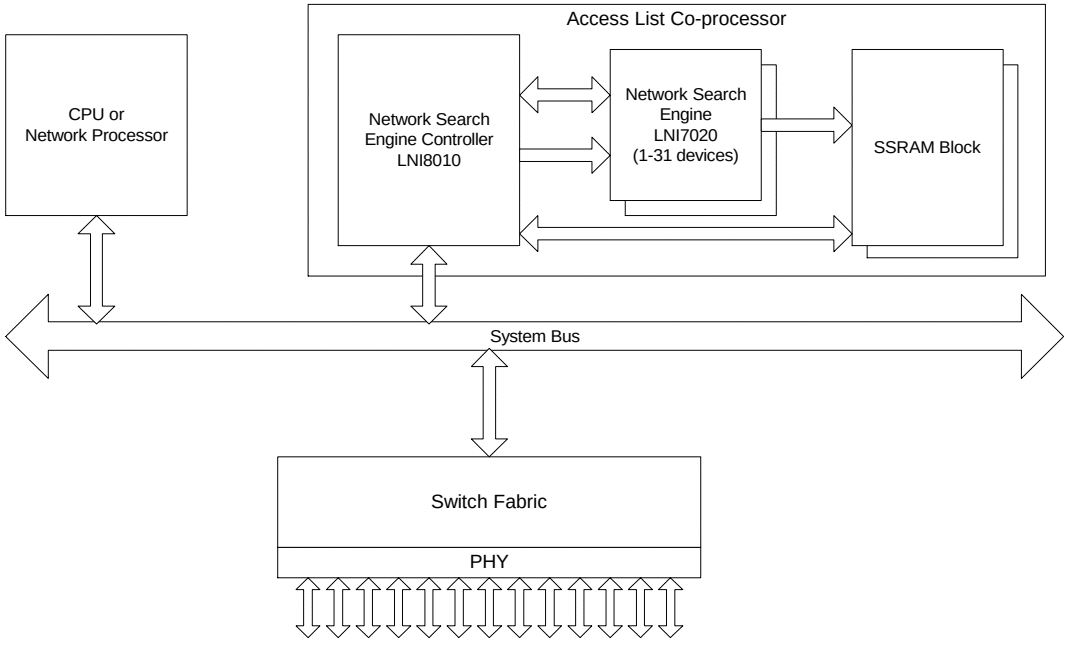

Figure 1

3. TABLE ORGANIZATION

3.1. Cisco IOS

In the Cisco IOS[i] implementation of access control lists, a number of lists can be configured for each protocol. These lists are numbered based on the protocol, i.e. for IPv4 numbers 1-199 may be used. These access lists can then be applied to specified ports, such that an incoming packet first has the protocol identified, and then is checked against all of the access lists that have been applied to that port for the protocol. In practice it is unusual for more than one access list to be applied to a port for any given protocol. This means that we need to implement multiple tables within the 'Network Search Engine'. Each LNI7020 device can segregated into four partitions, each of which can be configured to 34, 68, 136, or 272 bits wide. Multiple tables can be held within each partition, provided they are the same width, simply by using a few of the spare data bits to identify the table number. A 16-bit table identifier is used in this design. The structure of the records in each table will depend on the protocol and if the table is provided basic or extended control. Examples of how lists could be structured are shown below:

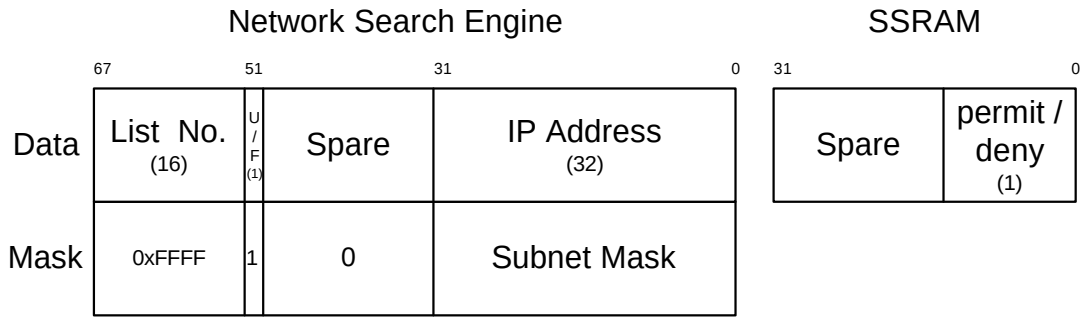

Figure 2. Basic IP

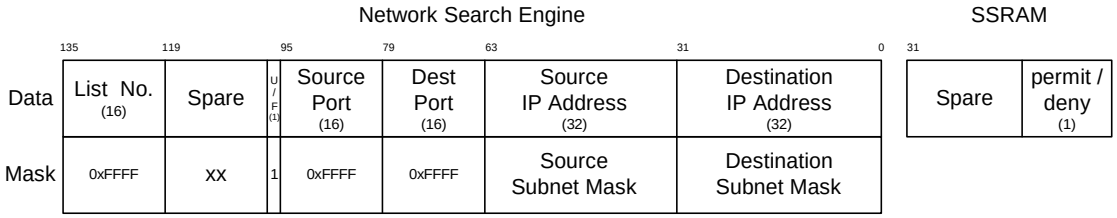

Figure 3. Extended IP

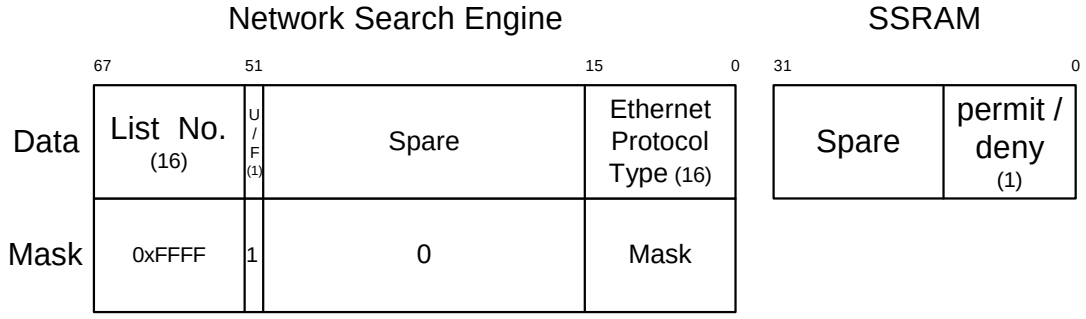

Figure 4. Ethernet Protocol Type

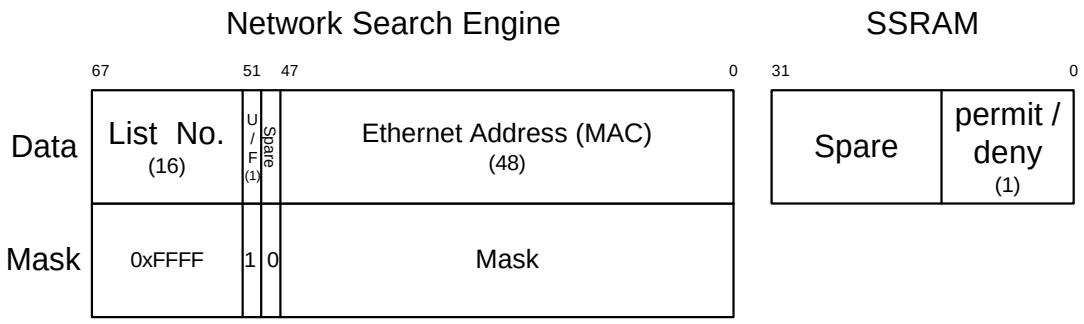

Figure 5. Ethernet Address

The U/F bit in the NSE is used to indicate if the entry is used or free. Only one bit is required in the SSRAM block to indicate the Permit / Deny result. The additional bits could be used to extend the functionality of the access list co-processor to CoS and ToS applications.

Additional protocols could be supported by using the same record structure. The “List No” and “Used / Free” fields must be retained in the same bit positions, but the remainder of the record format is determined by the protocol.

3.2. Table Management

Access lists are processed in order so that the first match found is used. When implementing the access list in an NSE, we need to also add the entries to the table in order so that the first match will be returned. New entries can only be added to the end of the table. If it is necessary to insert a record at a position in the table, then the access list must be deleted and re-entered with the new entry in the correct order. It is not normal for access lists to change dynamically, so this does not cause a problem. To support the differing word widths required for the different type of access lists, the NSE will be partitioned into tables of 68 and 136 bits wide. Initially, each table size will be allocated a single block in one of the LNI7020 devices. This will allow 4K, 136 bit entries and 8K, 68 bit entries before the tables will need to expand. When either of the tables requires more space, an additional block will be configured to the correct size and allocated to that table.

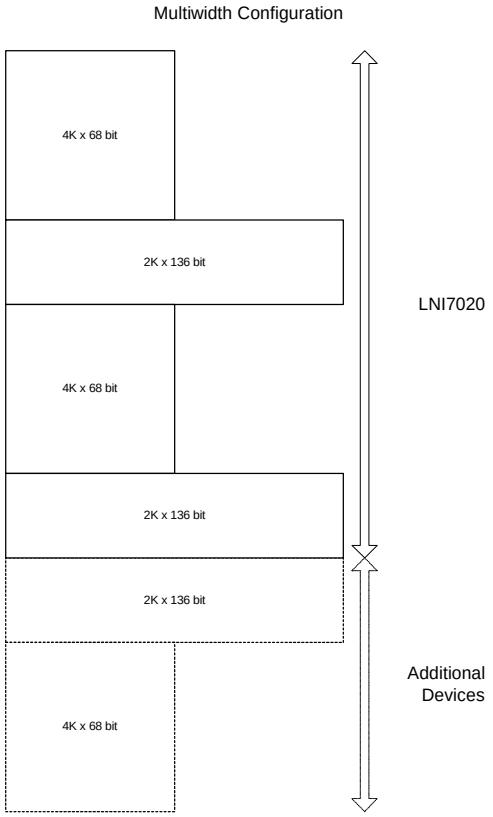

Figure 6

Each of the tables will hold multiple access lists which will be added as required. When resolving against an access list, the access list number will be used and this will ensure that only records from the required access list are compared. It is not necessary that an access list be added to the table fully before another list may be added. The records will simply be interleaved and as the access list number is used in all table queries, this will not affect the result. An example is shown in Table 1. Records shown in gray will be masked when a query is made using Access List number 1.

Table 1

Record	List No	Address	Mask	Permit / Deny
1	1	192.168.3.0	255.255.255.0	D
2	1	192.168.4.0	255.255.255.0	P
3	2	192.168.5.0	255.255.255.0	P
4	1	192.168.5.0	255.255.255.0	D
5	2	192.168.3.0	255.255.255.0	P
…	…	…	…	…

When an access list is deleted, the access list number portion of the record should be written with the value of 65535. This is an invalid number and will prevent the values from being used in future searches.

Periodically, after a large degree of update, the access list table should be dumped out, cleared and re-loaded to recover the entry positions taken up by access lists that have been deleted.

3.3. Alternative Organization

The Cisco IOS implementation of access control lists is designed specifically for use within routers and switches. Vendors may choose to implement their own access list layouts without using the IOS specification. In particular, an “Access List Coprocessor” used in the implementation of a high speed IP firewall product would only need to support IP, and would need to provide greater flexibility and improved management than the IOS specification allows.

3.4. Performance

In order to determine the number of concurrent ports that the “Access List Co-Processor” will support, it is necessary to look at the query throughput that can be achieved by the design. The LNI8010 Network Coprocessor may be connected to a 64 bit, 100MHz system bus. Three bus cycles are required to perform a 68 bit search as is used in the Basic IP and Ethernet tables, and four cycles are required to perform the 136 bit search used in the Extended IP table.

Table 2

Operation	Basic IP / Ethernet	Extended IP
Write Search Word	10ns	20ns
Write Command	10ns	10ns
Read Result Word	10ns	10ns
Total response time	**30ns**	**40ns**

The Permit/ Deny decision is performed for each packet that is transmitted through the device and so the time available to evaluate the parameters is dependant on the size of the packet. Average packet sizes are determined by the mix of applications that are supported by the network. The table shown in Table 3 details packet sizes for a number of applications and therefore the number of ports that this “Access List Coprocessor” could support at each data rate.

Table 3

	Basic IP / Ethernet			Extended IP		
		# 2.4 Gbps (OC-48) ports	# 9.6 Gbps (OC-192) ports		# 2.4 Gbps (OC-48) ports	# 9.6 Gbps (OC-192) ports
64 bytes (i.e. VoIP)	16	8	1	12	6	1
128 bytes	32	16	3	24	12	2
512 bytes (i.e. www)	128	64	13	96	48	10
1K bytes (i.e. FTP)	256	128	27	192	96	21

In practical applications, network devices deliver a range of different application services and thus a range of different packet sizes. The average achieved will be different for each network depending on the profile. The above table shows the number of ports that could be implemented without requiring packet buffering and if the data rate is sustained. In a practical firewall design, this technology should be able to support a general purpose IP gateway with up to four OC-192 ports. Or alternatively a single through port firewall running a sustained data rate of up to OC-768.

4. CONCLUSION

Access Lists are a key component of most network devices and provide the framework for preventing un-authorized access to network systems as well as performance features such as load balancing. As port data rates increase to Gigabit Ethernet and up to OC-192, conventional software technique for resolving access control will not be able to sustain wire speed transport, particular for small packet technologies such as VoIP. Additionally, in security applications, the access lists are becoming more complex requiring extensive lists to be evaluated to resolve access control. The addition of an “Access List Coprocessor” can provide a mechanism to offload the required processing, thus making the central CPU or Network Processor available to control the overall process. The “Access List Coprocessor” also provides significant performance enhancement thus allowing devices to support either higher data rates up to OC-768, or a greater number of ports at lower data rates.

CONTACT

Lara Networks, Inc.

110 Nortech Parkway

San Jose, CA 95134

www.laranetworks.com

Tel: (408) 519-6300

Fax: (408) 519-6399

---

[i] Cisco Systems, Inc. – IOS Software Specification.