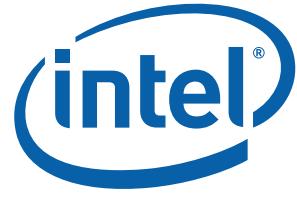


1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13	13
14	14
15	15
16	16
17	17
18	18
19	19
20	20
21	21
22	22
23	23
24	24
25	25
26	26
27	27
28	28
29	29
30	30
31	31
32	32
33	33
34	34
35	35
36	36
37	37
38	38
39	39
40	40
41	41
42	42
43	43
44	44
45	45
46	46
47	47
48	48
49	49
50	50
51	51
52	52
53	53
54	54
55	55
56	56
57	57
58	58
59	59
60	60
61	61
62	62
63	63
64	64
65	65
66	66
67	67
68	68
69	69
70	70
71	71
72	72



# Lincoln Tunnel L3 Software

## Application Note

*July 2009*

**Intel Confidential**

Reference Number: 426941-0.5

Modified on: 7/30/09

Draft Copy—Do Not Distribute



1	
2	
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	
15	
16	
17	
18	
19	
20	
21	
22	
23	
24	
25	
26	
27	
28	
29	INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.
30	
31	
32	
33	
34	Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.
35	
36	The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
37	
38	Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.
39	Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site.
40	
41	<a href="#">[When the doc contains software source code, include a copy of the software license or a hyperlink to its permanent location.]</a>
42	Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See <a href="http://www.intel.com/products/processor_number">http://www.intel.com/products/processor_number</a> for details.
43	
44	Code Names are only for use by Intel to identify products, platforms, programs, services, etc. ("products") in development by Intel that have not been made commercially available to the public, i.e., announced, launched or shipped. They are never to be used as "commercial" names for products. Also, they are not intended to function as trademarks.
45	
46	BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurstable, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel StrataFlash, Intel Vivid, Intel vPro, Intel XScale, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skoool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.
47	
48	
49	
50	*Other names and brands may be claimed as the property of others.
51	Copyright © 2009, Intel Corporation. All rights reserved.
52	
53	
54	
55	
56	
57	
58	
59	
60	
61	
62	
63	
64	
65	
66	
67	
68	
69	
70	
71	
72	



## Contents

---

1.0	Overview	5	10
1.1	Motivation	5	11
1.2	Packet Flow using Lincoln Tunnel	6	13
2.0	Flow Processing Code	7	14
2.1	Transmit Queues Example	8	15
2.2	Descriptor Bunching	8	16
3.0	Performance	8	18
A	e1000e-0.5.20 Modifications	9	19

## Figures

1	Packet Flow	6	24
2	Flow Processing Code	7	25
3	Transmit Queues Example	8	26

## Tables

30		30	28
31		31	28
32		32	28
33		33	28
34		34	28
35		35	28
36		36	28
37		37	28
38		38	28
39		39	28
40		40	28
41		41	28
42		42	28
43		43	28
44		44	28
45		45	28
46		46	28
47		47	28
48		48	28
49		49	28
50		50	28
51		51	28
52		52	28
53		53	28
54		54	28
55		55	28
56		56	28
57		57	28
58		58	28
59		59	28
60		60	28
61		61	28
62		62	28
63		63	28
64		64	28
65		65	28
66		66	28
67		67	28
68		68	28
69		69	28
70		70	28
71		71	28
72		72	28



## Revision History

Date	Revision	Description
July 2009	0.5	Initial release

6 

## 1.0 Overview

  
78 Lincoln Tunnel software modifies an existing Ethernet\* driver to implement a flow  
9 processing pipeline to efficiently classify and process network traffic at very high rates.  
10 Intel® architecture is well suited for this type of operation due to its multi-core  
11 architecture and large caches, as long as packets are delivered to the processing core  
12 efficiently.  
13  
1415 The software comprises two parts:  
1617 • **Driver Core**  
1819 The driver core contains a modified version of the e1000e driver (version 0.5.20).  
20 Modifications were made to e1000.h, netdev.c, param.c, and e1000\_82571.c files  
21 and two new files were added for flow processing (flows.c and flows.h). Refer to  
22 [Appendix A, “e1000e-0.5.20 Modifications”](#) for additional information.  
2324 • **Flows Setup**  
2526 The flows setup is a user space tool for rules configuration. It parses a policies file  
27 and builds a policy table based on the “HiCuts” algorithm. HiCuts is a decision tree-  
28 based packet classification algorithm. Refer to the “Packet classification using  
29 hierarchical intelligent cuttings” whitepaper<sup>1</sup> for additional information. The policy  
30 table is copied to kernel space for policy lookups.  
3132 Lincoln Tunnel software has the following characteristics:  
33

- 34 • Leverages existing e1000e driver:
- 
- 35     — Reuse of NIC configuration/communication code
- 
- 36     — Reuse of interrupt handling code
- 
- 37 • Use of Linux\* kernel for interrupts/NAPI and memory management:
- 
- 38     — Bypass of Linux kernel for packet processing:
- 
- 39         — Flow lookup/processing
- 
- 40         — Multiple per-CPU, per-interface lockless transmit queues
- 
- 41         — Descriptor bunching
- 
- 42

43 

## 1.1 Motivation

  
4445 One basic motivation for implementing Lincoln Tunnel software was the excessive time  
46 spent waiting for locks. When the RX interface receives a packet, the kernel locks the  
47 RX queue on the incoming interface. The kernel also locks the corresponding TX queue  
48 on the outgoing interface.  
4950 The other basic motivation was the lack of TX descriptor bunching. Without descriptor  
51 bunching, the driver updates the TX descriptor ring tail pointer for every packet, which  
52 requires a costly PCI configuration space write on each packet.  
5354 Lincoln Tunnel software solves these issues by:  
55

- 56 • Using separate TX queues
- 
- 57 • Bypass the Linux stack (complex “One size fits all” solution)
- 
- 58 • Doing RX and TX descriptor bunching
- 
- 59

60

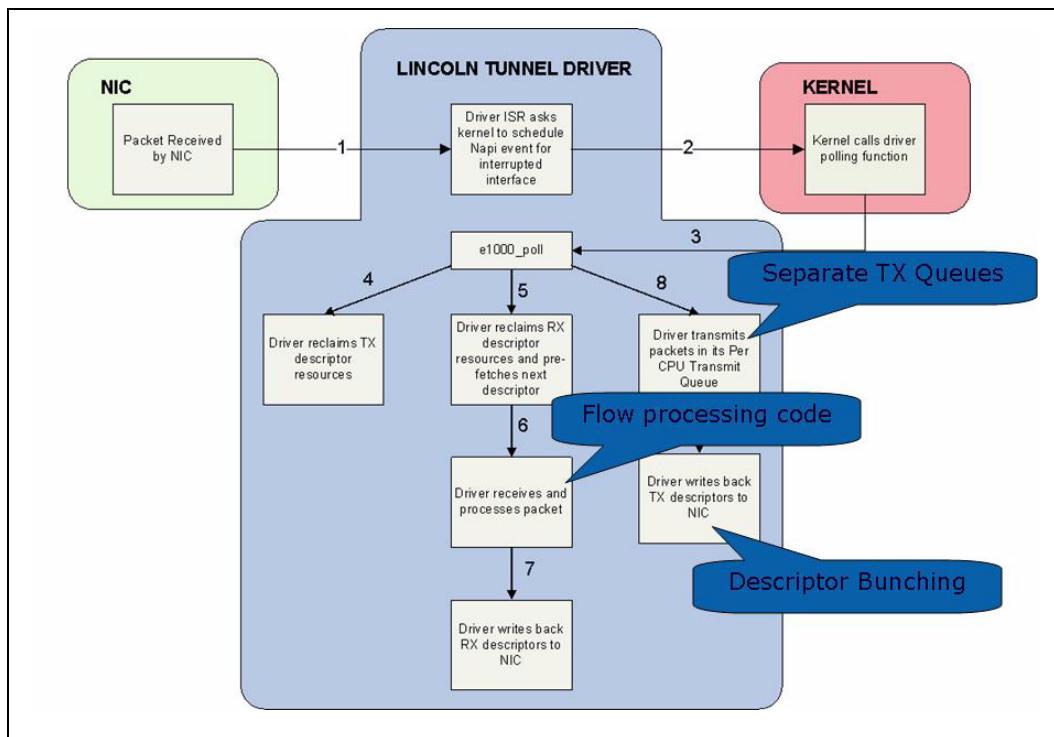
---

61 1. HiCuts: P. Gupta and N. McKeown, “Packet classification using hierarchical intelligent cuttings”, in  
62 Proc. Hot Interconnects, 1999.  
63  
64

## 1.2 Packet Flow using Lincoln Tunnel

Figure 1 shows the packet flow used by the Lincoln Tunnel software. The packet flow is described as a sequence of steps following the figure.

Figure 1. Packet Flow



### NIC

1. A packet is received by the NIC.

### Lincoln Tunnel Driver

2. The driver Interrupt Service Routine (ISR) asks the kernel to schedule the NAPI event for an interrupted interface.

### Kernel

3. The kernel calls the driver polling function.

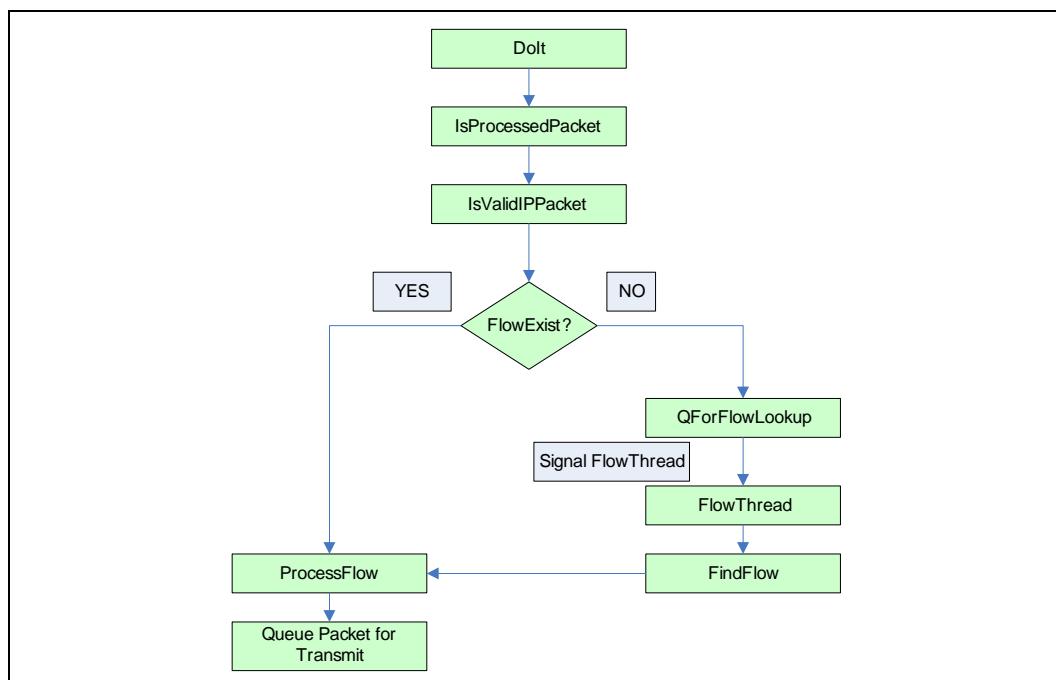
### Lincoln Tunnel Driver

4. The driver reclaims the TX descriptor resources.
5. The driver reclaims the RX descriptor resources and pre-fetches the next descriptor.
6. The driver receives and processes the packet. This is the flow processing code that is described in more detail in [Section 2.0, "Flow Processing Code" on page 7](#).
7. The driver writes back RX descriptors to the NIC.
8. The driver transmits packets in its per-CPU transmit queue.
9. The driver writes back TX descriptors to the NIC. Descriptor bunching is used to optimize packet flow. See [Section 2.2](#) for additional information.

## 2.0 Flow Processing Code

Figure 2 is a flowchart of the flow processing code. The flow processing code is described as a sequence of steps following the figure.

**Figure 2.** Flow Processing Code



1. The Dolt function is called by e1000\_receive\_skb
2. IsProcessedPacket checks to ensure the packet is a valid IP packet and extracts the 5tuple for hash computation. The 5tuple is based on: SRC/DST IP, SRC/DST Port, and protocol (TCP, UDP and so on).
3. FlowCheck first checks to see if there is a cached flow associated with the RX interface. If a cached flow is found, the flow exists. If no cached flow is found, the hash value is computed and a flow lookup in the flow table is performed. If flow is found in flow table, the flow exists. If it is not found, the flow does not exist in the flow table and it is not cached. If flow is found, proceed to ProcessFlow (Step 6.)

*Note:*

The flow does not exist in the flow table and it is not cached.

4. QForFlowLookup signals a separate kernel thread (FlowThread) to do the FlowLookup.
5. FlowThread calls FindFlow to look up the flow in the policies HiCuts tree. If found, the flow is added to the flows table and the receiving interface's cache. ProcessFlow is called.

*Note:*

The flow exists in the flow table.

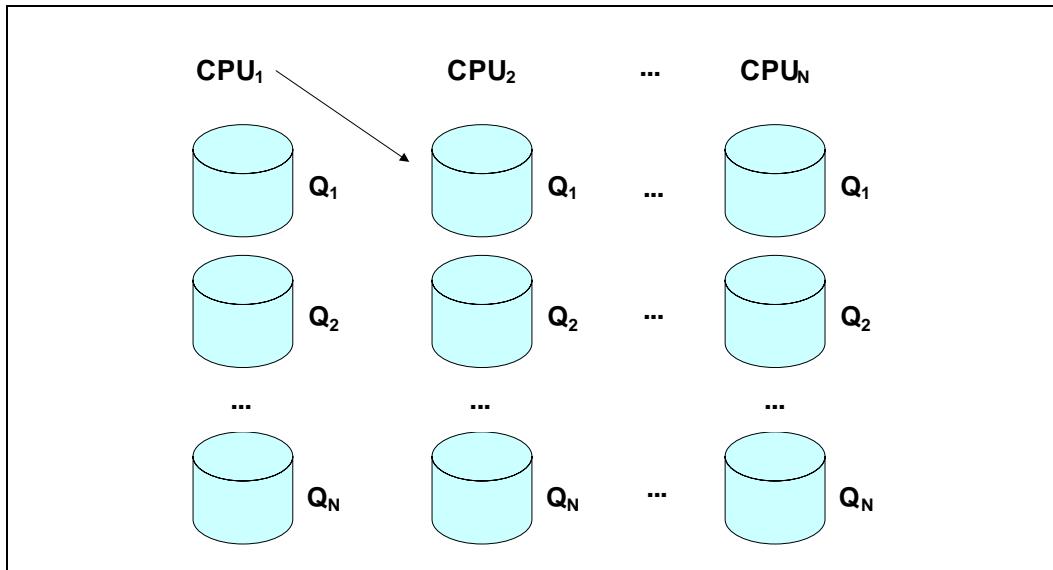
6. ProcessFlow performs actions on the packet based on the rules specified in the policies file. The policies file is created by user and loaded using the user-space tool.
7. ProcessFlow adds the packet to be forwarded to the per-CPU transmit queue. Transmit queues are described in Section 2.1.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72

## 2.1 Transmit Queues Example

Figure 3. shows an example of a transmit queue.

**Figure 3. Transmit Queues Example**



Consider the flow from eth1 -> eth2. eth1 was assigned to CPU<sub>1</sub> using IRQ affinity and receives a packet. eth1 determines the output device for eth2 using policies specified by the user. CPU<sub>1</sub> writes the packet to Q<sub>1</sub> belonging to CPU<sub>2</sub>.

Meanwhile, CPU<sub>2</sub> is polling all of its TX queues. CPU<sub>2</sub> sees that CPU<sub>1</sub> and transmits the packet.

## 2.2 Descriptor Bunching

The standard e1000e driver writes TX descriptors back to the NIC after transmitting every individual packet. This requires a costly PCI config space write.

To avoid this, the Lincoln Tunnel software uses a modified version of the driver to wait until many packets have been transmitted before updating the TX descriptor tail pointer on the NIC. (TODO: How many packets? Is this configurable?) This has the advantage of balancing the overhead of many packets and reduces the per-packet load.

*Note:* The standard e1000e driver performs RX descriptor bunching already, therefore the modification was implemented for TX descriptors only.

## 3.0 Performance

Do we want to include any performance information?

If so can we use contents from Jeff's presentation?

Do these results need legal review?



1  
2 Category—Lincoln Tunnel L3 Software  
3  
4  
5  
6

## 7 Appendix A e1000e-0.5.20 Modifications 8

9 A number of changes have been made to the standard e1000e driver to support the  
10 Lincoln Tunnel software. The following is a list of files that were created/modified and a  
11 brief description of the changes.  
12

- 13 • Files Changed:  
14     — e1000\_82571.c – Register configuration required for descriptor bunching  
15     — e1000.h – Added members to e1000\_adapter structure  
16     — netdev.c – Added code to use Lincoln Tunnel TX queues, descriptor bunching,  
17         and NUMA awareness  
18     — param.c - Added code to setup NUMA node assignment  
19
- 20 • Files Added:  
21     — flows.c – Flow searching/processing/queuing and condition checking  
22     — flows.h – Defines structures for in flow processing  
23

24                          § §  
25