algotronix
130-10 Calton Road
Edinburgh EH8 8JQ
United Kingdom
Phone: +44 131 556 9242
E-mail: cores@algotronix.com
URL: www.algotronix.com

IPSEC Core for Secure Communications

Algorithm	IPSEC
Test Method	Algotronix provided testbench
Performance/Area	Tradeoff via compilation options.
Deliverables	VHDL with Testbench. Targets all modern Xilinx FPGA families

Features

- Implements RFC 4303 (IP Encapsulating Security Protocol (ESP)) and RFC4106 (IPSEC with AES-GCM)
- Initial product offers 1Gbit/Sec and will scale to 10Gbit/sec. Higher performance is possible using our 100Gbit/sec AES-GCM core
- Based on the Algotronix AES-GCM-10G product
- Supports 128 and 256 bit keys
- Targets all modern FPGA families from Xilinx
- Supplied as easily customizable portable VHDL to allow customers to conduct their own code review in high-security applications
- Supplied with comprehensive test bench containing a behavioral model of IPSEC

Potential Applications

- Government/Military secure communications
- Secure communications for Internet of Things (IoT)
- Bridges and routers
- IPSEC Gateways
- Test equipment

General Description

Internet Protocol (IP) networking is the foundation of the internet and IP Security (IPSEC) is the standard way to provide security on IP networks. The Encapsulating Security Protocol (ESP) encapsulates IP packets within a security envelope providing both authentication and privacy for communications across untrusted public networks.

Where MACSEC offers security on ethernet networks normally with a relatively simple one-to-many topology where an endpoint has a single security association for transmit but receives using many security associations IPSEC offers security across the internet in a many-to-many environment where an endpoint maintains multiple security associations with different keys on both transmit and receive.

The IPSEC core is a high performance pipelined implementation of the Encapsulating Security Protocol mode of IPSEC: those components of the standard which need to operate at line rates are implemented in hardware. Elements such as key exchange which occur relatively infrequently are better implemented in software. The core is built on Algotronix' pipelined implementation of the AES-GCM encryption algorithm which itself builds on our G3 AES core. This release of the IPSEC core supports operation at 1Gbit/sec and a future release will operate at

10Gbit/sec. The core is intended to eventually scale up to data rates of 100GBit/sec on FPGAs using our 100Gbit AES-GCM implementation which uses parallel compute units as well as pipelining within the encryption circuitry.

The Algotronix IPSEC core is supplied with a VHDL testbench which generates a sequence of test packets and compares the responses of the IP core to the output generated by a behavioral model of IPSEC. It is supplied as VHDL source code and can be configured using a number of VHDL generic parameters to select only those features which are required in order to conserve area. The IPSEC core provides both transmit and receive channels. The core is an easy to use fully synchronous design with a single clock and separate flow control on the transmit and receive channels. The core has been designed for efficiency in modern FPGAs and makes full use of FPGA specific features such as dual port memory blocks.

Performance and Area

The core provides many options to allow the user to trade off area against throughput and latency and the FPGA architecture and speed grade also has a strong bearing on the results achieved. The block RAM usage within the core depends strongly on the number of entries required in the IPSEC SAD and SPD databases and whether the AES-GCM engine is configured to implement SBoxes using block RAM or LUTs. In addition, both the core itself and the FPGA manufacturer's design tools are regularly updated resulting in area and maximum clock frequency results rapidly becoming out of date. A clock frequency of 125MHz is needed for 1Gbit/sec performance and this is achievable on lower cost modern FPGA families such as Zynq.

For these reasons, rather than provide area and performance information in the data sheet Algotronix prefers to generate these estimates on demand for our customers. Contact us with the desired throughput, target FPGA family and speed grade, and we will work out the best core configuration and use the latest FPGA design tools and core source code to provide estimates of the attainable clock frequency and area in the required configuration.

Functional Description

This implementation of IPSEC is designed to offer 1 Gbit and 10 Gbit throughput on high performance FPGA devices.

There are two main functional blocks in the IPSEC design: the receive and transmit data paths. Both blocks provide a variety of memory mapped resources such as tables and content addressable memories for the security databases required by IPSEC. These resources are accessed through a processor interface with a 32 bit data bus.

The Receive path takes IPSEC encapsulated packets, captures the header to determine the security association information and decrypts and authenticates the packet using AES-GCM. It then removes the IPSEC encapsulation and provides a secure copy of the original non-IPSEC packet to the system. The receive operation reduces the size of incoming packets, therefore the bandwidth requirement on the output side is lower than that on the input side.

The transmit data path adds IPSEC encapsulation to an incoming packet, encrypts it using AES-GCM, which adds an Integrity Check Value (ICV), and outputs the data (Figure 2). The transmit operation increases the size of incoming packets, therefore the output side bandwidth requirement is higher than that on the input side. The throughput number specified for the IPSEC core refers to the network side of the system which deals with encapsulated packets. For example, a 10Gbit/sec IPSEC unit will not be able to cope with a sustained 10Gbit/sec stream of minimum sized un-encapsulated packets on the system input since this would require more than 10Gbit/sec on the network interface.

Within the receive data path the core inspects the Destination Address within the packet to determine whether to route it to the bypass port or process it as an IPSEC ESP packet. If it is determined to be an ESP packet the Security Parameter Index (SPI) field is looked up to find the corresponding entry in the Security Association Database which will provide the key and other parameters required for decryption.

Within the transmit datapath the packet source and destination address are looked up in the Security Policy Database (SPD) CAM in order to find the SPI which should be used. The SPI is then looked up in the SAD memory to find the key and other parameters required to encrypt the packet.

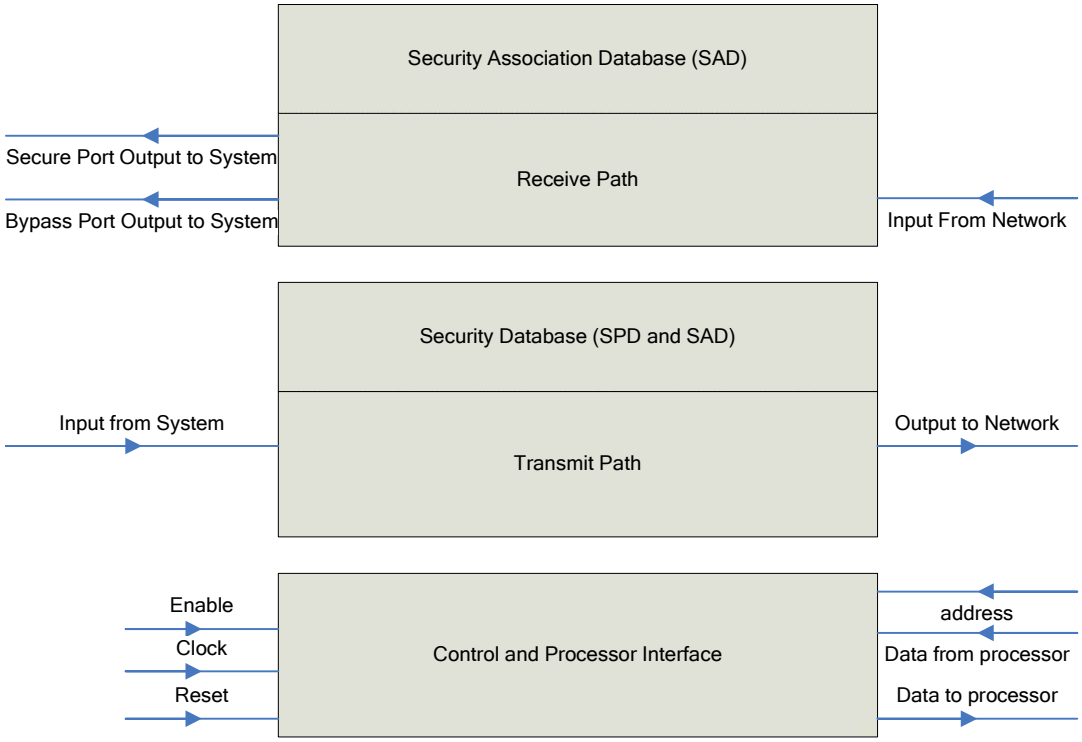

Figure 1, IPSEC Block Diagram

IPSEC Memory Map

Address	Resource
0x0000 - 0x00FF	Receive SAD SPI CAM (1 word per SAD entry)
0x0400 - 0x04FF	Receive SAD Flags Memory (1 word per SAD entry)
0x0800 - 0x0BFF	Receive SAD Data Memory (4 words per SAD entry)
0x1000 - 0x17FF	Receive SAD Key Memory (8 words per SAD entry)
0x2000 - 0x2FFF	Transmit SAD Memory
0x3000 - 0x37FD	Transmit SPD
0x3800	Receive Secure Packet Destination Address Register
0x3801	Receive Secure Packet Destination Address Wildcard Register
0x3802	IPSEC ID Register

Global Resources

The IPSEC ID Register is a read-only register with a constant value of X"49505365" ("IPSe" in ASCII). This register is useful for initial debug of a system using the IPSEC core since finding this value at the expected address indicates the interface between the core and the host processor is operating correctly.

Receive Channel Resources

Apart from two registers the resources in the receive channel are associated with the Security Association Database (SAD). The addressing scheme provides for an 8 bit field to specify the SAD entry index, thus a maximum of 256 entries are supported in the SAD database. If more entries are required changes to the core source code to provide a wider address bus would be required and this can be done on request. A user modifiable compilation parameter specifies how many of the 256 addressable entries should actually be implemented. This allows users to set a lower value if they do not need the full 256 entries so that FPGA resources are used efficiently.

Receive Secure Packet Destination Address Register

Packets to be processed by IPSEC entering the receive channel are identified by their Destination Address matching the value in the Secure Packet Destination Address register and a protocol type of Encapsulating Security Protocol (decimal 50). Packets which do not meet these criteria are forwarded immediately to the Bypass port rather than passing through IPSEC processing and emerging on the Secure Port.

Receive Secure Packet Destination Address Wildcard Register

The Secure Packet Destination Address Wildcard Register specifies bit positions in the Destination Address which will be regarded as always matching independent of their value allowing wildcard ('*') subfields in the destination address for flexibility. A '1' in a bit position of the Destination Address Wildcard register results in that bit in the Destination Address being regarded as always matching.

Receive SAD SPI CAM

The Receive SAD SPI CAM is a content addressable memory which contains the SPI value for the corresponding SAD entry. The SPI field in incoming ESP protected packets whose destination address matches the values specified by the Secure Packet Destination Address Register (taking account of the Secure Packet Destination Address Wildcard) is extracted and matched against these values to determine which SAD entry should be used in decoding

the protected packet. If no matching entry for the packet SPI is found the no_valid_security_association status indicator will be asserted and the packet dropped.

Receive SAD Flags Memory

SAD Flags Bit Field (4 : 0)	Function
(1 : 0)	Key Length ("10" - 256 bits, "00" - 128 bits)
2	Use Extended Sequence Number
3	Enable Anti-Replay
4	In Use

The Receive SAD Flags memory is a 5 bit wide memory providing information required to decode packets. When an incoming packet's SAD Index is determined by the SAD SPI CAM the flag values are looked up in this memory. If the in_use flag is false the CAM match is invalid and the no_valid_security_association status indicator will be asserted and the packet dropped.

The Key Length field uses the same two bit codes as Algotronix' AES products. In theory a single bit would suffice for IPSEC since only 128 and 256 bit keys are supported but the two-bit encoding also provides for 192 bit keys which are allowed in some variants of AES. The core behavior for the illegal key length values "10" (which would indicate 192 bit keys) and "11" is not specified.

The Use Extended Sequence Number bit specifies that the SA should use 64 bit extended sequence numbers rather than 32 bit Sequence Numbers. The packet itself only contains the least significant 32 bits of an extended sequence number, the most significant bits being inferred using the Top of Window register in the replay detect logic.

The Enable Anti Replay bit specifies that packets should be dropped with an Anti-Replay Detect status indication when the anti-replay detection logic deduces that the packet has already been seen based on the replay detect bitmap or is outside the allowable range based on the Top of Window Register. The anti replay logic updates the top of window and bitmap registers regardless of whether or not the Enable Anti-Replay bit is set (the Top of Window register is required for determining the Extended Sequence Number as well as replay detection): only the replay detect indication and packet drop signals on packet egress are controlled by the bit. If a packet is marked to be dropped as a result of a bad header, invalid security association or integrity check failure it will be ignored by the anti-replay logic and will not update the Top of Window or Bitmap registers or cause an Anti-Replay status indication.

Receive SAD Data Memory

Address (1 : 0)	Resource
"00"	Top of Window (63 : 32)
"01"	Top of Window (31 : 0)
"10"	Salt
"11"	Window Bit Map

The Receive SAD Data Memory has four 32 bit words for each SAD entry. The words within the entry are selected using bits (1:0) of the address bus and the entry index is provided on bits (9 : 2).

The Top of Window register is a 64 bit register containing the highest Extended Sequence Number of packets received so far on this Security Association. It is an approximation for the value of the sequence number counter in

the transmit channel of the IPSEC implementation on the other end of the link. The most significant word (bits (63:32)) is only significant if the Use Extended Sequence Number flag for the SA is true. The Top of Window register is used to determine the most significant word of the Extended Sequence Number for received packets (since this is not transmitted with the packet) using an algorithm specified in RFC 4303. It is also used in the replay detect calculation.

The Salt is a 32 bit value used in the construction of the Nonce for AES-GCM decryption of the incoming packet as specified in RFC 4106. Note that there is potential for confusion in terminology since the RFC 4106 IPSEC with AES-GCM standard uses the term Nonce to indicate the quantity the AES-GCM standard (NIST SP800-38D) calls Initialization Vector (IV) and the RFC 4106 uses the term Initialization Vector (IV) to indicate the non-Salt part of the Nonce.

The Window Bit Map tracks whether the packet corresponding to the 32 sequence numbers up to the sequence number in the Top of Window register have been received. A '1' in the register indicates a packet with the corresponding sequence number has already been received. If a packet arrives where the corresponding entry in the bit map is '1' then the packet is a duplicate and (if the replay_detect_enable flag for the SA is set) it will be dropped with a replay detect status indication. Received packets which are to be dropped for other reasons such as a bad packet header or integrity check failure are not considered when updating the Top of Window or Window Bit Map registers and will not be given a replay detect status indication even if their sequence number is a duplicate.

Receive SAD Key Memory

The Receive SAD Key Memory has eight 32 bit words for each SAD entry to allow for 256 bit AES keys. The words within the entry are selected using bits (2:0) of the address bus and the entry index is provided on bits (10 : 3).

The AES standards regard the most significant word of the key (i.e. bits (255: 224) for a 256 bit key or bits (127: 96) for a 128 bit key) as word 0 because normally the most significant bits will be loaded first into the encryption engine. The key is stored with the more significant words in the lower address locations (i.e. address(2:0) is word 0 or bits (255:224) for a 256 bit key). For a 128 bit key only address locations 0 to 3 are used.

Transmit Channel Resources and Detailed Functional Description

The transmit channel contains two major tables : the SPD is a content addressable memory used to lookup Source and Destination Address information from the packet to determine the index of the entry in the SAD which should be used to process that packet. The SAD Index (SADI) stored in the SPD CAM provides the location within the within the SAD containing the information needed to create the ESP envelope and encrypt the data.

Transmit SPD CAM Memory

Address (2 : 0)	Resource
"000"	Destination Address
"001"	Destination Address Wildcard
"010"	Source Address
"011"	Source Address WIldcard
"100"	In Use
"101"	Matching SAD Entry ((Bits 7:0) = SAD Index)

The SPD CAM matches the Source and Destination Address fields in the incoming packet with the values in the CAM entries. A '1' in a bit position of the Source Address Wildcard register results in that bit in the Source Address being regarded as always matching. A '1' in a bit position of the Destination Address Wildcard register results in that bit in the Destination Address being regarded as always matching.

If the In Use flag within the CAM entry is '0' the entry will not match independent of the source and destination address and wildcard registers.

If a match is found the Matching SAD Entry field is used to lookup information needed to process the packet in the SAD memory.

If multiple CAM entries match and have different values in the Matching SAD Entry field any of these values may be presented to the SAD memory. There is no guarantee of which one will be selected and this may vary from packet to packet. The expectation is that the user will configure the SPD memory to ensure a unique match.

Transmit SAD Memory

Address (3 : 0)	Resource
"0000"	Flags
"0001"	Reserved - Not currently used
"0010"	Salt
"0011"	SPI
"0100"	Extended Sequence Number (MSW)
"0101"	Extended Sequence Number (LSW)
"0110"	Tunnel Source Address
"0111"	Tunnel Destination Address
("1000" : "1111")	Key - "1000" contains Most Significant word. Same format as Receive Key.

The SAD memory provides the information required to transmit a packet using the Security Association determined from the SPD CAM lookup.

Flag bits	Function
(1:0)	Key Length ("10" - 256 bits, "00" - 128 bits)
2	Use Extended Sequence Number
3	Enable Anti-Replay
4	In Use
5	Tunnel Header DF Flag Source Select ('1' - Copy from inner header, '0' - Use Default Value).
6	Tunnel Header DF Flag Default Value
(15 : 7)	Reserved - not currently used
(31 : 16)	Tunnel Header Identification Field Counter

The key length for the transmit channel SAD is encoded in the same way as the key length in the receive SAD.

The Use Extended Sequence Number field is '1' when the transmit channel sequence number counter for outgoing packets is 64 bits wide, otherwise it is 32 bits wide. Although a 64 bit extended sequence number is maintained and used in the encryption process only the least significant 32 bits are actually transmitted with the packet. The core detects when the sequence number counter is about to overflow and all subsequent packets will be marked to be dropped until the counter is re-initialized. This is done to enforce a security requirement of the AES-GCM algorithm that the Nonce provided is never used twice with the same key. In IPSEC the nonce is derived from the sequence number. Driver software must ensure that the key is changed when resetting the sequence number counter to clear this condition.

The Enable Anti-Replay field should be '1' when anti-replay checking is enabled in the matching receive channel at the other end of the IPSEC link. When Enable Anti Replay is '0' the standard specifies the core should not check the transmit sequence number counter for overflow.

When the In Use field of the SAD is '0' the core will report that there is no matching SAD and the packet will be dropped.

The 'Do not fragment' or DF bit in the header of the outgoing ESP packet can either be copied from the header of the packet entering IPSEC transmit or a specified default value can be used.

The Identification field in the IP header of outgoing ESP packets is determined by the Tunnel Header Identification Field Counter. Normally this field would be initialized to 0 when the SAD entry is set up, the core will increment this field for each transmitted packet using this SAD entry and it will roll-over to 0 when it overflows.

Verification Methods

Verification of the AES-G3 core which forms the basis of the pipelined implementation of AES-GCM within IPSEC is through a comprehensive VHDL testbench which supports the standard AESAVS test suite with additional vectors from the SP800-38A publication to test the various AES modes. The testbench allows simulation of the design source code and also post place and route timing simulation. The testbench can be used in Regression mode to confirm the functionality of the core against known 'golden' test vectors provided by Algotronix or in Qualification mode to generate response files from vectors supplied by a NIST approved certification laboratory.

The AES-GCM core has its own testbench which runs standard vectors from the AES-GCM proposal and vectors from an Algotronix created software implementation of AES-GCM (based on the well known open-source implementation of AES by Brian Gladman).to check the functionality of the hardware.

The IPSEC testbench generates test packets programmatically for a wide variety of scenarios. A behavioral model of IPSEC is used to generate 'correct' values for output packets for comparison against the output from the hardware.

Recommended Design Experience

It is recommended that the user is familiar with the VHDL or Verilog language and with the Xilinx design flow and simulation tools. The coding standards used when creating the product allow automatic translation of the core into Verilog without loss of readability and the core can be supplied in Verilog on request.

It is recommended that the user has a background in data security or takes appropriate advice when considering how to implement IPSEC in a larger system.

Information on the IPSEC Algorithm

IPSEC Standards

IPSEC is specified by the Internet Engineering Task Force (IETF) in a series of RFC documents. RFC4303 specifies the overall architecture. RFC4301 and RFC4106 specify the performance critical parts which require implementation in hardware and are provided by the IPSEC core.

The AES and AES-GCM algorithms used in IPSEC are standardized by the Computer Security Division, National Institute of Standards and Technology (NIST), Gaithersburg MD. The relevant standard to this implementation is FIPS 197 which specifies the AES algorithm. The FIPS 197 document provides an excellent and concise description of the processing involved in implementing AES and therefore this basic information on the structure of the AES algorithm is not repeated here.

NIST Special Publication SP800-38D describes the AES-GCM algorithm, this document is derived from a proposal to NIST by Cisco..

These NIST documents can be downloaded free of charge from the NIST website ().

Customization Service

Algotronix can offer a cost effective customization service for this core in order to tune the implementation for easy integration into a larger system. It is also possible to produce variants with significantly higher performance at the expense of increased area and to create optimized variants of the core targeted at particular FPGA devices.

This cryptographic product is subject to export control. It is freely available within the European Union and can be supplied immediately to Australia, Canada, Japan, New Zealand, Norway, Switzerland and the United States under Community General Export Authorisation EU001.

Export to other countries requires an export licence. The UK Department of Business, Enterprise and Regulatory Reform publishes information on their website (www.berr.gov.uk) which gives an indication of average export licence processing times for various countries and the percentage of licence requests which are granted. For many countries obtaining an export licence can be done relatively quickly and with only a small amount of additional paperwork.

It is the the responsibility of the customer to comply with all applicable requirements with respect to re-export of products containing the AES technology.

Algotronix Ltd.
130-10 Calton Road
Edinburgh EH8 8JQ
Phone: +44 131 556 9242
E-mail: cores@algotronix.com

URL: www.algotronix.com

Version Control Information	
Subversion Revision Number	74
Date	2015/08/19 12:02:37
Document	IPSEC Product Brief
Status (blank field indicates OK/no warnings)	
	(Table auto-updates, do not edit field values by hand)