

**INTEL[®]
EXPERIENCE
DAY**



ETHERNET. ОБЗОР РЕШЕНИЙ INTEL ДЛЯ CLOUD/TELECOM

Иван Панков, специалист по продукции Intel® Ethernet в России и СНГ

LEGAL DISCLAIMER

All information provided here is subject to change without notice. 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.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at <https://www.intel.com>.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. .

ОБЗОР СЕТЕВЫХ АДАПТЕРОВ INTEL® ETHERNET

INTEL® ETHERNET 500 SERIES

X520 10GbE SFP+
10GbE CNA



X540 10GBASE-T
Single Chip 10GBASE-T




X550 10GBASE-T
2nd Generation
Single Chip 10GBASE-T



INTEL® ETHERNET 700 SERIES

X710 (10GbE SFP+)
Cloud and Network
Virtualization Overlays
(Dual- and Quad-Port)




XXV710 (25GbE SFP28)
Cloud and Network
Virtualization Overlays



XL710 (40GbE QSFP+)
Network Virtualization
Overlays Acceleration



X710-T4 (10GBASE-T)
Quad-Port 10GBASE-T



PHY Mezz Cards

iWARP

Lewisburg PCH

Intel® X722



OCP X527-DA4
OCP X527-DA2
OCP X547-T2



iWARP

X722 (10GbE SFP+)
Network Virtualization
Overlays Acceleration
Cloud and Network
iWARP



X710-T2L/T4L (10GBASE-T)
Dual- and Quad-Port



INTEL® ETHERNET 800 SERIES



Up to 100 GbE per port

Queue and Steering
Hardware Assists

- Application Device Queues (ADQ)

Fully Programmable
Pipeline

- Table definition with DDP profile packages

Storage

- RDMA (iWARP* & RoCE*v2)

ОБЗОР СЕТЕВЫХ АДАПТЕРОВ INTEL® ETHERNET

Campbell Pond



Legal Name	Intel® Ethernet Network Adapter X710-T2L/T4L		
Ports	Dual / Quad Port		
Product Codes	X710T2L, X710T2LBLK (Retail Dual) X710T4L, X710T4LBLK (Retail Quad)		
Connection Type	RJ45, Twisted-pair copper		
Connection Speed	100Mbps/1Gbps/2.5Gbps/5Gbps/10Gbps		
Cabling type/range	10GBASE-T 100m on CAT6A, 55m on CAT6 5GBASE-T, 2.5GBASE-T 100m on CAT5e, CAT6 or CAT6A 100BASE-TX, 1000BASE-T 100m on CAT5e, CAT6 or CAT6A		
Slot Specs	PCI Express* v3.0 x8 (8.0GT/s)		
Controller info	Intel® Ethernet Controller X710-AT2 (Carlsville Dual) Intel® Ethernet Controller X710-TM4 (Carlsville Quad)		
Power Consumption	10GbE Dual Port	Typical 8.2W	Max 9.6W
	10GbE Quad Port	Typical 13.6W	Max 14.2W
Virtualization acceleration	Intel® AVF / VXLAN, GENEVE, NVGRE, MPLS, and VXLAN-GPE with NSH Offloads		

INTEL® ETHERNET ADAPTERS FOR OPEN COMPUTE PROJECT (OCP)

OCP Mezz 2.0 slot

Intel® Ethernet Server Adapters **40GbE** QSFP+ for OCP

Single Port **XL710QDA1OCP**

Dual Port **XL710QDA2OCP**

Intel® Ethernet Server Adapters **25GbE** SFP28 for OCP

Single Port **XXV710DA1OCP**

Dual Port **XXV710DA2OCP1** OCP Spec v2.0 Type 1, 8 mm

Dual Port **XXV710DA2OCP2** OCP Spec v2.0 Type 2, 12 mm

Intel® Ethernet Server Adapters **10GbE** SFP+ for OCP

Single Port **X520DA1OCP**

Dual Port **X520DA2OCP** **X710DA2OCP**

OCP 3.0 Future Adapters

1/10/25/100 GbE per port

I350 / X710 / 800 Series based

(preliminary forecast)



НОВЫЕ ТРЕБОВАНИЯ ДЛЯ ВЫСОКОПРОИЗВОДИТЕЛЬНЫХ СЕТЕЙ




Одновременное использование сети множеством приложений

Эффективное распределение очередей для высокоприоритетных приложений



Поддержка высокопроизводительного доступа к СХД

Поддержка высокой производительности при переходе к конвергентной сети хранения и передачи данных



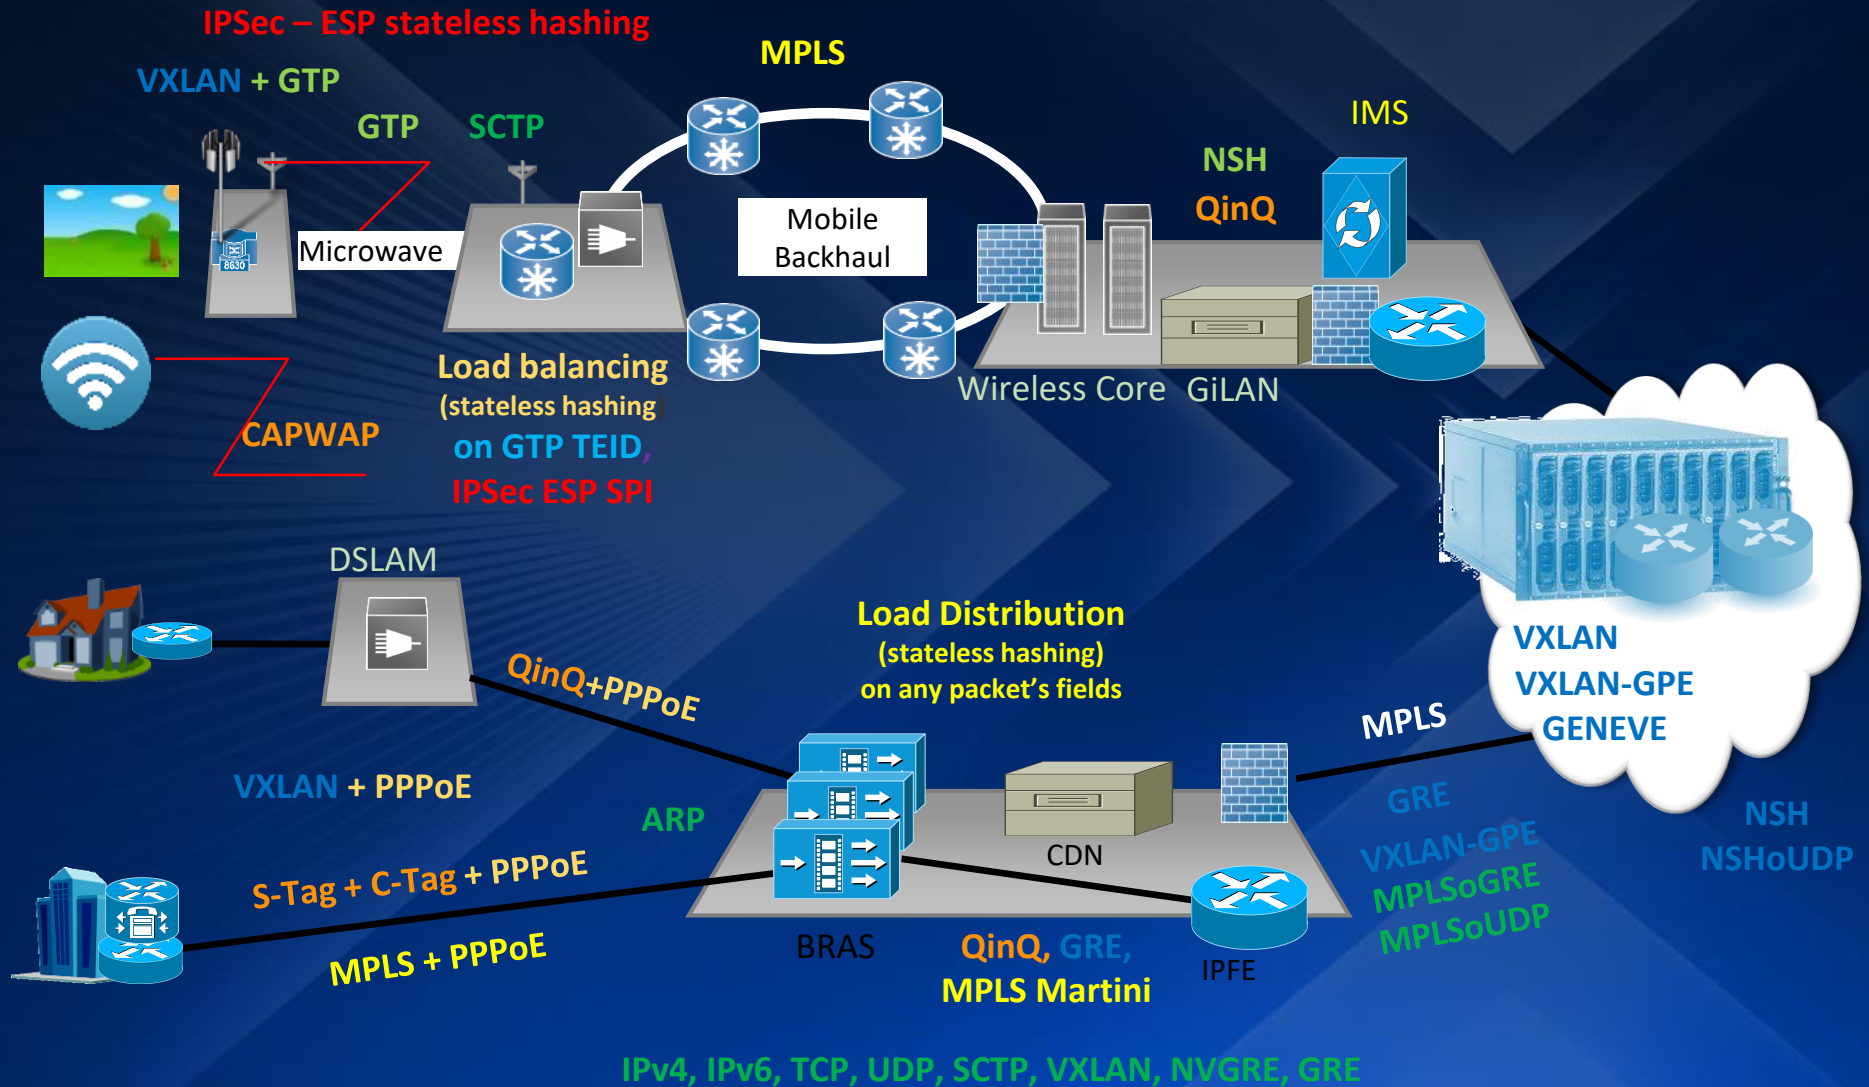
Поддержка множества сетевых протоколов

NV03 VXLAN NVGRE Geneve
C-VLAN S-VLAN Q-in-Q GPT
IPoE L2TP MPLS PPPoE NSH

Программируемый конвейер может адаптироваться к новым протоколам, снижая нагрузку на центральный процессор

Увеличение скорости делает выполнение этих требований более сложной задачей

РАЗНООБРАЗНЫЕ СЕТЕВЫЕ ЛАНДШАФТЫ – МНОЖЕСТВО СЕТЕВЫХ ПРОТОКОЛОВ



РАЗНООБРАЗНЫЕ СЕТЕВЫЕ ЛАНДШАФТЫ – МНОЖЕСТВО СЕТЕВЫХ ПРОТОКОЛОВ

Значительное расширение типов протоколов, с которыми способны работать сетевые адаптеры



Enterprise

Network Virtualization over Layer 3 (NVO3)

Virtual Extensible LAN (VXLAN) [RFC7348]

Generic Protocol Extension for VXLAN (VXLAN-GPE)

Network Virtualization using Generic Routing Encapsulation (NVGRE) [RFC7637]

Generic Network Virtualization Encapsulation (GENEVE)

Network Service Header (NSH)



Service Providers

C-VLAN Tag (C-Tag)

Customer VLAN (C-VLAN)

S-VLAN tag (S-Tag)

Service VLAN (S-VLAN)

Customized Protocols



Network Edge

GPRS Tunneling Protocol (GTP)

Internet Protocol over Ethernet (IPoE)

Layer 2 Tunneling Protocol (L2TP)

Multiprotocol Label Switching (MPLS)

Multi-Service BNG (MS-BNG)

Residential Gateway (RG)

Point to Point Protocol (PPP)

PPP over Ethernet (PPPoE)

Control and Provisioning of Wireless Access Points (CAPWAP)



SECURITY

Internet Protocol Security (IPsec)

Encapsulating Security Payloads (ESP)

Authentication Headers (AH)

Security Associations (SA)

Default Protocol Support + Programmability to Meet Segment Needs

DYNAMIC DEVICE PERSONALIZATION (DDP)

Broadband Remote Access Server (BRAS) Aggregated Forwarding Test

RUN-TIME PROGRAMMABILITY

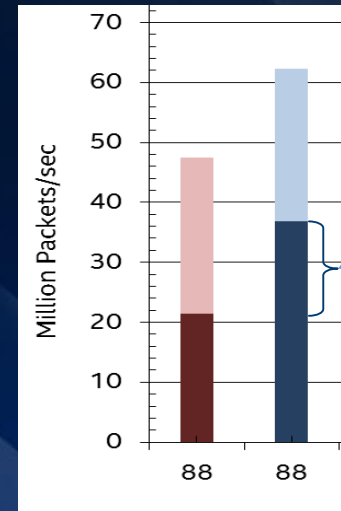
Packet pipeline customization to meet a wide range of customer deployment needs

Available on Intel® Ethernet 700 Series

OPTIMIZE PERFORMANCE

- Lower Latency
- Higher Throughput
- Improved CPU Utilization


Improved Packet Processing Efficiency



With PPPoE DDP Profile
70% MORE
UPSTREAM PROCESSING PERFORMANCE


without DDP Upstream (21.42 Mpps) Downstream (25.90Mpps)
with DDP Upstream (36.98Mpps) Downstream (25.18Mpps)

Unsupported protocols in the pipeline rely on the host to parse them



Dynamic Device Personalization (DDP) profile enabled

The pipeline parser can look deeper into the packets



AVAILABLE NOW

See user guide on the Intel® Developer Zone
<https://software.intel.com/en-us/articles/dynamic-device-personalization-for-intel-ethernet-700-series>

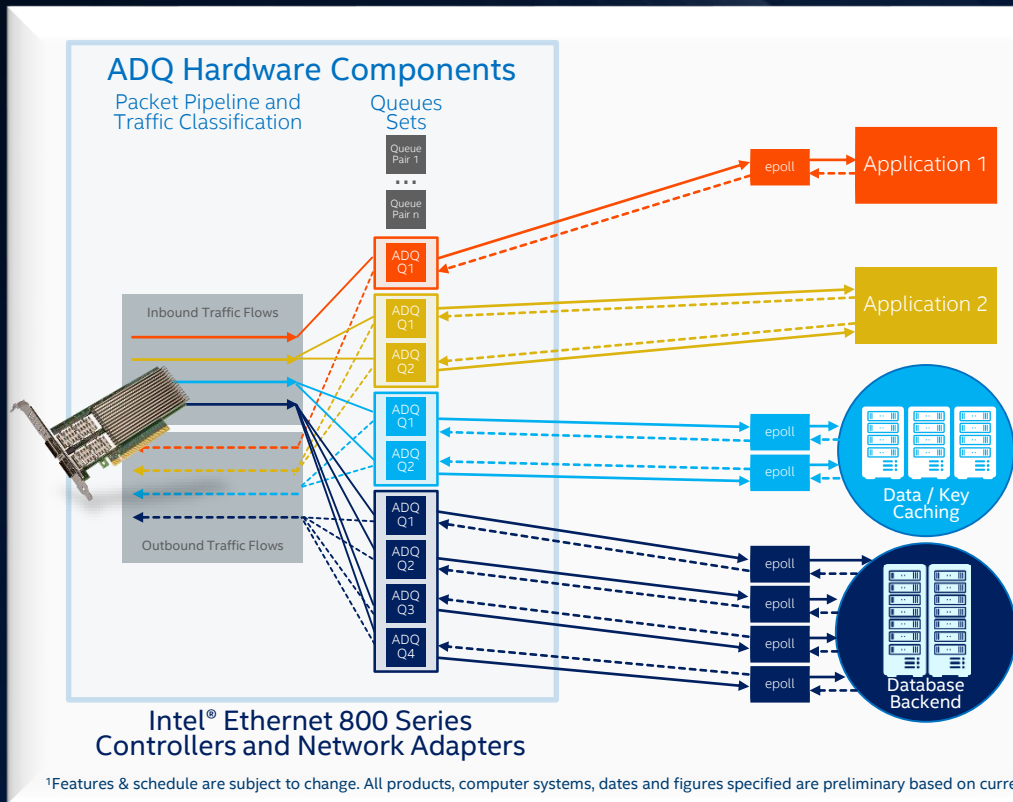
Intel® Ethernet 700 Series DDP Profiles

Publicly Released: GTPv1, PPPoE

Others: MPLSoGRE/MPLSoUDP, L2TPv3, QUIC, IPv4 Multicast, 4G Fronthaul, eCPRI, VXLAN-GPE, IPsec

*Other names and brands may be claimed as the property of others.

APPLICATION DEVICE QUEUES (ADQ)




ADQ Basics

- Выстраивание очередей трафика в зависимости от приоритета приложения
- Упорядочивание потоков данных приложения в соответствии с сценариями построения очередей ADQ
- Дополнительные возможности управления скоростью исходящего (Tx) трафика



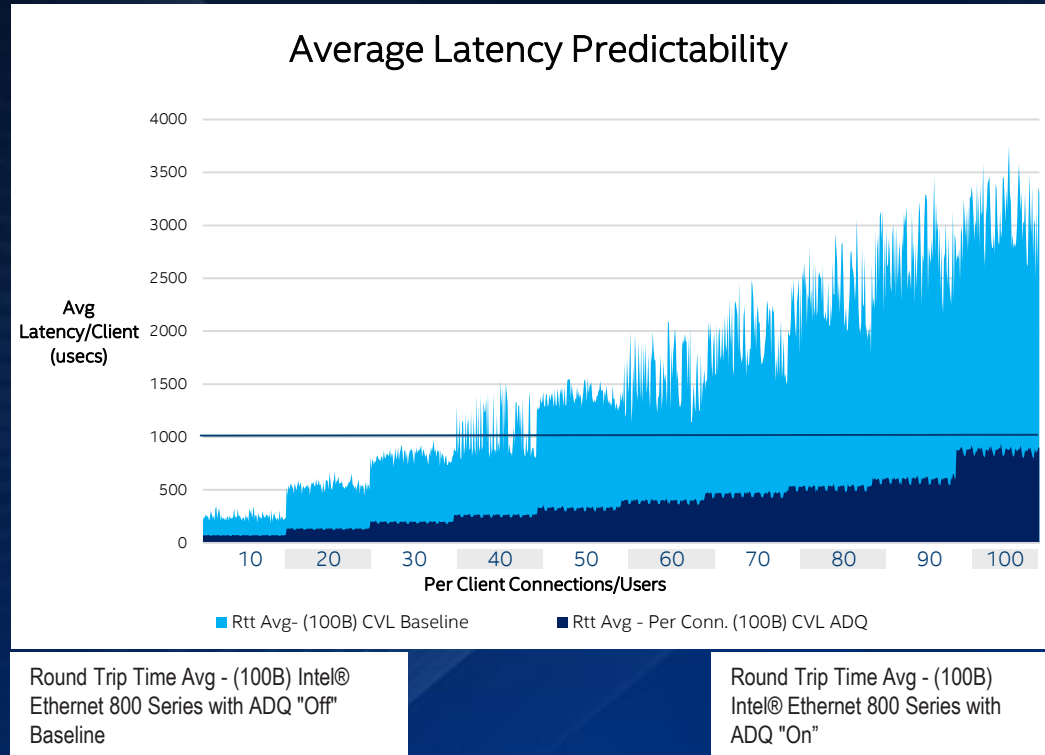
Intel® Ethernet 800 Series with ADQ on Redis* Solutions Brief



Существенное снижение разброса во времени отклика приложения – улучшение Application Latency Predictability

APPLICATION DEVICE QUEUES (ADQ) – REDIS* OPEN SOURCE DATABASE

**INCREASES
APPLICATION
PREDICTABILITY**



**>50%
PREDICTABILITY
INCREASE
WITH OPEN SOURCE REDIS**



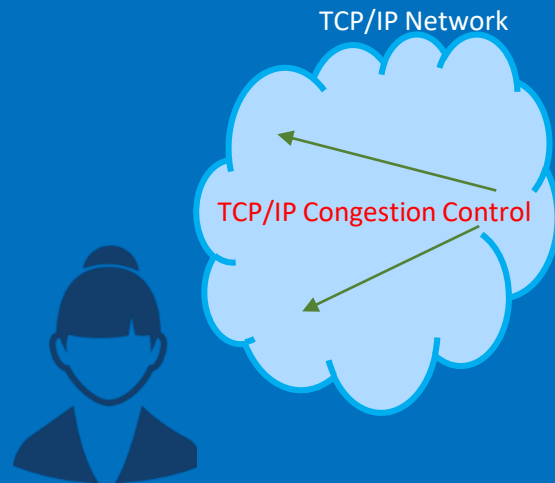
Lower
is Better



Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>. Source: Performance results are based on Intel internal testing as of February 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Tests performed using Redis Open Source on 2nd Generation Intel® Xeon Scalable processors and Intel® Ethernet 800 series 100GbE on Linux 4.19.18 kernel (see backup) Calculation: $(\text{new} - \text{old}) / \text{old} \times 100\%$ for reduction in variance of Standard Deviation of Rtt Average Latency across all runs (10 to 100) for baseline vs ADQ $(229-739)/739 \times 100\% = -69\%$ Reduction in Variance

INTEL® ETHERNET CONTROLLER E810 SERIES - iWARP* AND RoCE*

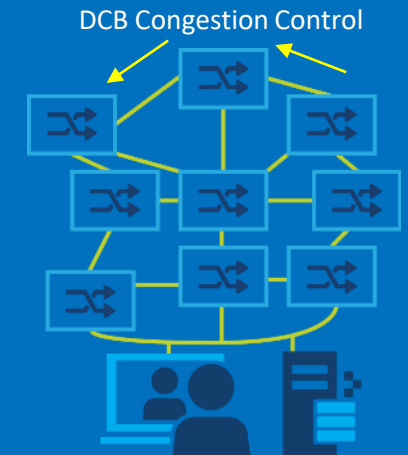
iWARP



iWARP utilizes proven TCP/IP congestion control technology

- Как iWARP*, так и RoCE* v2 потенциально могут предоставлять услуги RDMA при сравнимой стоимости и производительностью.
- iWARP построен на TCP/IP и сегодня является законченным решением, готовым к работе, готовым к масштабированию
- RoCEv2 требует использования коммутаторов с поддержкой Data Centre Bridging, с поддержкой таких сервисов как Enhanced Transmission Services и Priority Flow Control.
- iWARP представляется более простым и удобным с точки зрения настройки и масштабирования

RoCE



Requires fully deployed Data Center Bridging (DCB) technology

INTEL® ETHERNET ТЕХНОЛОГИИ ДЛЯ ОБЛАЧНЫХ РЕШЕНИЙ И ВИРТУАЛИЗАЦИИ

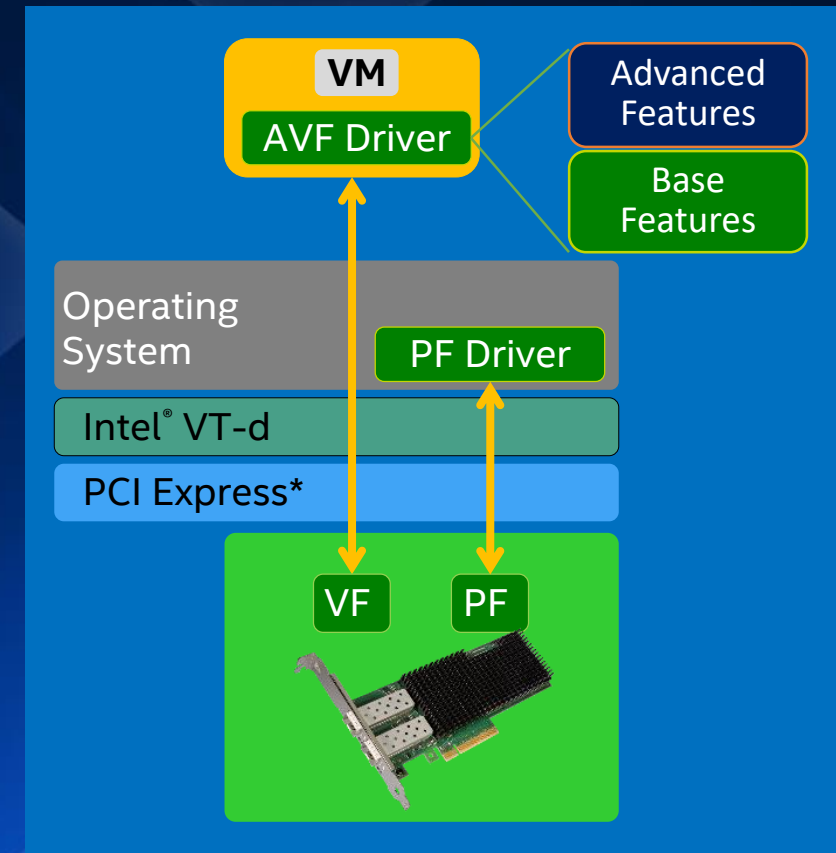
Intel® Ethernet AVF (Adaptive Virtual Function) – единый драйвер виртуальной функции для SR-IOV. При модернизации оборудования сохраняется базовый набор функций

- 4 Queue Pairs (QP) and associated Configuration Status Registers (CSRs) for Tx/Rx
- i40e descriptors and ring format
- Descriptor write-back completion
- 1 control queue, with i40e descriptors, CSRs and ring format
- 5 MSI-X interrupt vectors and corresponding i40e CSRs
- 1 Interrupt Throttle Rate (ITR) index
- 1 Virtual Station Interface (VSI) per VF
- 1 Traffic Class (TC), TCO
- RSS with 64 entry indirection table and key, configured through the PF
- 1 unicast MAC address reserved per VF
- 16 MAC address filters for each VF
- Stateless offloads - non-tunneled checksums
- Large Send Offload (LSO)/ TCP Segmentation Offload (TSO) for buffers up to 256K
- Common device ID
- HW mailbox is used for VF to PF communications (including on Windows)

Расширяемый функционал
Основано на Physical Driver
сетевого контроллера


Intel® Ethernet 700 series
Further product series

Linux Kernel 4.14 in




QUAD PORT ADAPTERS FOR FIBER OPTICS – FULL HEIGHT AND LOW PROFILE

X710-DA4 FH



X710-DA4 LP



Гарантийные обязательства

Intel® Ethernet Products

- Ограниченная пожизненная гарантия на выпускаемую продукцию
- 5-летняя гарантия на продукцию, снятую с производства

• Intel® Points

