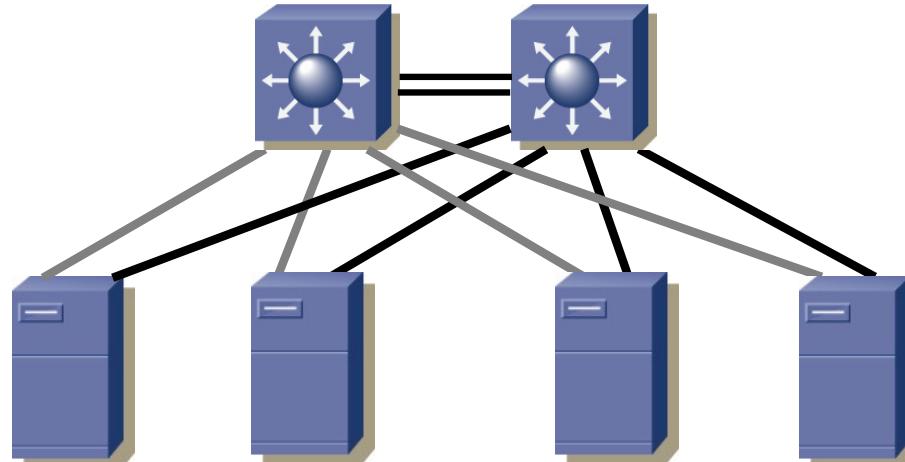


NICs Ethernet para servidor

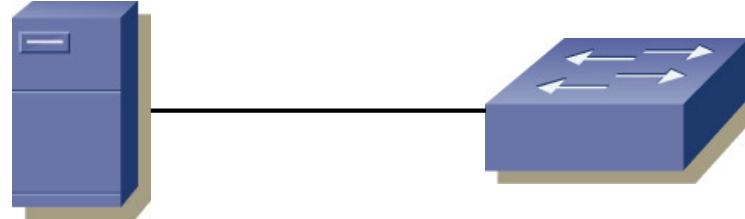
Requerimientos

- Soporte de algún tipo de agregación de interfaces
 - Gestionable a nivel de usuario
 - O de sistema operativo
- Alto rendimiento
 - Arquitectura PC limitada para altas tasas de paquete/s
- Virtualización
 - Soporte eficiente de VMs en el host (más adelante)
- Almacenamiento
 - Integración con almacenamiento en red (más adelante)



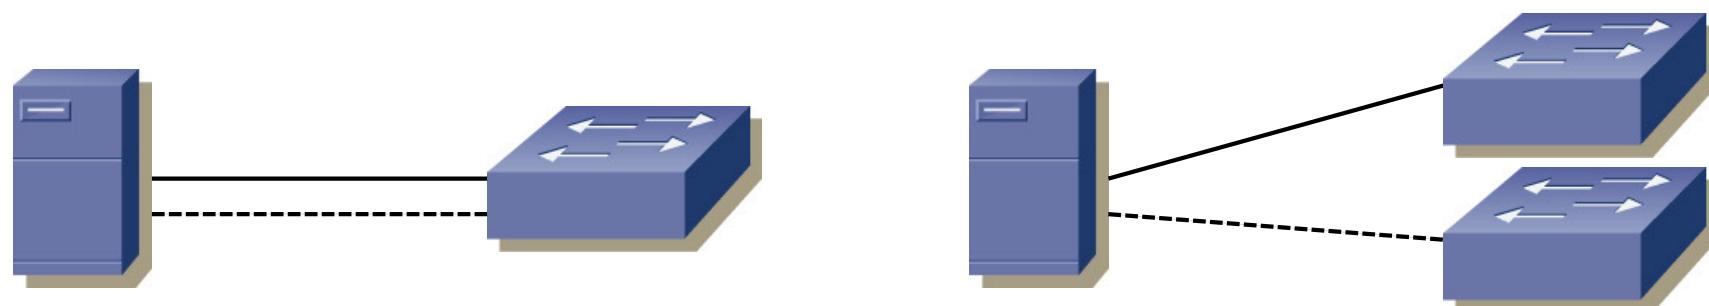
Server multihoming

- *NIC teaming / bonding / aggregation*
- Un servidor conectado a un conmutador presenta puntos únicos de fallo: la NIC, el cable, el conmutador
- Estas soluciones requieren colaboración del driver y normalmente también del sistema operativo
- Tenemos varias mejoras posibles (con una segunda o más NICs)
- (...)



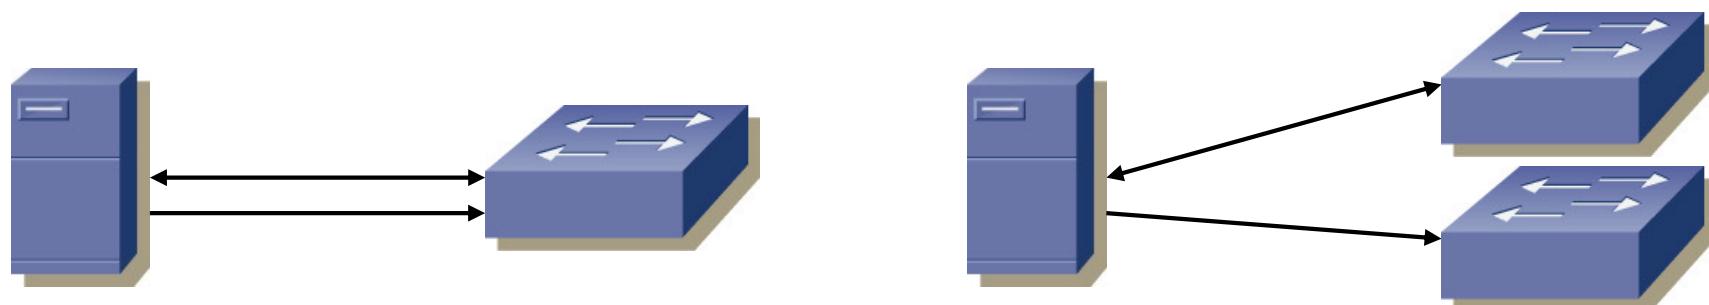
Server multihoming

- Un segundo enlace, modo activo-pasivo
 - Si falla el primero (la NIC, el commutador o el cable) se activa el segundo con la misma dirección MAC e IP
 - Se desaprovecha el segundo enlace



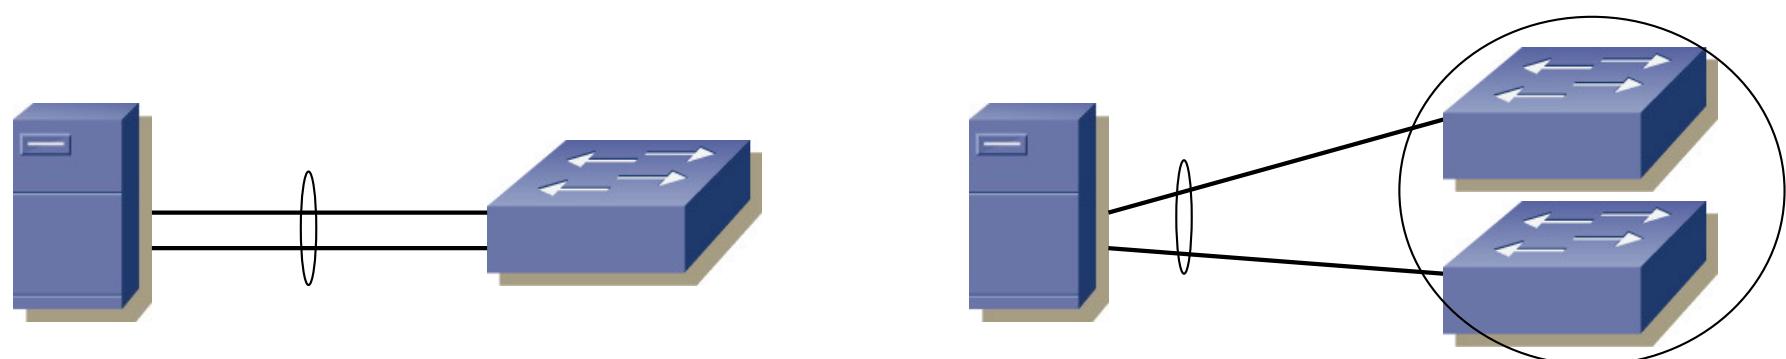
Server multihoming

- Un segundo enlace, modo activo-pasivo
- O se usan los dos enlaces para transmitir pero solo se recibe por uno
- Cada interfaz suele enviar con diferente dirección MAC origen para no tener *MAC flapping* en el conmutador



Server multihoming

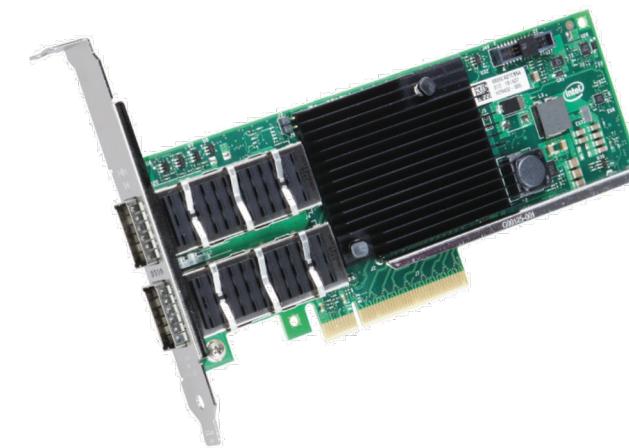
- Un segundo enlace, modo activo-pasivo
- O se usan los dos enlaces para transmitir pero solo se recibe por uno
- O se forma un LAG (802.3ad / 802.1AX)
 - Permite usar la capacidad de ambos enlaces
 - Normalmente requiere colaboración por parte del switch
 - Si se quiere redundancia de switch hay que hacer una agregación en la que un extremo son 2 comutadores



Alto rendimiento

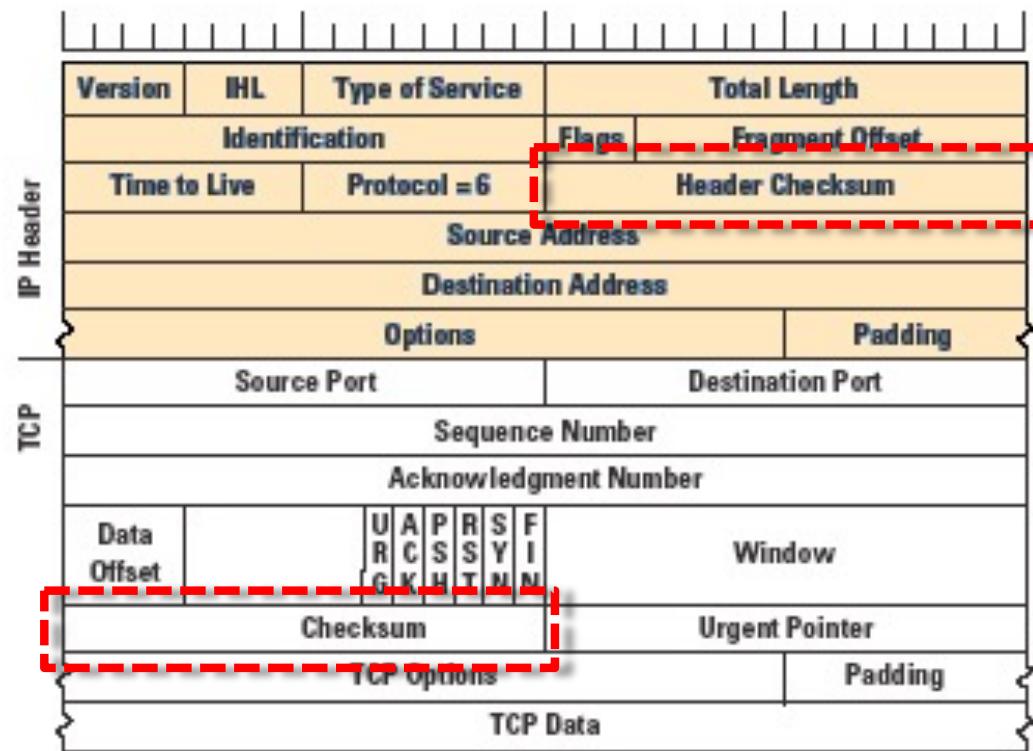
Tareas en la NIC

- Por un enlace 10GE pueden llegar en 1 segundo más de 14 millones de tramas de 64 bytes
- Eso da a la CPU unos 67ns para procesar cada una
- Las CPUs tienen serios problemas para procesar en ese tiempo cabeceras TCP/IP
- Una NIC puede incluir electrónica para llevar a cabo ciertas tareas de TCP/IP descargando a la CPU
- La NIC puede incluir ASICs, Network Processors o un procesador con un sistema operativo de tiempo real
- A 400Gbps una trama cada 1,67ns lo cual está en el rango de los mejores tiempos de acceso a memoria



Checksum offload

- La NIC descarga del cálculo a la CPU
- En transmisión y recepción
- Checksum IP (v4 y v6), UDP y TCP

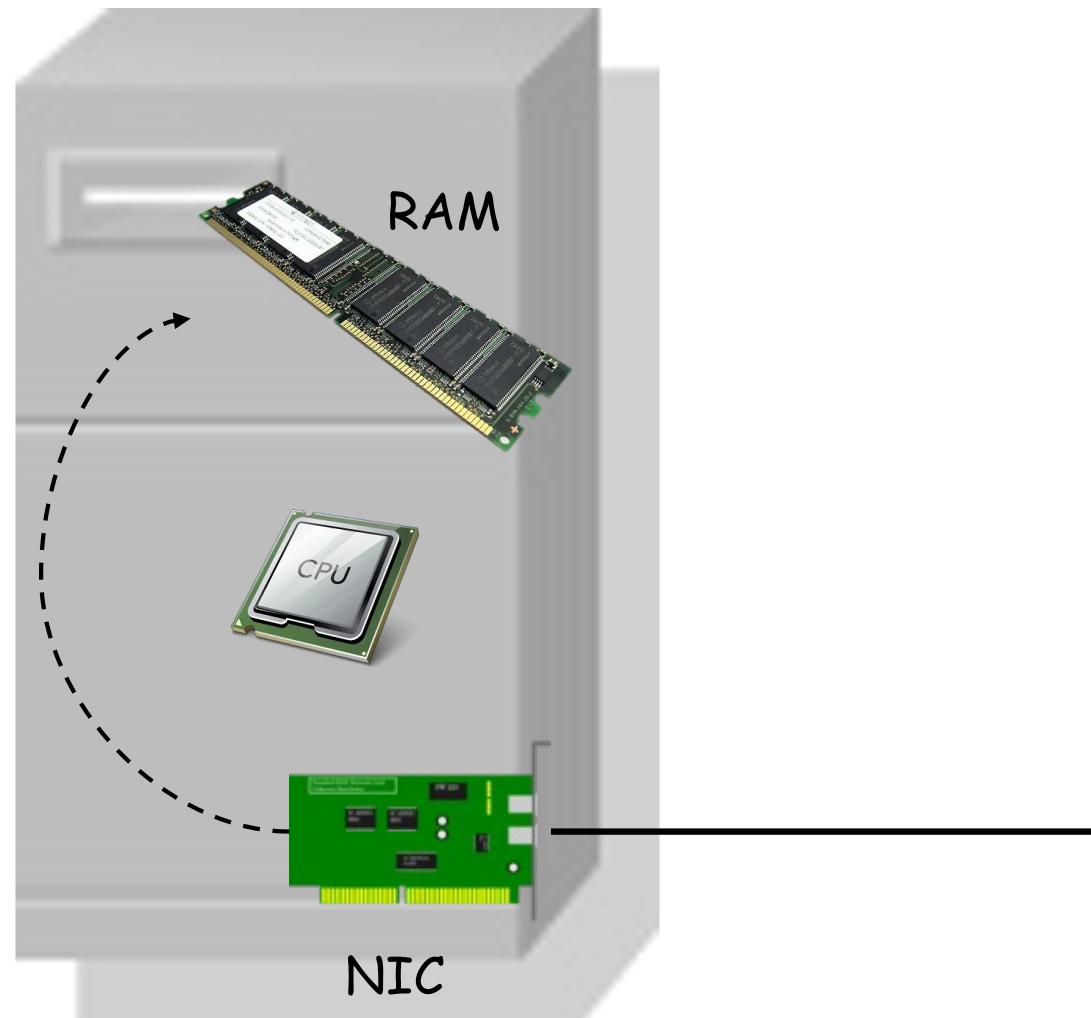


Integración en el bus

- Coalescencia de interrupciones
 - Las NICs solían generar una interrupción por paquete
 - Alto coste para la CPU
 - Por ejemplo los mainframes tienen CPUs dedicadas a atender I/O
 - La coalescencia hace que la NIC genere una interrupción para un grupo de paquetes en vez de por cada uno
 - También puede hacer *polling* la NIC

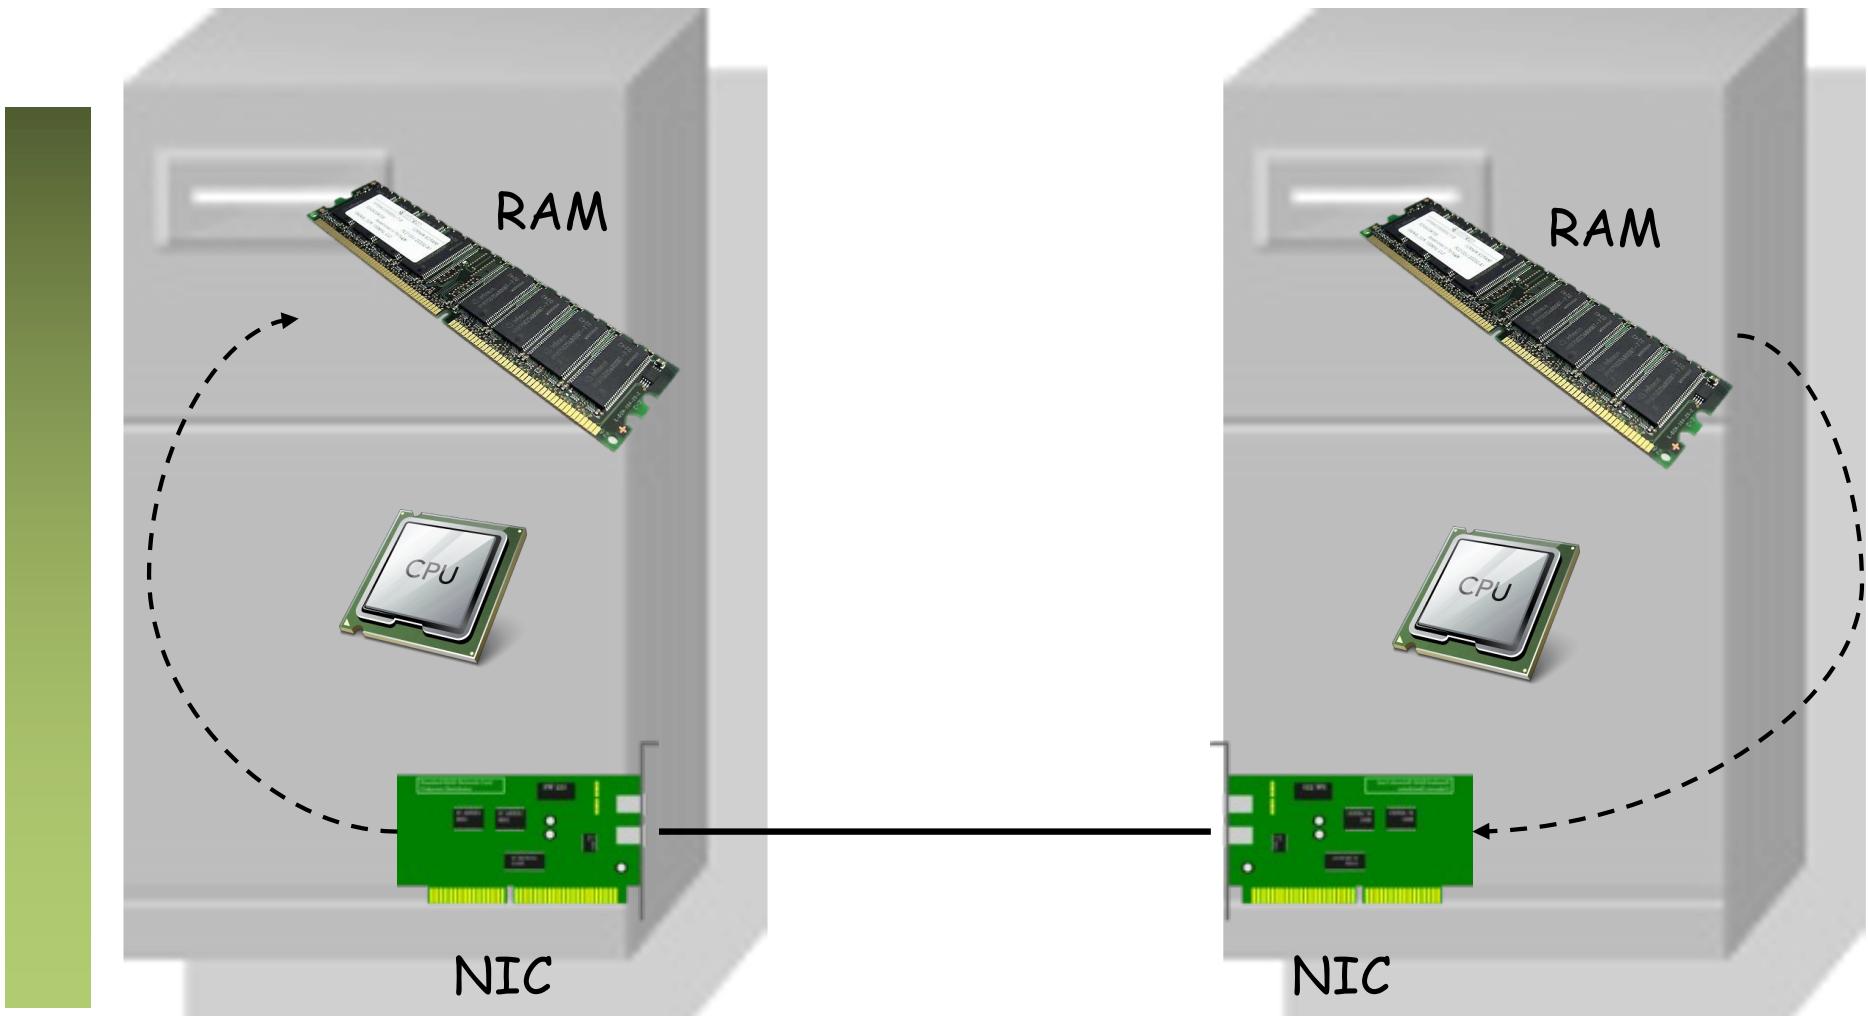
Integración en el bus

- DMA
 - *Direct Memory Access*
 - Transferencia desde la NIC a memoria sin requerir a la CPU



RDMA

- Remote Direct Memory Access
- Copias entre RAM de hosts diferentes sin involucrar a la CPU
- Latencia de pocos microsegundos



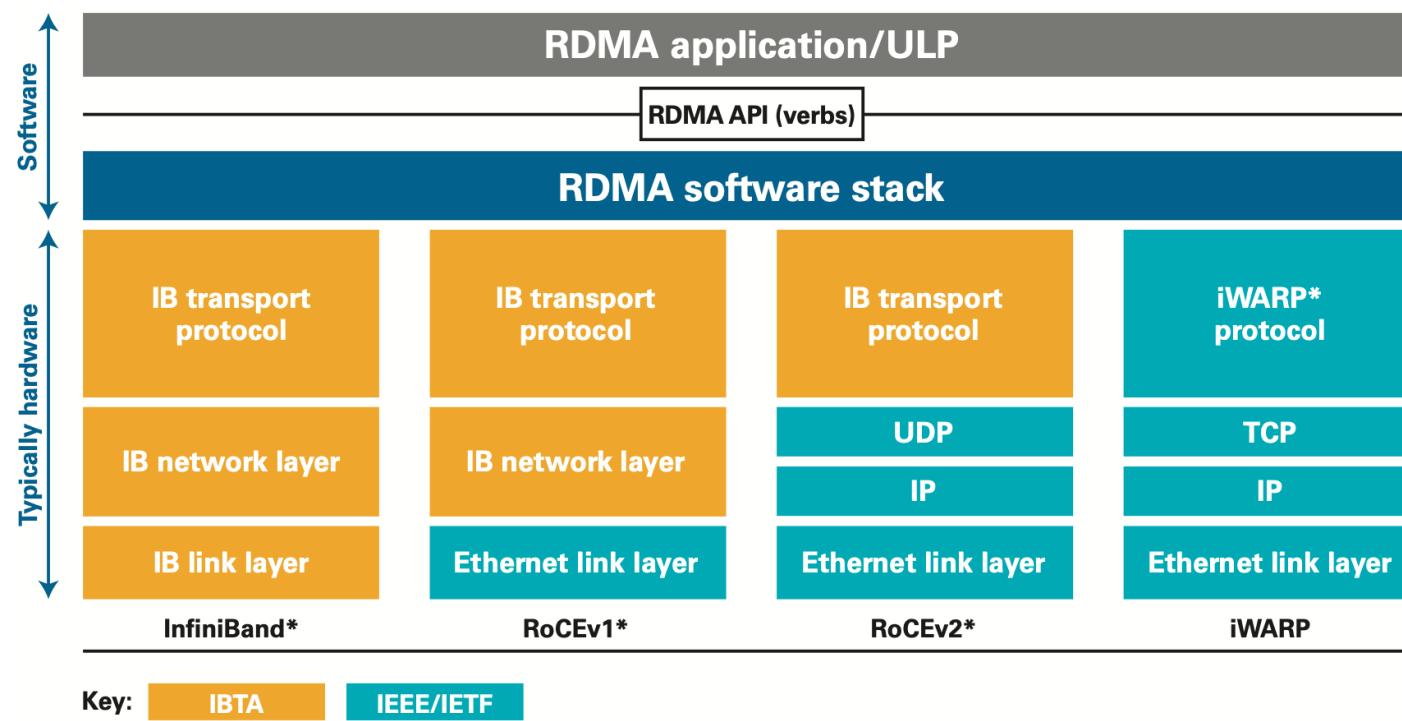
RDMA

iWARP

- RFCs 5040, 5041, 5044
- Sobre TCP o SCTP

RoCE

- RDMA over Converged Ethernet (DCB, Data Center Bridging)
- RoCE v1 sobre Ethernet, v2 sobre UDP
- RoCE v1 mecanismos de control de flujo y congestión de DCB
- RoCE v2 emplea control de congestión basado en ECN



Jumbo frames

- Tramas Ethernet con MTU superior a 1500bytes
- No están estandarizadas, la MTU estándar sigue siendo de 1500bytes
- Motivos para limitarlo
 - NICs tenían memoria limitada
 - Se quería limitar el tiempo que una estación tenía capturado el medio transmitiendo
 - El CRC es menos efectivo cuanto más grande es la trama
- Hoy en día no son problemas reales:
 - Decenas o centenares de Megabytes en la NIC
 - No tenemos medio compartido (ni coaxial ni hubs)
 - El CRC de Ethernet soporta más de 11 Kbytes de trama



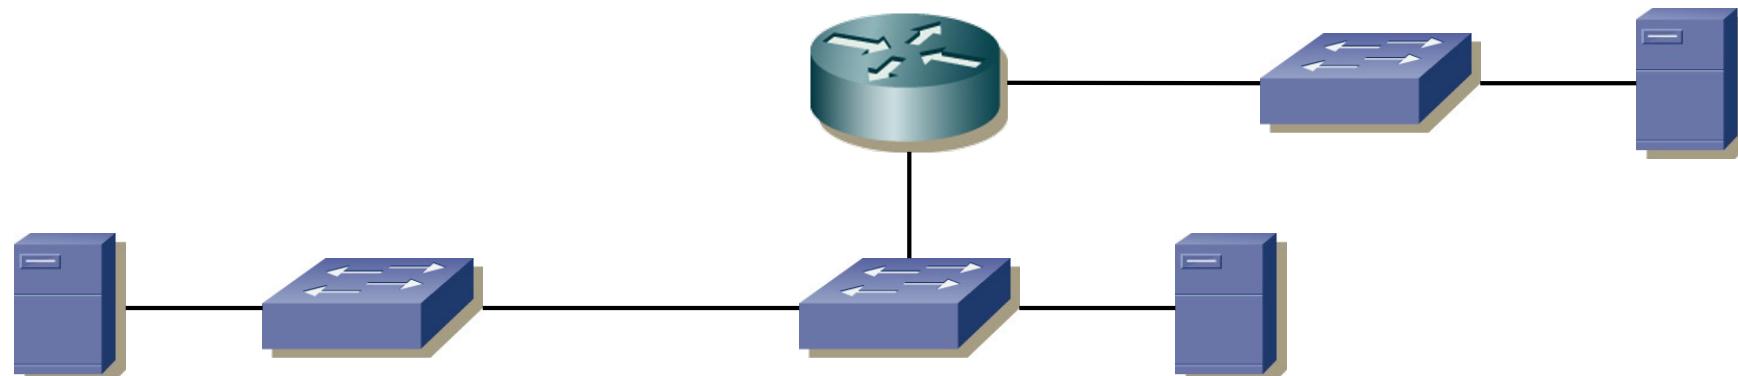
Jumbo frames

- Diversos estándares han ido aumentando el tamaño máximo de la trama (802.1Q, 802.1ad, MPLS, FCoE, etc)
- A estas últimas en ocasiones se las llama “Baby Giant”
- Jumbo frames suelen estar cerca de los 9 Kbytes (que se puedan transportar bloques de datos de 8Kbytes + encapsulados varios)
- ¿Postivo?
 - Cuanto más grandes menor ratio de cabeceras y menos interrupciones
 - Menos carga de procesado de cabeceras en equipos de red y hosts
- ¿Negativo?
 - (...)



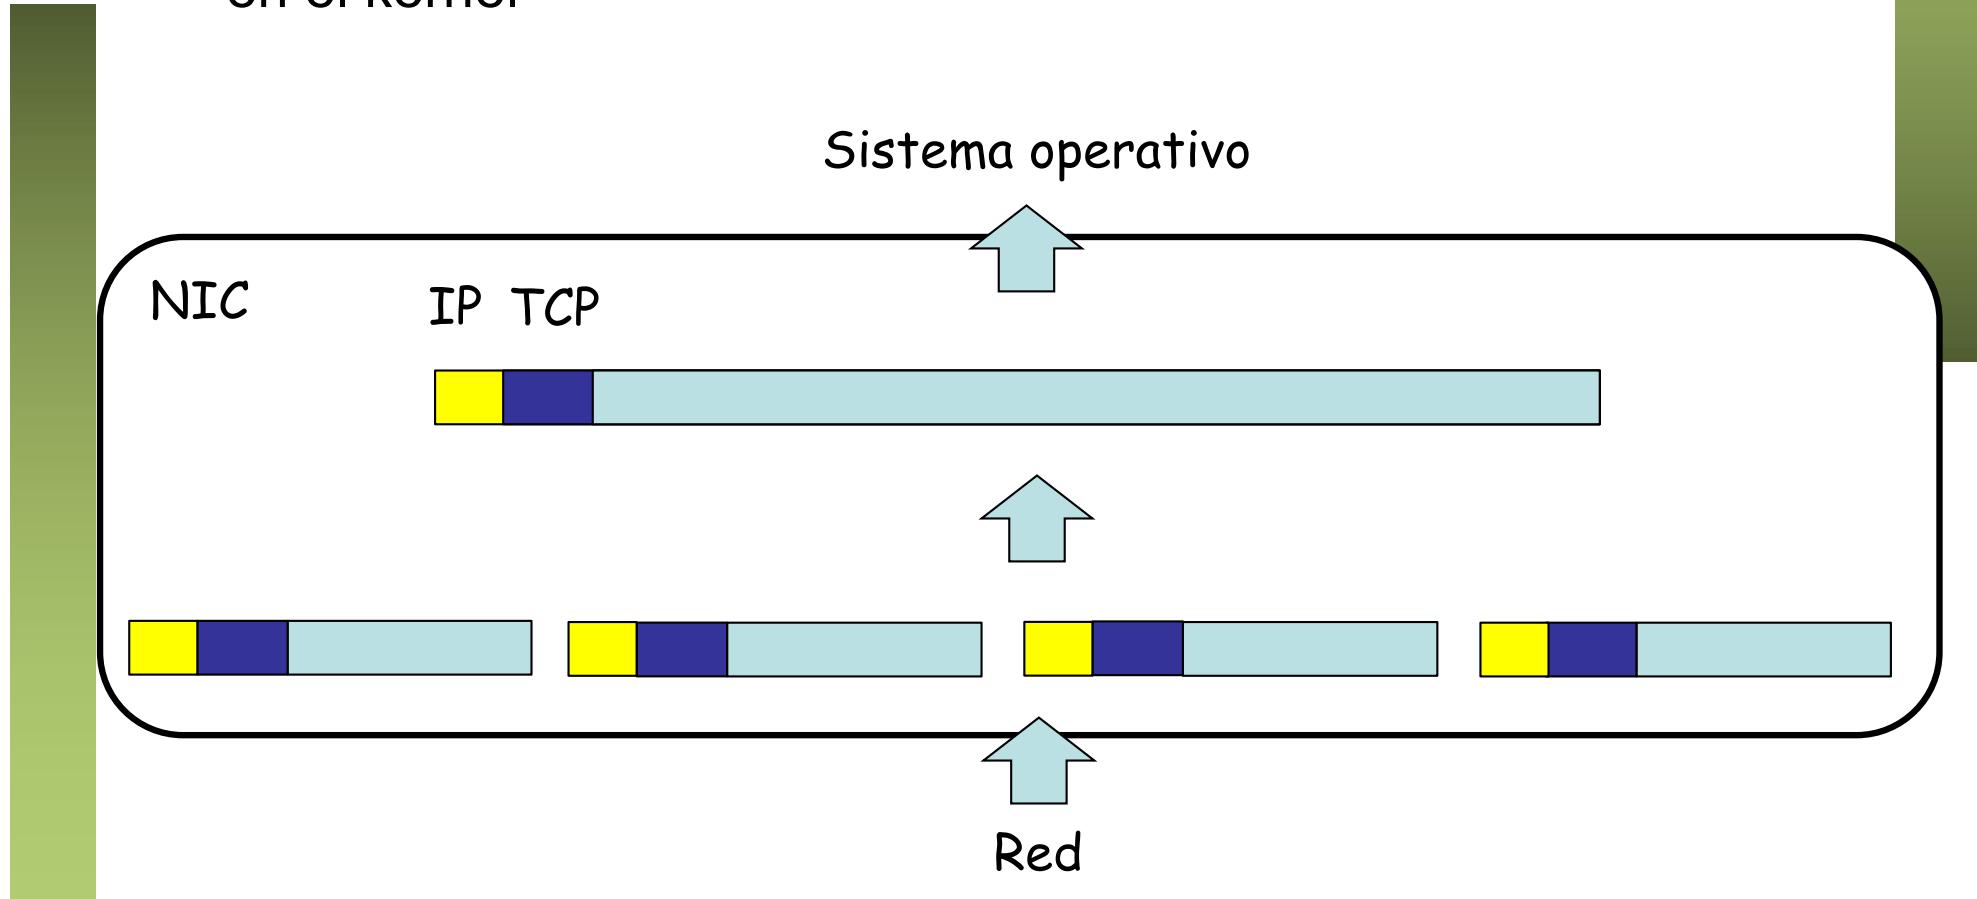
Jumbo frames

- Diversos estándares han ido aumentando el tamaño máximo de la trama (802.1Q, 802.1ad, MPLS, FCoE, etc)
- A estas últimas en ocasiones se las llama “Baby Giant”
- Jumbo frames suelen estar cerca de los 9 Kbytes (que se puedan transportar bloques de datos de 8Kbytes + encapsulados varios)
- ¿Postivo?
- ¿Negativo?
 - Todos los equipos del camino deben soportarlas
 - Posibles problemas con implementaciones que esperan 1500 bytes
 - Mayores tramas sufren mayor retardo así que no son adecuadas para todos los servicios
 - Mayores tramas pueden llenar antes los buffers de los conmutadores



LRO

- *Large Receive Offload, Receive Segment Coalescing*
- La NIC une varios segmentos TCP en uno solo
- Crea unas cabeceras TCP e IP para ese nuevo segmento
- Reduce el número de interrupciones y procesado de cabeceras en el kernel



LRO: Ejemplo

h2-1-images-beginning_cnx_60260.pcap

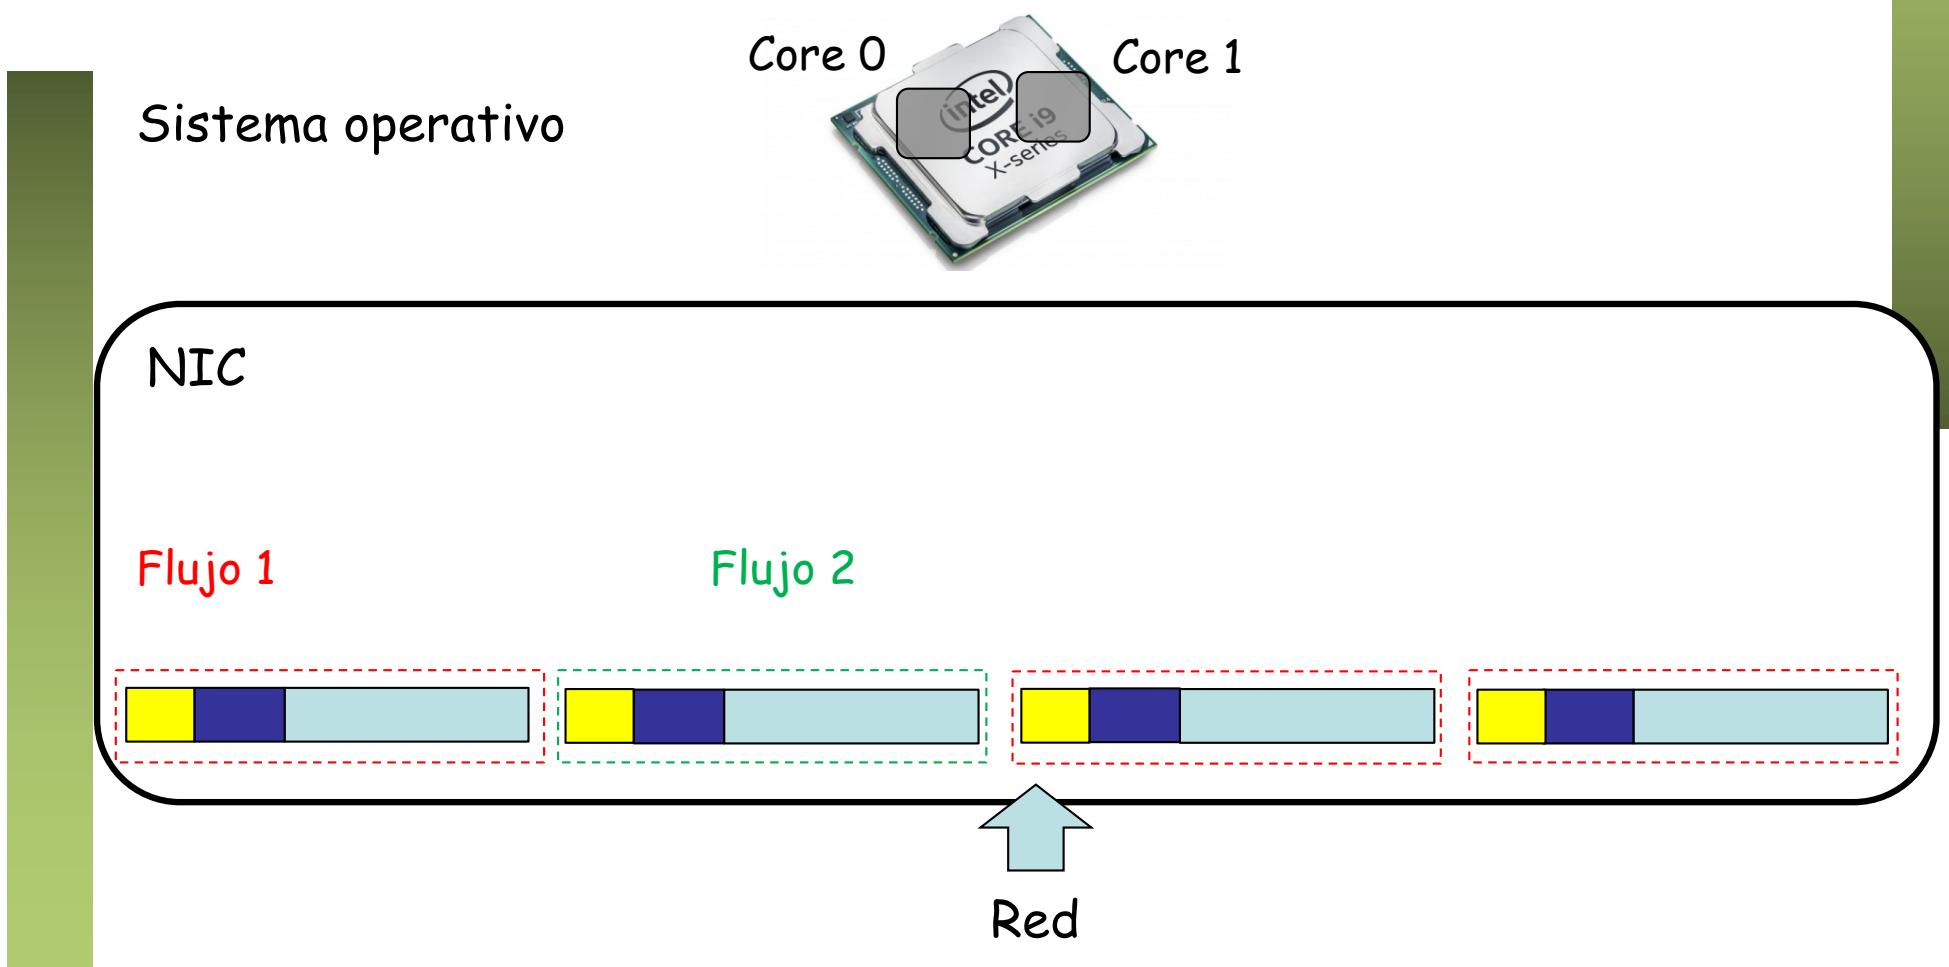
Apply a display filter ... <⌘/>

No.	Time	Source	Destination	tcp.len	frame.len	Info
24	1612366802.319898	192.168.1.3	192.168.1.2	158	224	Application Data
25	1612366802.319977	192.168.1.2	192.168.1.3	0	66	60260 → 443 [ACK] Seq=1077 Ack=2978
26	1612366802.320132	192.168.1.3	192.168.1.2	2896	2962	Application Data, Application Data
27	1612366802.320133	192.168.1.3	192.168.1.2	1448	1514	Application Data [TCP segment of a connection]
28	1612366802.320191	192.168.1.3	192.168.1.2	944	1010	Application Data
29	1612366802.320193	192.168.1.3	192.168.1.2	2896	2962	Application Data, Application Data
30	1612366802.320194	192.168.1.2	192.168.1.3	0	66	60260 → 443 [ACK] Seq=1077 Ack=2978
31	1612366802.320194	192.168.1.2	192.168.1.3	0	66	60260 → 443 [ACK] Seq=1077 Ack=2978
32	1612366802.320196	192.168.1.3	192.168.1.2	2392	2458	Application Data, Application Data

▶ Frame 26: 2962 bytes on wire (23696 bits), 2962 bytes captured (23696 bits)
▶ Ethernet II, Src: Universa_2c:dc:32 (00:1e:37:2c:dc:32), Dst: Universa_2c:dc:6c (00:1e:37:2c:dc:6c)
▼ Internet Protocol Version 4, Src: 192.168.1.3, Dst: 192.168.1.2
 0100 = Version: 4
 0101 = Header Length: 20 bytes (5)
▶ Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
 Total Length: 2948
 Identification: 0x12b6 (4790)
▶ Flags: 0x40, Don't fragment
 Fragment Offset: 0
 Time to Live: 64
 Protocol: TCP (6)
 Header Checksum: 0x9968 [validation disabled]
 [Header checksum status: Unverified]
 Source Address: 192.168.1.3
 Destination Address: 192.168.1.2
▶ Transmission Control Protocol, Src Port: 443, Dst Port: 60260, Seq: 2978, Ack: 1077, Len: 2896
▼ Transport Layer Security
 ▼ TLSv1.3 Record Layer: Application Data Protocol: http-over-tls
 Opaque Type: Application Data (23)
 Version: TLS 1.2 (0x0303)
 Length: 1317

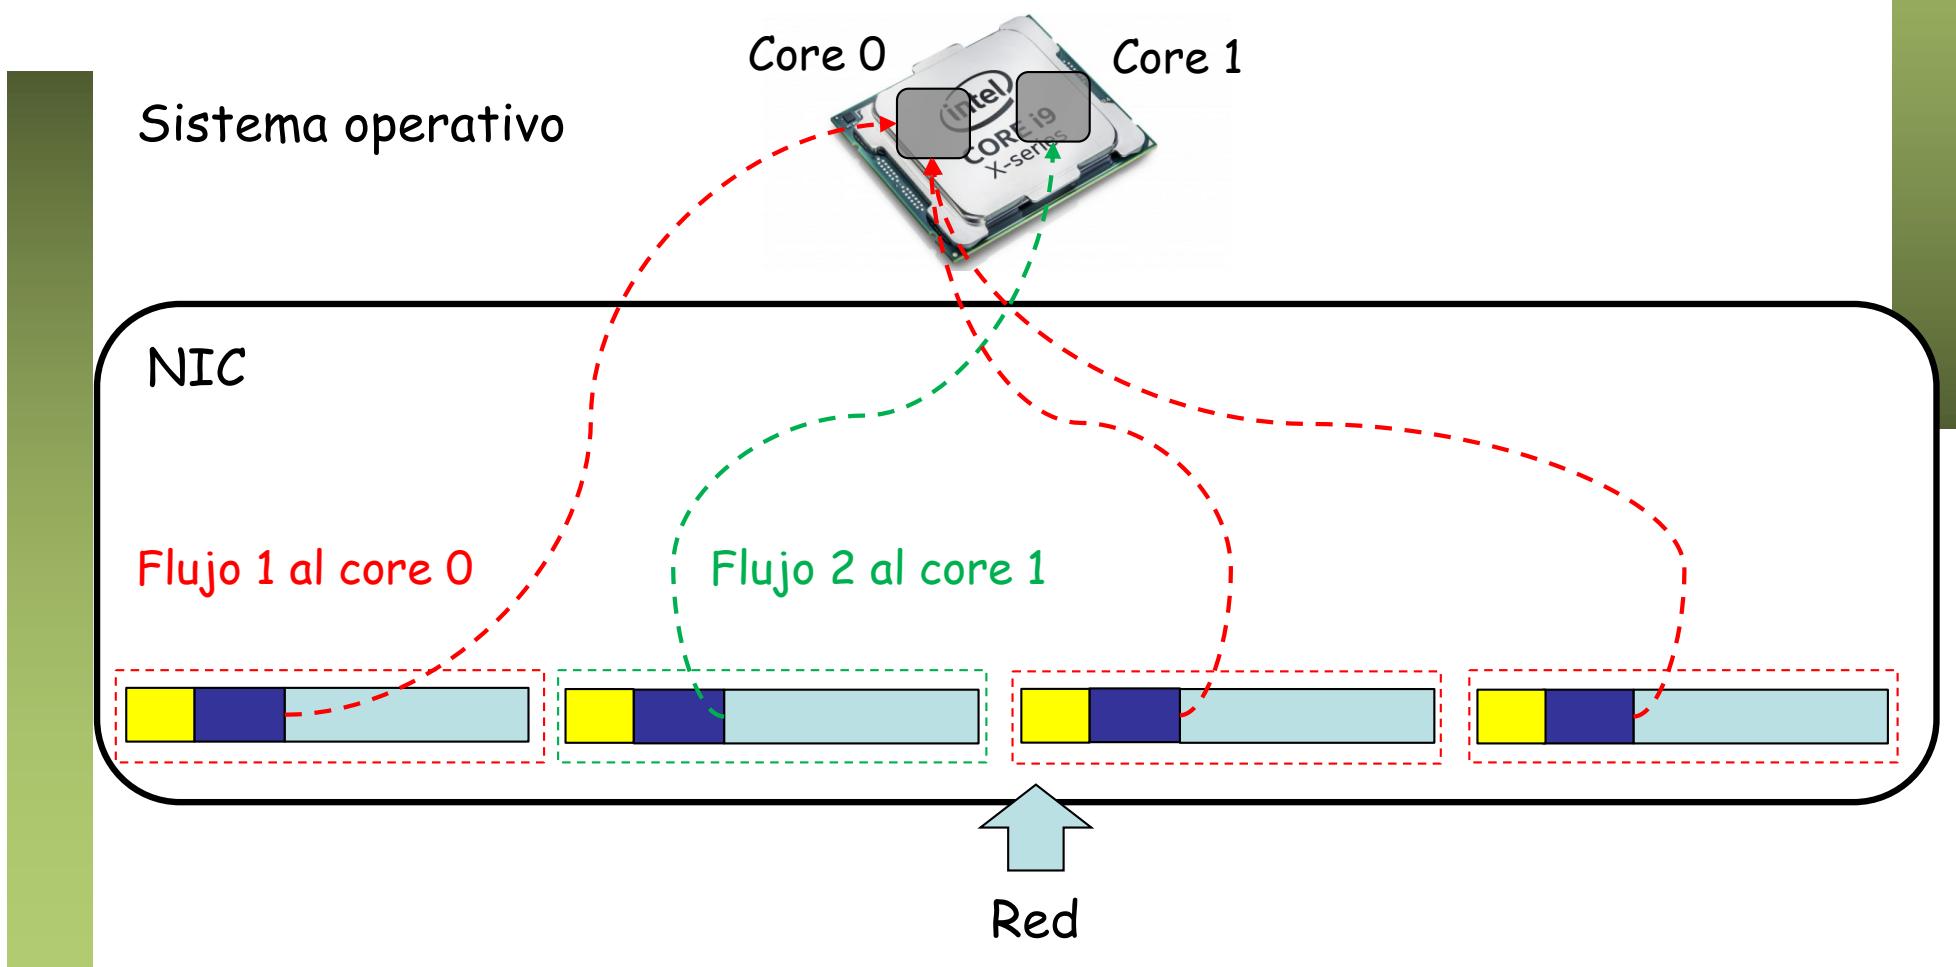
RSS

- *Receive Side Scaling*
- Multi-CPU o CPU multi-core



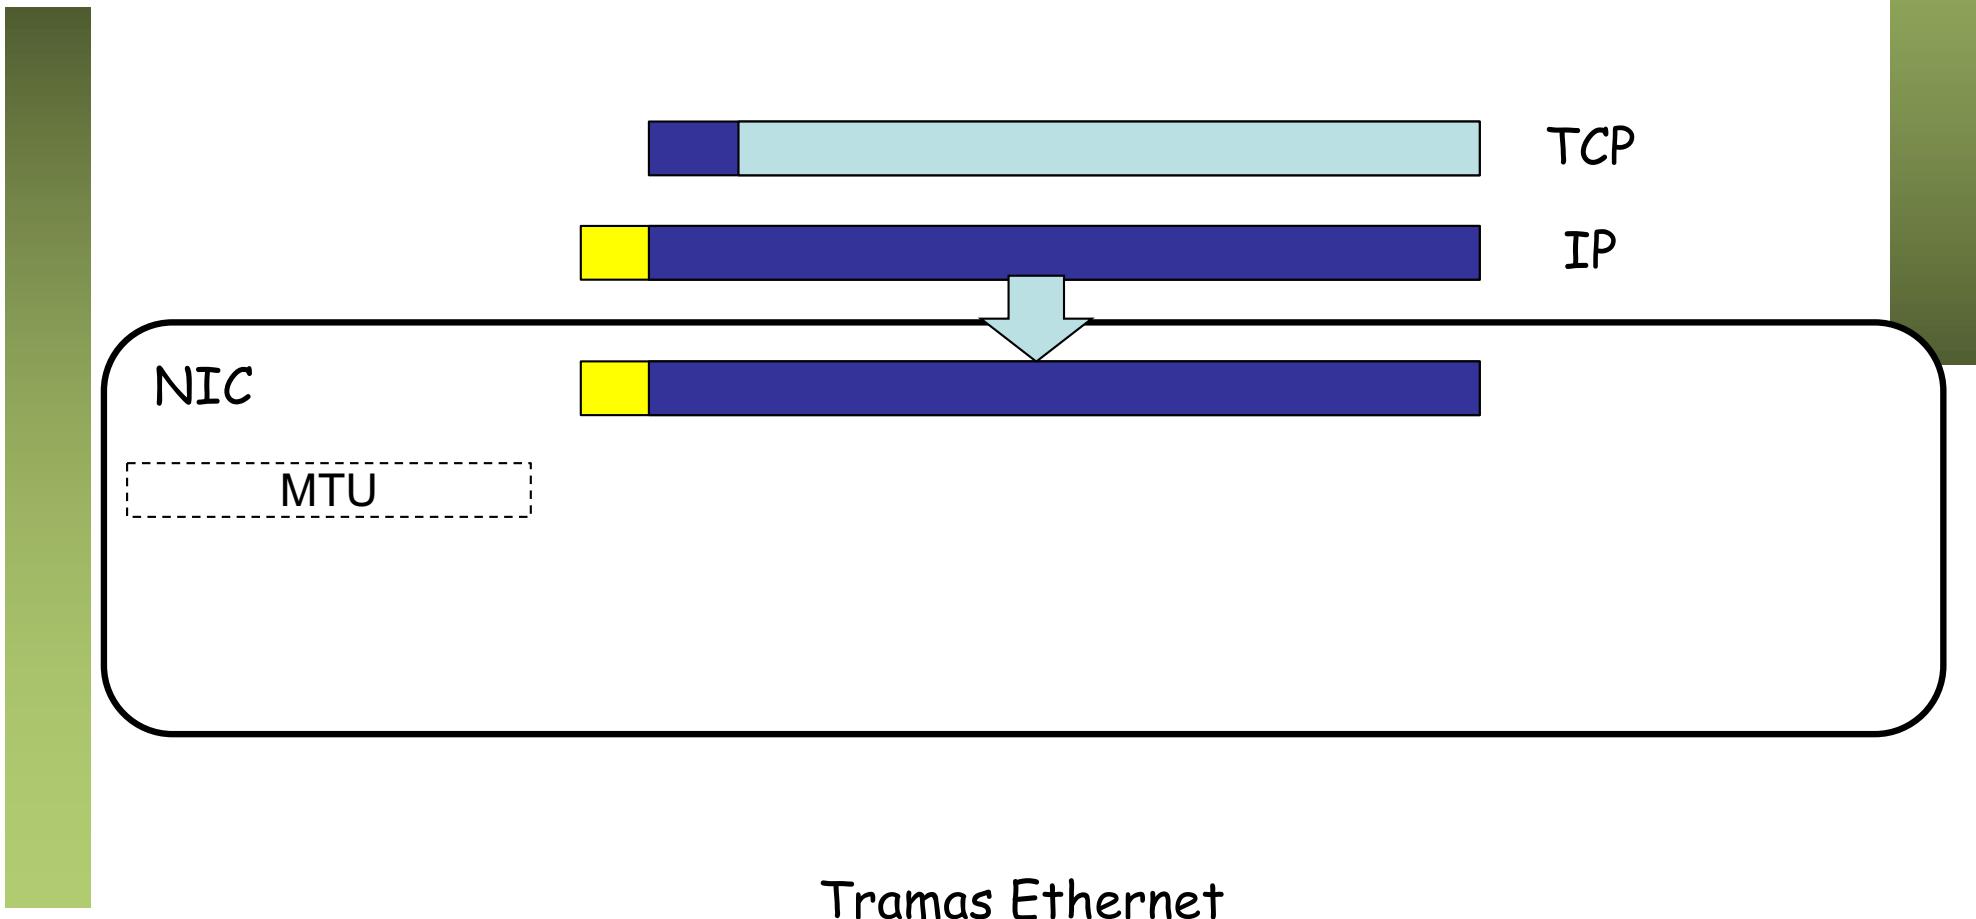
RSS

- NIC calcula un hash sobre el paquete recibido y con él decide a qué CPU manda la interrupción
- Permite paralelizar entre varias CPUs el procesado del tráfico recibido



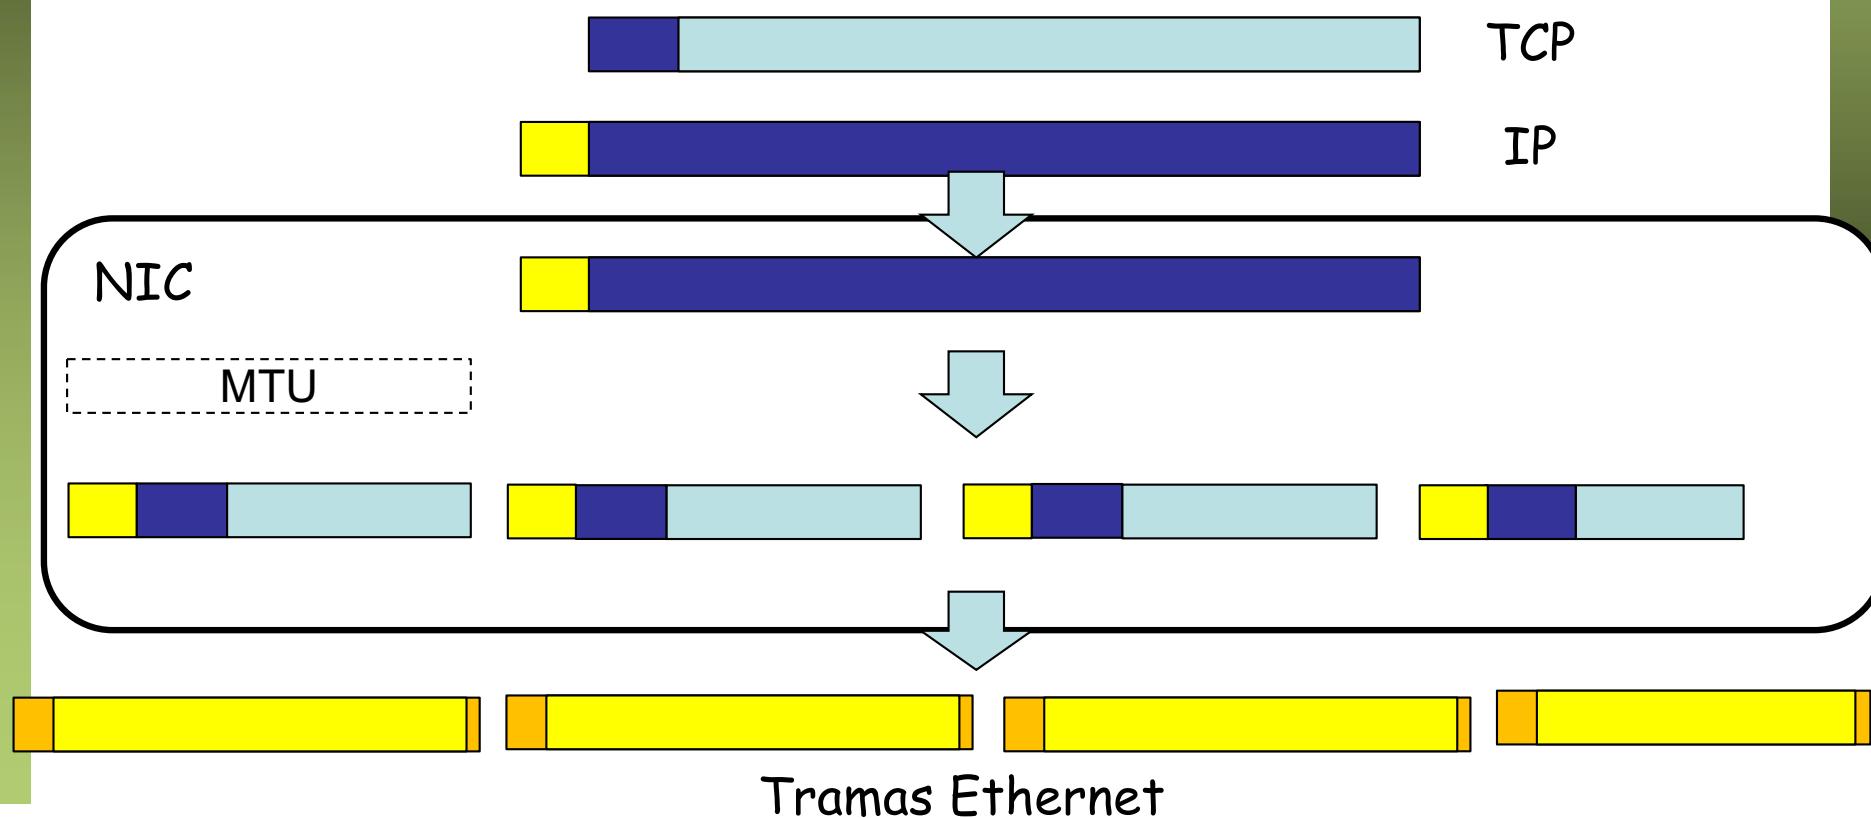
LSO

- *Large Segment Offload, TCP Segmentation Offload*
- Busca reducir carga de trabajo a la CPU en transmisión
- TCP entrega a la NIC paquetes más grandes que la MTU
- (...)



LSO

- La propia NIC hace la segmentación de nivel TCP
- Eso le obliga a crear nuevas cabeceras TCP e IP, descargando de ello a la CPU
- Requiere que la NIC sepa segmentar el protocolo (solo TCP)
- Problemas con encriptación (IPSec)
- Genera ráfagas de tráfico





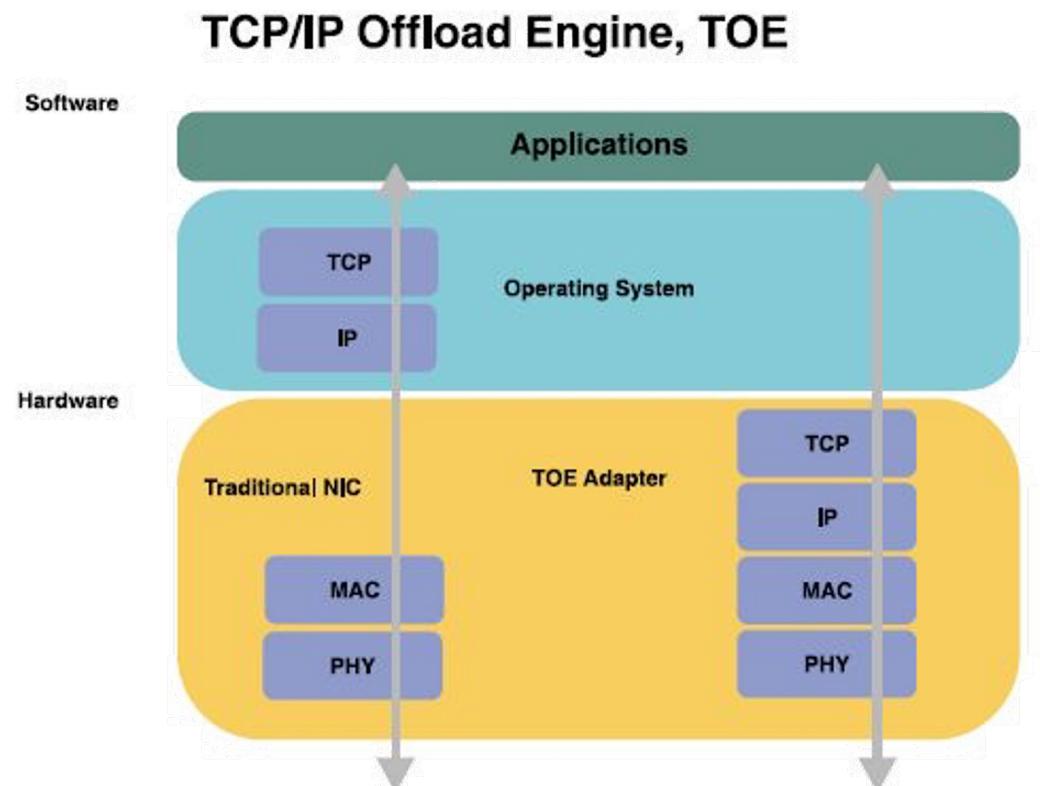
Universidad Pública de Navarra
Nafarroako Unibertsitate Publikoa

Redes de Nueva Generación
Área de Ingeniería Telemática

TOE

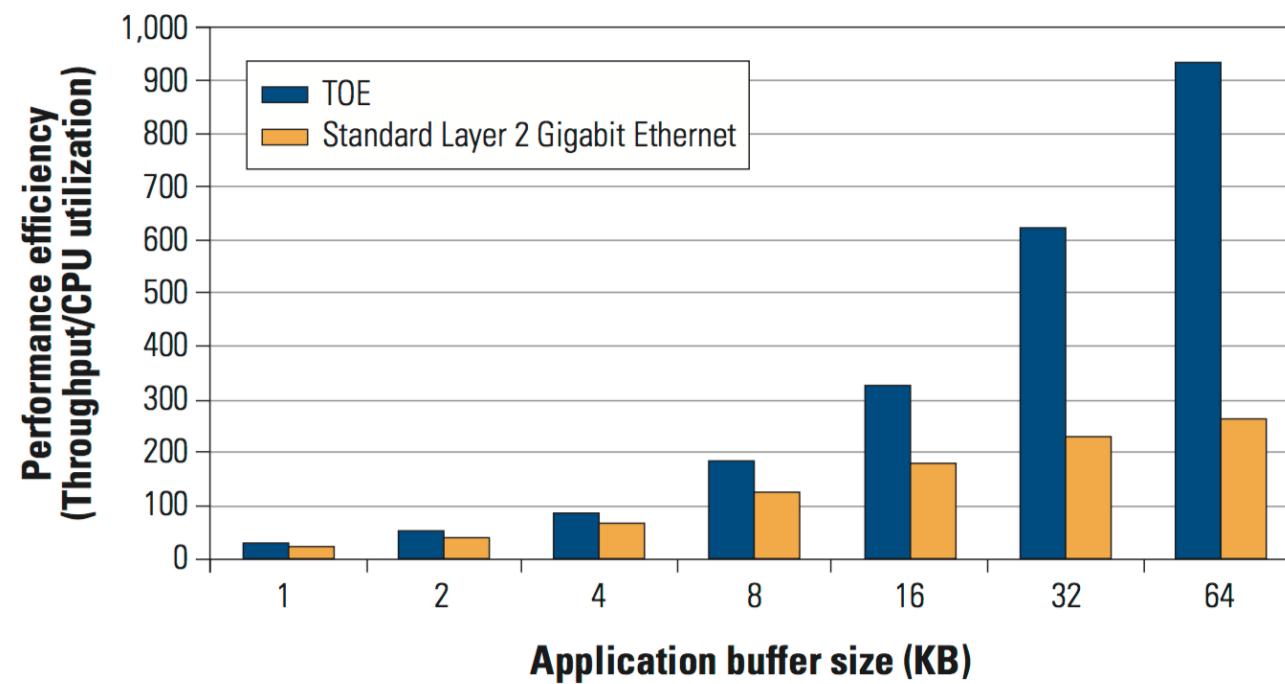
TOE

- *TCP/IP Offload Engine*
- Los datos pueden pasar directamente de la aplicación a la NIC
- La NIC puede emplearse para todas las tareas de la fase de transferencia y emplear la CPU para el establecimiento y terminación
- O se puede emplear la NIC para todo
- Requiere soporte del sistema operativo



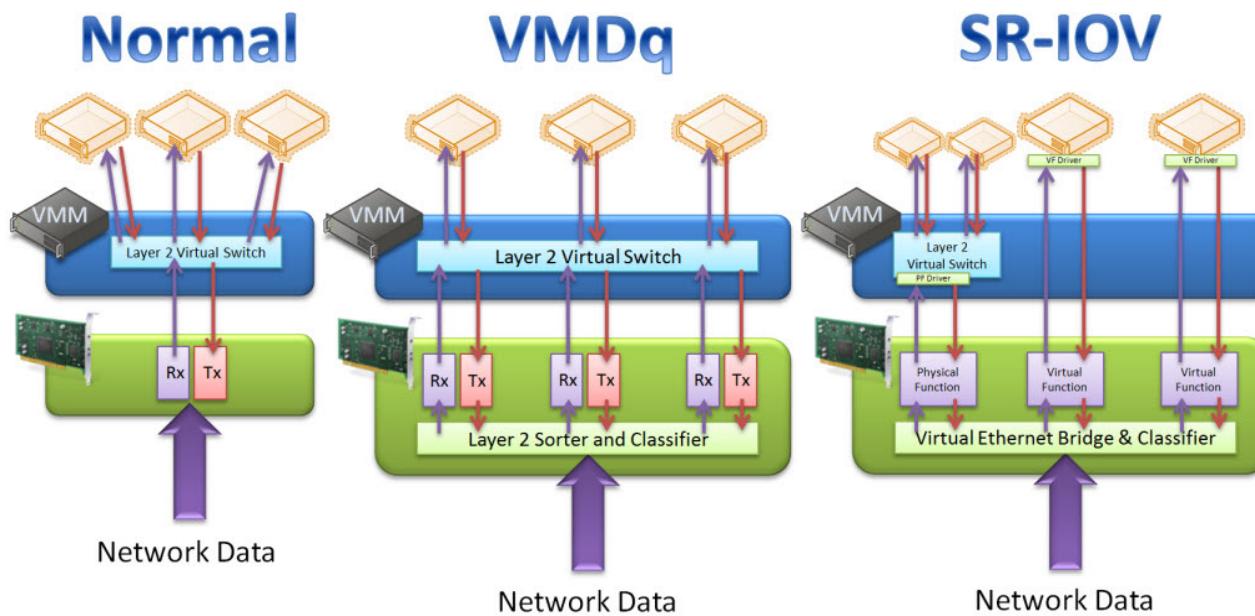
TOE

- Puede mejorar el throughput
- Reduce la carga sobre la CPU



Otras funcionalidades

- VMDq, SR-IOV, etc, asociadas a la presencia de máquinas virtuales



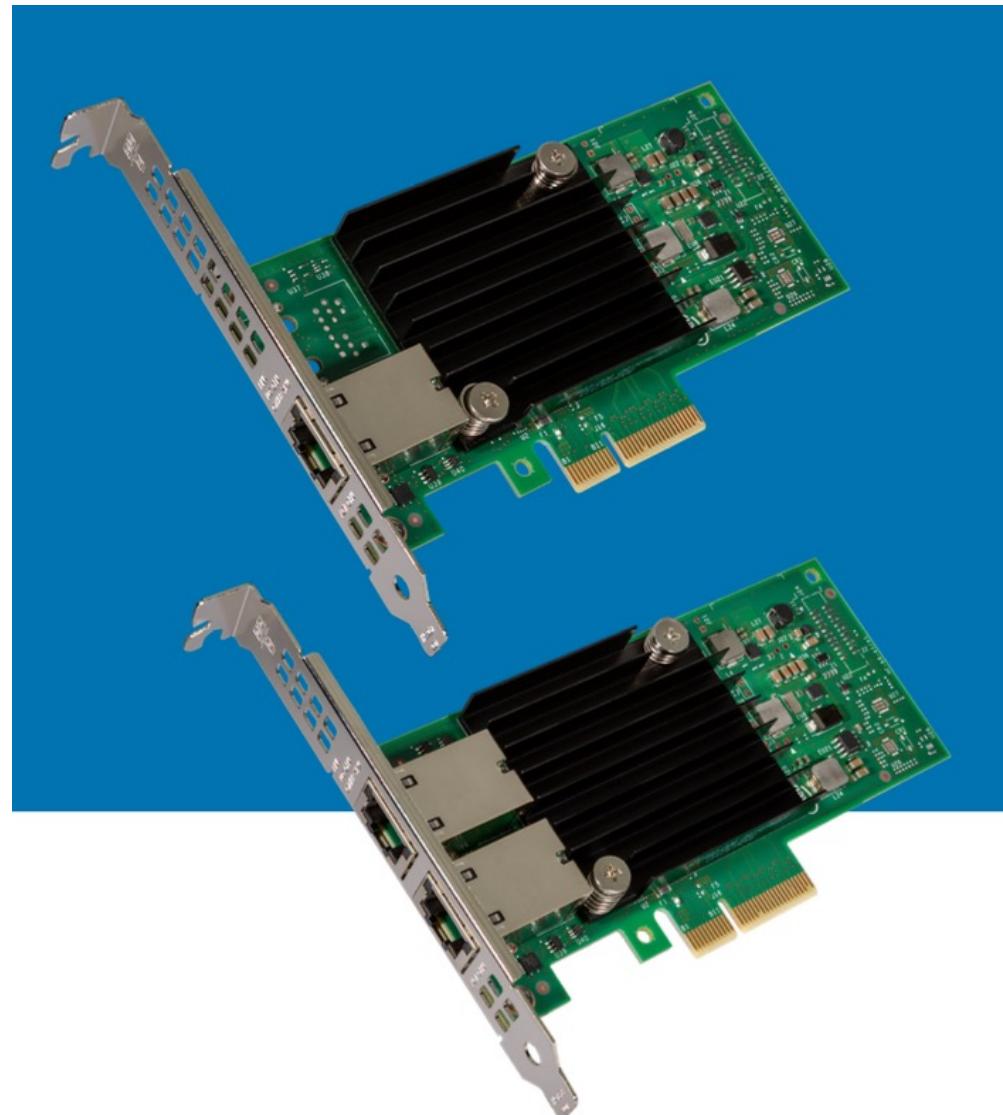
Alto rendimiento

Ejemplos

Intel® Ethernet X550-T2

Key Features

- Backward compatible with existing 1000BASE-T networks
- Supports NBASE-T technology (2.5 and 5GbE over CAT5e)
- Standard CAT6a cabling with RJ45 connectors
- Low cost, low power, 10GbE performance for the entire data center
- Flexible I/O virtualization for port partitioning and quality of service (QoS) of up to 64 virtual ports
- Single-chip solution with integrated MAC + PHY
- PCIe 3.0 with up to 8.0GT/s



<https://www.intel.es/content/www/es/es/products/sku/88209/intel-ethernet-converged-network-adapter-x550t2/specifications.html>

<https://cdrdv2.intel.com/v1/dl/getcontent/333369>

Intel® Ethernet X550-T2

Features	Description
General	
RJ45 connections over CAT6A cabling	<ul style="list-style-type: none">Ensures compatibility with cable length up to 100 meters.
RoHS-compliant, lead-free technology	<ul style="list-style-type: none">Complies with the European Union (EU) directives to reduce the use of hazardous materials.
I/O Features for Multi-core Processor Servers	
MSI-X support	<ul style="list-style-type: none">DMA engine – Enhances data acceleration across the platform (network, chipset, processor), lowering CPU usage.
Low latency	<ul style="list-style-type: none">Based on the sensitivity of the incoming data, the adapter can bypass the automatic moderation of time intervals between interrupts.
Header splits and replication in receive	<ul style="list-style-type: none">Helps the software device driver focus on the relevant part of the packet without the need to parse it.
Multiple queues – 64 Tx and Rx per port	<ul style="list-style-type: none">Network packet handling without waiting or buffer overflow providing efficient packet prioritization.
Tx/Rx IP, SCTP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities	<ul style="list-style-type: none">Checksum and segmentation capability extended to a new standard packet type.
Tx TCP segmentation offload (IPv4, IPv6)	<ul style="list-style-type: none">Increased throughput and lower processor usage.Compatible with large-send offload feature (in Microsoft Windows Server operating systems).
IPsec	<ul style="list-style-type: none">Offloads IPsec capability onto the adapter instead of software to significantly improve throughput and CPU usage.
Compatible with x4, x8 and x16 standard and Low-profile PCIe slots	<ul style="list-style-type: none">Enables each PCIe slot port to operate without interfering or competing with other PCIe slot port.
Receive Side Scaling for Windows Environment and Scalable I/O for Linux Environments (IPv4, IPv6 and TCP/ UDP)	<ul style="list-style-type: none">Enables the direction of the interrupts to the processor cores in order to improve CPU use rate.

Specifications

General	
Connections	RJ45 copper
Cable Distance	10GBASE-T: 100 m using CAT6A, 55 m using CAT6 1000BASE-T: 100 m using CAT5e, CAT6 or CAT6A

Intel® Ethernet X550-T2

Virtualization Features

Multi-mode I/O virtualization operations	<ul style="list-style-type: none">Supports two modes of operation of virtualized environments:<ul style="list-style-type: none">Direct assignment of part of the port resources to different guest operating systems using the PCI SIG SR-IOV standard (also known as native mode or pass-through mode)Central management of the networking resources by hypervisor (also known as software switch acceleration mode)A hybrid model, where some of the VMs are assigned a dedicated share of the port and the rest are serviced by a hypervisor is also supported
VxLAN stateless offloads	<ul style="list-style-type: none">A framework for overlaying virtualized layer 2 networks over layer 3 networks. VxLAN enables users to create a logical network for VMs across different networks.
NVGRE stateless offloads	<ul style="list-style-type: none">Network Virtualization using Generic Routing Encapsulation. The encapsulation of an Ethernet layer 2 Frame in IP that enables the creation of virtualized layer 2 subnets that can span physical layer 3 IP networks.
Virtual Machine Device Queues (VMDq)	<ul style="list-style-type: none">Offloads data sorting from the hypervisor to silicon, improving data throughput and CPU usage.QoS feature for Tx data by providing round-robin servicing and preventing head-of-line blocking.Sorting based on MAC addresses and VLAN tags.
64 Transmit (Tx) and receive (Rx) Queue pairs per port	<ul style="list-style-type: none">Supports VMware NetQueue and Microsoft VMQ.MAC/VLAN filtering for pool selection and either DCB or RSS for the queue in pool selection.
FPP – 64 VFs per port	<ul style="list-style-type: none">VFs appear as Ethernet controllers in Linux operating systems that can be assigned to VMs, Kernel processes or teamed using the Linux bonding drivers.
Support for PCI-SIG SR-IOV specification	<ul style="list-style-type: none">Up to 64 VFs per port.
IEEE 802.1Q VLAN support with VLAN tag insertion, Stripping and packet filtering for up to 4096 VLAN tags	<ul style="list-style-type: none">Ability to create multiple VLAN segments.Filtering packets belonging to certain VLANs.

Intel E830-CQDA2

Key Specifications

- Speeds up to 200GbE
- Supports multiple port configurations using EPCT
- PCIe 5.0 x8 or PCIe 4.0 x16
- Dynamic Device Personalization (DDP)
- Data Plane Development Kit (DPDK) enabled
- IEEE 1588 Precision Time Protocol v2
- Precision Time Measurement (PTM) v1.0a
- Commercial National Security Algorithm (CNSA) 1.0 compliant
- Modern security with signed firmware, secure boot, and hardware root of trust (RoT)

Programmable Pipeline / Dynamic Device Personalization (DDP)

DDP improves packet processing performance by using the E830 Controller's programmable pipeline to classify frames instead of the CPU. DDP increases throughput, lowers latency, and reduces host CPU overhead in both network functions virtualization (NFV) workloads and cloud-native architectures.



Open vSwitch (OVS) Acceleration

The E830 is optimized for Intel® Xeon® processors to minimize packet parsing overhead and flow table search. DPDK integration with OVS increases performance by eliminating extra layers in the architecture and native OVS stack.

Precision Time Synchronization and Measurement

Growth in 5G RAN and edge deployments is driving demand for high-precision timing synchronization across the network.

Intel® Ethernet E830 Network Adapters enable service providers to build open, disaggregated vRAN solutions with off-the-shelf components to meet unique customer needs, including system size and budget.

This adapter also features an Ethernet Port Configuration Tool (EPCT). In addition to the default configuration of 2x100, four other configurations are possible: 1x200 (port 1 only), 4x50, 2x50, and 8x25 (via breakout cables).

Intel E830-CQDA2

Adapter Features

Data Rate Supported	200/100/50/25/10GbE
Bus Type/Bus Width	PCIe 5.0 x8 or 4.0 x16
Form Factor	Standard PCIe; ships with both low-profile and full-height brackets
Controller	Intel Ethernet Controller E830
Supported Operating Systems	Linux
Hardware Certifications	BSMI, CE, CMIM, FCC, ICES, KCC, RCM, UKCA, cURus, and VCCI
Compliance	RoHS and BMSI RoHS compliant. Product is compliant with Taiwan Bureau of Standards, Metrology and Inspection (BMSI) and EU RoHS Directive 2 2011/65/EU (Directive 2011/65/EU) and its amendments.

Ethernet Media Supported

100GbE QSFP28 – 2 Ports

100GBASE-CR4, 100GBASE-SR4, 100GBASE-LR4, 100GBASE-FR,
100GBASE-DR, 100GBASE-PSM4 (Optical Breakout)

200GbE QSFP56 – 1 Port*

200GBASE-CR4, 200GBASE-SR4, 200GBASE-LR4, 200GBASE-FR4

50GbE SFP56 – Up to 4 ports with breakout cables

50GBASE-CR, 50GBASE-KR, 50GBASE-LR, 50GBASE-SR

25GbE SFP28 – Up to 8 ports with breakout cables (4 ports per cage)

25GBASE-CR (802.3by 25G twinax), 25GBASE-CR1 (Consortium 25G twinax), 25GBASE-SR, 25GBASE-LR, 25G-AUI C2M, CA-25G-N (DA Breakout), CA-25G-S (DA Breakout), CA-25G-L (DA Breakout)

10GbE SFP+ – Up to 8 ports using breakout cables (4 ports per cage)

10Gb SFI-DAC (SFP+ twinax); 10Gb SFI Limiting (SFP+ optics/AOC)

*Use only Port 1 to configure the adapter as a single 200GbE port; using this configuration, Port 2 will not have connectivity.



Broadcom 5720 Quad-Port 1GbE

Features

- Quad-port GbE Network Daughter Card for Dell PowerEdge 12G rack servers
- Two x1 PCI Express® (PCIe™) v2.0 (5 GT/s)
- Energy Efficient Ethernet (EEE)
- Full line-rate performance across all ports
- Broad OS and hypervisor support
- iSCSI remote boot support
- Preboot eXecution Environment (PXE) support
- Support for VMware® NetQueue™ and Microsoft® VMQ
- Link aggregation and automatic load balancing
- Wake-on-LAN support
- MSI and MSI-X support
- IPv4 and IPv6 offloads
- Stateless offload
- TCP, UDP, and IP checksum
- Large Send Offload (LSO)
- TCP Segmentation Offload (TSO)
- Receive Side Scaling (RSS)
- Transmit Side Scaling (TSS)
- VLAN support with VLAN tagging
- Jumbo frame support for frames larger than 1500 bytes
- Precision Time Protocol (PTP)
- Broadcom Advanced Control Suite (BACS) management application and integration into Dell's embedded management framework (iDRAC7 and Lifecycle Controller)

OS Support

Microsoft	Windows Server® 2008, 2008 R2, all editions
Linux	Red Hat® Enterprise Linux (RHEL) 5.7/5.8, 6.1/6.2
VMware	Novell® SUSE® Linux Enterprise Server (SLES) 10 SP4, 11 SP2
Citrix	vSphere™ 4.1 and 5.0
	XenServer 6.0



IEEE 802.3x—Flow control

IEEE 802.3 (Clause 30)—Statistics for SNMP MIB II, Ethernet-like MIB, and Ethernet MIB

IPv4 and IPv6 offload

Teaming support

Intel Ethernet Controller XL710-BM1/BM2

Performance

40Gb throughput	Wire-rate down to 128 bytes
10Gb throughput	Wire-rate down to 64 bytes
Standard Linux Stack Latency	~8 μ s

Additional Features

- Enhanced Transmission Selection (draft IEEE 802.1az)
- Priority Flow Control (draft IEEE 802.1Qbb)
- Data Center Bridging (DCB/DCB-X) Support; up to eight traffic classes
- Jumbo Frame Support—Up to 9.5 KB (9728 Bytes)
- VLAN Support

TCP/IP/L2 Features

- Receive Side Scaling (RSS) for TCP and UDP traffic
- Large Send Off-load (LSO) / Generic Send Off-load (GSO) including encapsulated traffic
- TCP/UDP/IP/SCTP Checksum Off-load including encapsulated traffic
- IPv4, IPv6



Virtualization Interface Features

Features	Implementation
Emulated Support	Driver Optimizations and VMDQ enablement
Direct Assignment Support	PF and VF assignment with SR-IOV
Virtual Bridging Support	VEPA/802.1Qbg
Virtual Functions	Up to 128 per device
Network Virtualization	VxLAN, MACinUDP, NVGRE, IPinGRE