ARISTA

reinventing data center switching

Arista Data Center Portfolio

vEOS	7048	7100 ‘S’	7100 ‘T’
Manages VMware VSwitches	48-port GigE Data Center Switch	24/48 port 1/10Gb SFP+ Low Latency Data Center Switches	24/48-port 1/10GBASE-T Cat5e/6a Data Center Switches

Ultra-low Latency Switching

2009

Received Numerous Industry Awards

NETWORKWORLD® CLEAR CHOICE

HPC wire Readers' Choice Awards 2009

LIGHT READING

TiE 50 2009 TOP STARTUPS WINNER

Arista Data Center Portfolio

Extensible Operating System (EOS)

vEOS

Manages VMware Virtual Switches

Provides visibility and consistency between physical, virtual, and cloud networks

7048

48-port Data Center Class Gigabit Ethernet Switch

Advanced congestion control and avoidance for bursty and mission-critical traffic

7100 'T'

Dense 24/48-port 1/10GBASE-T Data Center Switches

Prepares the network for future 10Gb deployments with multi-speed 10Gb on existing cable plants

7100 'S'

Ultra Low Latency 24/48-port 1/10Gb SFP+ Leaf Switches

Best possible performance for latency sensitive HPC and Trading Applications

7500

Lossless, High Density, Modular Switching System supporting up to 384 Wirespeed 10Gb Ports

10 Terabit Switching capacity for 40GbE and 100GbE. Designed for high-capacity network aggregation with deep buffering

Network World's Clear Choice Winner

NETRESULTS

Product	Arista DCS-7124S	Cisco Nexus 5010	Summit X650-24x	HP ProCurve 6600-2
Vendor	Arista Networks	Cisco	Extreme Networks	HP ProCurve Networking
Price	$26,080	$67,030	$46,665	$63,594
Pros	Very low latency and jitter; extensible Linux operating system; standards-based.	Full Fibre Channel/FCoE support; extensive virtualization features.	Highest IGMP group capacity.	Largest MAC address capac
Cons	Relatively low MAC address capacity.	High latency; no Layer-3 support; some leakage in multicast tests.	Relatively high power consumption; uneven distribution in some link aggregation tests.	Low unicast and multicast throughput; relatively high late frames forwarded out of sequ
Score	4.29	3.68	3.7	3.36

Outperformed Cisco, Extreme, HP, Dell and others

Network World's Clear Choice Review

"Arista's EOS also runs on Linux, and does more than any other switch tested to make Linux features available to users. The command set also allows network managers to drop into a Bash shell and run virtually any Linux command – including applying bug fixes without a reboot, a unique feature in this test"

"Arista's 7124S was more consistent across the board, with the least variation between average and maximum join and leave times. This is largely a function of control-plane processing power, and reflects Arista's use of a dual-core 1.8-GHz x86 CPU, a powerful processor for a top-of-rack switch"

- David Newman, Network Test

Latency: the Race to Zero

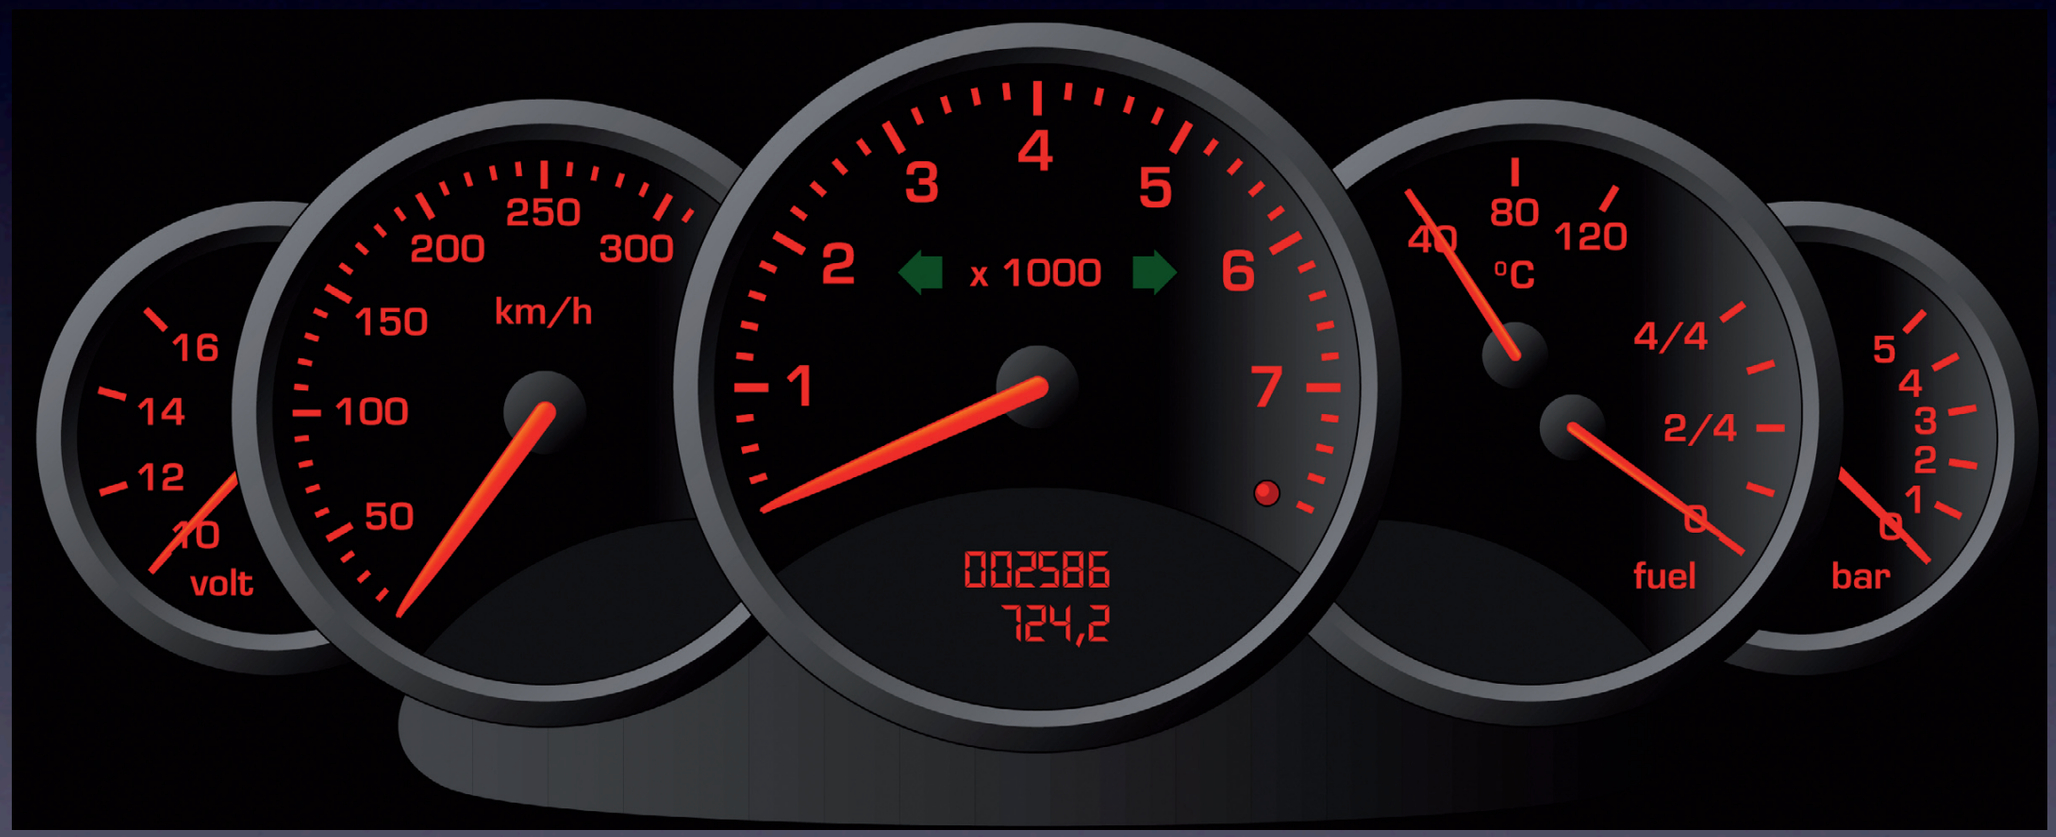

Every Microsecond Counts™

High Frequency Trading

Network Requirements:

1. Lowest Latency
2. Highest Throughput
3. Unbreakable Software

Arista's Switches Leading the Industry

World Record Performance

LLM over RDMA latency				
Rate [msgs/sec]	median [usec]	average [usec]	max [usec]	std [usec]
1,000	5.00	5.50	20.00	1.30
10,000	4.00	4.50	15.50	0.90
50,000	4.00	4.50	26.00	1.10
100,000	4.00	4.50	41.00	2.10
250,000	4.50	5.00	88.00	7.30
500,000	4.50	6.00	80.00	7.50
1,000,000	6.00	7.50	94.00	9.60

IBM x3650, Intel X5570 CPU, Mellanox ConnectX
IBM Websphere LLM, OFED 1.5, Arista 7124S

Electronic Trading Sample Flow

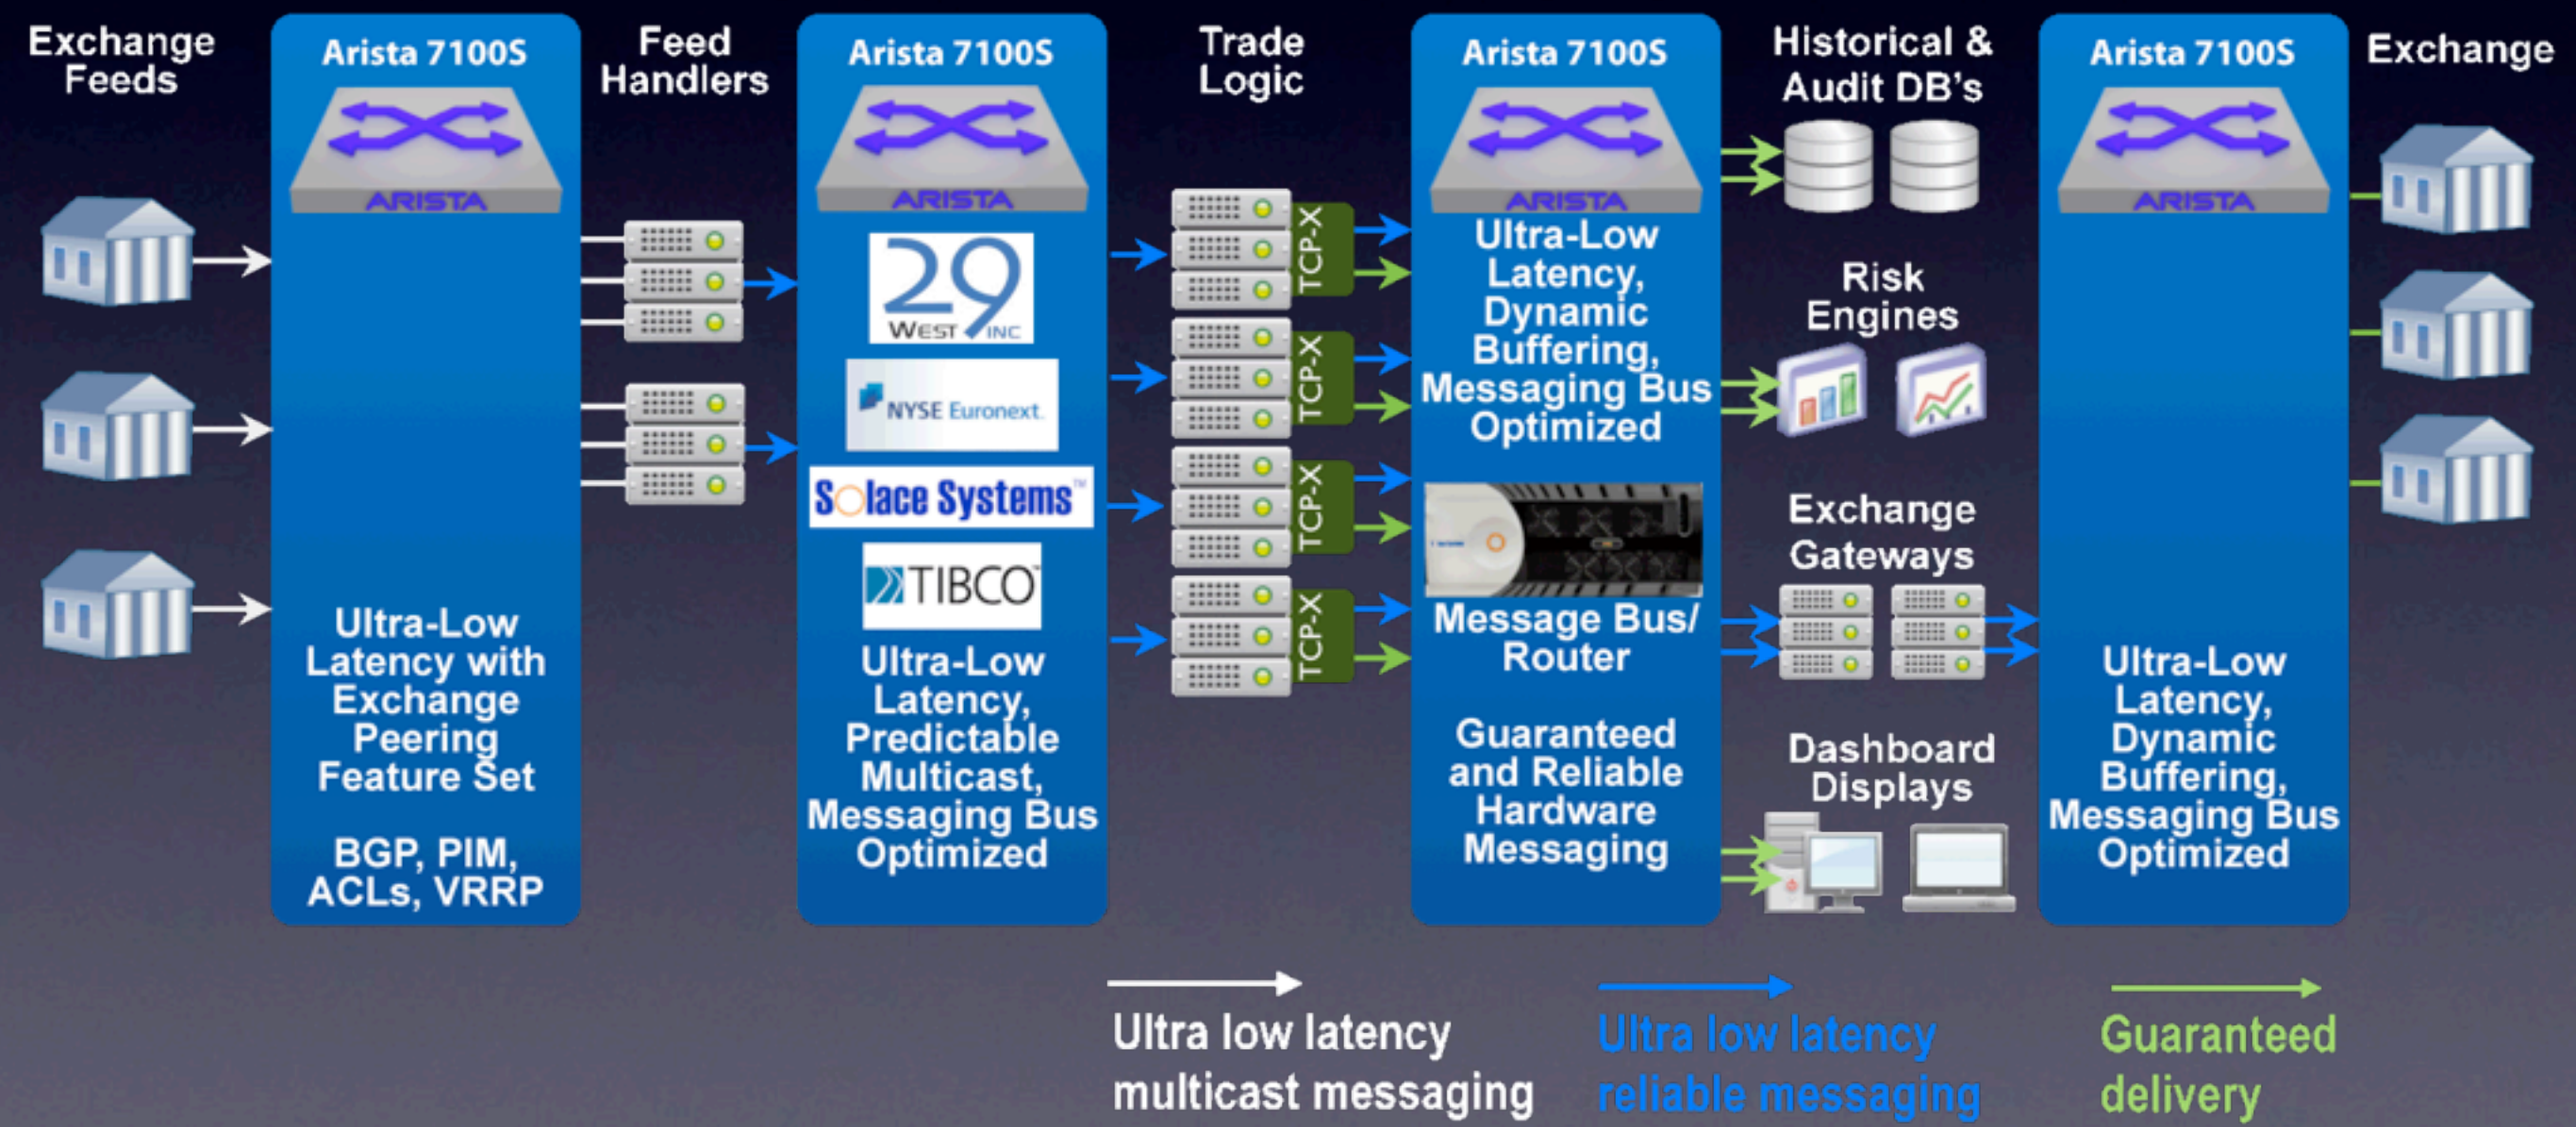

Real World Transaction Volume Tests

100% hardware data path for market data distribution

- **1 million msgs/sec**
  24 µsecs mean latency
  27 µsecs 99.9th percentile
- **2 million msgs/sec**
  25 µsecs mean latency
  30 µsecs 99.9th percentile
- **4 million msgs/sec**
  31 µsecs mean latency
  39 µsecs 99.9th percentile

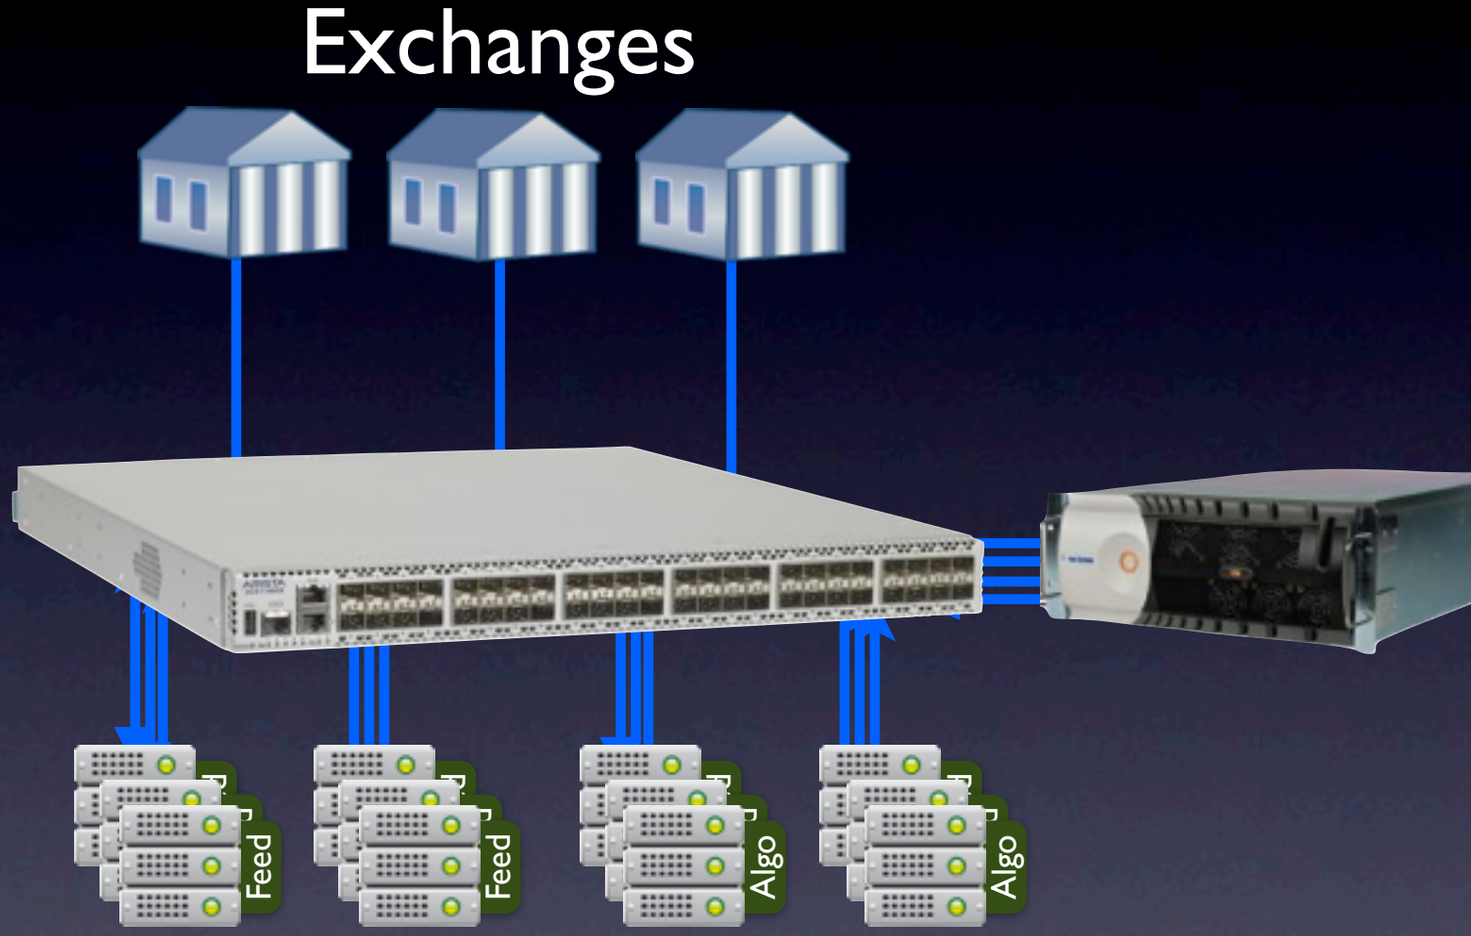

Less than 7µsecs end to end mean latency variance from 1M to 4M messages per second

Arista 7500 Switch

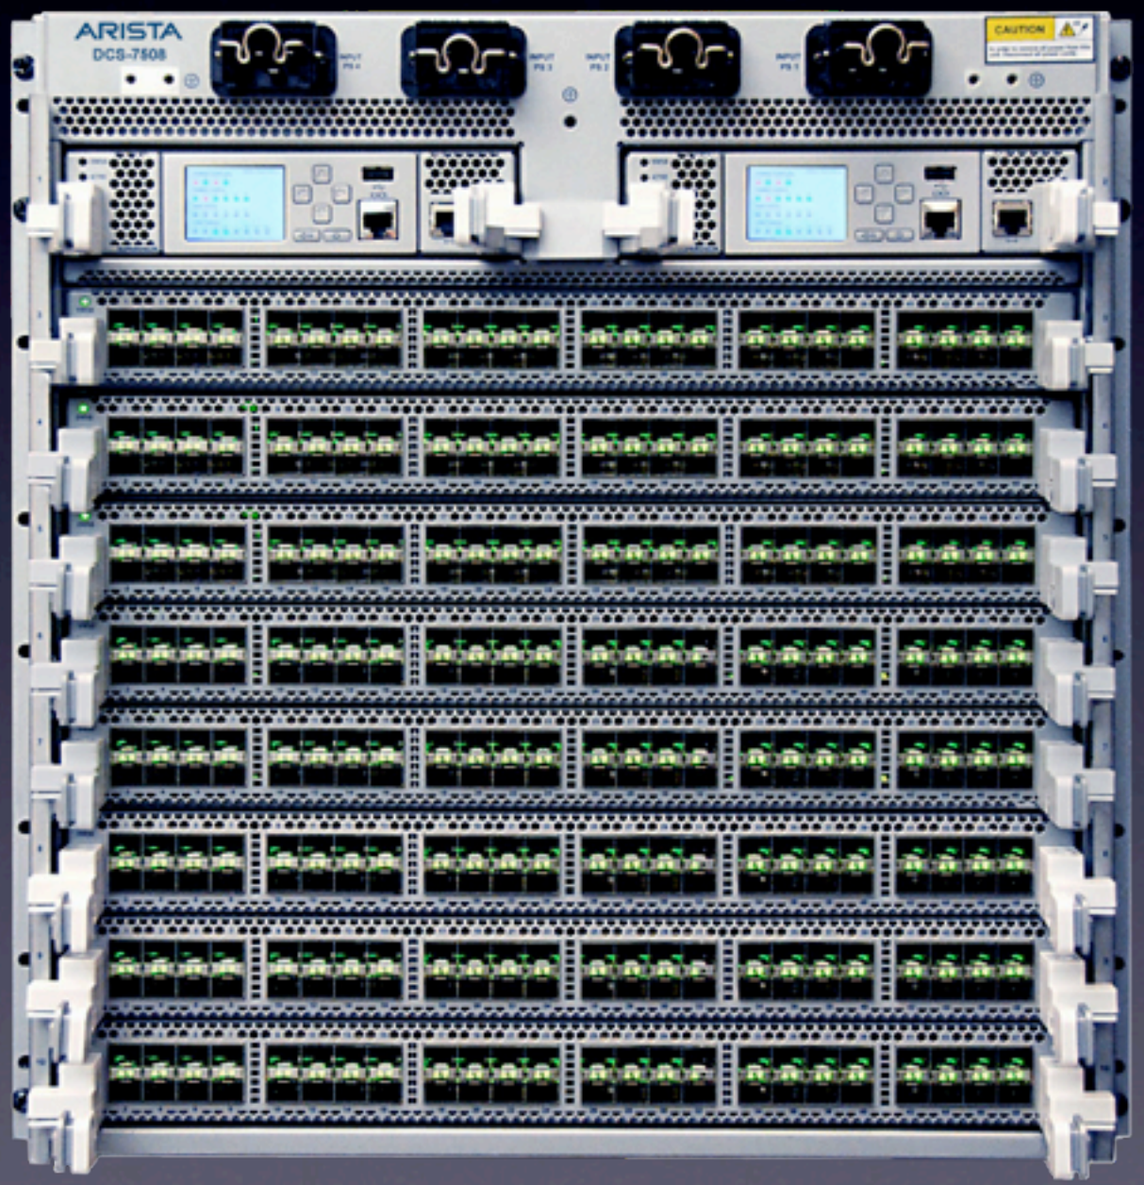

World's fastest 10G Switch

10 Terabits/sec Non-blocking Fabric

384 Wirespeed 10G ports per chassis

5.7 BPPS (Billion Packets per Second)

3.5 usec Latency (64 Byte packets)

3-5X the throughput of other products

World's greenest 10G Switch

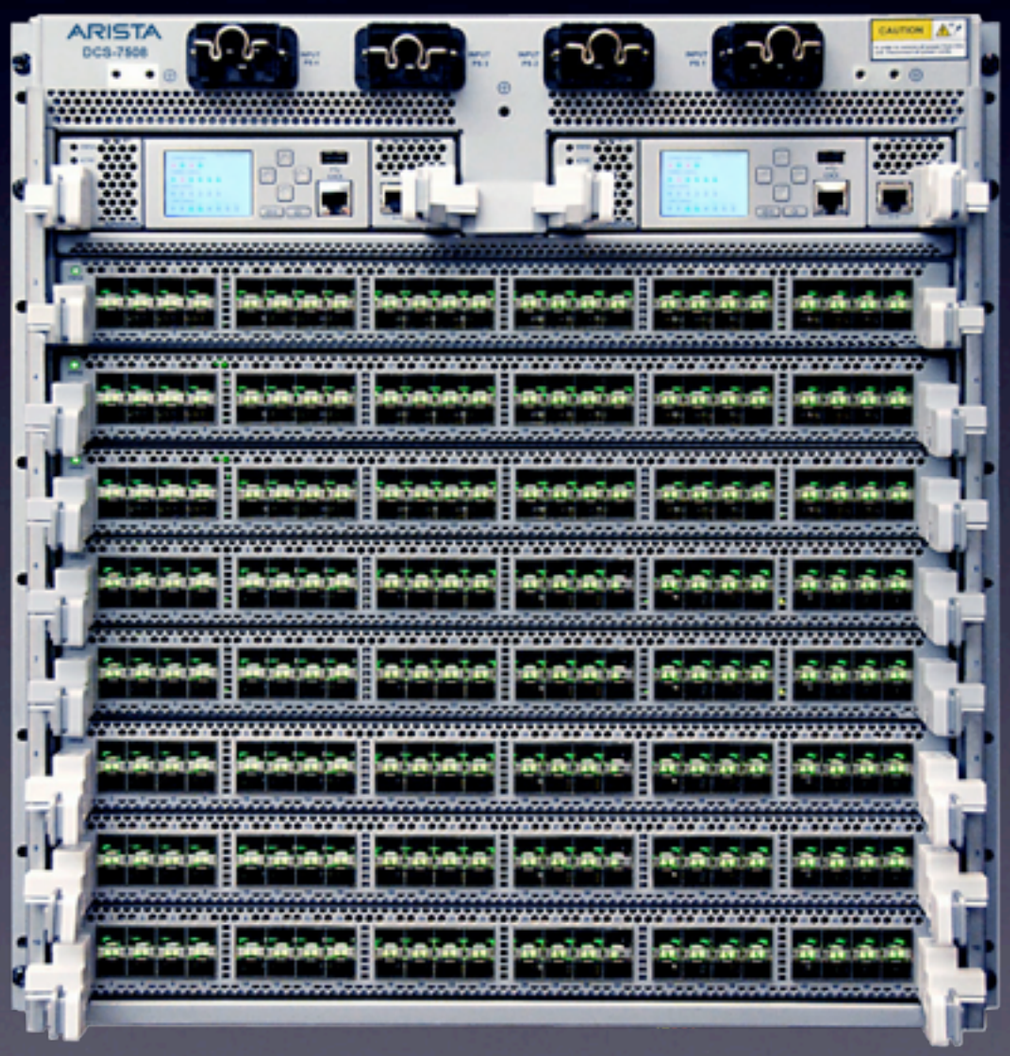

10W typical power/port

3.8KW for 384 port chassis

5-10X greener than others

Power per wirespeed 10G Port

10G Switch	10G Ports	Watts/Port	OPEX*
Juniper T1600	80	105W	$126,000
Juniper MX960	48	106W	$127,200
Cisco 12816	64	156W	$187,200
Cisco Nexus 7010	64	139W	$166,800
Force10 e1200	54	106W	$127,200
Force10 X-series	140	50W	$60,000
Foundry MLX-32	128	84W	$100,000
Arista 7500	**384**	**13.2W**	**$15,840**

*Annual Operating Cost Per 1000 10G ports, 12c/KWh

World's densest 10G Switch

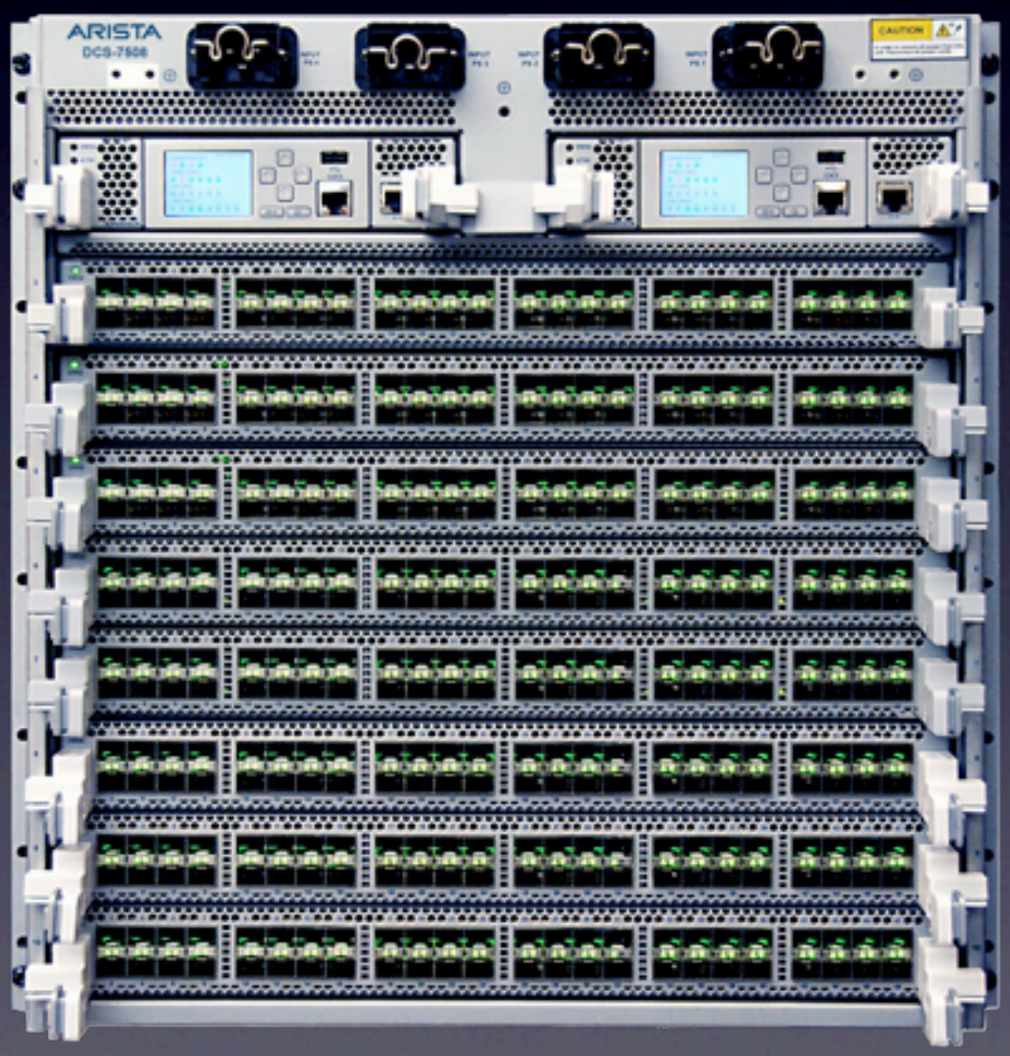

384 ports per 11U chassis

1536 ports per 44U rack

5-10X greener than others

Designed for the Data Center

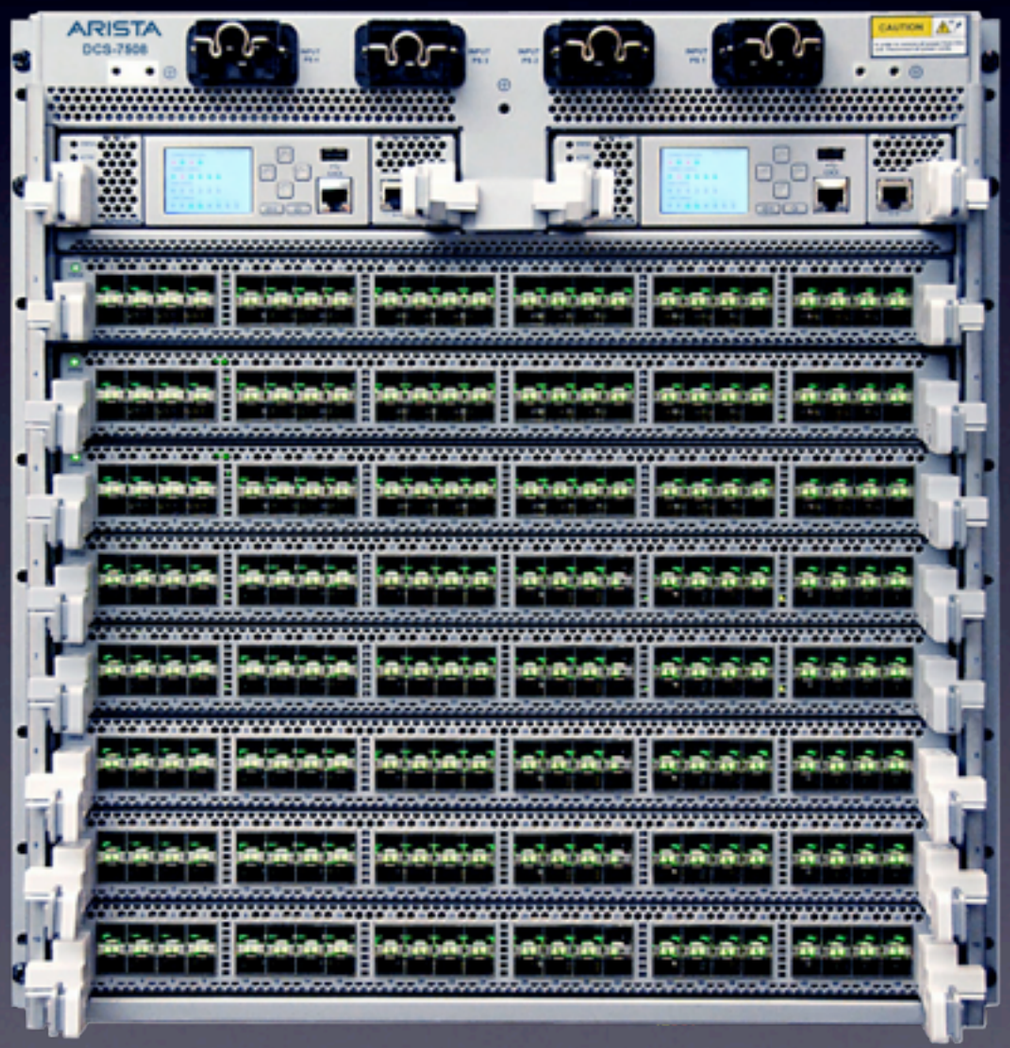

True front-to-rear cooling

All Connectors Face Front

Cordless Power Supplies

40 and 100 Gigabit Ready

Long Roadmap Ahead

Advanced Architecture

VOQ Architecture eliminates Head-of-line blocking

Non-blocking Fabric ready for 40 and 100 GbE

Large packet buffers avoid packet drops under load

Distributed Packet Scheduler Scales Performance

Low Latency Fair Access to Fabric Bandwidth

Arista 7500 Summary

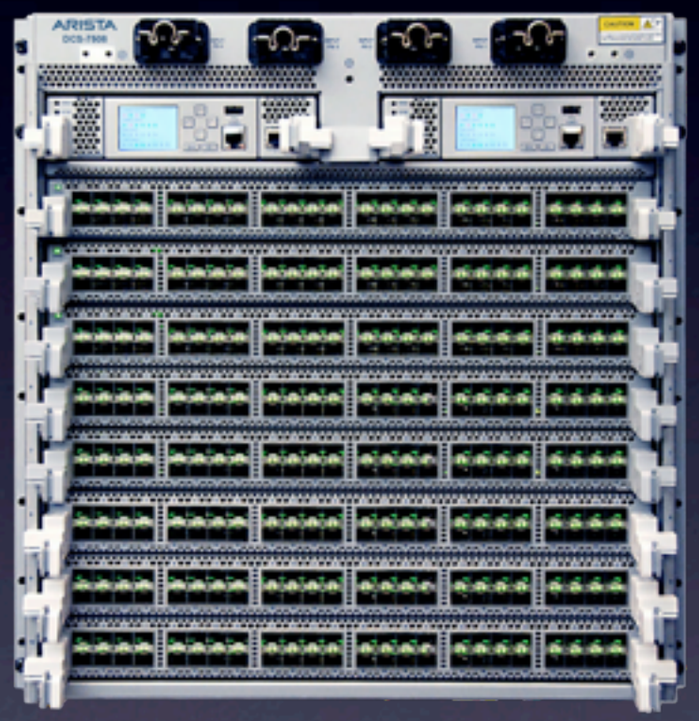

Outstanding Throughput

Outstanding Power Efficiency

Outstanding Cost-performance

Designed for large-scale data centers

Scalable Networks for Cloud Data Centers

Scalable Network Design

- Dual Core Switch
  - 768 10G ports
  - Wire-speed Performance
- Leaf Switch
  - 48 1G or 10G ports
  - 4 or more 10G uplinks
- Overall Capacity
  - 10,000+ ports
  - 10 + Tbps throughput

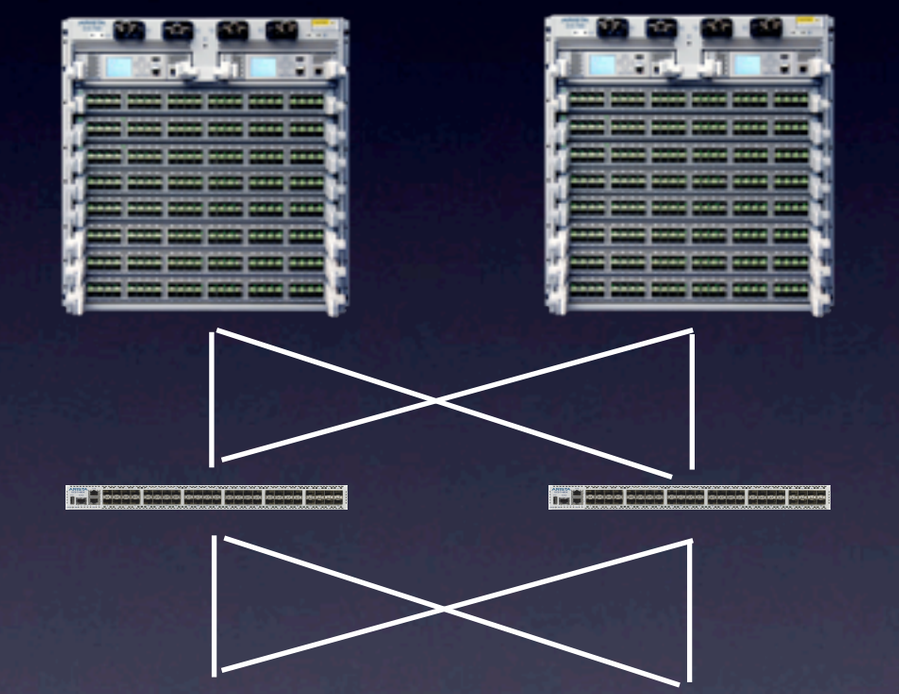

Cost-effective Scalability

- Large-scale Data Centers require large-scale networks
- Wirespeed fabrics key to minimize latency
- Low latency key to avoid locality/placement issues
- Flat L2 networks key for large-scale VMotion
- Legacy Network cost prohibitive for this solution

MLAG: Multi-Chassis LAG

- Active/Active Load Sharing
- Uses standard LAG protocol
- No proprietary Lock-in
- Works with any Server O/S
- Enables L2 Multipathing
- Automatic Failover

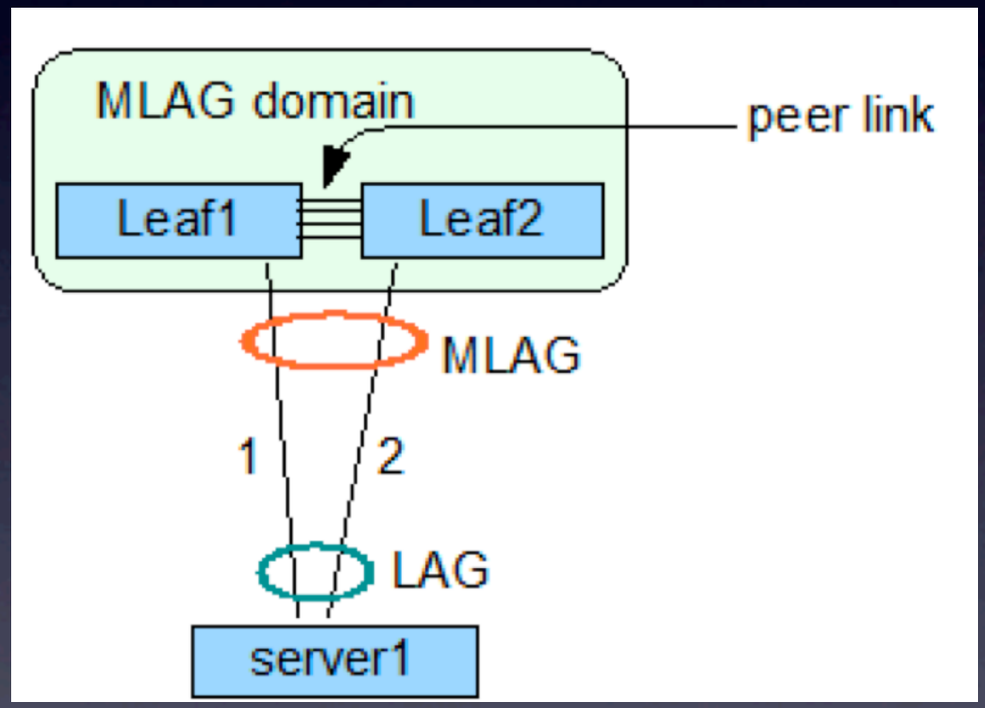

Multi-stage MLAG

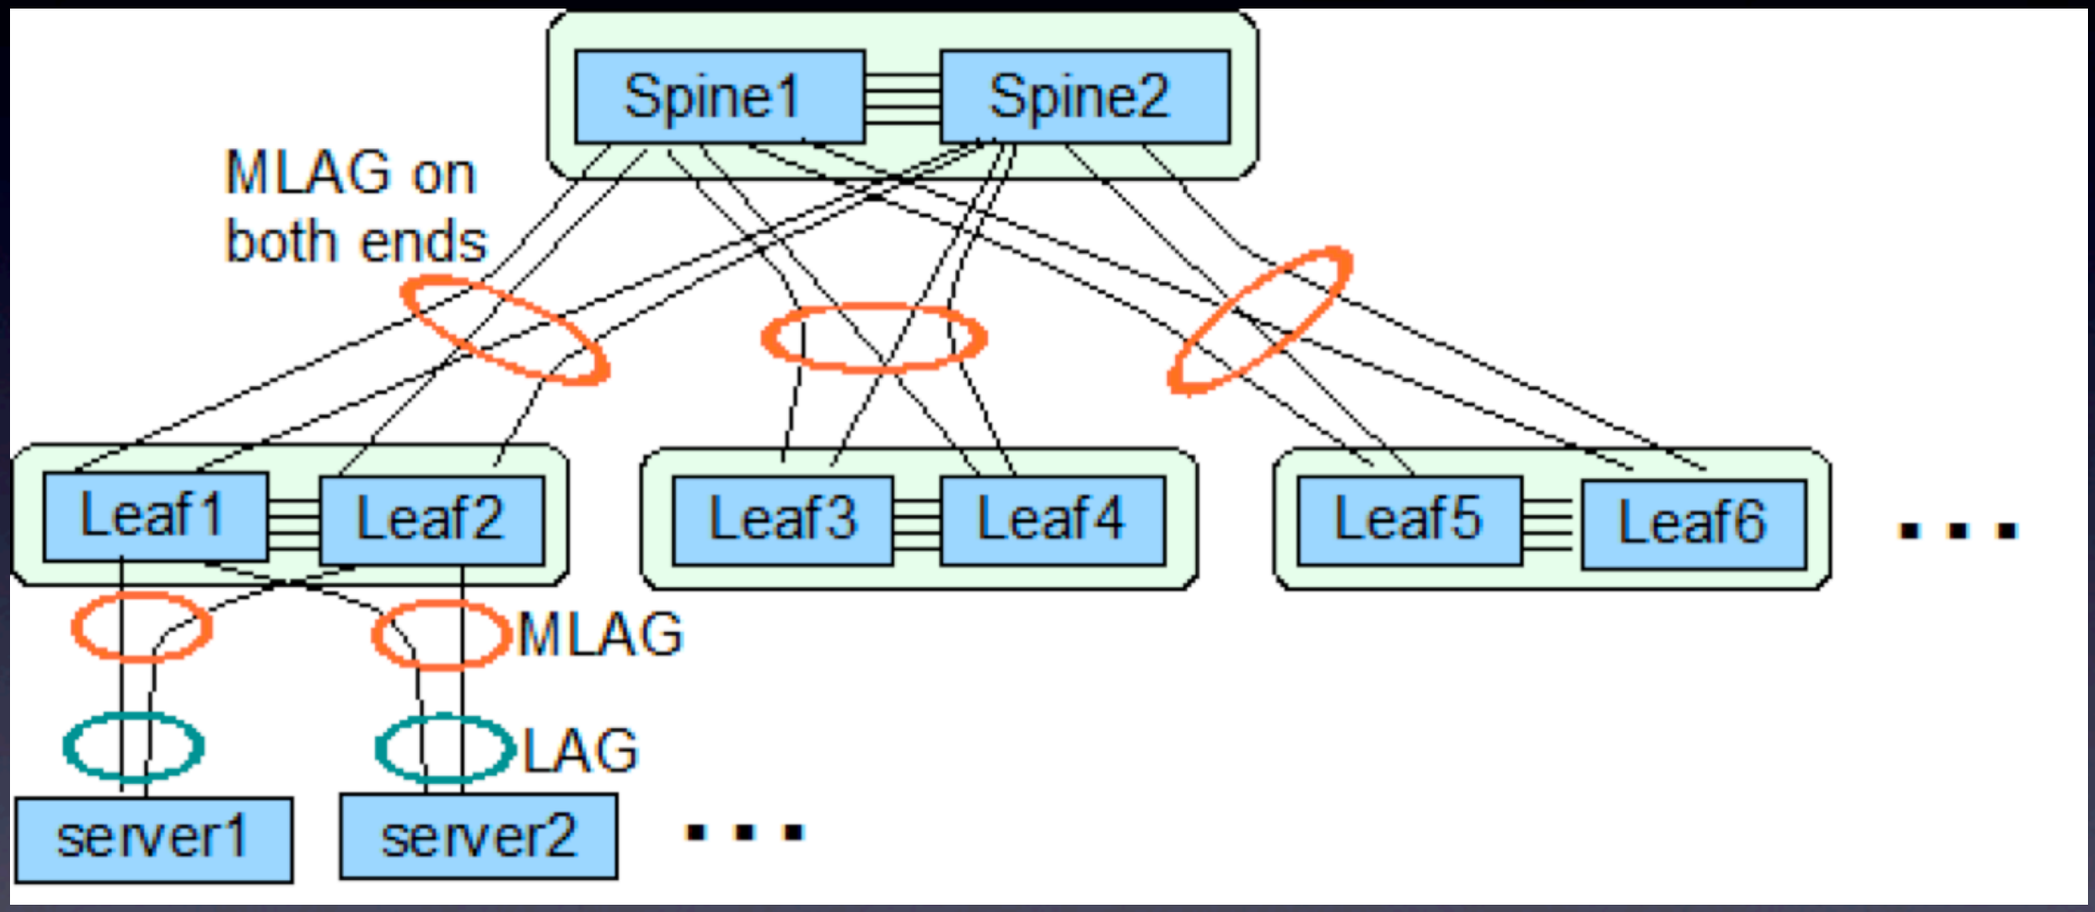

MLAG can be extended from leaf to core

Arista EOS

Extensible Operating System

Network OS Attributes

Network O/S	Legacy	Arista EOS
3rd Party Apps	N/A	Yes
3rd Party Mgmt	N/A	Yes
ISSU	N/A	Yes
Self Healing	N/A	Yes
Fault Tolerant	N/A	Yes

EOS Modular Architecture

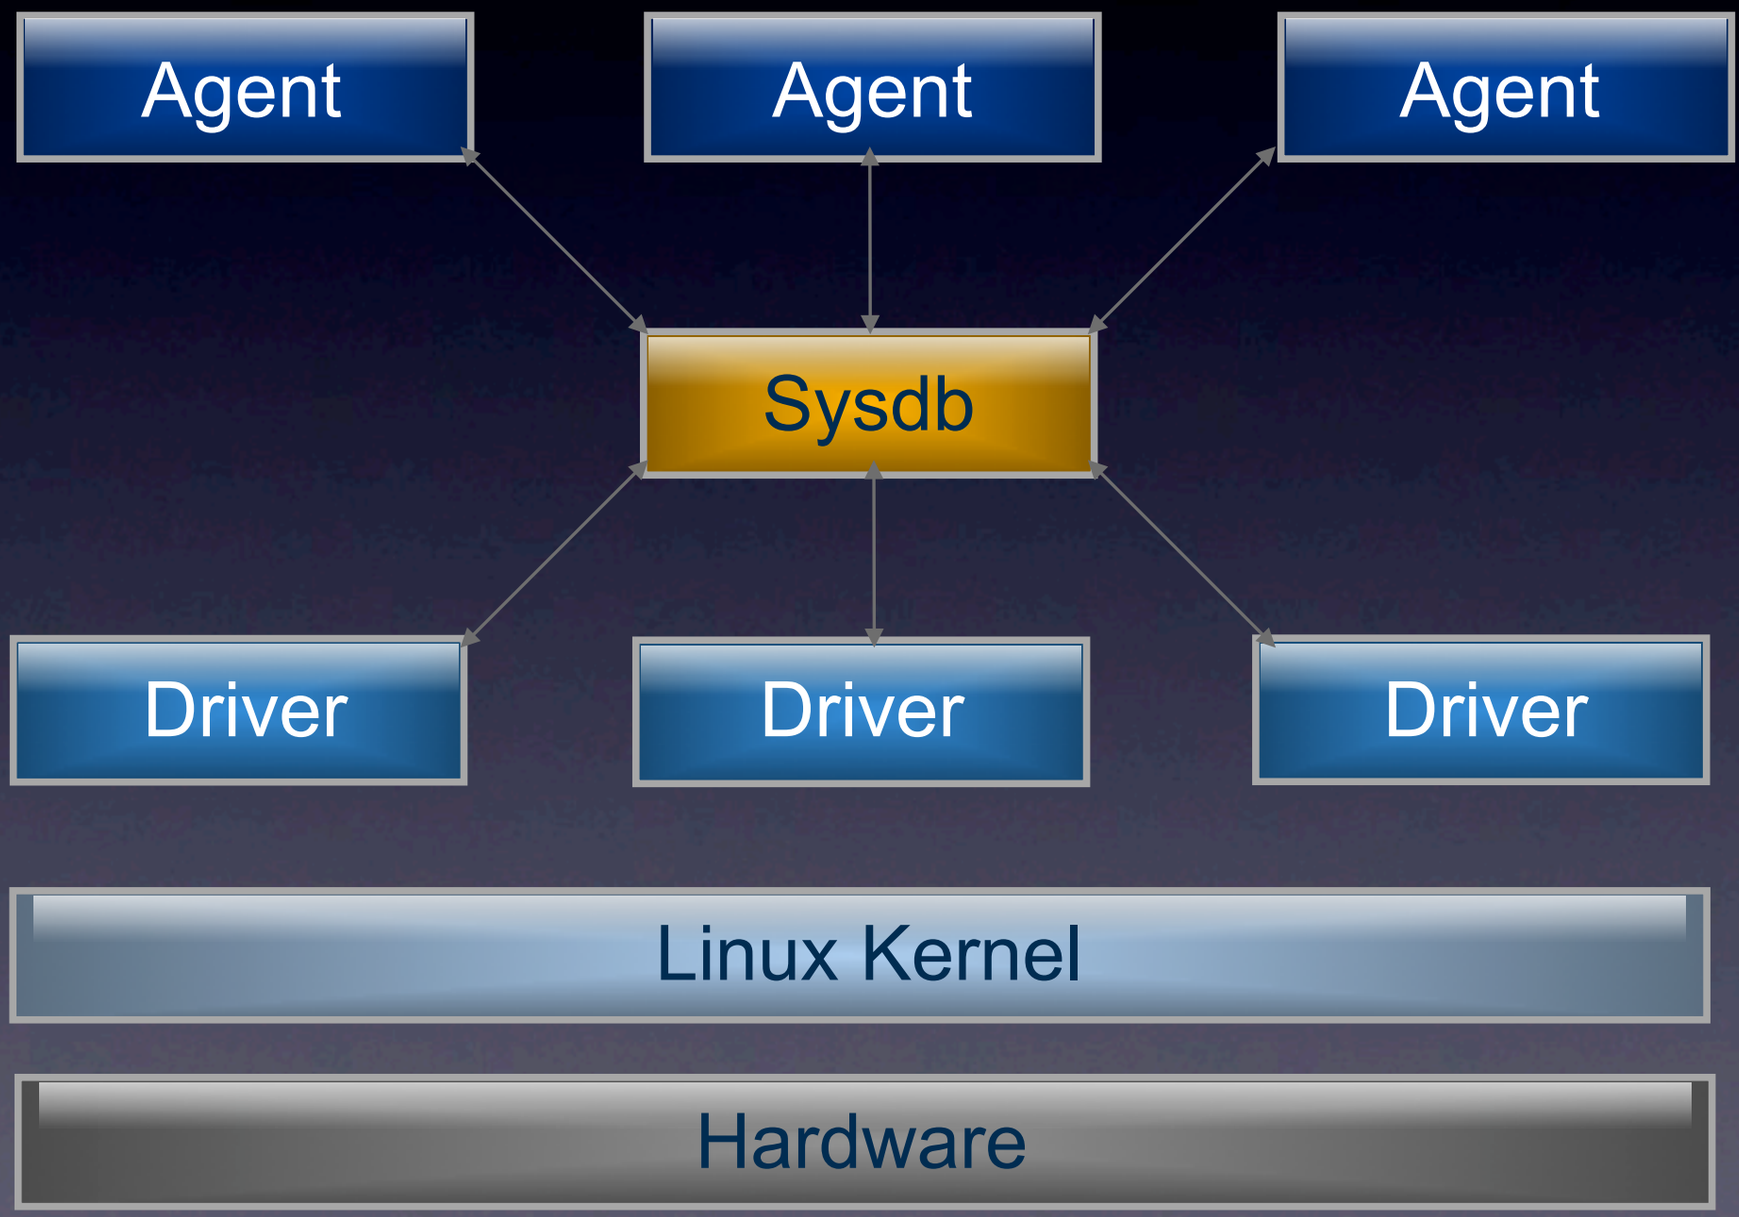

Event Driven State Updates

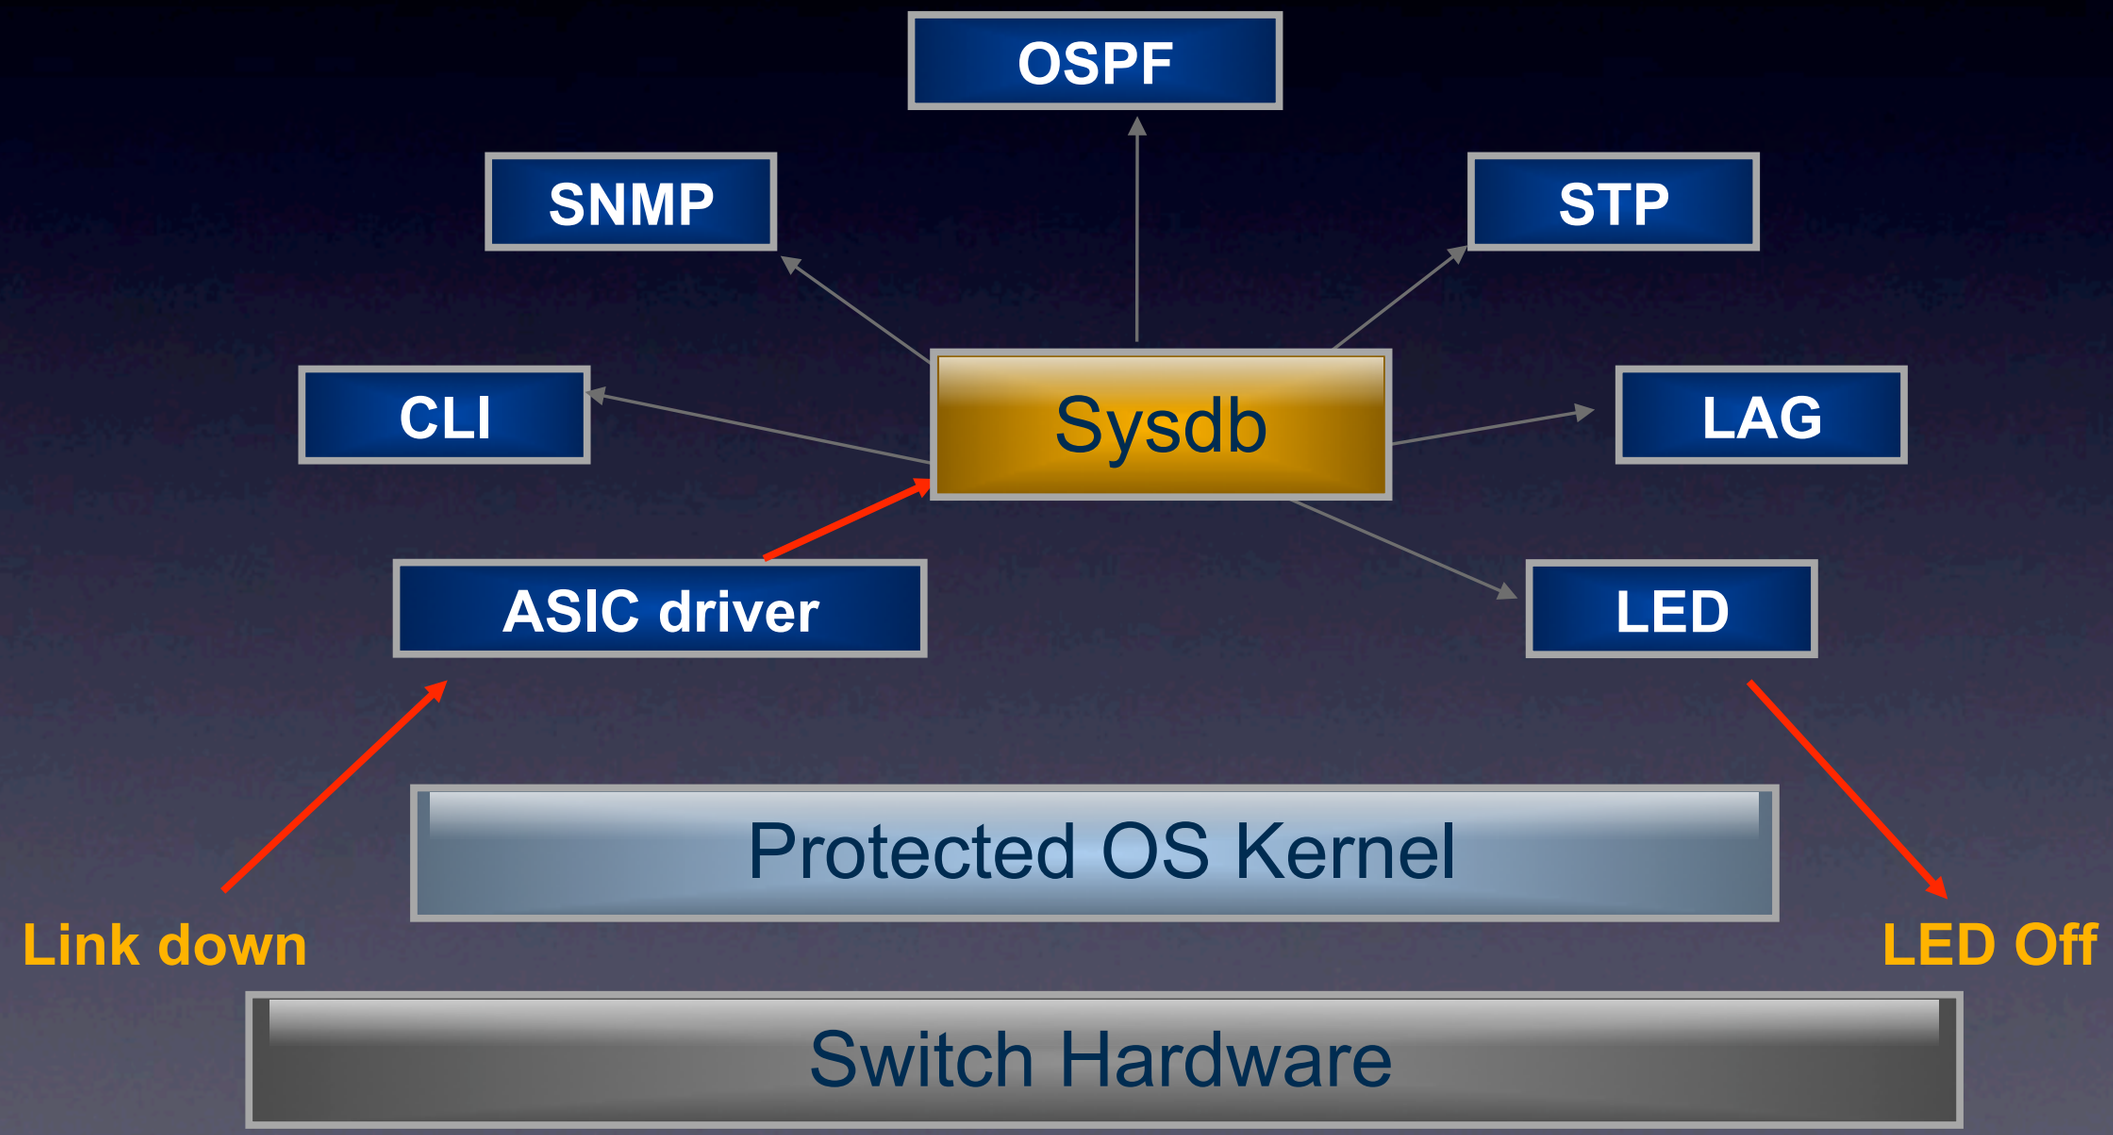

It is a Switch - it is a Server!

```
ge203>enable
ge203#show interfaces status
Port      Name                 Status       Vlan       Duplex  Speed Type
Et1                            notconnect   1            full  10000 Not Present
Et2                            notconnect   1            full  10000 Not Present
...
Et47                           connected    1            full  10000 10GBASE-CR
Et48                           connected    1            full  10000 10GBASE-CR
ge203#bash
[admin@ge203 ~]$ cd /mnt/flash
[admin@ge203 flash]$ ls
EOS.swi  boot-config  persist  startup-config
[admin@ge203 flash]$ head startup-config
! device: ge203 (DCS-7148S, EOS-4.1.0-135910 (engineering build))
!
hostname ge203
ip name-server 172.17.0.22
...
[admin@ge203 flash]$
```

» **EOS provides server capabilities on a switch**

Full Access to all Linux Tools

```
[admin@st321 ~]$ uptime
 06:21:23 up 48 min,  4 users,  load average: 0.36, 0.24, 0.21
```

```
[admin@st321 ~]$ vmstat 1
procs -----------memory---------- ---swap-- -----io---- --system-- -----cpu------
 r  b   swpd   free   buff  cache   si   so    bi    bo   in   cs us sy id wa st
 0  0      0  15560  61852 505448    0    0    38    26  405  487 13  1 86  0  0
 0  0      0  15548  61852 505500    0    0     0     0  731  582  2  0 98  0  0
```

```
[admin@st321 ~]$ free
             total       used       free     shared    buffers     cached
Mem:       1035976    1020484      15492          0      61916     505436
-/+ buffers/cache:     453132     582844
Swap:            0          0          0
```

- **Linux system management tools standard with EOS**

Example: TCPdump

Debugging a Spanning Tree Issue:

```
[admin@ge203 ~]$ sudo tcpdump -n -i eth48 -v -vv
tcpdump: listening on eth48, link-type EN10MB (Ethernet), capture size 96 bytes
06:30:52.907720 STP 802.1w, Rapid STP, Flags [Proposal, Learn, Forward, Agreement],
        bridge-id 8000.00:1c:73:04:1a:27.8030, length 36
        message-age 0.00s, max-age 20.00s, hello-time 2.00s, forwarding-delay 15.00s
        root-id 8000.00:1c:73:04:1a:27, root-pathcost 0, port-role Designated
06:30:54.910031 STP 802.1w, Rapid STP, Flags [Proposal, Learn, Forward, Agreement],
        bridge-id 8000.00:1c:73:04:1a:27.8030, length 36
        message-age 0.00s, max-age 20.00s, hello-time 2.00s, forwarding-delay 15.00s
        root-id 8000.00:1c:73:04:1a:27, root-pathcost 0, port-role Designated
06:30:56.911804 STP 802.1w, Rapid STP, Flags [Proposal, Learn, Forward, Agreement]
        bridge-id 8000.00:1c:73:04:1a:27.8030, length 36
        message-age 0.00s, max-age 20.00s, hello-time 2.00s, forwarding-delay 15.00s
        root-id 8000.00:1c:73:04:1a:27, root-pathcost 0, port-role Designated
```

- **Customer reaction: “You can do that?!”**

Install Extensions in Minutes

```
ge203#copy http://mirrors.reflected.net/fedora/linux/releases/9/Everything/i386/os/Packa
ges/fping-2.4b2-8.fc9.i386.rpm extension:
ge203#show extensions
Name                           Version/Release Available     Status          RPMs
------------------------------ --------------- ------------- --------------  ----
fping-2.4b2-8.fc9.i386.rpm     2.4b2/8.fc9     available     not installed      1
ge203#config
ge203(config)#extension fping-2.4b2-8.fc9.i386.rpm
ge203(config)#bash fping -g 172.17.0.32/29
172.17.0.32 is alive
172.17.0.34 is alive
172.17.0.35 is alive
172.17.0.36 is alive
172.17.0.37 is alive
172.17.0.38 is alive
172.17.0.33 is unreachable
172.17.0.39 is unreachable
ge203#
```

» **nc, sendmail, mrtg, nagios, dhcpd, named, ptpd, tftpd, pdsh …**

ARISTA

reinventing data center switching