

Product Highlights

System Scale and Performance

- Over 27k x 800Gbps Interfaces
- 22 Pbps (44 Pbps FDX) total capacity
- Up to 8 Trillion packets per second
- Up to 24 TB of system-wide buffer

Optimized for AI/ML and HPC Clusters

- Fully scheduled lossless interconnect
- 100% efficient cell based load balancing
- Distributed hardware scheduler
- Per port Virtual Output Queuing to eliminate head of line blocking
- Hardware accelerated link health management
- Under 4 microsecond latency (64 bytes)
- Ultra Ethernet ready

Arista Etherlink for AI

- AI Analyzer powered by AVA
- Advanced RDMA load balancing
- Optimized DCQCN, ECN, PFC congestion management
- AI workflow integration

High Availability

- Hot swappable redundant power
- Field serviceable supervisors
- Integrated over provisioning
- Real-time hardware health checking
- Graceful uplink capacity degradation

Power Efficiency

- High efficiency power and cooling design
- ORv3 54V power support
- High performance SerDes enable
 - 4m passive copper (DAC)
 - Linear-drive pluggable optics (LPO)

Resilient Distributed Control Plane

- Fully distributed architecture
- No single point of failure
- High performance 8-core x86 CPU

Arista Extensible Operating System

- Single 64-bit binary image
- Fine-grained truly modular network OS
- Stateful Fault Containment & Repair
- Full access to Linux shell and tools
- Open APIs and Real-time telemetry
- Extensible platform - bash, python, C++

Overview

The Arista 7700R4 Distributed Etherlink Switch (DES) is the first in a new series of ultra-scalable, smart distributed systems that take the architectural foundations of the 7800R beyond the confines of a chassis to deliver thousands of 800GbE ports and Petabits per second of capacity combined with perfect fairness and lossless packet delivery - imperative for supporting large scale, open, AI and ML applications built on any flavor of XPU.

The Ultra Ethernet ready DES system is designed exclusively to support the needs of leading edge, very large scale Ethernet based generative AI clusters which depend on the interconnect between accelerators for the efficient execution of thousands of simultaneous high bandwidth transactions. Congestion, packet loss and failure can dramatically reduce the efficiency of a workload or potentially lead to a stall in processing, all of which reduce the useful working time and therefore return on investment (ROI) of high value compute farms.

The unique scalability and fully co-ordinated logical single-hop behavior is made possible by an automated and fully hardware accelerated architectural framework. Massively parallel, distributed, end-to-end scheduling combined with deterministic traffic spraying across all available leaf-spine interfaces delivers 100% efficient path utilization and equality to all traffic flows with no tuning required. These are critical attributes for supporting both single large workloads as well as mixed multi-tenant and multi-generational jobs in parallel.

As a distributed system, DES is designed for pay-as-you-grow scaling starting from hundreds of ports, with a predictable, linear CapEx and OpEx trajectory to maximum capacity.

Arista Etherlink for AI adds layers of smart, workload specific functionality improving, traffic management, configuration consistency and visibility - critical factors in building, operating and troubleshooting very large clusters and minimizing downtime.



Arista 7700R Distributed Etherlink Switch

Arista EOS

All Arista products including the 7700R Series run the same Arista EOS software, simplifying network administration with a common standard across all switches. Arista EOS is a modular switch operating system with a unique state sharing architecture that cleanly separates switch state from protocol processing and application logic. Built on top of a standard Linux kernel, all EOS processes run in their own protected memory space and exchange state through an in-memory database. This multi-process state sharing architecture provides the foundation for in-service-software updates and self-healing resiliency together with stateful switchover without the loss of data plane forwarding.

Arista EOS enables advanced monitoring and automation capabilities such as AI Analyzer, LANZ, Zero Touch Provisioning, and enables Linux based tools to be run natively on the switch.

System Overview

The 7700R4 Distributed Etherlink Switch series builds on seven previous generations of Arista 7500 and 7800 series systems, leveraging the fundamental system architecture that is proven by hundreds of thousands of systems in the world's largest and most critical networks and is acknowledged as the most efficient solution for scaling back-end networks for AI. Where the 7500 and 7800 Series offer modularity within a single system, the 7700R series enhances scaling beyond single systems, enabling scalability to thousands of ports.

Separating the traditional line card and fabric functions into discrete devices enables higher scalability through substantially larger topologies than is possible in a single physical chassis while also distributing power consumption, cooling needs and the control plane across multiple devices, and enabling hosts to connect to distributed leaves using cost effective, low power passive copper cabling.

The fully distributed control-plane runs locally on each system and combines with hardware based topology configuration, integrated end to end scheduling, distributed queueing and automatic link health monitoring to minimize deployment time and maximize performance and reliability in service.

The 7700R4 system consists of two major components, which are deployed in similar ways to a typical leaf-spine configuration that creates single system topologies that scale to over 27,000 x 800GbE or over 31,000 x 400GbE host ports.

- **7700R4C-38PE** - 38 Port 800GbE leaf switch with 18 host ports and 20 fabric ports
- **7720R4-128PE** - 128 Port 800GbE spine switch with 128 fabric ports

Arista 7700R4 systems can be combined in various topologies to provide a choice of density, scale and speed, with linear CapEx and OpEx. 7700R4 Series systems are designed with maximum efficiency in mind, based on an advanced 5nm silicon geometry combined with optimized power and cooling design. Support for passive copper cables up to 4m and Linear-drive Pluggable Optics dramatically reduces the power and cooling footprint of very large scale networks.



7700R4C-38PE: 38 port 800GbE Distributed Leaf Switch

- 18 x 800GbE host facing ports
- 20 x 800GbE spine facing fabric ports
- 14.4 Tbps of forwarding and 5.4 Bpps with 16GB of buffer
- Field serviceable 8-core x86 supervisor



7700R4-128PE: 128 port 800GbE Distributed Spine Switch

- 128 x 800GbE leaf facing fabric ports
- 102.4 Tbps non-blocking fabric bandwidth
- Field serviceable 8-core x86 supervisor



Rear view of 7700 Series systems

- 7700R4C-38PE is shown with F-R airflow and PWR-2411-MC
- 7720R4-128PE is shown with F-R airflow and PWR-2421-HV

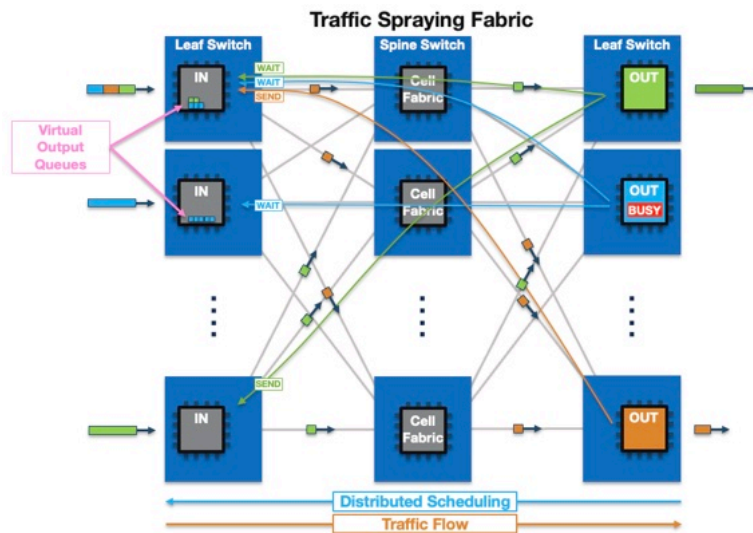


7700 DES Fully Scheduled Lossless Network Performance

The Arista 7700R Series Distributed Etherlink Switch is optimized to deliver massively parallel end to end scheduling and coordination - even in the largest topologies, the platform operates like a single extremely large switching chip that is 100% internally lossless and fair. Four advanced technologies converge to this highly desirable architecture:

- **Cell-based traffic spraying fabric** - elegantly mitigates elephant and mice flows. Uniform traffic spraying for 100% efficiency
- **Virtual Output Queues (VOQ)** - invokes ingress virtual queues to every egress port, eliminating head of line blocking (HOLB)
- **Distributed Credit Scheduling** - ensures all egress ports are independent eliminating HOLB and noisy neighbors
- **Deep Buffering** - easily handles in-cast, bursts and speed mismatches without packet loss, keeping TCP running efficiently

This advanced architecture enables the Arista 7700R4 to handle the most demanding cluster computing workloads with ease. Generative AI clusters, ML and HPC environments all benefit from the 7700R4 Series' ability to linearly scale bisectional bandwidth and logical radix while handling high bandwidth low entropy traffic and mixed traffic loads without increasing latency or introducing unnatural congestion.



7700R4 Series Fully Scheduled Architecture

7700 DES Topology and Workload Flexibility

AI training workloads typically involve synchronized, high throughput traffic flows that require careful network design to maximize XPU efficiency and performance. Building a non-blocking, fully-connected fabric is a common starting point but optimal design depends on a variety of factors. Traffic patterns and their appropriate topologies are deployment specific and may require specific connectivity plans which require specialized physical topologies. Planar (rail based) topologies, for example, require a unique physical cabling plan between accelerators and leaf switches which makes running mixed workloads on the same cluster more complicated. Factors that influence network design may include:

- The number, type and behavior of workload(s) and model(s) in use
- How workloads are sub-divided (e.g. sharding)
- Number of accelerators per workload and how they are distributed across the cluster
- The selection of accelerator processor architecture and the requirements for network connectivity
- Optimizations provided by the Collective Communications Libraries (CCL) for each workload
- Levels of resilience / redundancy required and tolerance to reduced performance

DES' innovative logical single-hop architecture, robust traffic management and inherently fair forwarding paradigm eliminates the complexity of workload-specific architectures and tuning, abstracting logical and physical topologies while enabling multiple unique workloads in parallel. In many cases, DES removes the need for unique, job specific, cabling plans, as every accelerator is logically connected to the same single distributed switch. This allows the use of low cost, short reach cables and optics, saving significant expenditure and power consumption on cross-facility optical interconnections and eliminating the need to re-cable for different workload topologies.

Deploying the Distributed Etherlink Switch


A typical deployment of the 7700R4 Distributed Etherlink Switch follows the traditional, well understood and widely implemented, 2-tier leaf-spine physical topology model. Each leaf 7700R4 switch provides 18 x 800GbE ports for host connectivity and 20 x 800G fabric ports which are designed to be distributed evenly across a set of 7720R4 spines. Consistent with other leaf-spine designs, spines only connect to leaf switches, not to each other or to other networks.

In a simple 2-tier configuration, DES scales to support over 4600 x 800GbE hosts (4x the size of the largest 7816R4 modular system), in a single logical system with single-hop logical forwarding. By introducing an orthogonal super-spine layer, DES can continue to expand to support tens of thousands of 400GbE or 800GbE connected accelerators.

DES Leaf switches are built with an inherent over-provisioning ratio (18 host ports : 20 spine ports) which serves as an integrated redundancy mechanism to avoid congestion in the event of either an uplink or spine failure, allowing workloads to continue unaffected. It is typical to connect up to 16 accelerators to each leaf switch with the remaining two ports per leaf being used for storage nodes or inter-cluster connectivity.

Starting with the total number of accelerator ports required, the number of leaf switches can be determined and subsequently the number of spine switches necessary to support equal connectivity to all leaf devices. In typical deployments, each leaf would be connected to each spine with total bandwidth of 1.6T (i.e. 2 ports), 800G or 400G (i.e. a 50:50 breakout of a port across two spines).

Despite physically resembling a two or three tier leaf-spine network, DES deployed at any scale continues to provide fully co-ordinated, logical single-hop forwarding, eliminating the traditional congestion management protocols and tuning required by a leaf-spine constructed of independent devices. With fully automated fabric configuration and operation, overheads and time to operation are minimized while providing deterministic performance for any workload.



7700R4 Series Typical Deployment Topology

Scaling Examples	1152 x 800G Hosts	2304 x 800G Hosts	4608 x 800G Hosts	Over 4608 x 800G Hosts
Number of Spine Nodes	10	20	40	Up to 800 inc. Super Spine
Number of Leaf Nodes	64	128	256	Over 1500
Leaf-Spine Interconnect	16 Tbps (1600G per spine)	16 Tbps (800G per Spine)	16 Tbps (400G per Spine)	16 Tbps (400G per Spine)

Ideal Architecture for AI/ML and HPC Networks

Maximizing the ROI of high value Generative AI clusters requires a different approach to infrastructure than a typical data center. Back-end training networks need to be highly performant, delivering lossless and deterministic connectivity under 100% load conditions in order to minimize the completion times of each of the hundreds of thousands of inter-accelerator transactions that occur through the life of a training job. At high scale and for long running jobs, hardware and software quality and reliability are critical to ensuring jobs reach completion in one pass and expensive accelerators are not under utilized and jobs do not need to be repeated.

The AI ecosystem is innovating rapidly, with continuous development in accelerator hardware, supporting libraries and underlying communications topologies. The 7700 Series is accelerator agnostic and Ultra Ethernet ready, enabling innovation in the compute layer, while providing unparalleled performance with its unique fully scheduled cell based architecture and industry leading hardware and software quality.

CloudVision® for Accelerated Computing

CloudVision is a network-wide approach for workload orchestration and workflow automation as a turnkey solution for Accelerated Computing. CloudVision extends the EOS publish subscribe architectural approach across the network for state, topology, monitoring and visibility.

In large scale high performance environments, such as back-end networks for AI/ML training, it is critical to apply best-practise tuning parameters consistently across all devices and every interface to maximise performance. Due to the large number of devices and interfaces involved, this must be achieved in a fully automated fashion to avoid human-error.

CloudVision combined with Arista Validated Designs (AVD) enables a template driven automated common configuration model to be deployed deterministically across all network elements, implementing best-practise configuration parameters with minimal user input. When combined with the Arista AI Agent for compute hosts, configuration consistency and visibility is extended into the compute platform improving cluster deployment time, operational stability and end to end telemetry

Precise Data Analysis

Arista's Streaming Telemetry, Latency Analyzer (LANZ), AI Analyzer and Precision Data Analyzer (DANZ) are integrated features of EOS. Together these capabilities provide a complete solution to the monitoring and visibility challenges at speeds up to 800Gbps giving IT operations the ability to proactively deliver feedback on congestion events, filter, replicate, aggregate and capture traffic without affecting production performance. EOS monitoring features include both event triggered monitoring for real-time micro-burst and congestion tracking as well as high rate counter polling down to 1ms granularity.

Model Comparison	7700R4C-38PE	7720R4-128PE
800G Host Ports	18 x OSFP800	—
800G Fabric Ports	20 x OSFP800	128 x OSFP800
Additional Ports	1 x QSFP100	—
Throughput (FDX)	14.4 (28.8) Tbps	102.4 (204.8) Tbps
Packets/Second	5.4 Bpps	—
End to End Latency		From 3.8 us
CPU	Eight-Core x86	Eight-Core x86
System Memory	64 Gigabytes	64 Gigabytes
Packet Buffer Memory	16 GB	—
Supervisor Module	7001-SUP-EL	7003-SUP-EL
USB and RS232 Ports		1 USB, 1 RS232 (RJ45)
100/1000 Mgmt Ports	1	2
Hot-swap Power Supplies	2 (1+1 redundant)	4 (N+N redundant)
Hot-swappable Fans	4 (N+1 redundant)	12 (N+1 redundant)
Airflow Direction (Fan Tray)		Front to Rear
Rack Units	2 U	7 U
Size (WxHxD) inc. handles	17.32 x 3.46 x 23.06 in (44 x 8.8 x 58.6 cm)	17.32 x 12.11 x 22.06 in (44 x 30.8 x 56 cm)
Typical/Max Power Draw ¹	593 W / 1840 W	1032 W / 3848 W
Weight	47.2 lbs (21.4 kg)	110.3 lbs (50 kg)
Fan Tray		FAN-7021H
Power Supplies		PWR-2421-HV or PWR-2411-MC
EOS Feature Licenses	LIC-DL	—
Minimum EOS		TBD

1. Typical power consumption measured at 25C ambient with 50% load on all ports. Excludes optics power as this is a significant variable.

Supervisor Module	7001-SUP-EL	7003-SUP-EL
Processor	Eight-Core x86	Eight-Core x86
System Memory	64 GB	64 GB
SSD Storage	512 GB	512 GB
Default Secure Boot	Yes	Yes
Size (WxHxD)	7.05 x 3.27 x 12.57 in (17.91 x 8.31 x 31.93 cm)	7.05 x 1.16 x 10.29 in (17.91 x 2.96 x 26.13 cm)
Weight	3.35 lbs (1.52 kg)	3 lbs (1.36 kg)
Typical/Max Power ¹	TBD	TBD
System Compatibility	DCS-DL-7700	DCS-DS-7720

Power Supply Specifications

Power Supply	PWR-2421 HV	PWR-2411 MC
Input Voltage	200-277VAC 240-380VDC	48 to 72 VDC
Max Input Current	13.5 A (200V AC)	70 - 40 A Max (43 - 72 V)
Input Frequency	50/60 Hz or DC	DC
Output Power	2400 W	2400 W
Input Connector	SAF-D	6-Pin M-CRPS
Efficiency (Typical)	96%	96%

Standards Compliance

EMC	FCC Class A, ICES-003, EN 55032, EN IEC 61000-3-2:2019, EN 61000-3-3
Immunity	EN 55035 EN 300 386
Safety	EN 62368-1:2014 + A11:2017 IEC 62368-1:2014
Certifications	BSMI (Taiwan) CE (European Union) KCC (South Korea) NRTL (North America) RCM (Australia/New Zealand) UKCA (United Kingdom) VCCI (Japan)
European Union Directives	2014/35/EU Low Voltage Directive 2014/30/EU EMC Directive 2012/19/EU WEEE Directive 2011/65/EU RoHS Directive 2015/863/EU Commission Delegated Directive

Environmental Characteristics

Operating Temperature	0 to 40°C (32 to 104°F)
Storage Temperature	-40 to 70°C (-40 to 158°F)
Relative Humidity	5 to 95%
Operating Altitude	0 to 10,000 ft, (0-3,000m)

Further Information [Product Certification Portal](#)

¹. Typical power consumption measured at 25C ambient with 50% load on all ports. Excludes optics power as this is a significant variable.

Systems and Optional Components

DCS-DL-7700R4C-38PE-F	7700R4 Series Distributed Leaf Switch with 18 x 800G OSFP ports and 20 x 800G OSFP fabric ports, Sup-EL, front to rear air, 2 x AC and 2 x SAF-D cords
DCS-DL-7700R4C-38PE#	7700R4 Series Distributed Leaf Switch with 18 x 800G OSFP ports and 20 x 800G OSFP fabric ports, Sup-EL, configurable fans and PSU
DCS-DS-7720R4-128PE-F	7720R4 Series Distributed Spine Switch Bundle with 128 x 800G OSFP fabric ports, Sup-EL, front to rear air 4 x AC and 4 x SAF-D cords
DCS-DS-7720R4-128PE#	7720R4 Series Distributed Spine Switch Bundle with 128 x 800G OSFP fabric ports, Sup-EL, configurable fans and PSU
DCS-7001-SUP-EL	7001 series Supervisor-EL module with 8c/16t 2.3 GHz CPU, 64GB RAM, 512GB NVMe (spare, -F airflow, secure boot enabled)
DCS-7003-SUP-EL	7003 series Supervisor-EL module with 8c/16t 2.3 GHz CPU, 64GB RAM, 512GB NVMe (spare, -F airflow, secure boot enabled)
FAN-7021H-RED	Spare Front-to-Rear High Speed Fan module for select Arista 7000 Series. Refer to product datasheets for compatibility.
FAN-7021H-RED#	Configurable Front-to-Rear High Speed Fan module for select Arista 7000 Series. Refer to product datasheets for compatibility.
FAN-7021H-BLUE	Spare Rear-to-Front High Speed Fan module for select Arista 7000 Series. Refer to product datasheets for compatibility.
FAN-7021H-BLUE#	Configurable Rear-to-Front High Speed Fan module for select Arista 7000 Series. Refer to product datasheets for compatibility.
PWR-2411-HV-RED	Spare Arista PSU, 1RU, HVAC/DC, 2400W, SAF-D, FORWARD, 73.5MM
PWR-2411-HV-RED#	Configurable Arista PSU, 1RU, HVAC/DC, 2400W, SAF-D, FORWARD, 73.5MM
PWR-2411-MC-RED	Spare Arista PSU, 1RU, +48VDC, 2400W, 6-PIN Modular Connector, FORWARD, 73.5MM
PWR-2411-MC-RED#	Configurable Arista PSU, 1RU, +48VDC, 2400W, 6-PIN Modular Connector, FORWARD, 73.5MM

Warranty

The Arista 7700 Series switches come with a one-year limited hardware warranty, which covers parts, repair, or replacement with a 10 business day turn-around after the unit is received.

Service and Support

Support services including next business day and 4-hour advance hardware replacement are available. For service depot locations, please see: <http://www.arista.com/en/service>

Headquarters

5453 Great America Parkway
Santa Clara, California 95054
408-547-5500

Support

support@arista.com
408-547-5502
866-476-0000

Sales

sales@arista.com
408-547-5501
866-497-0000