Bandwidth drivers for 100 G Ethernet

Vik Saxena, Ph.D.

Sr. Director, Network Architecture

CTO Office, Comcast Cable

Outline

- Comcast Profile & Triple Play Offering
- Network Overview
- What is driving network bandwidth?
- Where will we see high bandwidth requirements?
- Why 100GE?

comcast.

Comcast Profile

Largest U.S. Cable Company
> 40 Million homes passed
24 Million Basic Cable subs
12 Million Digital Cable subs
11 Million HSD subs
2 Million CDV subs

A Portfolio of Growing Cable Channels


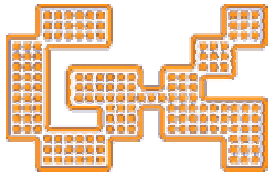


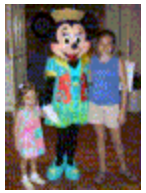


Services Currently Supported by Network Infrastructure

Current Spectrum Utilization

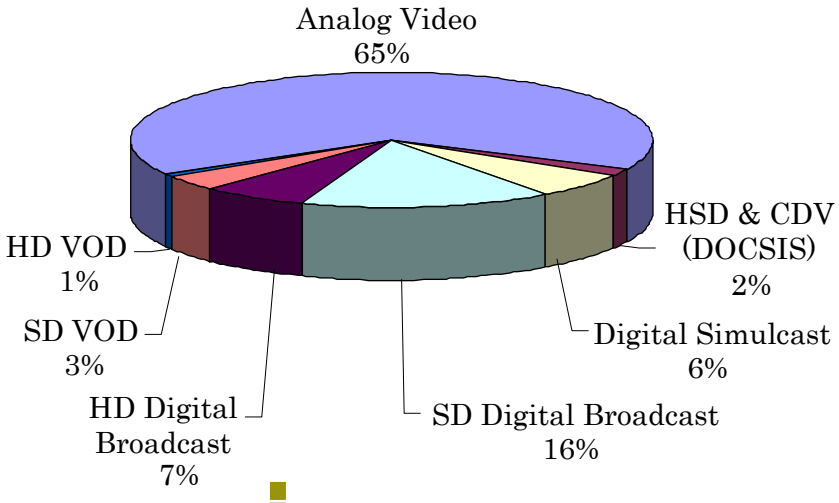

Shift from Broadcast to Unicast services drives core network capacity!

Unicasting for personalized & interactive services

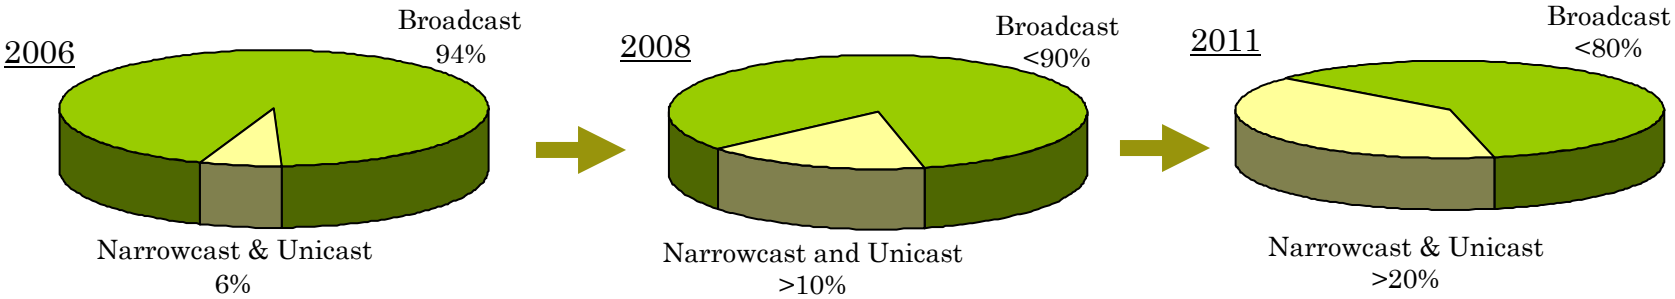

Broadcast services: Analog video, digital broadcast, digital simulcast

Narrowcast and Unicast services: DOCSIS, VOD, SDV

New Unicast / Interactive Services on Roadmap

Video	High-Speed Internet	Voice
VRN Guide		Single Address Book
Video Search Tools		
TV Recommendations		Communications Hub Voice, text, IM, video
Online TV Planner		
"Follow-me" On Demand	Comcast Dashboard	Personalized/ Celebrity Greetings
	E-Mail Reminders	
	Click to Record	CDV Quick Info news, weather, sports, traffic
TiVo	PowerBoost Upstream	

Growth in On-Demand Services

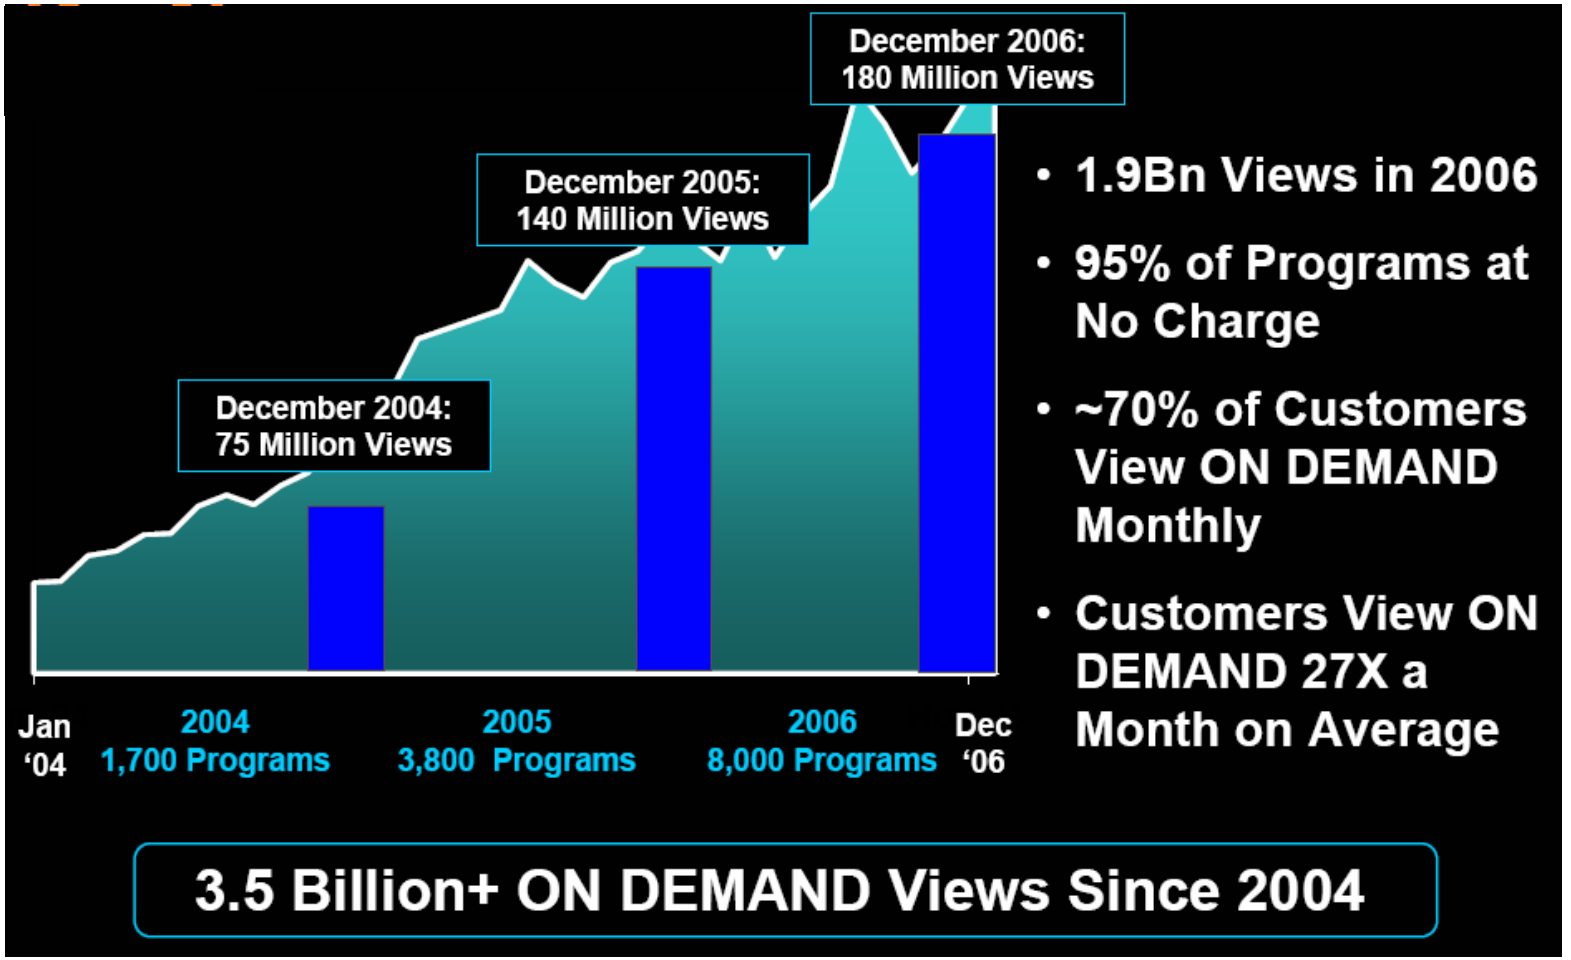

Growth in On-Demand Services

A Superior Hi-Def Experience

- **A Growing High-Def Offering**
  - Up to 20 Linear HD Channels Today… and Growing
- **Leveraging ON DEMAND**
  - 100+ Hours of HD ON DEMAND
  - HD VOD Hours to Double in 2007 and Again in 2008
- **Blockbuster HD ON DEMAND**

2 Million HD-VOD Views Since Launch

Network Infrastructure Segments: Backbone, Regional and Access

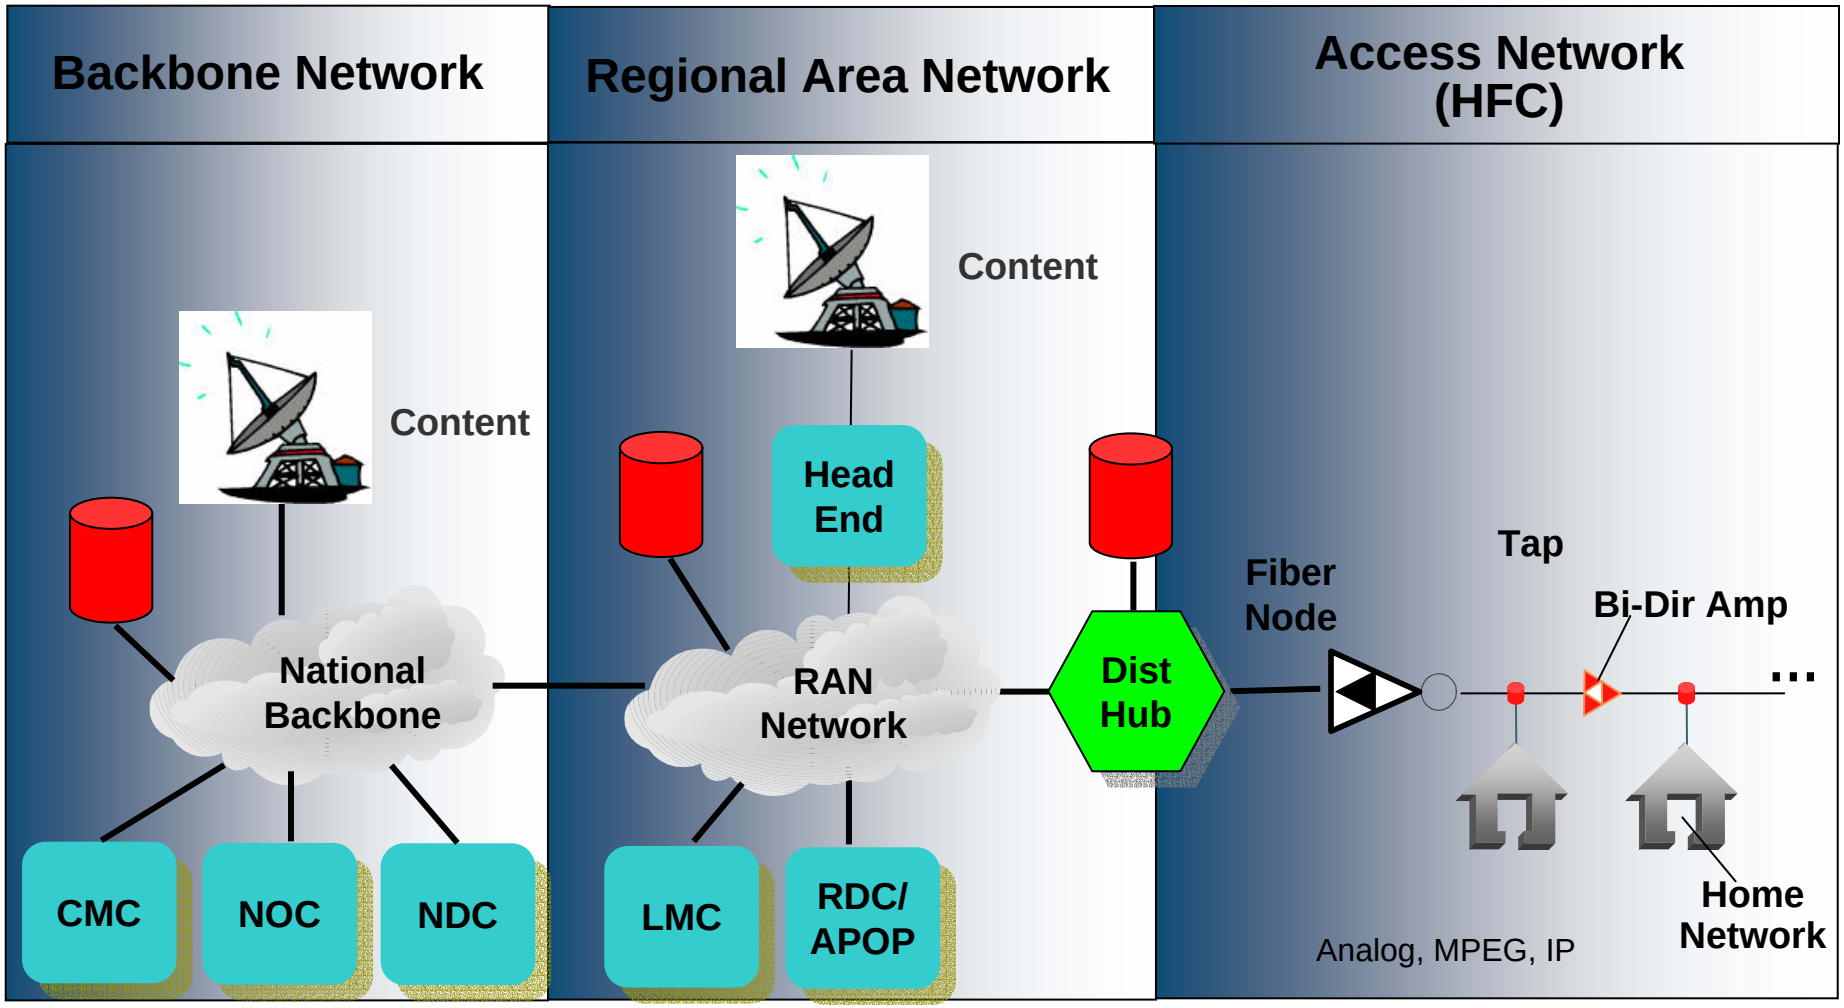

The Converged Regional Area Networks (CRAN) supports the delivery of multiple services over a common IP infrastructure

Able to support very high bandwidths and low price points required for video delivery

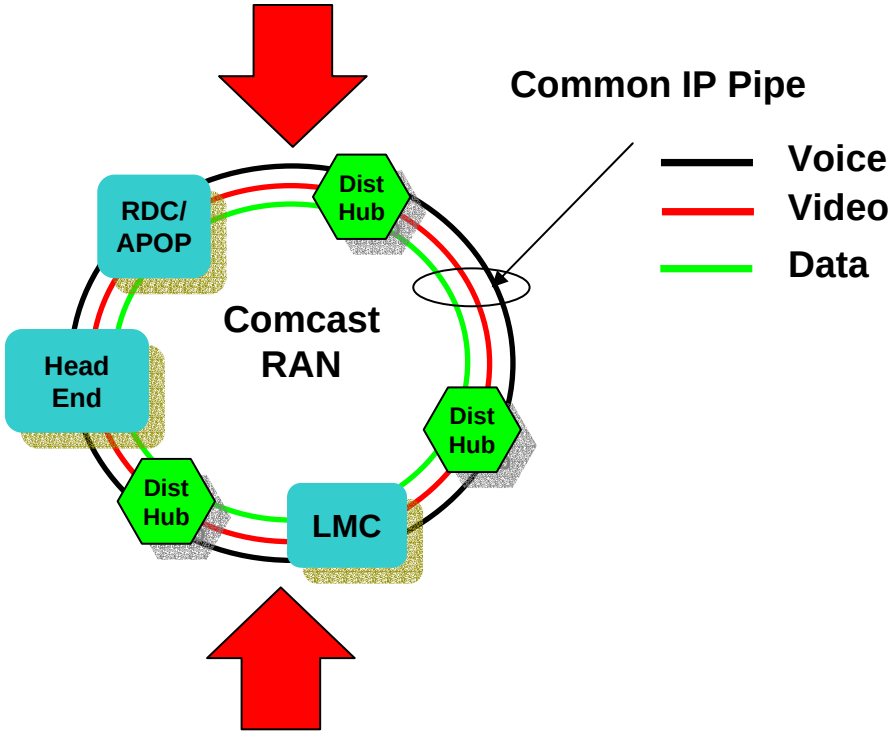

Able support low latency, minimal jitter and high reliability required to delivery quality voice services

Features

- ❑ Provides a common pipe for the delivery of all services
- ❑ Bandwidth can be allocated dynamically across services to accommodated demand changes
- ❑ Supports a range of quality of services levels
- ❑ Allows servers to be consolidated at a regional level

Benefits

- ❑ Reduces the cost of transporting all services by up to 40% (In Greenfield deployment)
- ❑ Allows Comcast to offer a variety of end customer service levels
- ❑ The network architecture more scalable and future-proof minimizing future capital expenditure
- ❑ Enhances operational efficiency by unifying operations

The backbone interconnects regional networks to create a unified national network

Comcast National Backbone Network

- ***Over 19,000 route miles of national fiber***
- ***Covers 95+% of Comcast Homes Passed***
- ***All major peering points***
- ***Transport capability of >2500+ Gbps in 2006***
- ***40Gbps IP links in 2006***
- ***QoS based voice, video and data over IP***

Features

- ❑ Links all the regional networks into a unified network
- ❑ Consolidates peering and interconnection with other operators
- ❑ Managed QoS service delivery end-to-end with no third-parties involved
- ❑ Supports centralized management functions

Benefits

- ❑ Very flexible and low cost linear and OnDemand video distribution
- ❑ Backbone transport payments to transit providers reduced
- ❑ Reduce overall operational costs and complexity by supporting centralization
- ❑ Increased revenue opportunities by providing high quality, national end-to-end services
- ❑ Opens up wholesale opportunities

comcast

VoD Adoption is Shaping New Network Needs

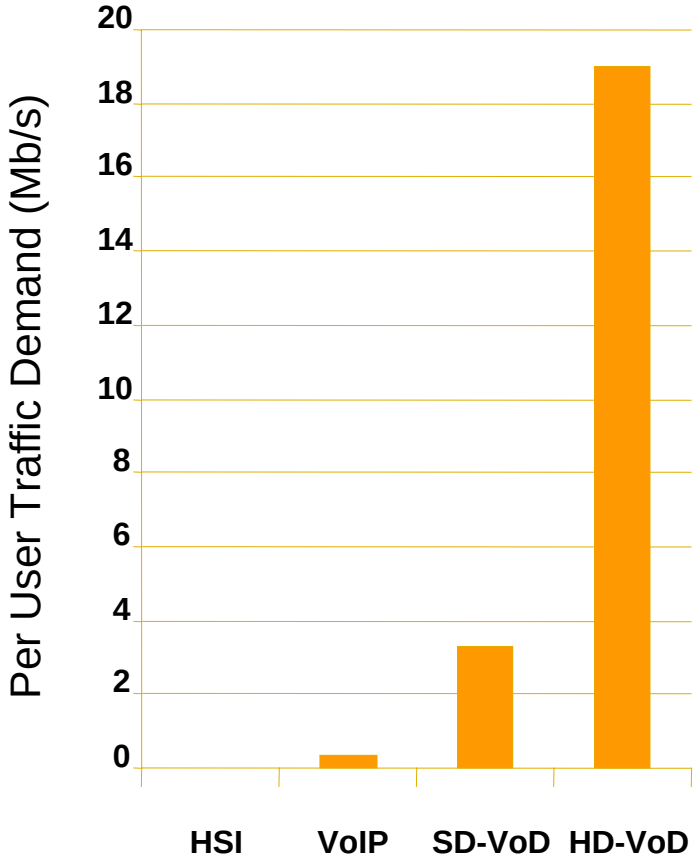

Architectural Drivers

- Rebuilds Push Fiber (FTTN)
- DOCSIS Enables Ethernet (ETTH)
- HSI Drives IP Foundation
- VoIP Drives Availability & dQoS
- VoD Sets Capacity & Economics
- TV & STB (DSG) Drive Huge Scale
- SIP/PCMM Shapes Future services

comcast.

Example of VOD Server Connectivity

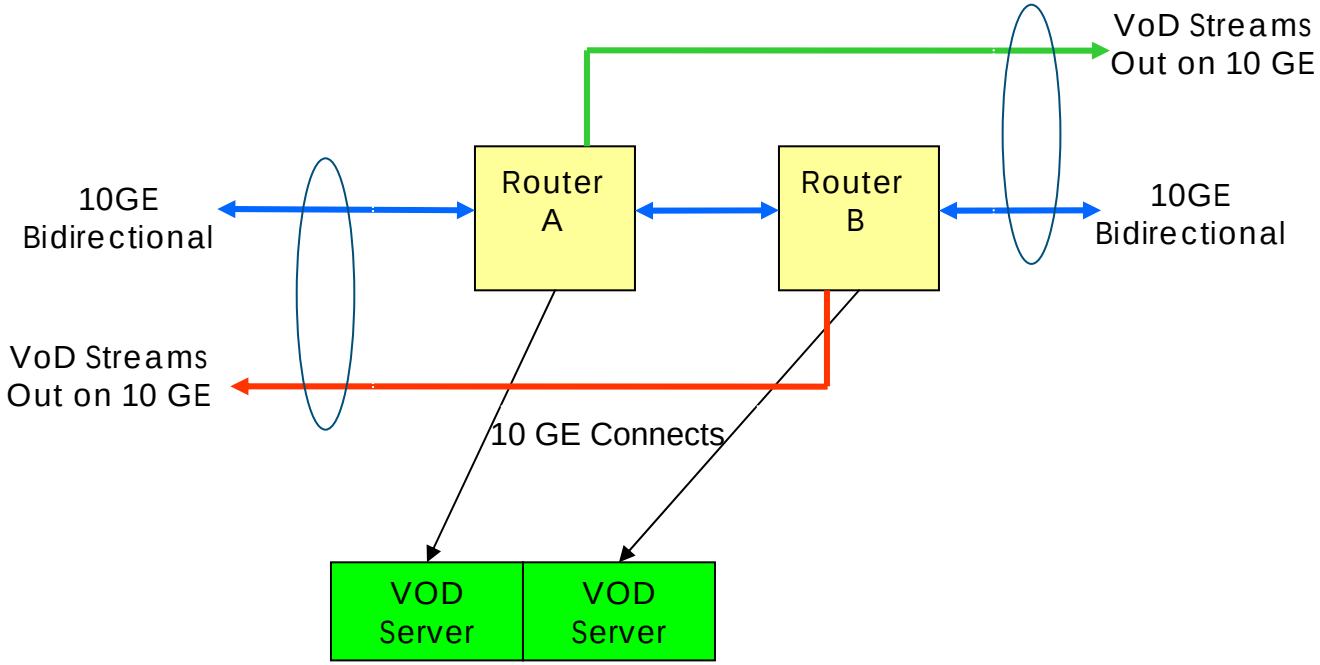

VOD streams for each Hub/OTN are distributed evenly between:

- Two downstream fibers (diverse paths) from HE to Hub/OTN
- Two routers at Hub/OTN or two distinct ports on the same router

Example of Edge Connectivity

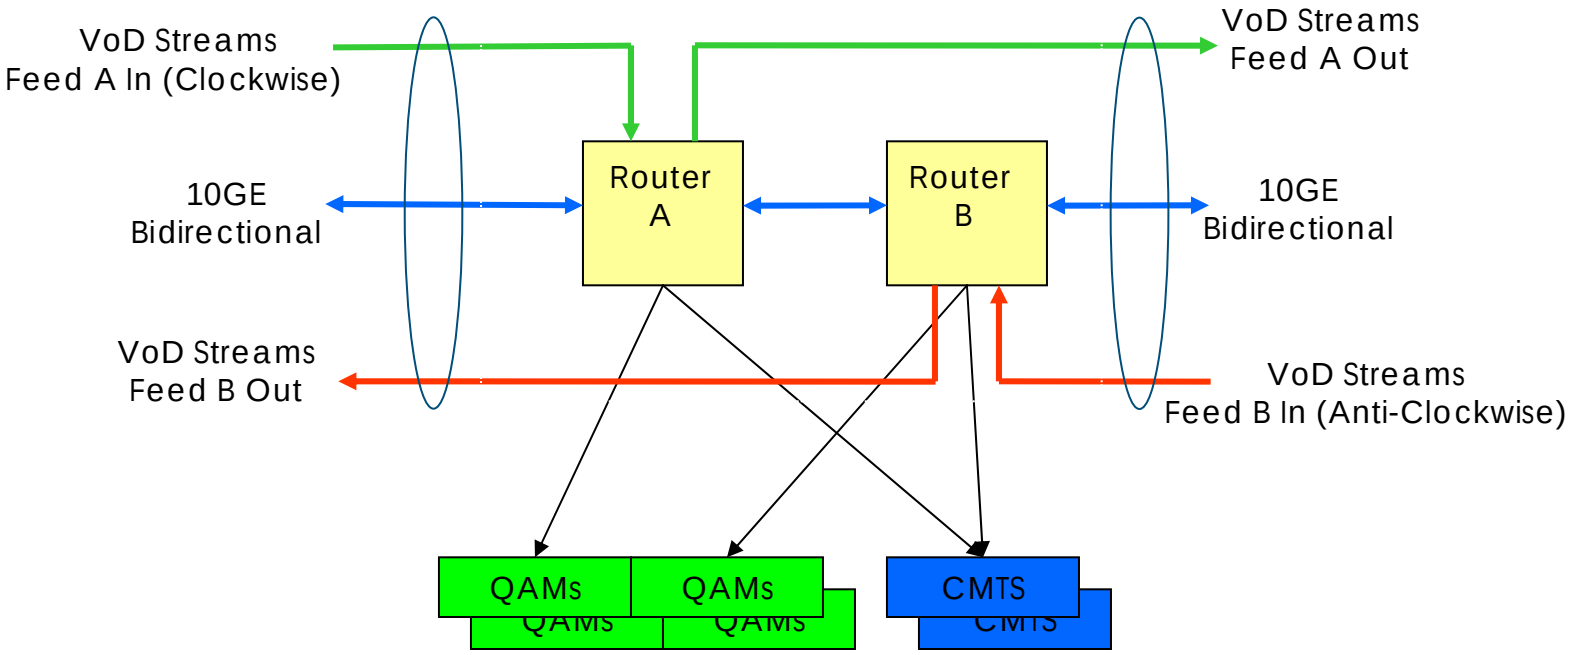

VOD streams for each Hub/OTN are distributed evenly between:

- Two downstream fibers (diverse paths) from HE to Hub/OTN
- Two routers at Hub/OTN (when available) or two distinct ports on the same router

Example of VOD Implementation

1M subscribers, 1 HE, 6 Dhubs, 100,000 Total Video Streams

265 SD streams per 1 Gbps, 56 HD streams per 1 Gbps

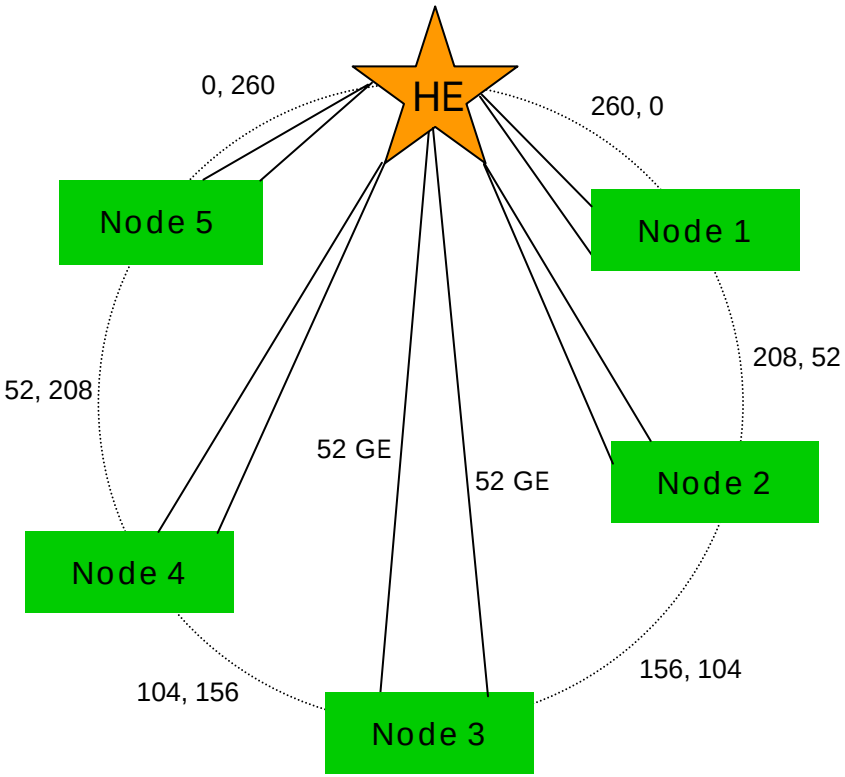

2007

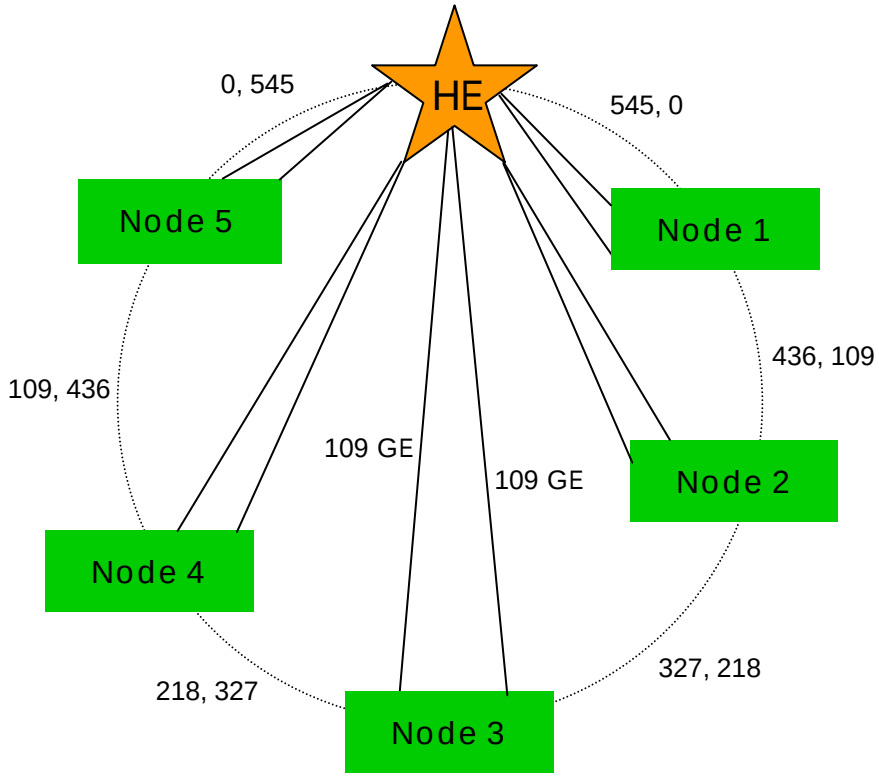

2010

90% SD ~ 18,000 SD/ Dhub ~ 68 GE/ Dhub

10% HD ~ 2,000 HD/Dhub ~ 36 GE/ Dhub

50% SD ~ 10,000 SD/ Dhub ~ 38 GE/ Dhub

50% HD ~ 10,000 HD/Dhub ~ 180 GE/ Dhub

comcast.

Limitations imposed by parallel nx10Gbps links

- Multiple ways to utilize parallel nx10G links
  - L3 ECMP
  - L2 Link-bundling
- Distribution of traffic over parallel links done via flow-based hash mechanism in both cases
  - Per-packet/round-robin distribution CANNOT be used due to packet re-ordering which results in significant drop in “goodput”
- Effectiveness of flow-based Hash distribution determined by traffic characteristics
  - Flow diversity - large number of flows
  - Average bandwidth per flow - determines number of flows that can be supported on any given link
- Flow-based hash mechanism CANNOT guarantee equal distribution of load
  - With ideal traffic characteristics it is statistically possible to uniformly distribute load over all links

comcast

Limitations imposed by parallel nx10Gbps links (continued)

- Effectiveness of load distribution depends on
  - Hash algorithm
  - Diversity of hash input (variability in SRC/DEST IP address)
  - Number of flows/size of flows
- Hash can result in un-equal load distribution
  - Caused by Non-ideal traffic characteristics such as:
  - Small distribution of src/dest IP addresses
  - High per-flow bandwidth
- Un-equal load distribution can results in under-utilization of available capacity
  - May potentially cause artificial congestion and packet loss
- 10x10Gbps is not the same as 1x100Gbps from a real throughput perspective
- How many parallel 10Gbps links do you need to match usable bandwidth on one 100Gbps link?
  - Depends on traffic characteristics…

comcast

Key Points

- There is a market need for 100GE
- Must standardize 100GE, pre-standard implementations are less accepted in the market
- Non-standard 40/80GE solutions will only slow 100GE development and adoption
- Providing a standard, cost effective solution is the best way to make 100GE successful

comcast.

Thank You!