

ARISTA



4K and Beyond – Is Your Network Ready?

Addison Chi, Technical Lead, addison@arista.com

Live Music
DIGITALCMAA N

360全景VR直播

王菲 幻乐一场 演唱会

Faye's Moments
Live 2016

天后全球首场VR直播

12月30日 19:30

腾讯视频 不负好时光



Augmented Reality AR Virtual Reality VR Online Live Show



零距离互动

用我的方式，表达我的态度



Business Drivers for Transition SDI to IP/Ethernet

Serial Digital Interface



Customized SDI Broadcast hardware
SDI Cabling – Die in Place
Limited distance
Parallel Infrastructure
Single Direction Per Stream
Forklift as formats change

Ethernet/IP



Ethernet/IP based Data Center
Scale Economics from communication
Flexible – formats, facilities, workflows
Flexible distance and cabling
Capacity Grows as Moore's Law 2X/2Y

JT-NM Market & Technology Drivers

- Take advantage of the marketplace **economics of IT** Technology;
- Make use of the extensive and well trained base of **design and maintenance personnel** available in this field;
- Deploy enterprise-class capabilities and **redundancy options**;
- Use any one of a number of monitoring, **diagnostic and troubleshooting tools** that currently exist for enterprise deployments of IT infrastructure;
- **Facilitate higher resolution and transport bit rates with 10/25/40/50/100G Ethernet**

Standard	Name	Introduced	Bitrates	Example video formats
SMPTE 259M	SD-SDI	1989 ^[2]	270 Mbit/s, 360 Mbit/s, 143 Mbit/s, and 177 Mbit/s	480i, 576i
SMPTE 344M	ED-SDI		540 Mbit/s	480p, 576p
SMPTE 292M	HD-SDI	1998 ^[2]	1.485 Gbit/s, and 1.485/1.001 Gbit/s	720p, 1080i
SMPTE 372M	Dual Link HD-SDI	2002 ^[2]	2.970 Gbit/s, and 2.970/1.001 Gbit/s	1080p60
SMPTE 424M	3G-SDI	2006 ^[2]	2.970 Gbit/s, and 2.970/1.001 Gbit/s	1080p60
SMPTE ST-2081	6G-SDI	2015 ^[4]	6 Gbit/s	4Kp30
SMPTE ST-2082	12G-SDI	2015 ^[5]	12 Gbit/s	4Kp60
SMPTE ST-2083*	24G-SDI		24 Gbit/s	4Kp120



SDI to IP Evolution - Video and Audio Streams

Video, Audio,
Ancillary



Serial Digital Interface
SMPTE 259M, 344M, 292M, ...



Video, Audio,
Ancillary



SDI / HD-SDI over IP
RFC 4175, ST 2022-6



Video, Audio,
Ancillary



HD-SDI over IP
ST 2022-6

Audio over IP
AES67



Separate Video &
Audio Streams
TR-04

Video, Audio,
Ancillary



Video over IP
RFC 4175

Audio over IP
AES67

ANC over IP
RFC RTP Anc.



Separate Ancillary
Data, Video & Audio
Streams
TR-03


Verticals Within Media & Entertainment

- Editing
- Rendering
- Transcoding
- Finishing
- All About Storage

File-Based
Workflows



Real-Time
Broadcast



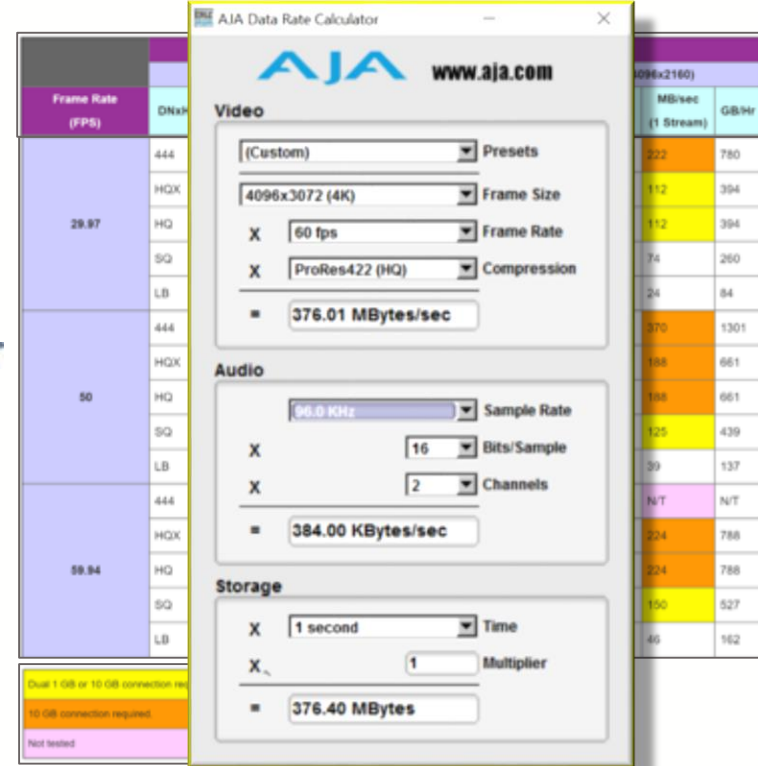
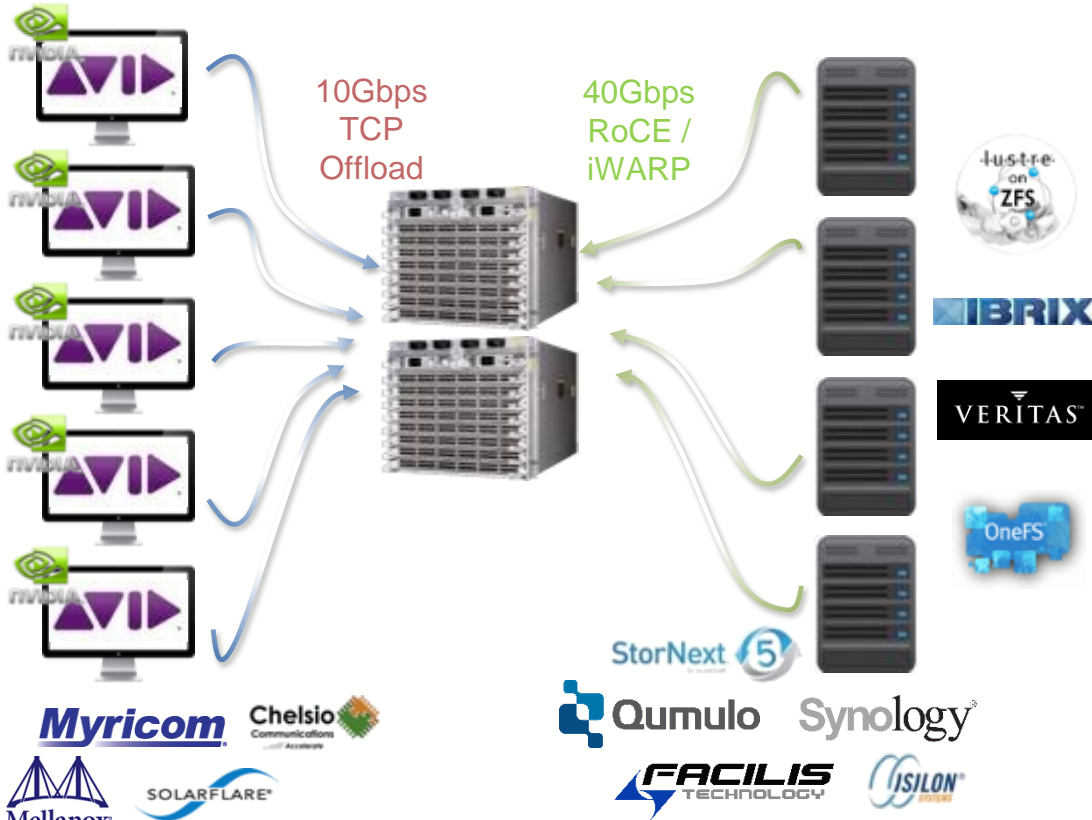
- Uncompressed
- Live
- Sports
- Broadcast Plant
- Best Quality
- SDN / Control Systems

Streaming /
OTT




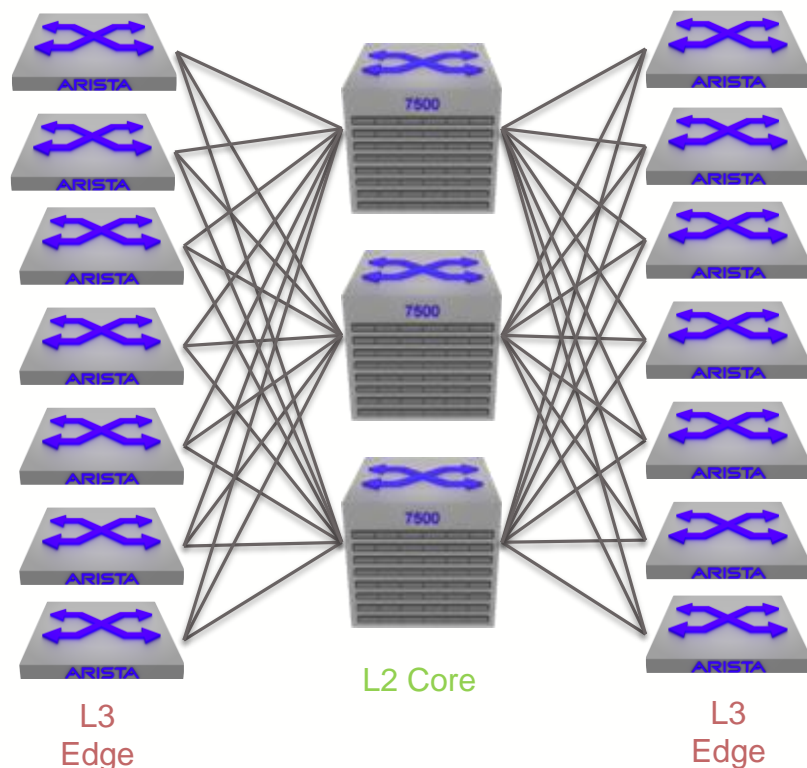
- CDN
- Time-Shift
- Compressed
- Internet Route Scale
- Any Device, Any Time

File Based Workflow – Bandwidth = Productivity



4K Storage: ~ 23GB per Minute ProRes422 (HQ)
4K Transport: ~ 3Gbps

Rendering Farm / VFX / Gaming




ARISTA

Buffering Challenges in Leaf-Spine Network

N/S Bottleneck:
Oversubscribed Ports
Non-Deterministic Latency
Limited Expansion



Serialization Bottleneck:
1G Servers lack capacity
to handle 10G traffic
bursts



Spine Bottleneck:
TCP request from host to
storage targets.



Leaf Bottleneck:
TCP reply from storage to
compute.



TCP Stack Bottleneck:
TX and RX buffer size,
number of threads,
packet size, PCIe lanes.



Small Buffer Switch – Output Discards

Ethernet1/2 is up

Dedicated Interface

```
Hardware: 1000/10000 Ethernet, address: 002a.6a25.22c9 (bia 002a.6a25.22c9)
Description: *** Connection to TC-E 6/2 (Test Point F) ***
MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec
reliability 255/255, txload 134/255, rxload 20/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s, media type is 10G
Beacon is turned off
Input flow-control is off, output flow-control is off
Rate mode is dedicated
Switchport monitor is off
EtherType is 0x8100
Last link flapped 22:08:06
Last clearing of "show interface" counters 18:11:27
0 interface resets
30 seconds input rate 1,438,945,632 bits/sec, 497,203 packets/sec
30 seconds output rate 9,479,333,624 bits/sec, 3,254,913 packets/sec
Load-Interval #2: 5 minute (300 seconds)
  input rate 803.26 Mbps, 277.51 Kpps; output rate 5.29 Gbps, 1.81 Mpps
```


RX

```
61210683 unicast packets 0 multicast packets 24 broadcast packets
61210698 input packets 22144318040 bytes
0 jumbo packets 0 storm suppression bytes
0 runs 0 giants 0 CRC 0 no buffer
0 input error 0 short frame 0 overrun 0 underrun 0 ignored
0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
0 input with dribble 0 input discard
0 Rx pause
```

TX

```
400221438 unicast packets 263648 multicast packets 209 broadcast packets
400485237 output packets 145885543572 bytes
2659898 jumbo packets
0 output error 0 collision 0 deferred 0 late collision
0 lost carrier 0 no carrier 0 babble 354790439 output discard
0 Tx pause
```

- Packet buffer availability at the moment a particular packet arrives.
- Poisson-like distribution of bandwidth per flow.
- Resulting in certain flows will use more bandwidth than others.
- Typical Storage or Server are small request, big reply.




4K Network Architecture Best Practice

Microbursts, Fan-in, Speed
changes
4K Needs Large Switch
Buffers
Multiple **40G** and / or **100G** Links
Non-Blocking Fabric



Large Buffers at Leaf
Virtual Output Queues
Smart System Upgrades
Latency Monitoring
Telemetry and Management



ECMP
Aggregation
Resiliency
Easy Maintenance

QoS: Priority Flow Control PFC/RDMA
Enhanced LACP Hashing Algorithm
Automated Interface Configuration
Distributed Gateway – vARP


Reduce Storage Cost, Increased Productivity and Throughput

Network Features for Broadcasters

- Low latency
 - Real time, Uncompressed
- Multicast
 - IP Video is Multicast
- AVB (802.1Q)
 - Bandwidth Reservation for Lossless Transport
- SMPTE 2022-6/2059-2
 - HD-SDI over IP & Synchronization
- VSF TR-03
 - Elemental Media Flows
- SDN Control
 - Broadcast Control Systems (BCS)
- 25G/50G/40G/100G



Audio Video Bridging (IEEE 802.1BA) is L2 As Now



gPTP - Generalized Precision Time Protocol / (IEEE 802.1AS)

Grand Master - gPTP master

AVB Domain - Layer 2 domain including AVB capable endpoints and one or more AVB capable Ethernet switches

MSRP - Multiple Stream Reservation Protocol (IEEE 802.1Qat)

MAAP - MAC Address Acquisition Protocol

MVRP - Multiple VLAN Registration Protocol

Queuing and Forwarding (802.1Qav)

AVB Traffic Classes

A = Max latency 2ms

B = Max latency 50ms

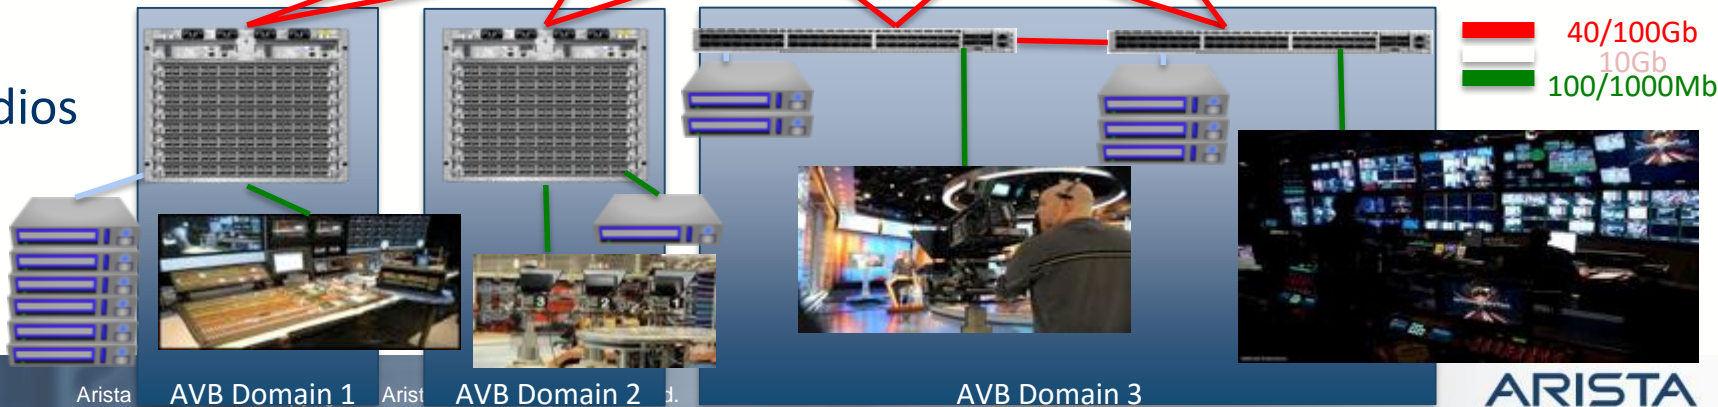
One Big Sports TV Customer

Broadcast Core

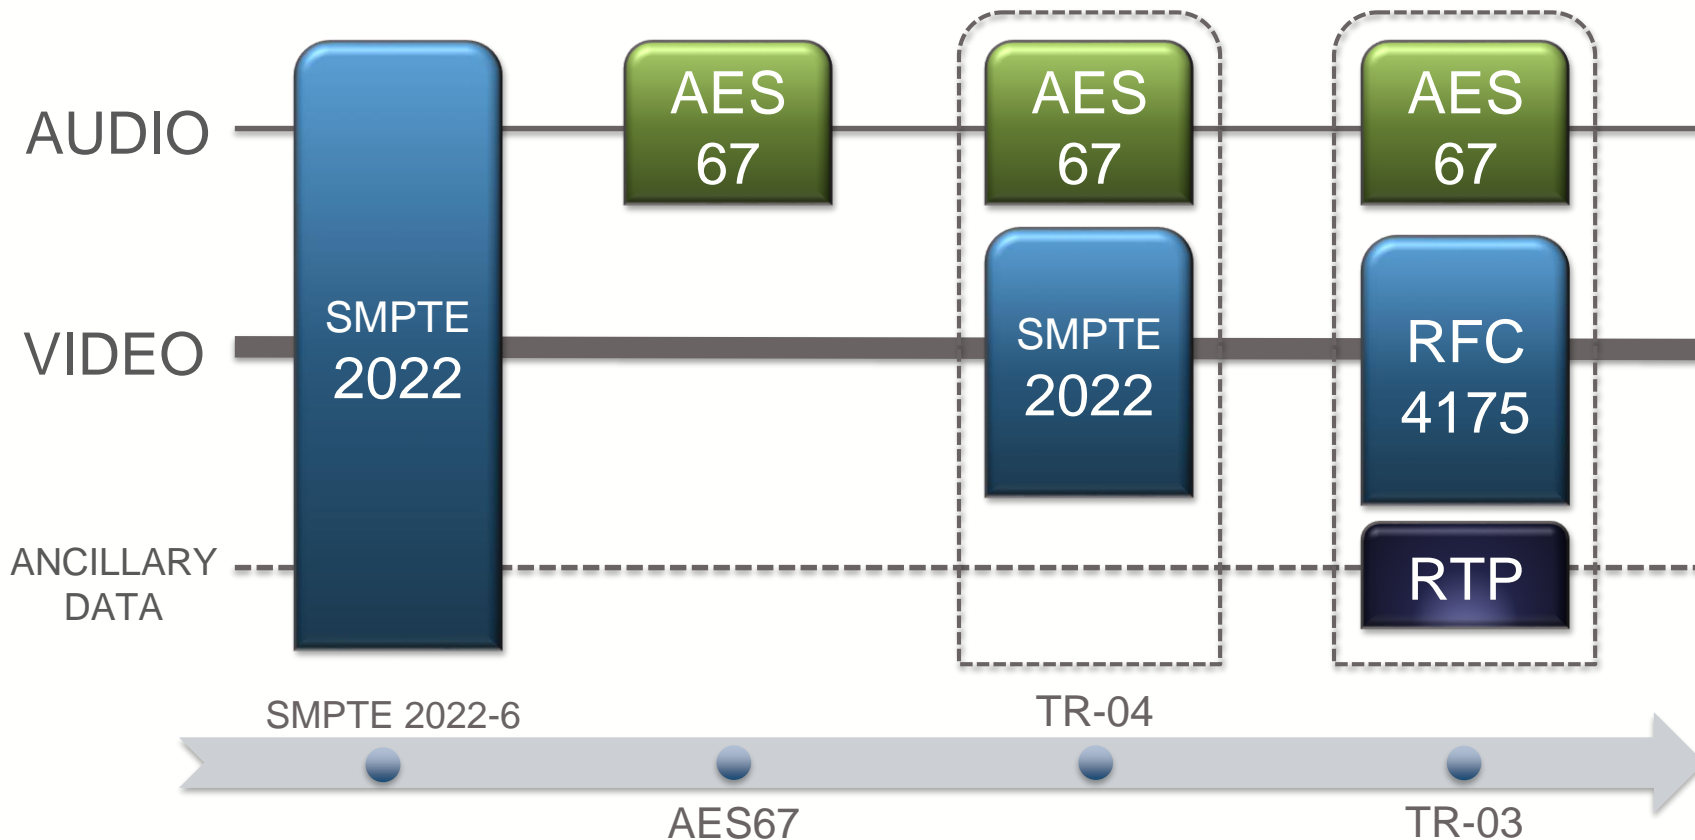
2304 x 10Gbps ports
46 Tbps, non-blocking
Supports > 50,000 signals
(traditional ~1150)
Format agnostic
SD, HD, 3G, 4K, 8K....
Any/All compression codecs
Ethernet MPEG2 -TS, multi-cast
All signals everywhere

Current
88 concurrent inbound recording paths
~2,500 hours of storage
Goal
200 concurrent paths initially
Easily expandable to 400+
20,000+ hours of on-line storage

Studios

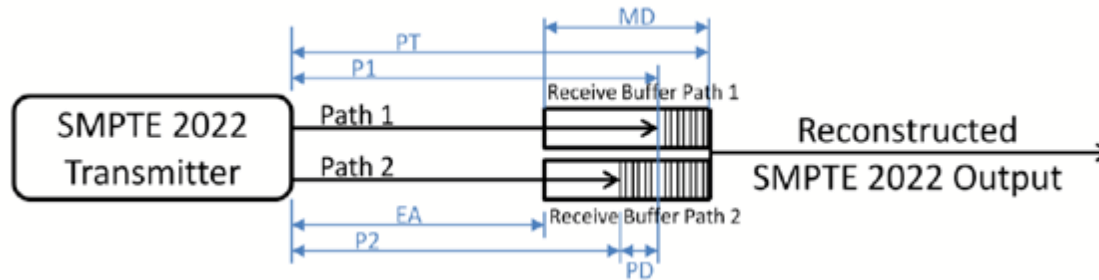


SMPTE Extends to L3 and is Driven by AIMS




SMPTE 2022-5/6/7 – Network Attributes

- SMPTE 2022-7 – Transmission of **two streams** of identical content over **potentially diverse paths**. (2022-6 Video Source ID Primary or Protected Stream)




- Ex: 1080p/60Hz HD-SDI bit rate 2.970 Gbps of 1376 byte packet, generating 270 packets per millisecond / 270,000 pps. 3ms protection is 810 packets.
- **Receiver Based Multicast or SDN path determinism to ensure diverse paths.**
- **Network L2 and L3 convergence, packet loss, and controller feedback mechanism.**

Euro 2016



Seamless Switching

- Channel change needs to feel immediate upon button push
- IGMP join/leave not as fast as analog
- Partners need a better mode of operation
- EOS extensibility to the rescue



For 720p/59.94, there are 2249 SMPTE 2022-6 RTP datagrams per frame, and the datagrams should be transmitted approximately every 7.4 μ s.

SDN Integration for Control

Lawo

EOS SDK Integration

- Rewrite multicast MAC table using our speedy SDK
- Direct integration with their control software
- Provides seamless switching without requiring IGMP



Imagine

EAPI Integration with Magellan Panel

- Arista EOS CS supplied extension with it's own API
- Programs static IGMP snooping entries
- Less overhead than eAPI, more batch-friendly, much faster than IGMP or eAPI




<http://sportv.globo.com/site/>

2016 Olympics

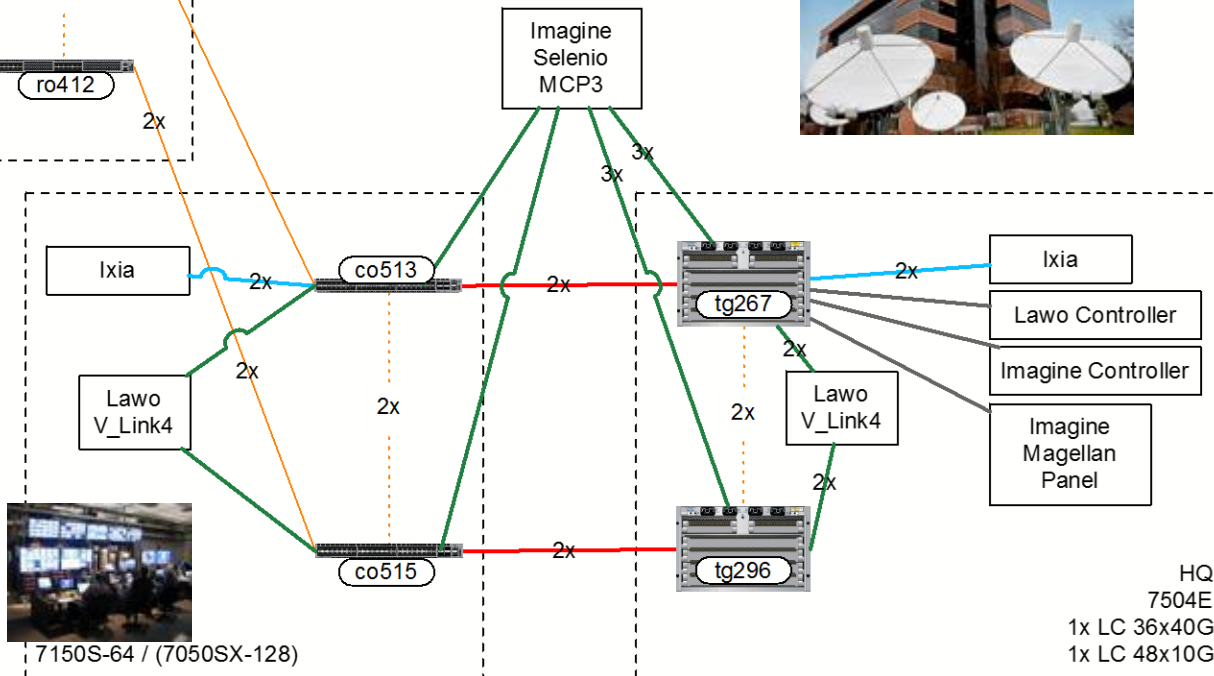
LOBOSAT



Imagine COMMUNICATIONS.



- 1G UTP —
- 10G SR —
- 10G DAC —
- 40G AOC —
- 40G SR —



Network Features for Webcasters


- Low latency
 - Video Streaming
 - East / West Traffic
 - Transcoding Engines
- Automated Application Burst
 - VMware and Docker Integration
 - Automated Network Reconfiguration
- High Throughput Storage
 - RoCE, iWARP, 40G/100G
- High Throughput Database
 - Big Data
 - Advertisement Placement
- Firewall Integration & Offloading
- Hitless Upgrades
- CDN / Internet Route Scale



NETFLIX CDN DC

Every POD

1. 1152 x 10G
2. 96x100G
3. Deep buffer
4. SDN Switch as Router



100G
or nx10G

Internet

NETFLIX

40G-PLR4

40G

40G-SR4
To Flash

Fiber Panel

Fiber Panel

Fiber Panel

Fiber Panel

Fiber Panel

Fiber Panel


4 x 10G-LR
LC Connector



Server / Storage

ARISTA

First SDI/IP Hybrid Switch Live at NAB 2016



Arista Focus for Media & Entertainment



Best-of-Breed Physical Characteristics

- Performance
- Throughput
- Low Latency
- High Port Density
- Multicast at Scale
- Uptime Reliability
- Hitless Upgrades

Managing Congestion / Speed Changes

- Deep Buffers
- VoQ Architecture
- Telemetry for Proactive Network Monitoring

Traffic Steering, Broadcast Control System Integration

- SDN
- OpenFlow
- DirectFlow
- Extensions
- SMPTE 2022-6
- AVB

Timing and Synchronization

- SMPTE 2059-1, 2
- AVB gPTP
- PTP IEEE 1588
- Boundary & Transparent Clock
- High Speed On-Board OCXO

Scalable Universal Cloud Network

Uniform, Single Image Control Plane

Data Center Automation with CloudVision