



Video Transcoding Over NVMe

August 8, 2019



Flash Memory Summit

Daniel Zhou

Director of Firmware Development
NETINT Technologies

Topics

- Video Transcoding Demand is Growing
- Video Transcoding Needs to Scale at Cloud and Edge
- ASIC-based Computational Storage
- Application of NVMe for Transcoding Control

Video Traffic is Large and Growing

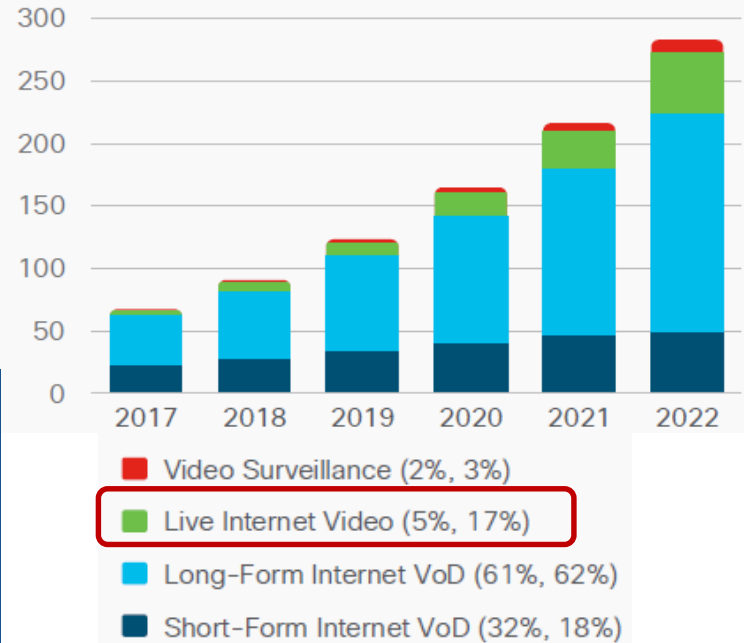
IP Video will be
82%
of Global IP Traffic
by 2022

33% CAGR
2017-2022

Exabytes
per Month

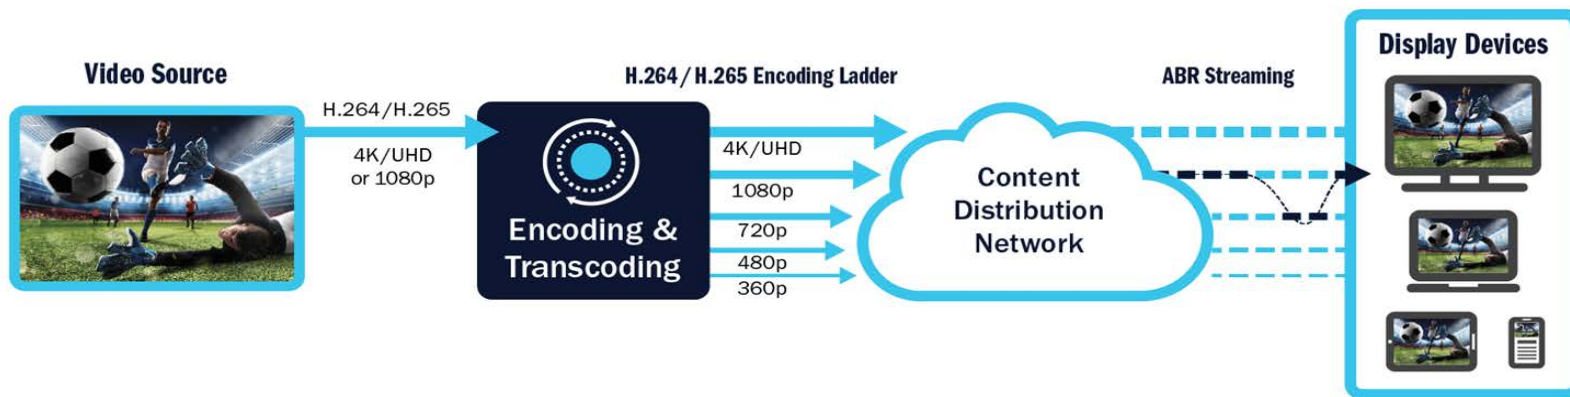
Live Video
will grow
15x
By 2022

Global Internet video by subsegment



Source: Cisco VNI Global IP Traffic Forecast. *Figures (n) refer to 2017, 2022

Video Transcoding in Cloud

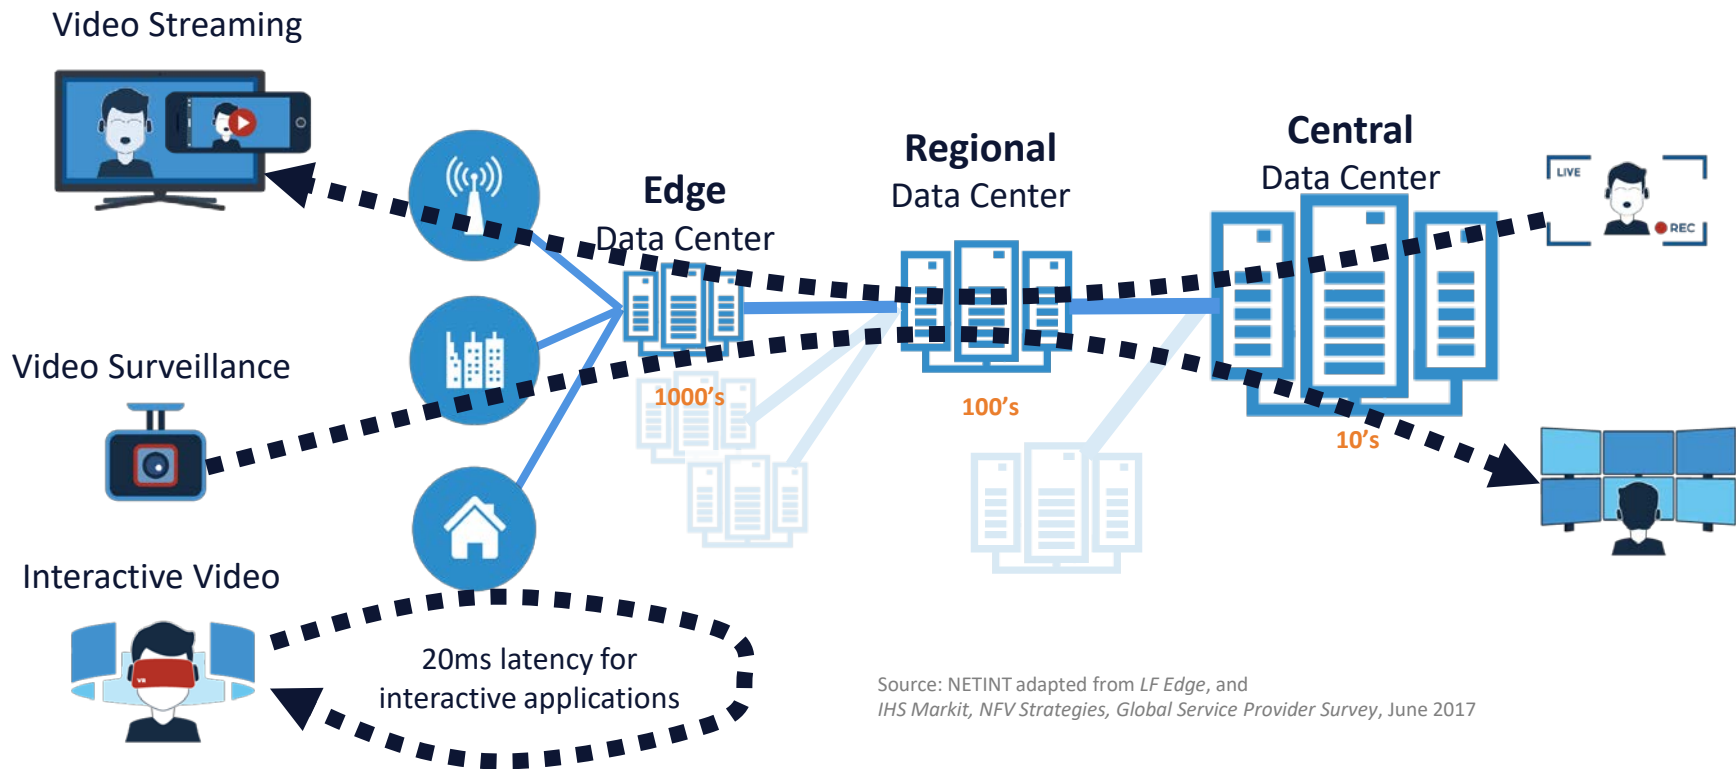


- Live Streaming
- Long-form video
- Short-form video
- Video Surveillance as a Service

- > Real-time, low-latency
- > Highest quality
- > High quality
- > Real-time

Video Edge Encoding and Storage in the Video Cloud

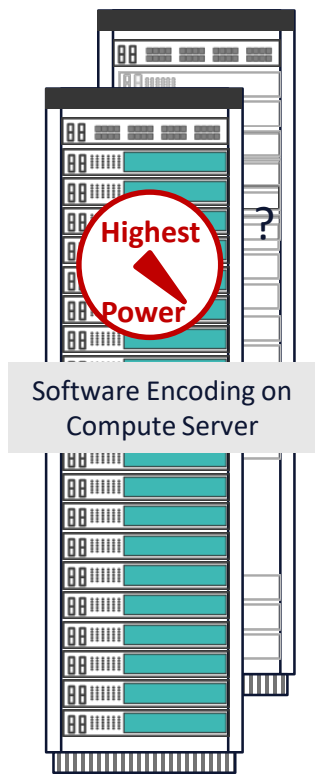
Use Cases with primary video flows




Source: NETINT adapted from *LF Edge*, and *IHS Markit, NFV Strategies, Global Service Provider Survey*, June 2017

Video Encoding Alternatives Compared: Density and Power

Approximated infrastructure required for
80x 1080p30 Encoding Streams, or
40x Typical Encoding Ladders.



Software Encoding on
Compute Server



GPU Accelerators
Hosted in Server



FPGA Accelerators
Hosted in Server

Lowest Power

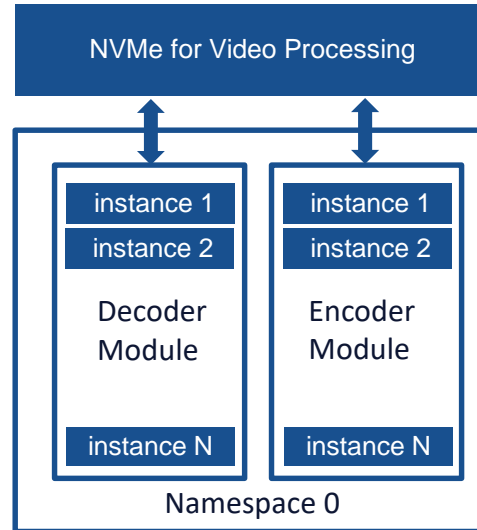
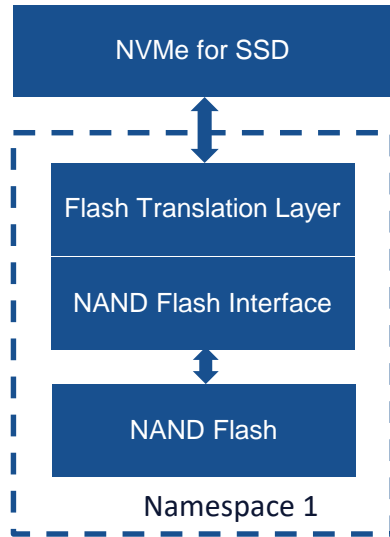
T408 Transcoders
Hosted in
1RU NVMe Server

Codensity T408 Video Transcoders

- NVMe Interface
- ASIC-based Transcoding
- Highest Density
- Least Rack Space
- Lowest Power


Codensity
T408 Video Transcoder
Powered by
NETINT

Application of NVMe to control SSD and video processing

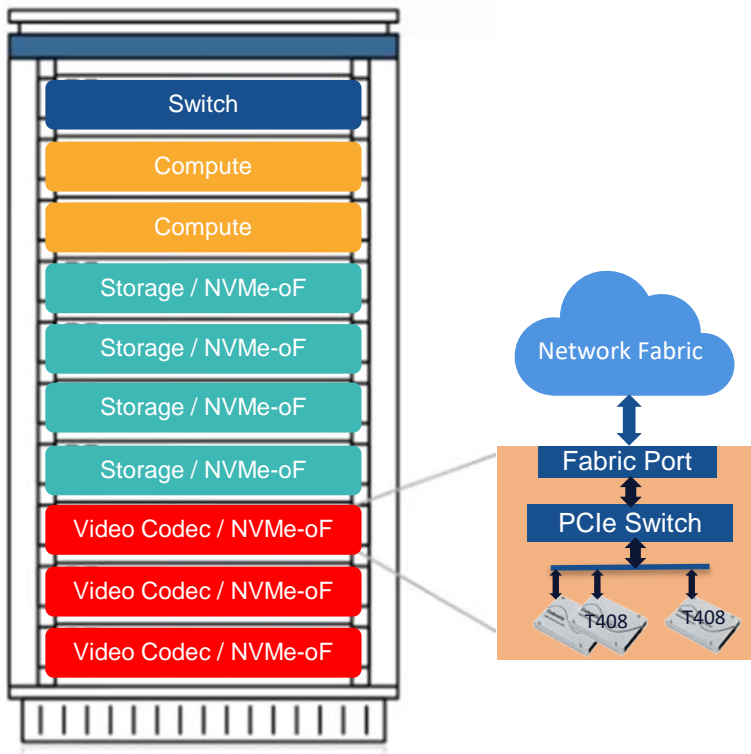


Scaling-up Video Transcoding within NVMe Servers

- Leverages standard NVMe drivers
- Transcoding U.2 modules plug into SSD slots of NVMe Server

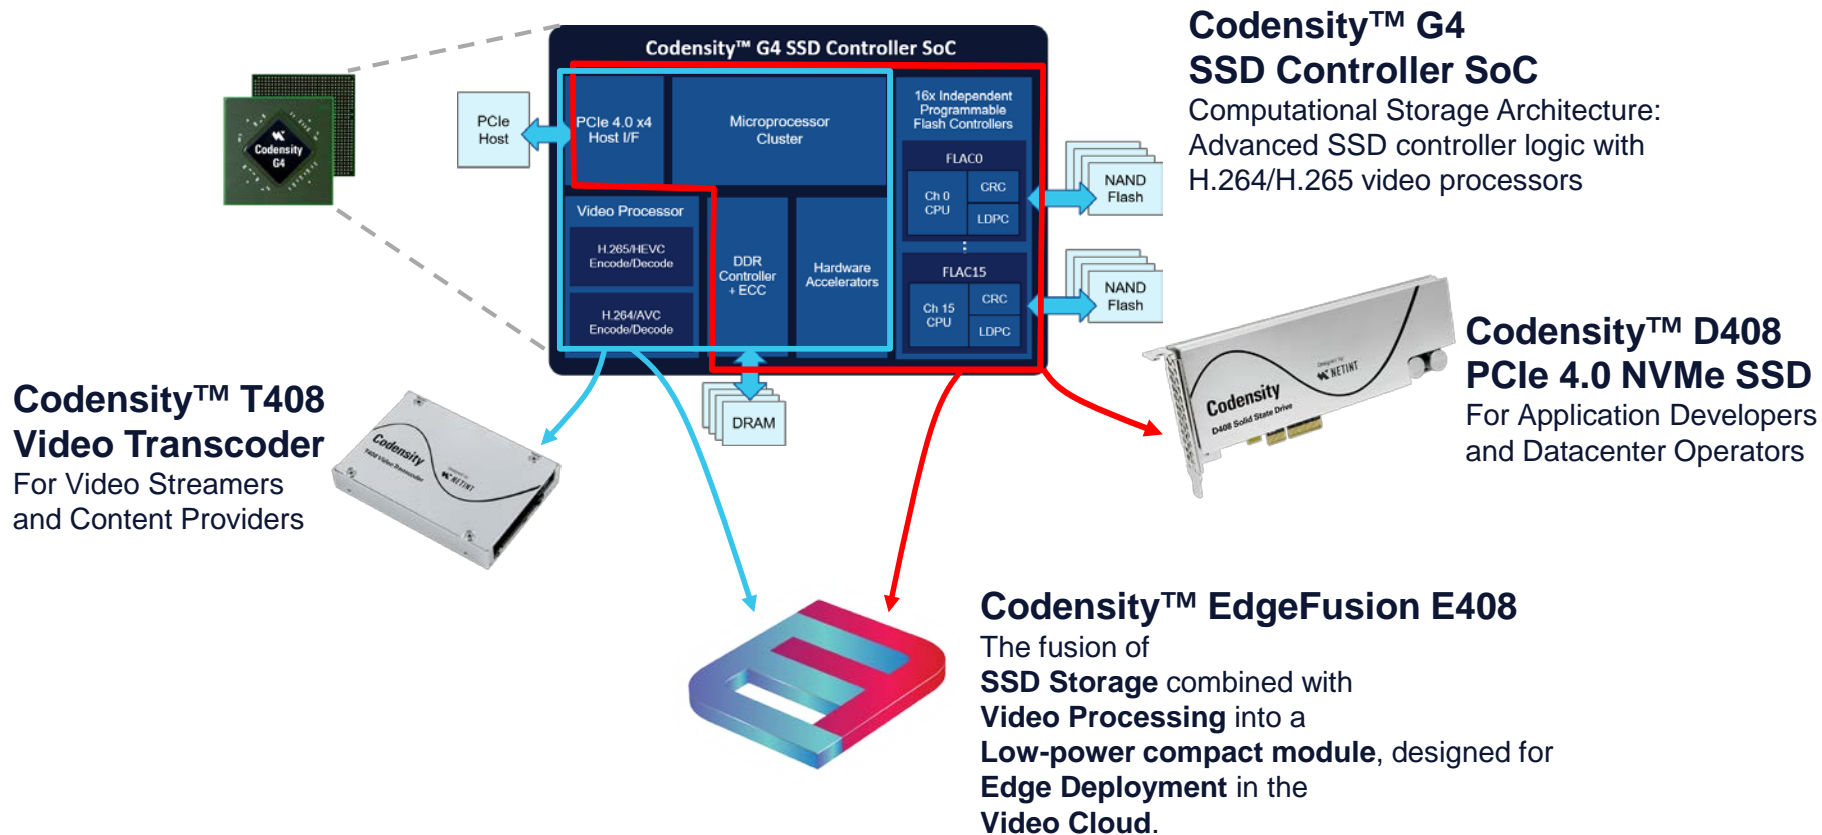


Scaling-out Video Transcoding with NVMe-Over-Fabrics



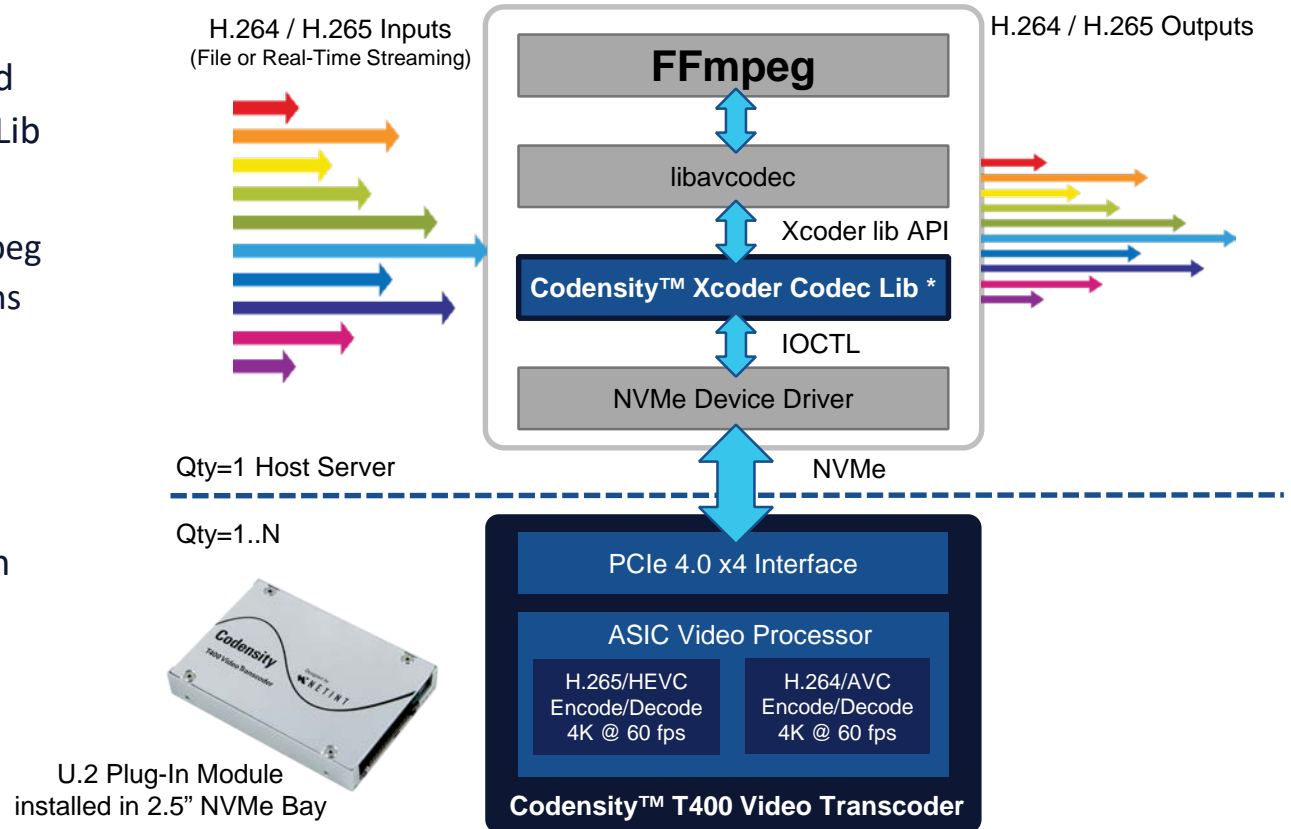
- Work with proven NVMe and NVMe-oF device drivers
- Scaling video transcoding resources outside servers
- Sharing video transcoding resources among servers

Scaling-up Video Transcoding Together With Storage



Codensity T408 Video Transcoder – Software Integration

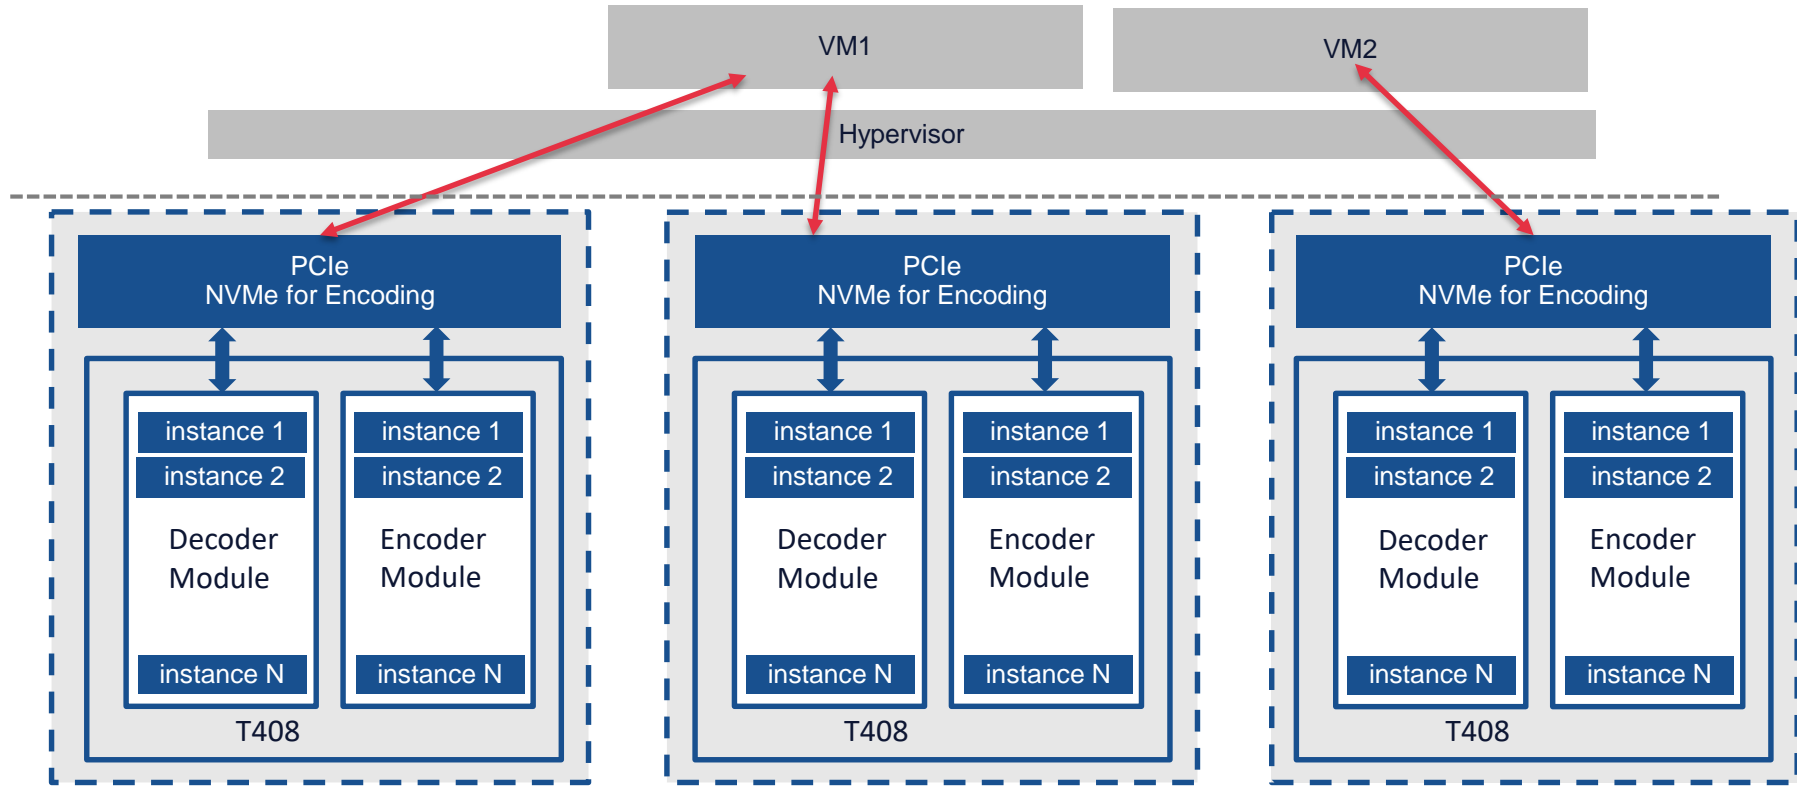
- FFmpeg integration achieved by installing FFmpeg Codec Lib and SDK into host server
 - Seamlessly abstracts FFmpeg video transcoding functions from 1 or more T400 transcoder modules
- T400 video transcode functions controlled through standard NVMe protocol



U.2 Plug-In Module installed in 2.5" NVMe Bay

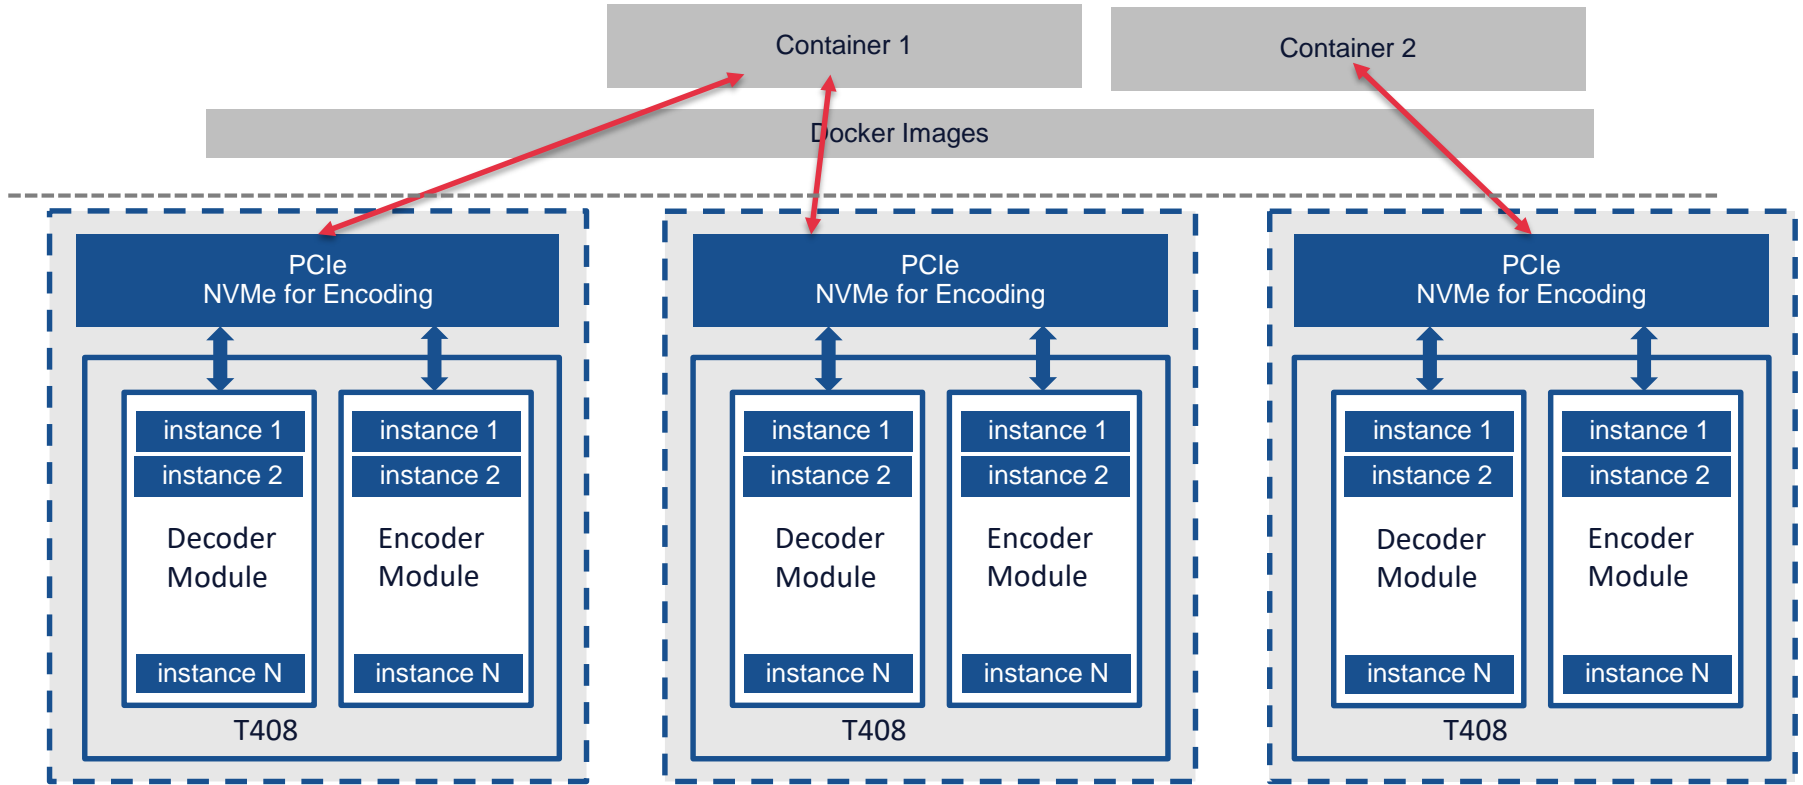
Virtualization for Cloud with Device-passthrough

Allocate Different T408s to Different Virtual Machines



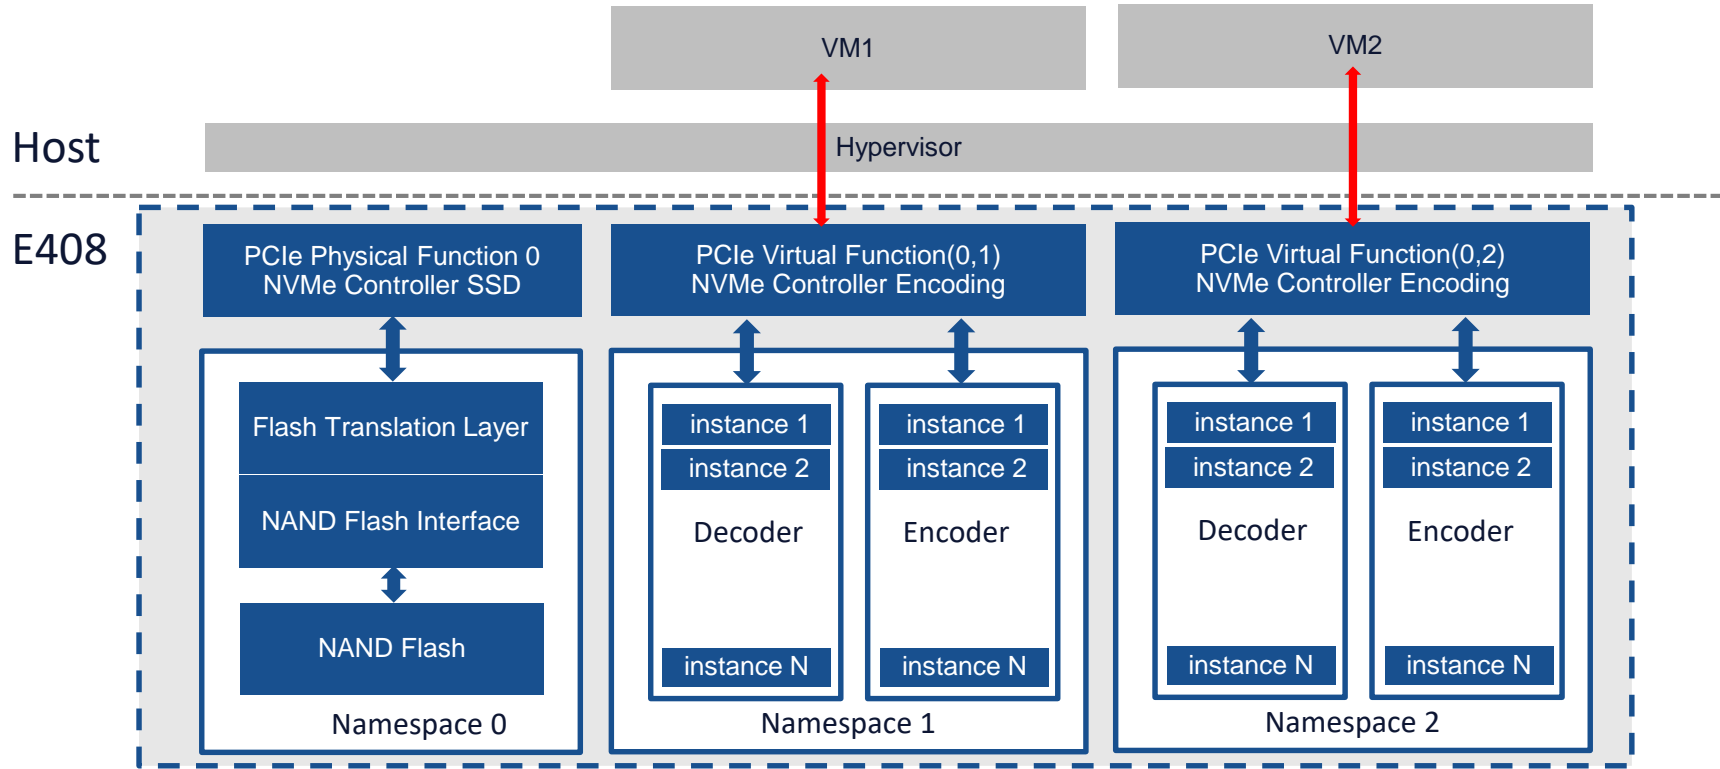
Virtualization for Cloud with Device-passthrough

Allocate Different T408s to Different Containers



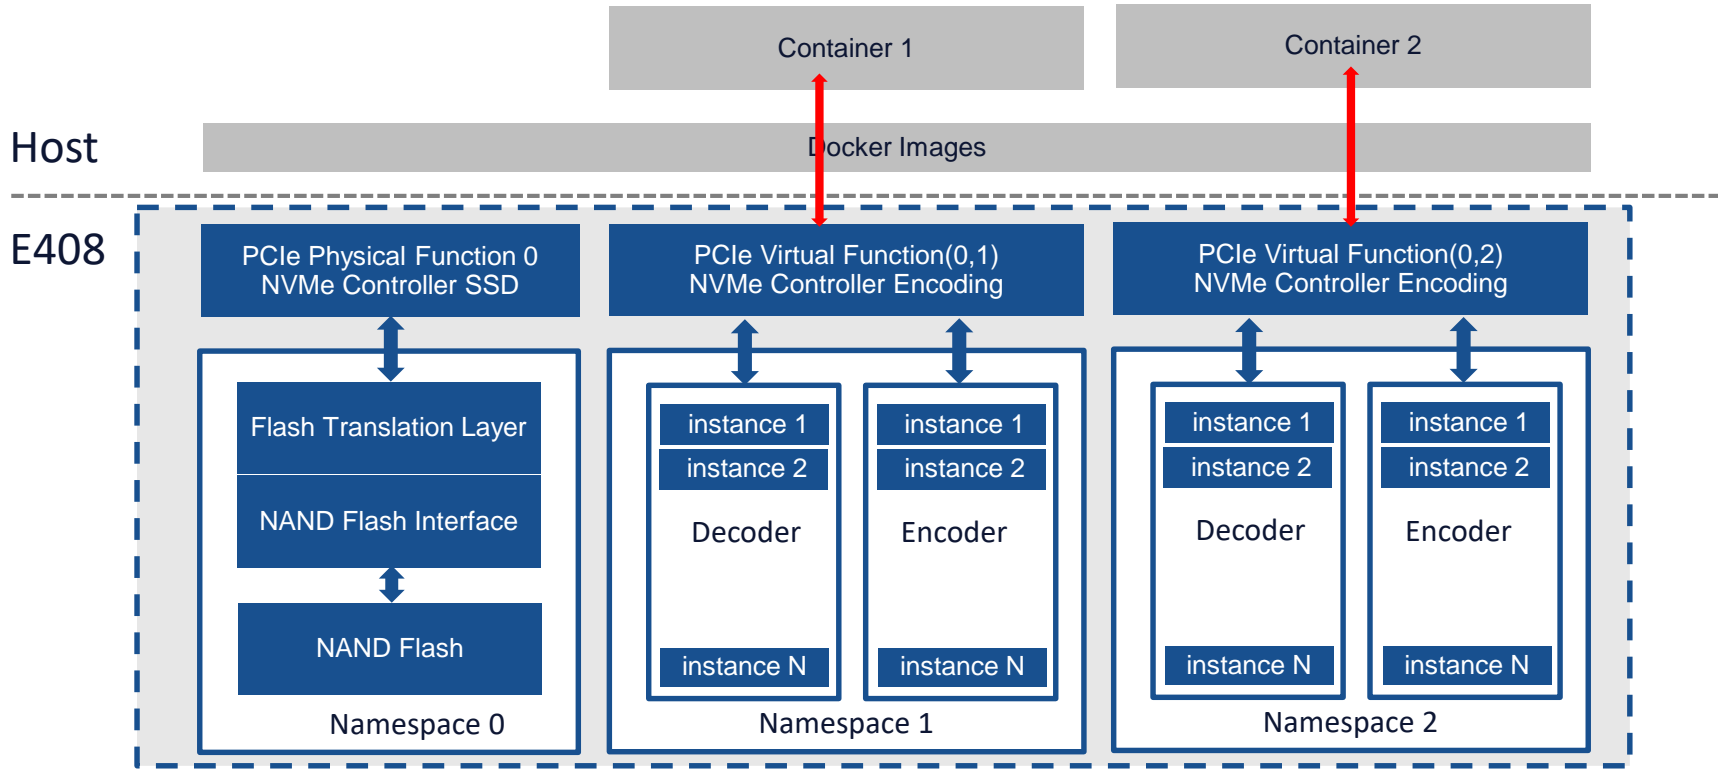
Virtualization for Edge with SR-IOV

Share One E408 among Virtual Machines



Virtualization for Edge with SR-IOV

Share One E408 among Containers



Summary

- Video content on Internet is dominant and growing.
- Video processing migrating to network edge.
- It is challenging to scale up video transcoding in the cloud or edge.
- Solution is ASIC-based NVMe-enabled Computational Storage
 - Highest Density, Economic, Scalable, Pragmatic
 - NVMe provides high-performance control for SSD and video transcoding
 - Leverages existing NVMe solutions for scaling-out (NVMe-oF) and Virtualization



Flash Memory Summit

Questions?

Visit our booth#724 during FMS 2019,
or www.netint.ca for more information

Daniel Zhou

Director of Firmware Engineering,
NETINT Technologies

daniel.zhou@netint.ca