



Scalable IP Architectures for Live Production and Playout

Gerard Phillips, Systems Engineer
Arista Networks




IP SHOWCASE THEATRE AT IBC – SEPT. 14-18, 2018



Scalability – Adding more stuff right?



- End points
- Bandwidth
- Geographic sites
- Formats and workflows
- Network functions





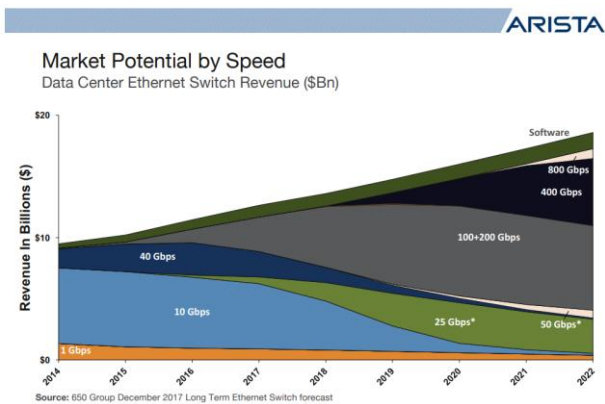
More Formats and workflows

- 4K, 8K, VR, eSports
- What else? We just don't know!!
- Enable scalability by building on cloud principles
 - Vanilla, standardized IP - no vendor lock-in
 - BGP – resilient, scalable, hardened, fast convergence
 - Extensible NOS that allows you to add functionality as you need it
 - Demand a rich API, designed for automation
 - Telemetry + monitoring



More Bandwidth

- More hosts. Higher resolution, high frame rate ...
- Tomorrow's unknown workflows
- The data heavy immersive viewing experiences of tomorrow
- Fiber is relatively inexpensive, bury more than you think you will need
 - You're bound to use it
 - Digging up the carpark is likely way more expensive than the fiber!
- Moores law for merchant silicon is still running....






Merchant Silicon Firsts

- 2008: First ultra-low latency 24-port 10G single chip
- 2010: First Large Buffer 10G Chip with VOQ Fabric
- 2011: First 64-port 10G single chip switch
- 2012: First 32-port 40G single chip
- 2013: First Large Buffer 40G Chip with VOQ Fabric
- 2015: First 32-port 100G single chip
- 2016: First Router 100G Chip with VOQ Fabric
- 2017: First 64-port 100G single chip
- 2018: First 32-port 400G single chip (forecast)
- 2019:


Bandwidth Improvement



Architectures for host scale

- One **BIG** switch
- Hub and spoke
- Leaf and spine

- What might we want to consider?
 - How many hosts can be accommodated?
 - Blocking or non-blocking infrastructure?
 - PTP distribution
 - Resilience options
 - Failure domain size - Physical, L2, PTP...





One BIG switch

- Non-blocking
- Simple IGMP works fine
- > 2k hosts @ 25Gbe
- Fixed or Modular chassis
- Modular – populate as you grow
- Maybe 100% non-blocking is not needed??




ARISTA



Hub and spoke

- >16k hosts @ 25Gbe, 1:10 provisioning
- 400Gbe around the corner....
- Scale at the rate you want to
- Just like traditional broadcast Tie-lines ☺
- Non-blocking no longer makes sense
 - Does that fit your workflow?
- Flow orchestration (SDN) is now required



ARISTA




Leaf and Spine

- A great architecture for future thinking converged network
- ...think Virtual Network Functionality

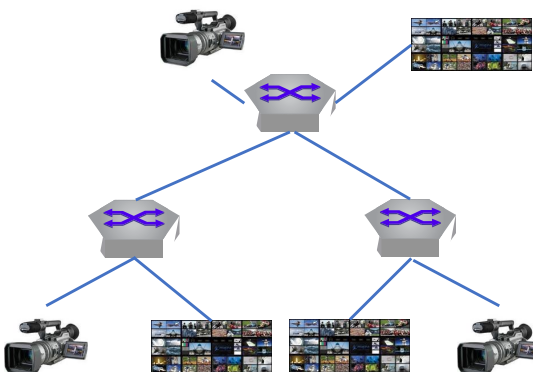
But....

- Multicast does NOT scale like unicast...
- Flow orchestration (SDN) is required



Expansion – if you want to start with a single switch (pair -7) and grow

- Start with a leaf
- Either add a sp(l)ine...
- Or promote the leaf pair to “spline”
- Your decision

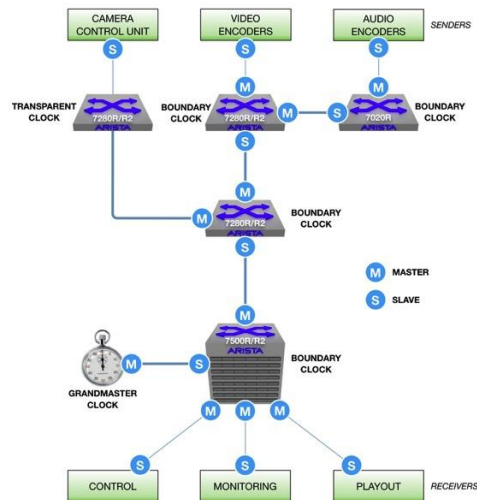




PTP Scaling

BC and TC provide useful ways to scale and protect the performance of your PTP distribution:

- A single switch may not need BC/TC
- Multiple switches – use BC/TC
- GM capability needs to be understood
- Unicast delay-request is preferred
- Slaves often have different requirements



Don't forget the Dev/Net Ops guys....

- Whether the network starts big, or grows, the effort to manage and sustain grows exponentially
- The key to being able to scale the management function is:
 - Automation, enabled by multiple flexible API's
 - Integration and support for the tools the server guys use - Puppet / Ansible / Chef / bash / python
 - Real-time, comprehensive telemetry
 - Powerful on switch scripting / programmability
- More on this subject, back here, Monday 4:30





Summary

- The cloud is all about increasing scale, and reducing OPEX
- Their principles are very applicable to our industry
 - Vanilla, standardized IP - no vendor lock-in
 - Extensible NOS that can allow you to add functionality as you need it
 - Multi-switch architectures that allow for growth, and add resilience
 - Rich API's designed for automation, reducing errors and cost
 - Telemetry + monitoring give forward notice of trends before they become issues

ARISTA



Thank You

Gerard Phillips, Arista Networks

gp@arista.com

+44 7949 106098

ARISTA

IP SHOWCASE THEATRE AT IBC – SEPT. 14-18, 2018