



# CDN Use Case Revisited

Marc Fiuczynski

A close-up photograph of a person's hands holding a tablet computer. The tablet screen displays a soccer game in progress, with a player in a blue jersey in the foreground and a crowd in the background. The person holding the tablet is wearing a light-colored, textured sweater. The background is blurred, showing a checkered pattern, possibly a shirt or blanket. The text "The ultimate in video quality over IP at scale" is overlaid in a bold, orange font across the center of the image.

**The ultimate in  
video quality over  
IP at scale**


## Five years from now... Akamai's CDN traffic **may** look like:

- Daily peaks of 500 Tbps from Global CDN with ~ 3x avg-to-peak ratio
- Event peaks >10 Tbps from Private CDNs with >> 3x avg-to-peak ratios
- “*CDN Interconnection*” between Global CDN and Private CDNs (optional)
- 90% of the Global CDN traffic as VoD over IP. The rest...
  - Extremely high-definition video calling and conferencing (most of the rest?)
  - E-commerce and retail traffic with very high security requirements
  - IoT (extremely low volumes but massive transaction rates)
  - VR-based massive multiplayer gaming
- 99%+ of the Private CDNs traffic as VoD/Live over IP.
- Most, if not all, traffic encrypted
- POP deployment sizes ranging from 1 Gbps to 10 Tbps
- Many tens of thousands of POP deployments of every possible kind
- Many different delivery protocols; multicast-on-demand finally possible?!




# Traditional Distributed CDN


*Deployed in Top-Level  
Datacenters*



# Dynamically Scaling Out / In CDN Capacity



# Scale Out/In Deep with NFV/SDN



## vCDN: Virtualized CDN Deployment Scenarios

- **Base vCDN**  
Static deployment onto a virtualized infrastructure → NFVI (KVM, ESXi, Containers)
- **Elastic vCDN**  
Bandwidth scaling out/in for peak events on a virtualized infrastructure → VIM (OpenStack, vSphere)
- **Deep vCDN**  
Deployments in operator network for backhaul/access net offload via NFV → NFVO

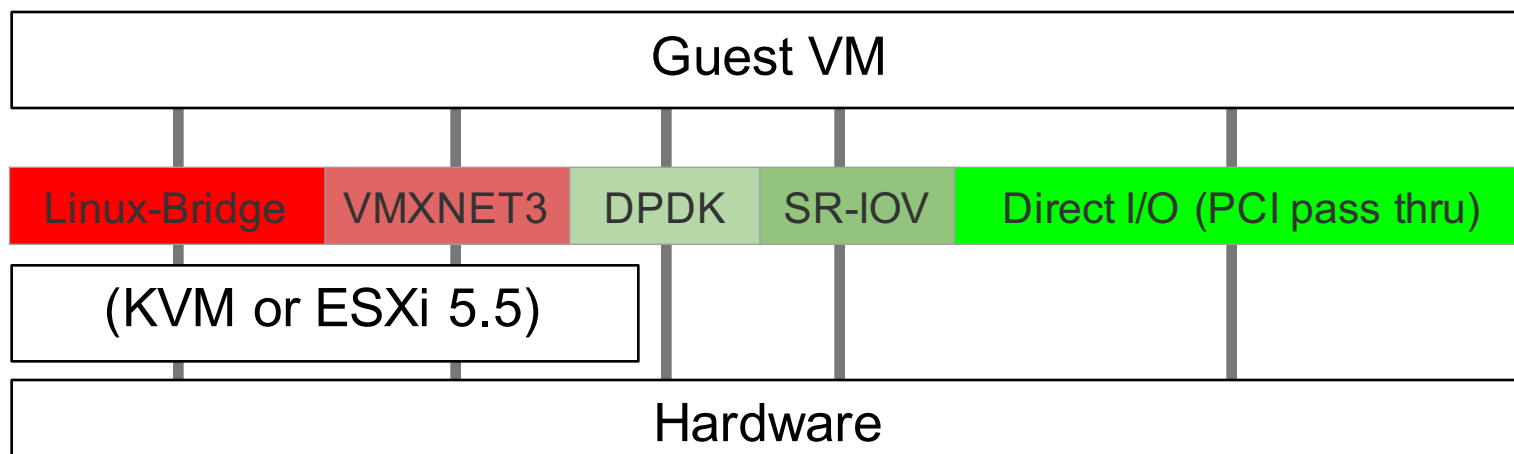


## Base vCDN Considerations

- High Performane Network I/O
- High performace Storage I/O
- Security

## NFVI: Networking performance

- TCP Bulk delivery issues -- e.g., HTTP Delivery of VoD/Live Content
  - send side ok: TCP GSO/TSO well supported by virtual machines and overall network stack
  - receive side challenges: pure ack TCP floods caused by *simple* consumer devices
    - consumer devices tend not to have TCP LRO/GRO support, ack every other segment
    - @ 10Gbps video delivery this produces ~ 500K PPS of pure ack TCP responses



## NFVI: Storage I/O performance

Storage Solution	OS Disk	Cache Disks
local disk (HDD or SSD)	1	2
iSCSI (e.g. HP StorageWorks P2000)	1	2
Fibre Channel over Ethernet (FCoE)	1	2
OpenStack Ceph	1,3	2,3
VMware vSAN	1,3	2,3
VMware NFS Datastore	1,3	2,3
EMC VNX or EMC Isilon	1,3	2,3
...		

### Legend:

1. *redundancy preferred (e.g., RAID 1)*
2. *direct disk **or** no replication*
3. *untested by Akamai, expecting suitability feedback from partners / customers*

## NFVI: Microbenchmarking to identify poor configurations


- Networking: apache bench (`ab`) with `nginx` server configuration
  - emulate *simple* consumer device video delivery
    - no TCP GRO/LRO to produce pure ack floods
- Storage: `fio` utility based with CDN read/write profiles
  - compare with locally attached SSD/HDD performance profiles

## Base vCDN: Security Considerations

- CDN has high value secrets that must reside at the edge
  - e.g., TLS certificates
- Trusted Compute Cloud solution stack
  - compute: TPM + Intel TXT
  - storage: self encrypting disks
  - trusted on-boarding: attestation
  - trust-based access control
  - compliance reporting
  - trusted compute pools
- *Operator must have the above in place for Global CDN traffic*

## Elasitc vCDN: VIM Considerations

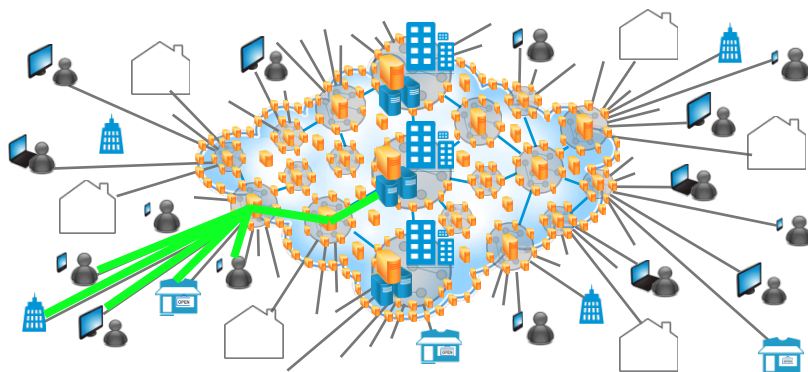
- Dozens to hundreds of PoP deployments... per operator
- Assuming *Trusted Compute Pools* across all these locations



- Ephemeral vCDN cache instantiations – e.g., for just 6-10 hours per day
- Interaction possibly with  $\geq 1$  distinct VIMs
- Resource management → stay&pay for as long as you like

## Deep vCDN: NFVO & VIM Considerations

- Hundreds to thousands of PoP deployments... per operator
- Trusted Compute Pools across all these locations?



- Per operator  $\geq 1$  NFVO,  $\geq 1$  VIMs,  $\geq 1$  administrative/network domains
- PoPs may be resource constrained; e.g., no storage  $\rightarrow$  RAM-only cache
- Ephemeral vCDN cache instantiations – e.g., for 6 hours to 6 days per month
- Zero conf techniques for vCDN cache at deeply deployed locations is critical

