Greening the Internet with Nano Data Centers

V. Valancius (Georgia Institute of Technology), N. Laoutaris (Telefonica Research), L. Massoulie, C. Diot (Thomson), P. Rodriguez (Telefonica Research)

Presented by Sean Barker

THOMSON
images & beyond

Gateway vs. Set-top-box

Home gateway

- Internet connection
- Traffic management

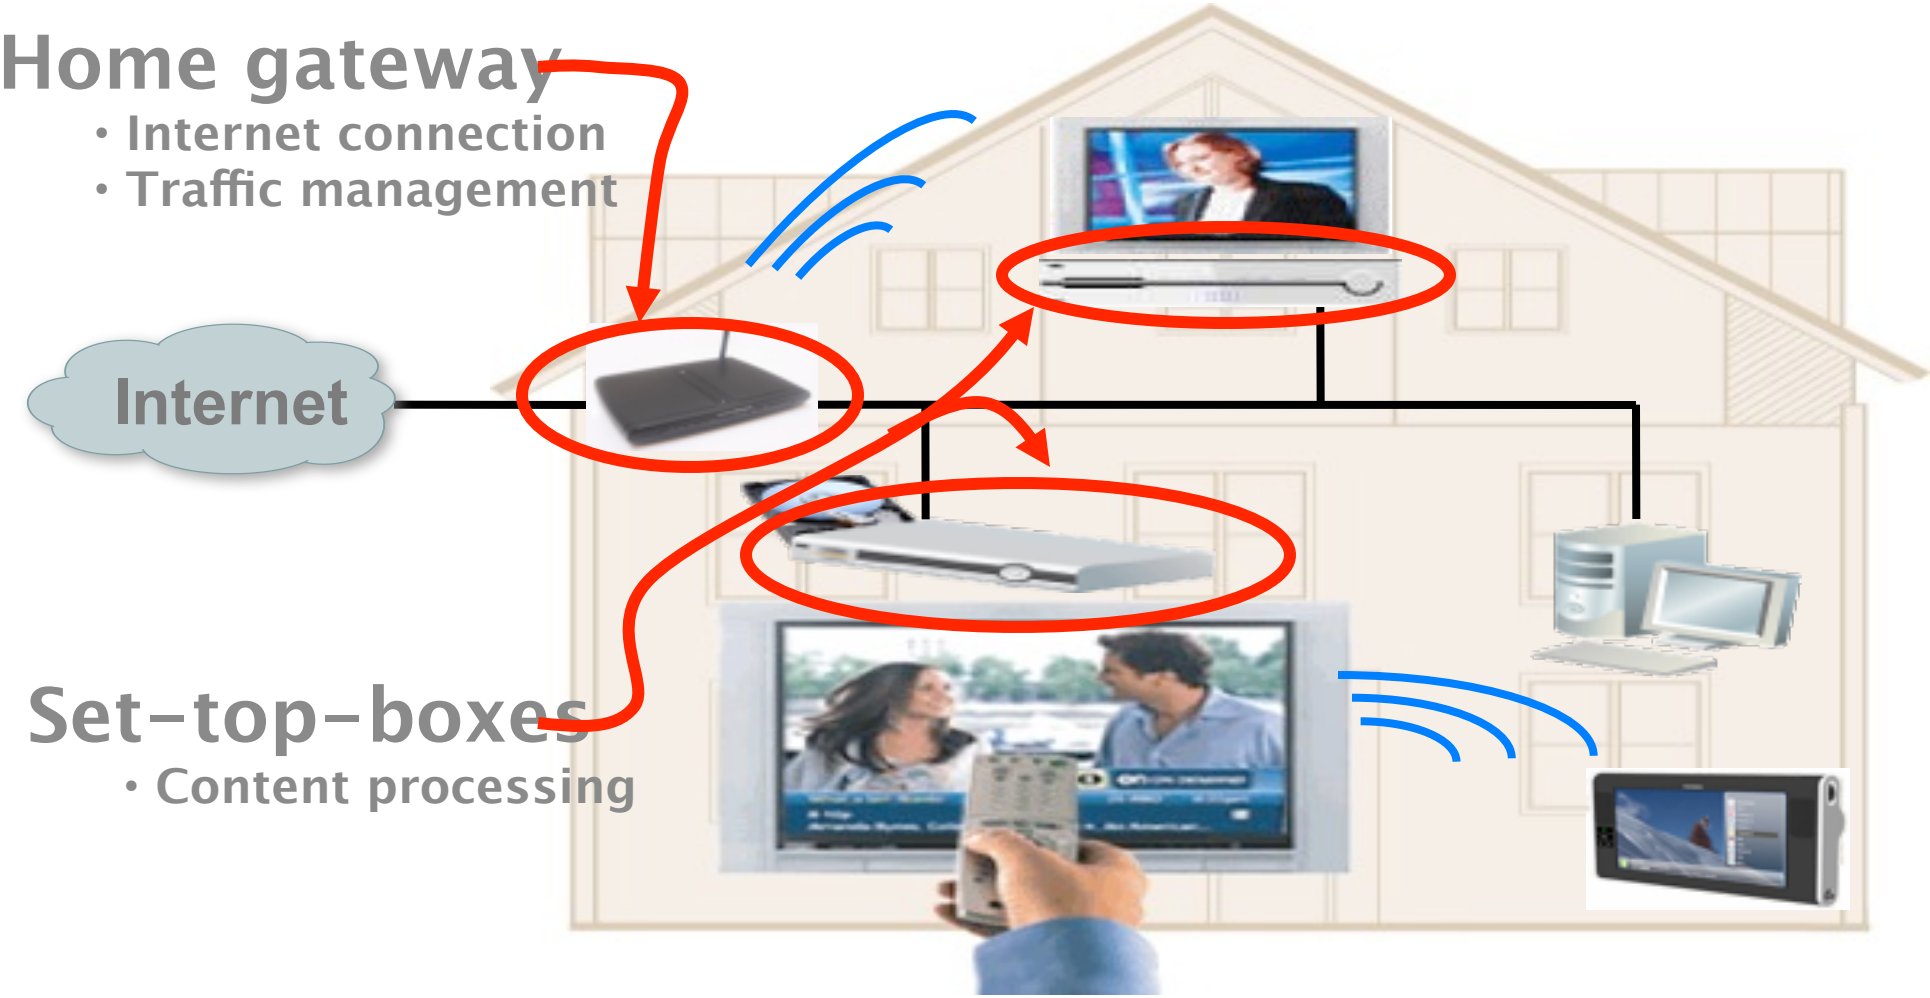

Set-top-boxes

- Content processing

Gateway as new center of media experience

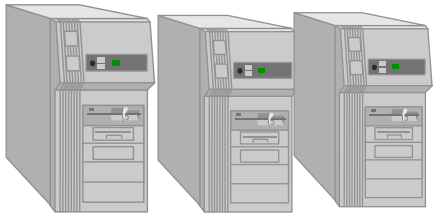

Content & Control Servers

Internet services

The Internet

Home network

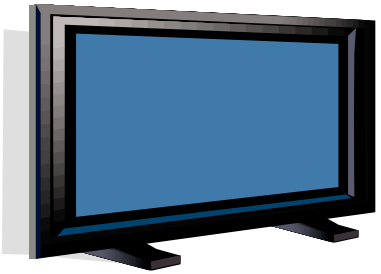

Extended home network

THOMSON
images & beyond

The current model: Data centers

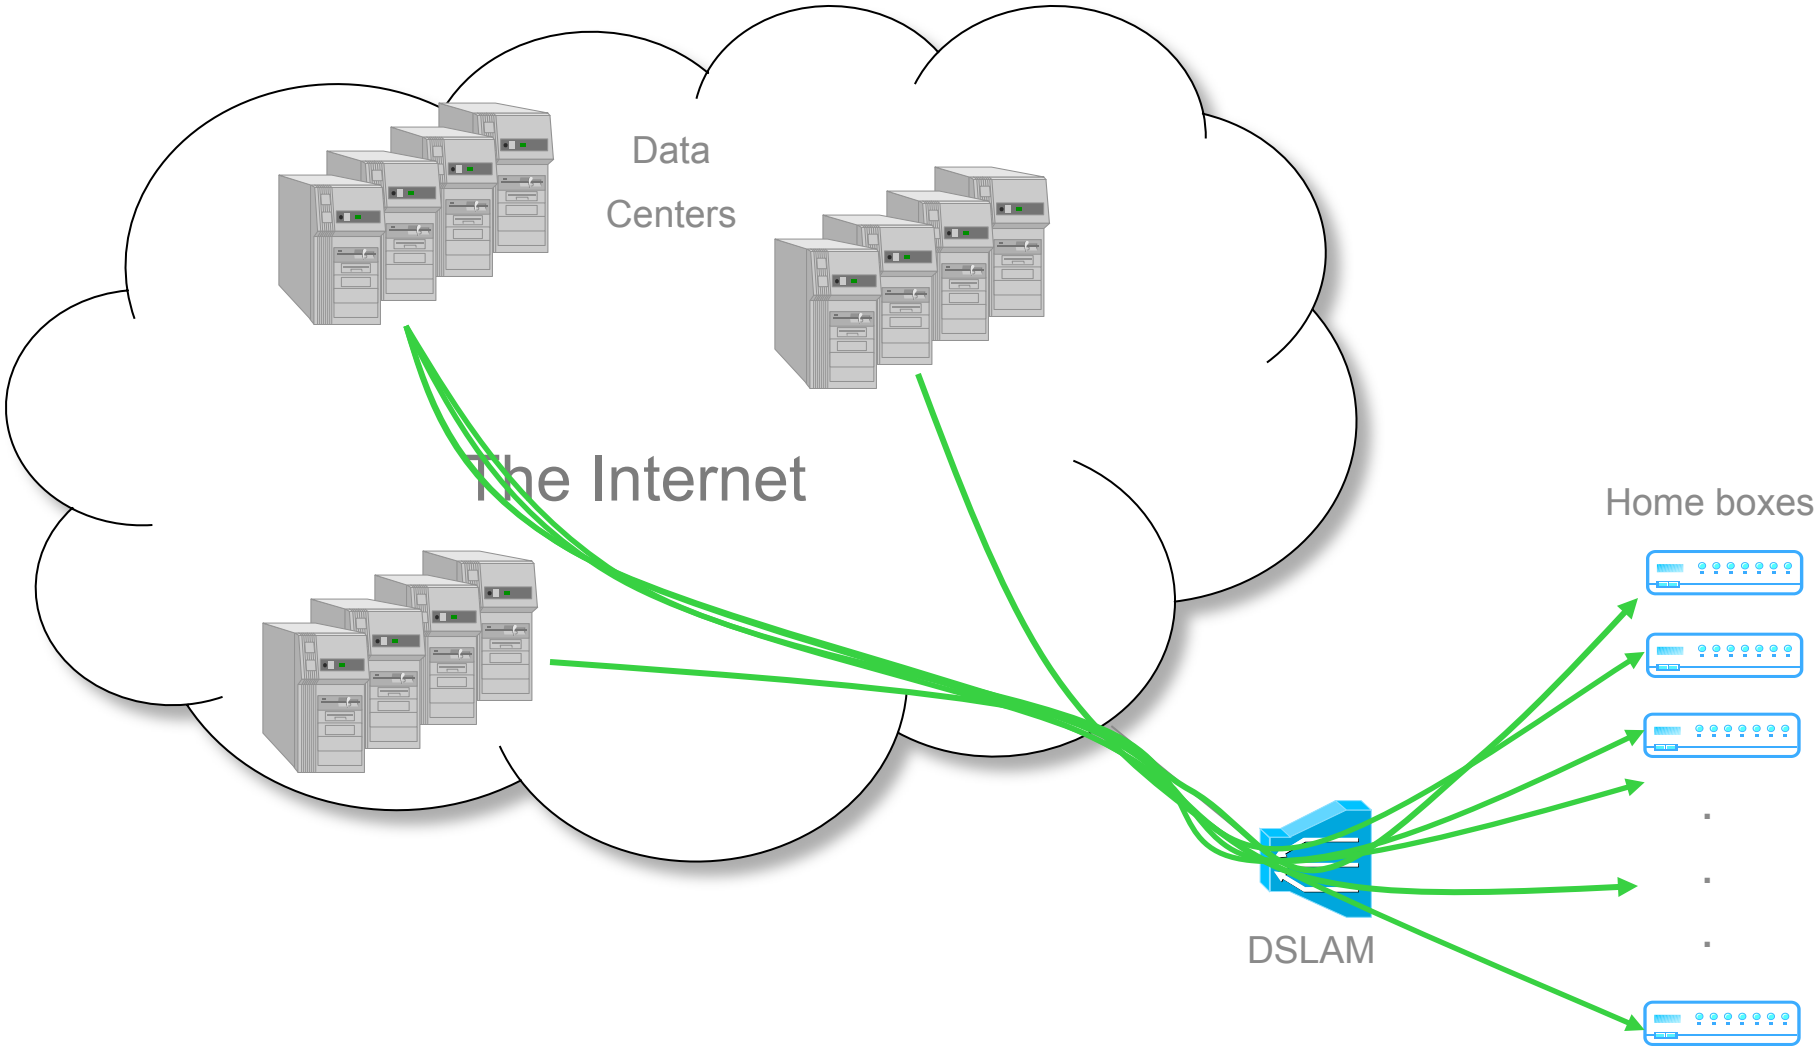

THOMSON
images & beyond

Limitations

- **Expensive**
  - High capital investment
  - Customer generally pays per byte
- **Location constraints in order to be “central”**
- **Requires a lot of redundancy to be robust**
  - Electricity shortage
  - Content availability
- **Power, power, power**
- **New service deployment is slow**
  - ISPs not encouraged to take risks, nor to deploy new services

The nano data center

- Take advantage of always-on gateways
- Add memory and stronger CPU to home gateways
- Push content to gateways when bandwidth is cheap
- Manage millions of gateways as a logical ‘single server’ using P2P infrastructure

The nano data center model

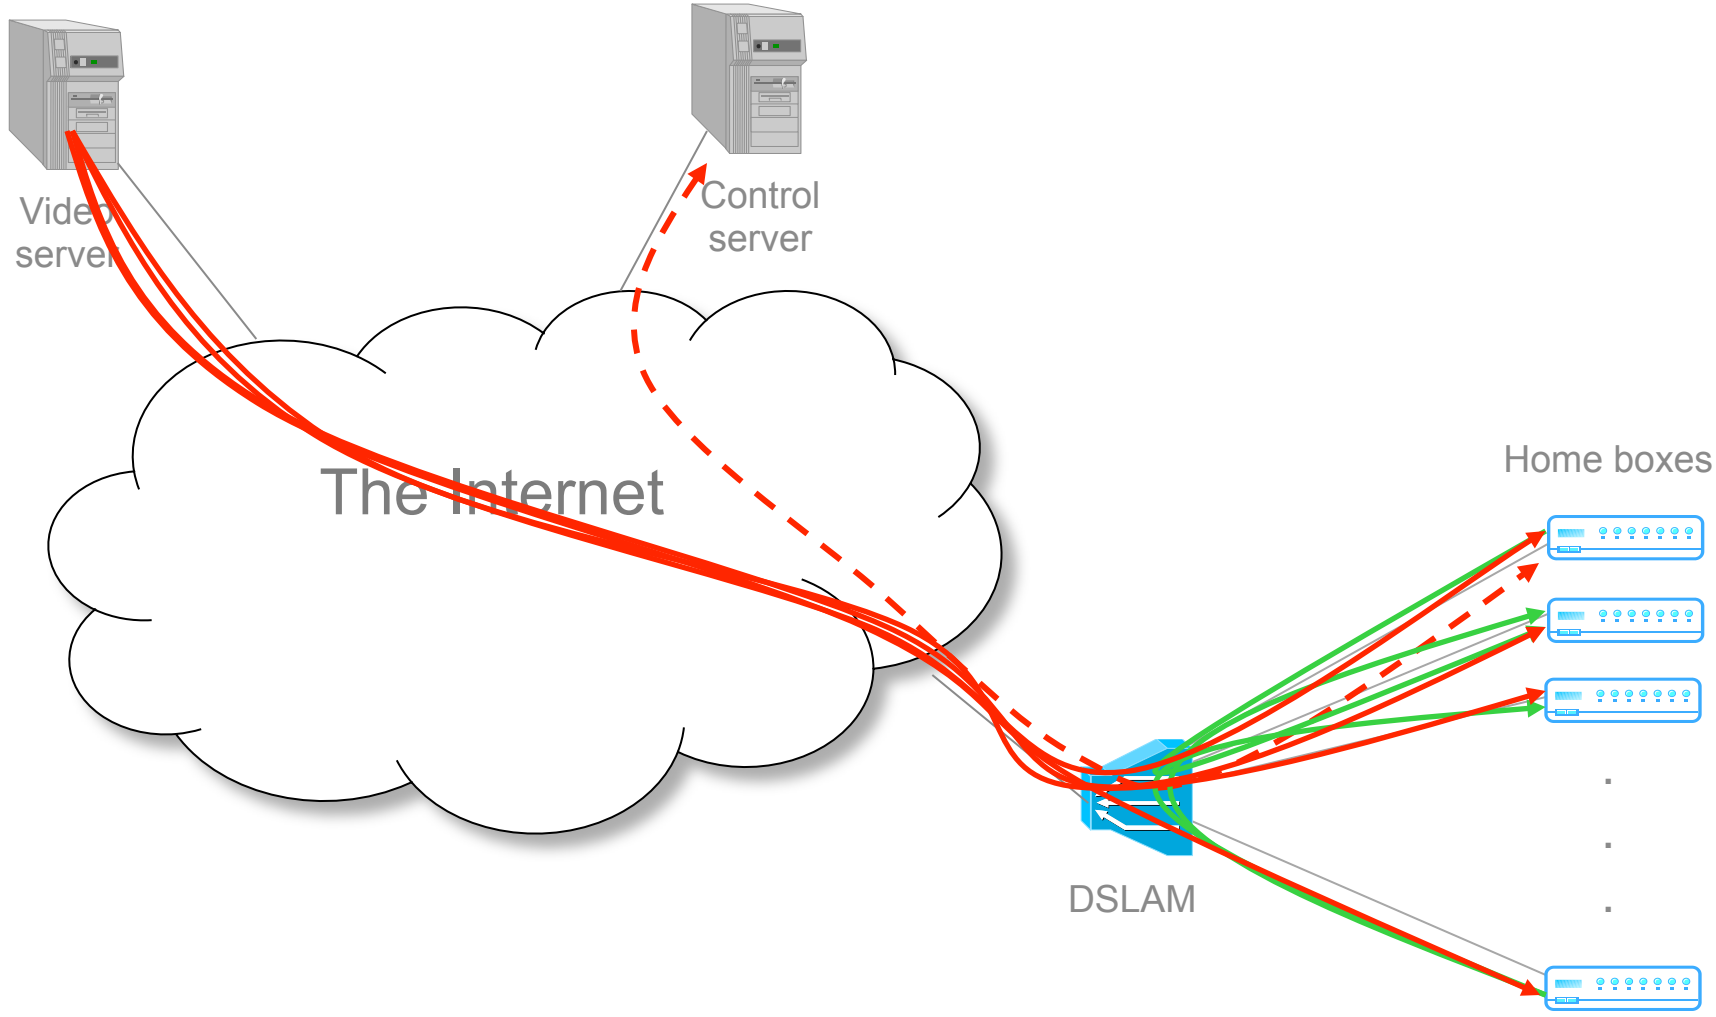

THOMSON images & beyond

The nano data center

- **PROS**
  - Multiple applications can take advantage of the model (VoD, gaming, catchTV, UGC)
  - ISP friendly
  - Reduces traffic volumes and variability on backbones.
  - Highly scalable and robust by design
  - Cheap for ISPs
  - Flexible for users
  - Localized & personalized services
- **CONS**
  - Uplink bandwidth often limited
  - Millions of boxes to manage using P2P
  - Cost of gateway
  - Incentive?
  - Privacy?
  - Always on?

Gateway uptime (*)

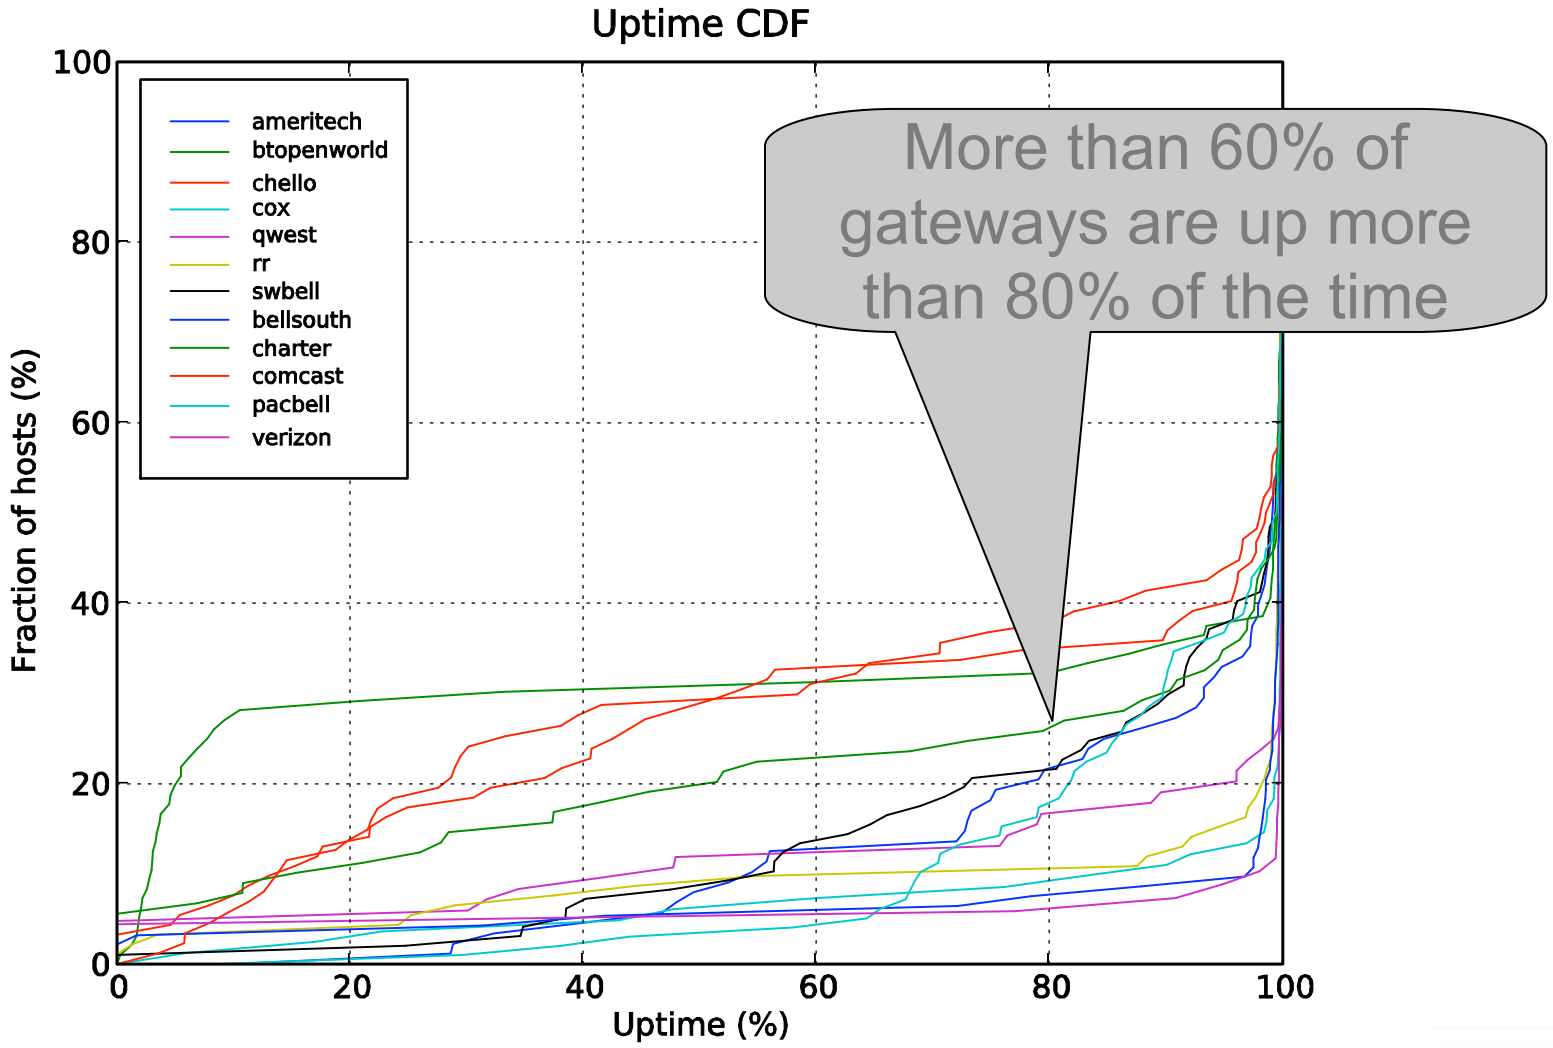

(*) Courtesy of Krishna Gummadi

Push phase

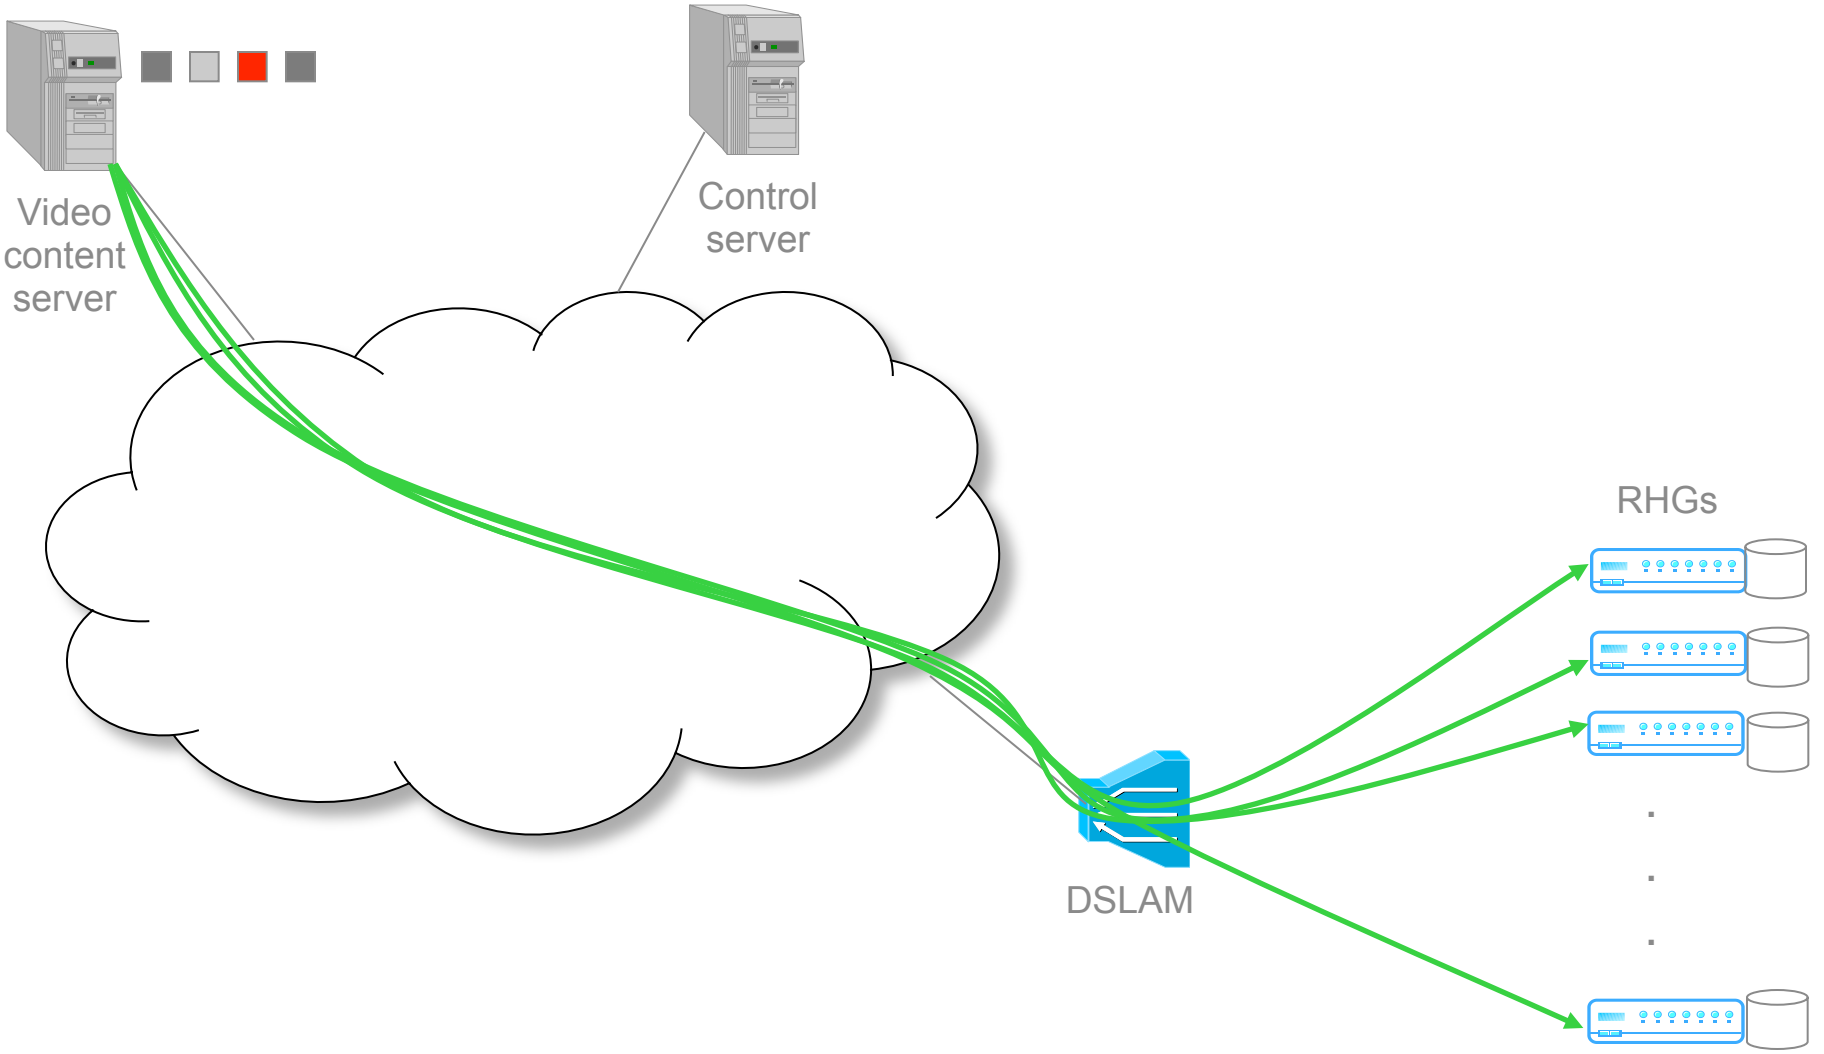

Pull phase

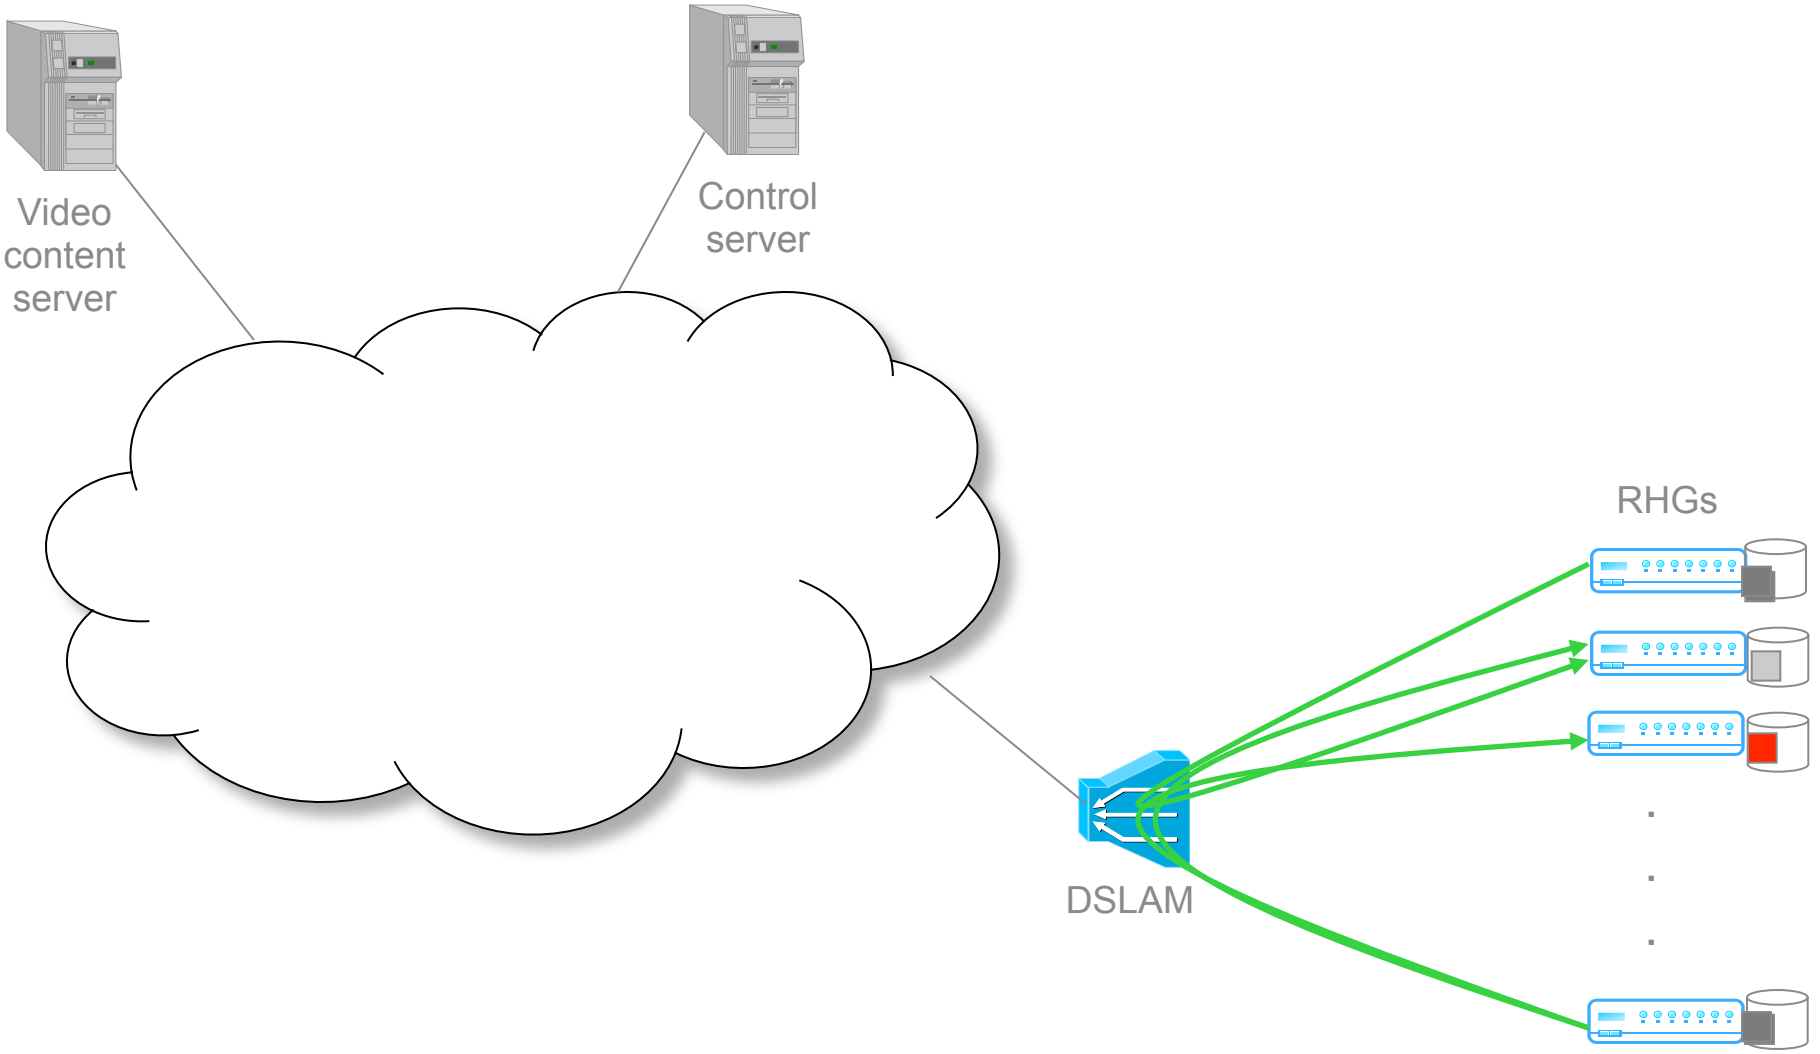

Placement strategy

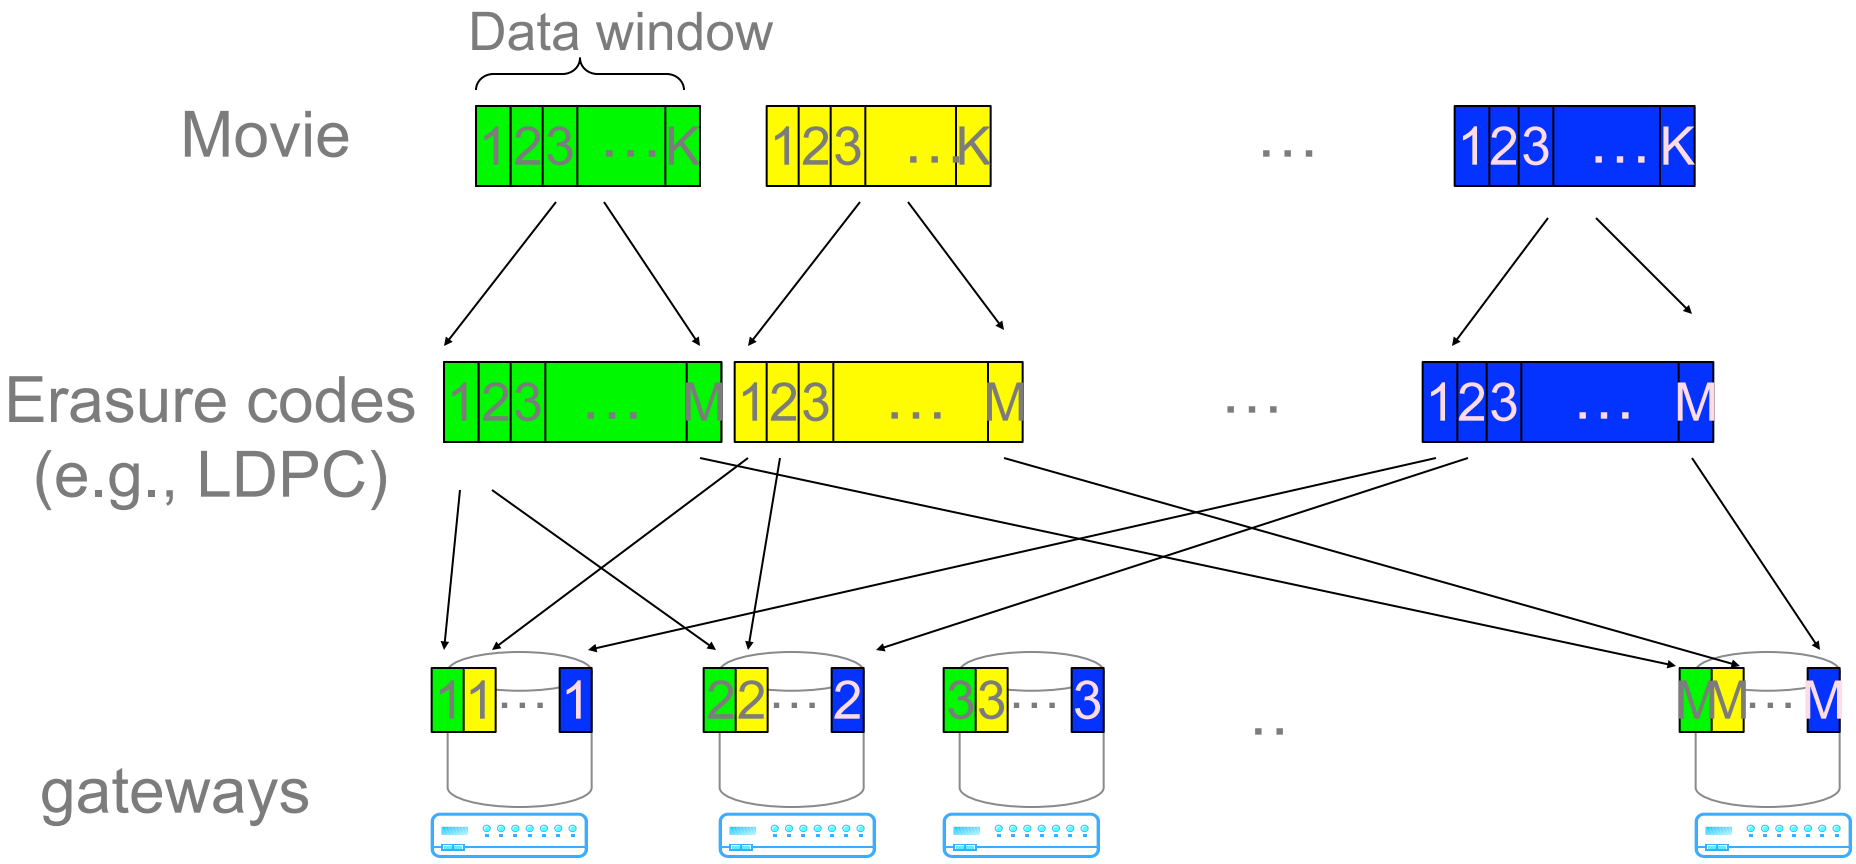

THOMSON
images & beyond

Placement strategy

- **Replicate content according to popularity**
- **Popular content served by gateways**
  - slack bandwidth from original content servers
- **Number of replicas determined by solving optimization problem**
  - Constraints on available upload and storage, number of clients, request rates, etc.

Popularity aware placement

- **Partition content into hot / warm / cold categories**
  - Hot: replicate on all gateways
  - Warm: use code-based placement
  - Cold: no proactive placement (stays on servers)

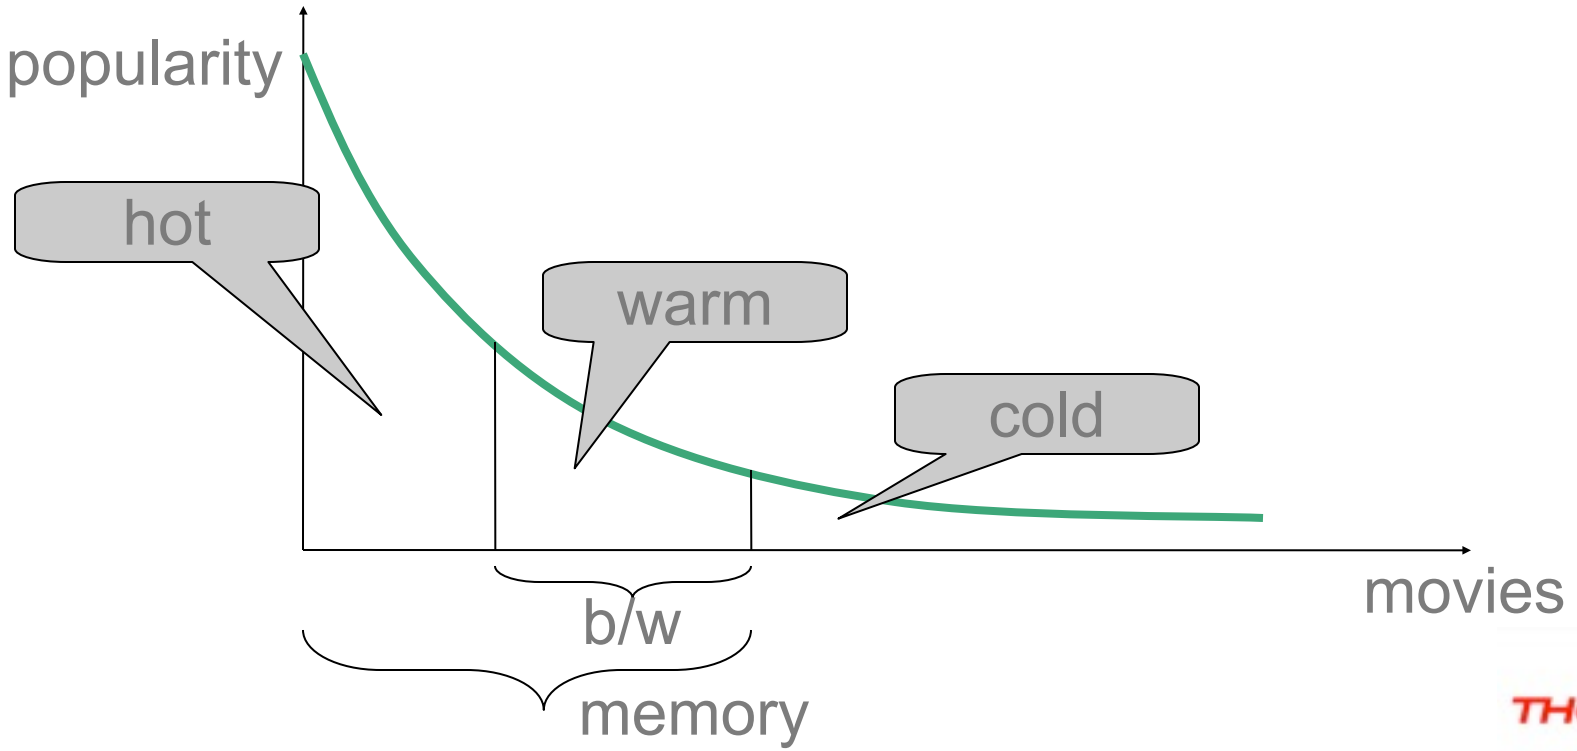

Where shou ve energy?

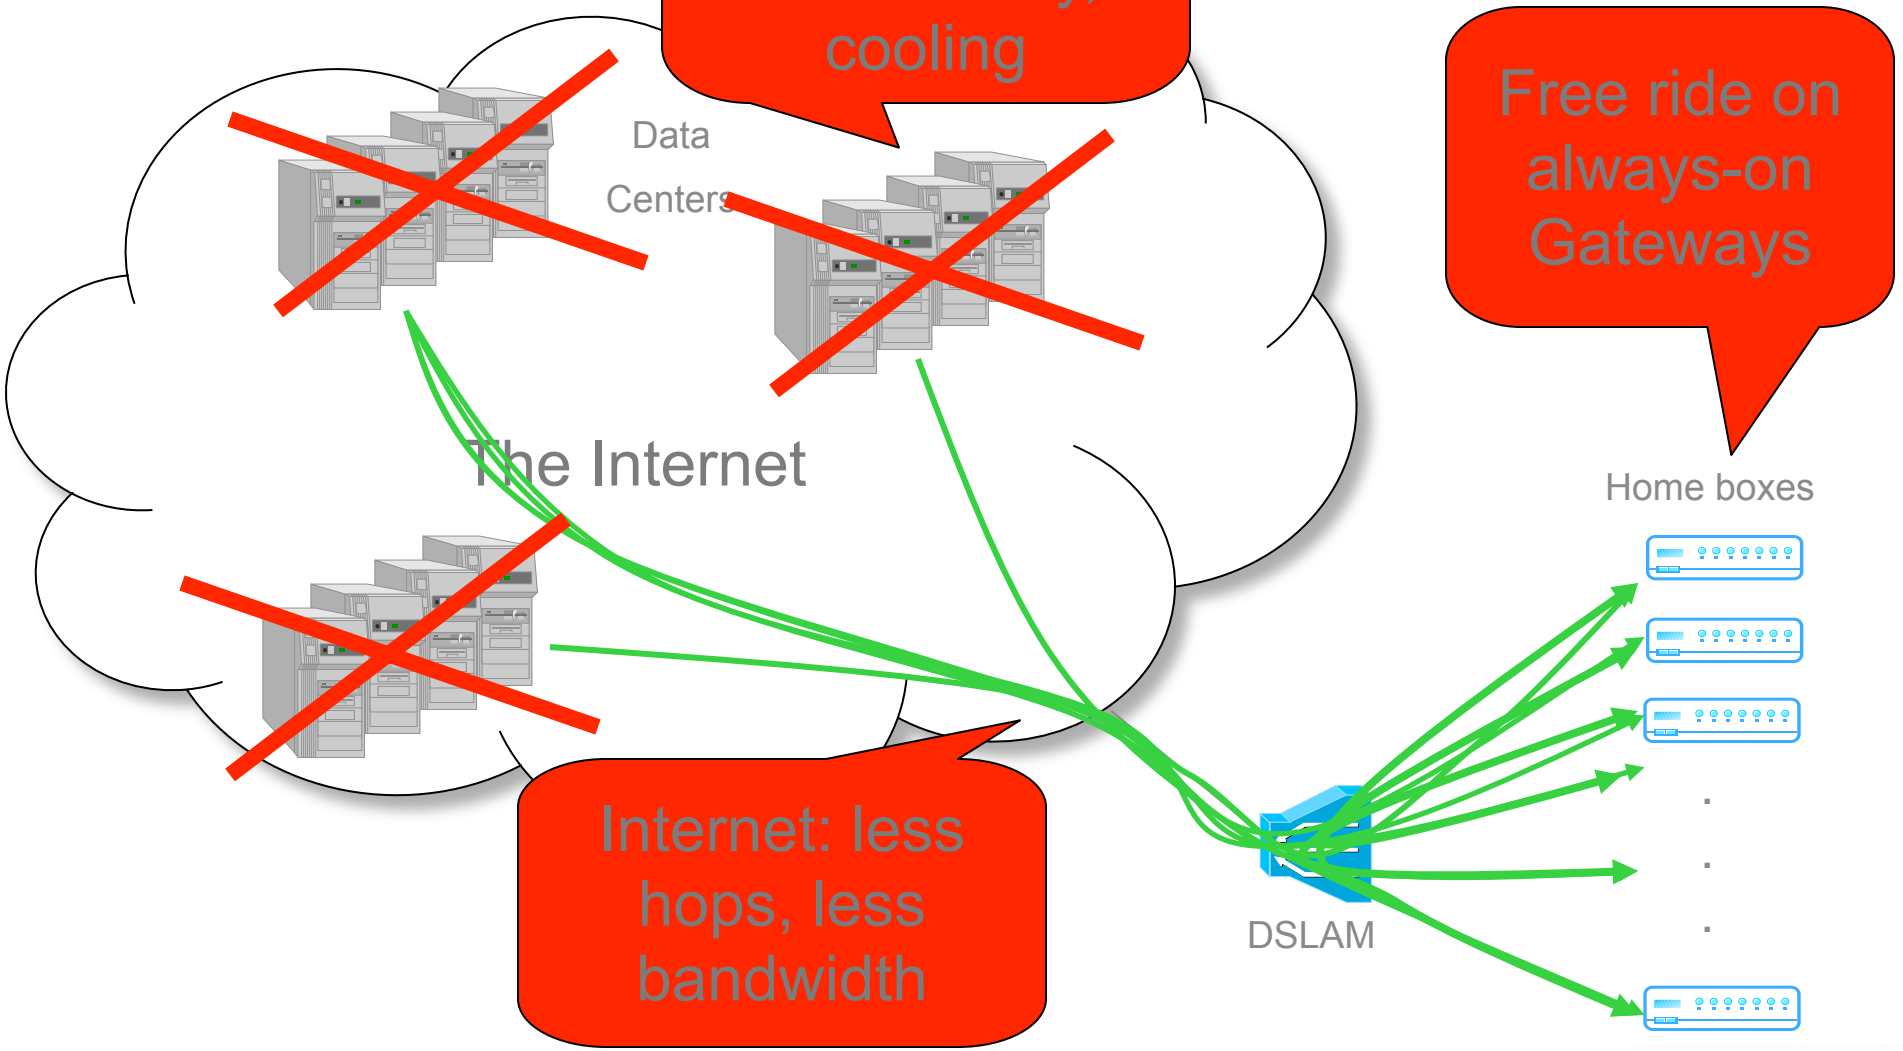

Energy issues

- **Variables**
  - Network topology
  - Hardware power consumption
  - Placement algorithms
  - Content popularity
  - User behavior
- **Data available**
  - DSL gateways and VoD servers power (Thomson)
  - Routers power (Cisco data)
  - Telefonica Spain and Peru network topologies
  - Imagenio VoD platform (Telefonica Spain)
  - Telefonica IPTV
  - Netflix movie popularity

VoD server power

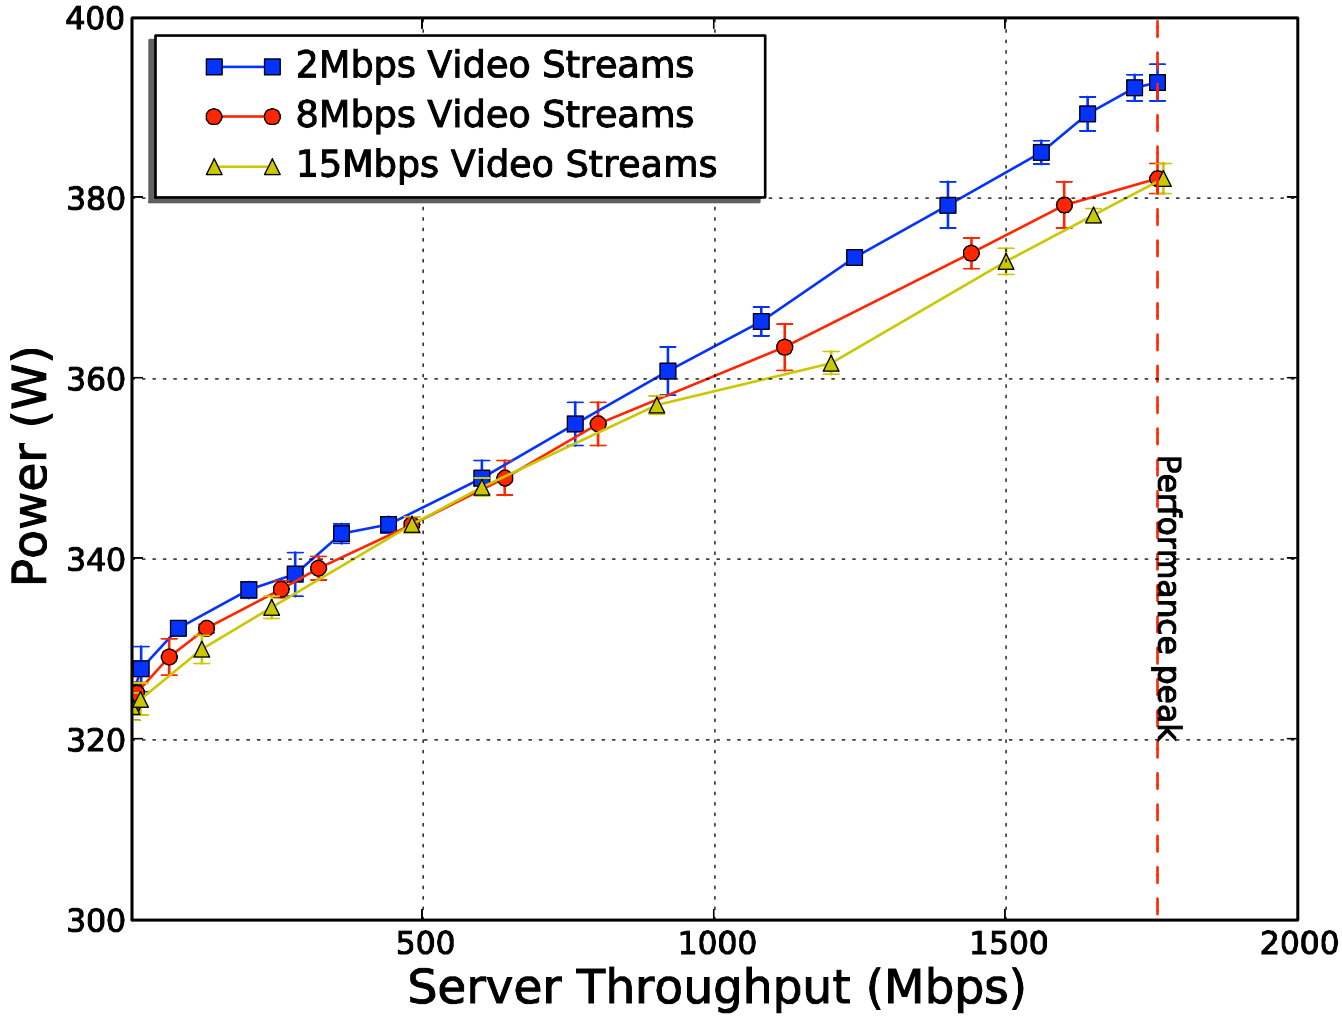

Gateway power

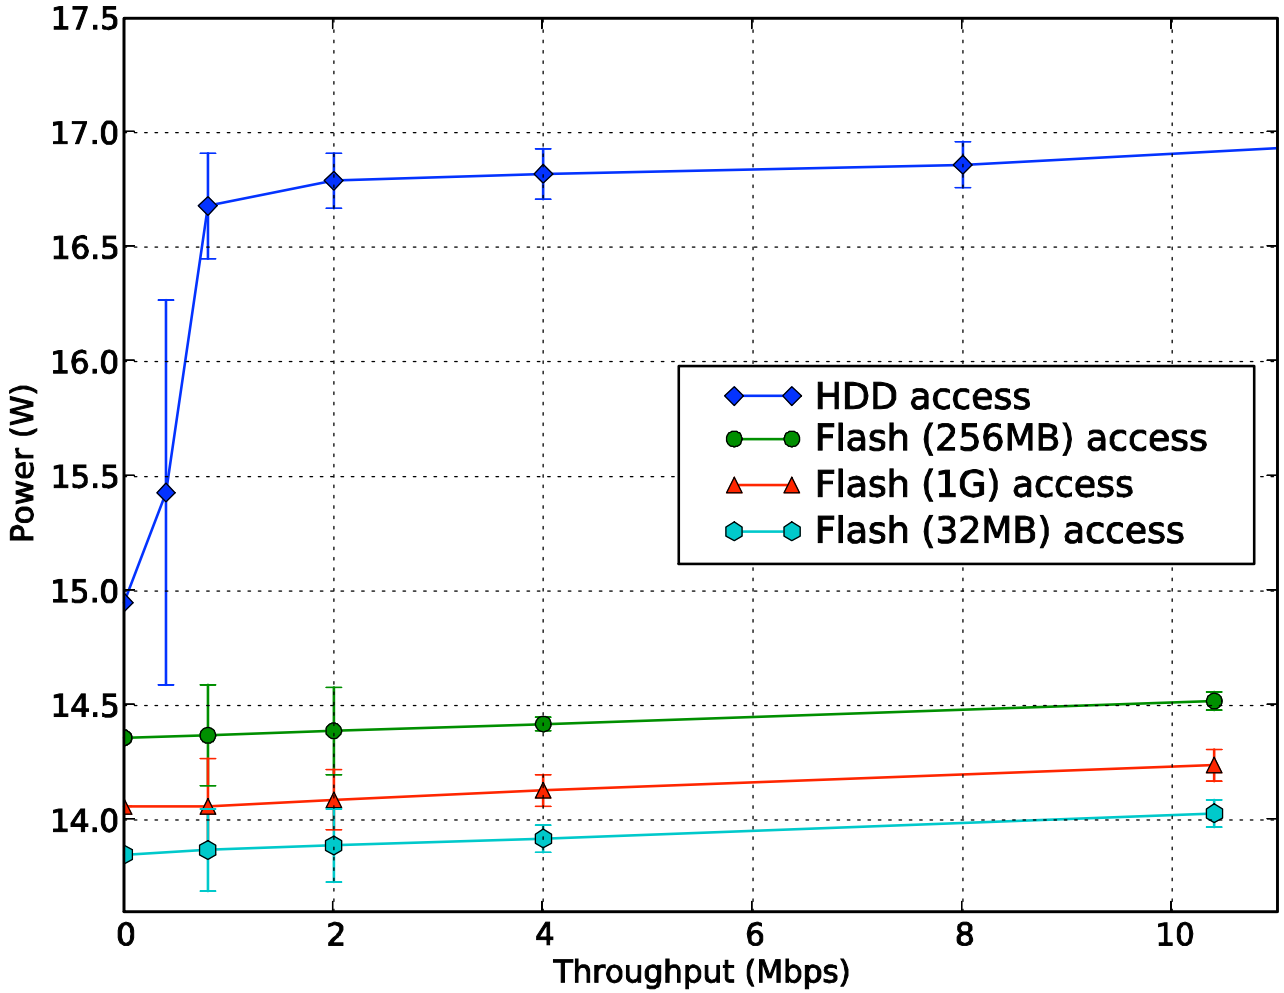

When would NaDa not work ?

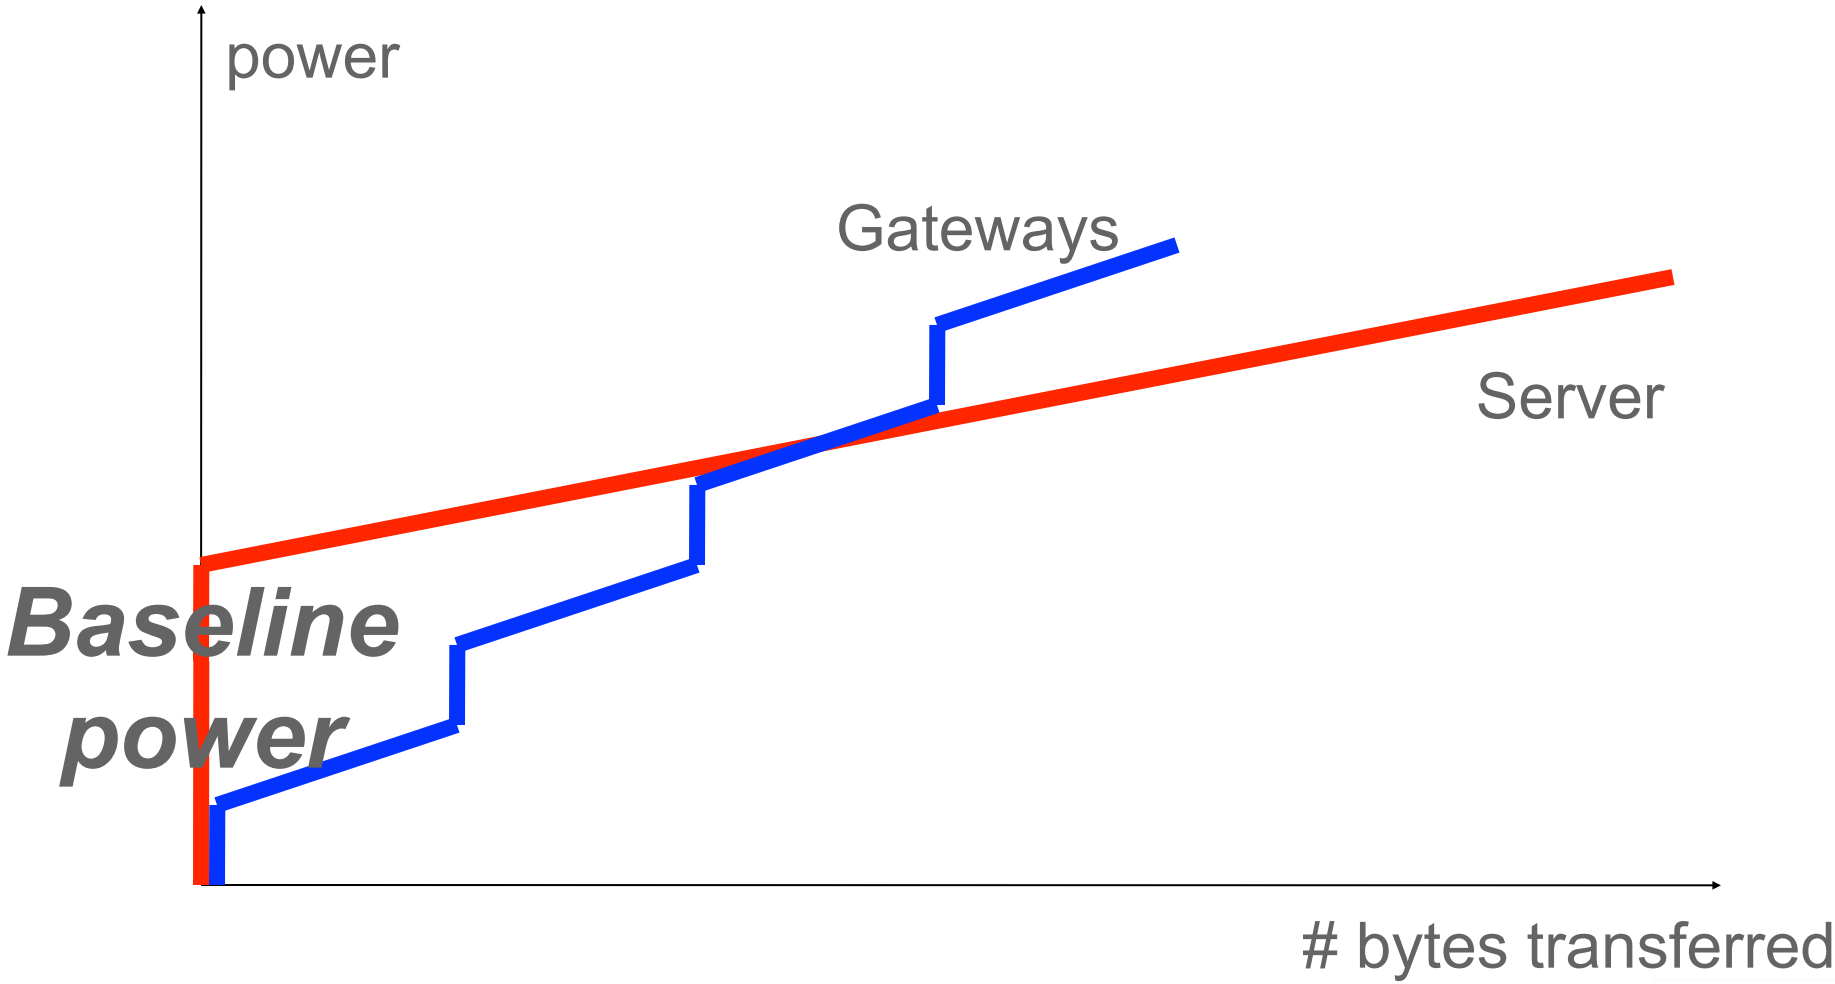

When does NaDa work ?

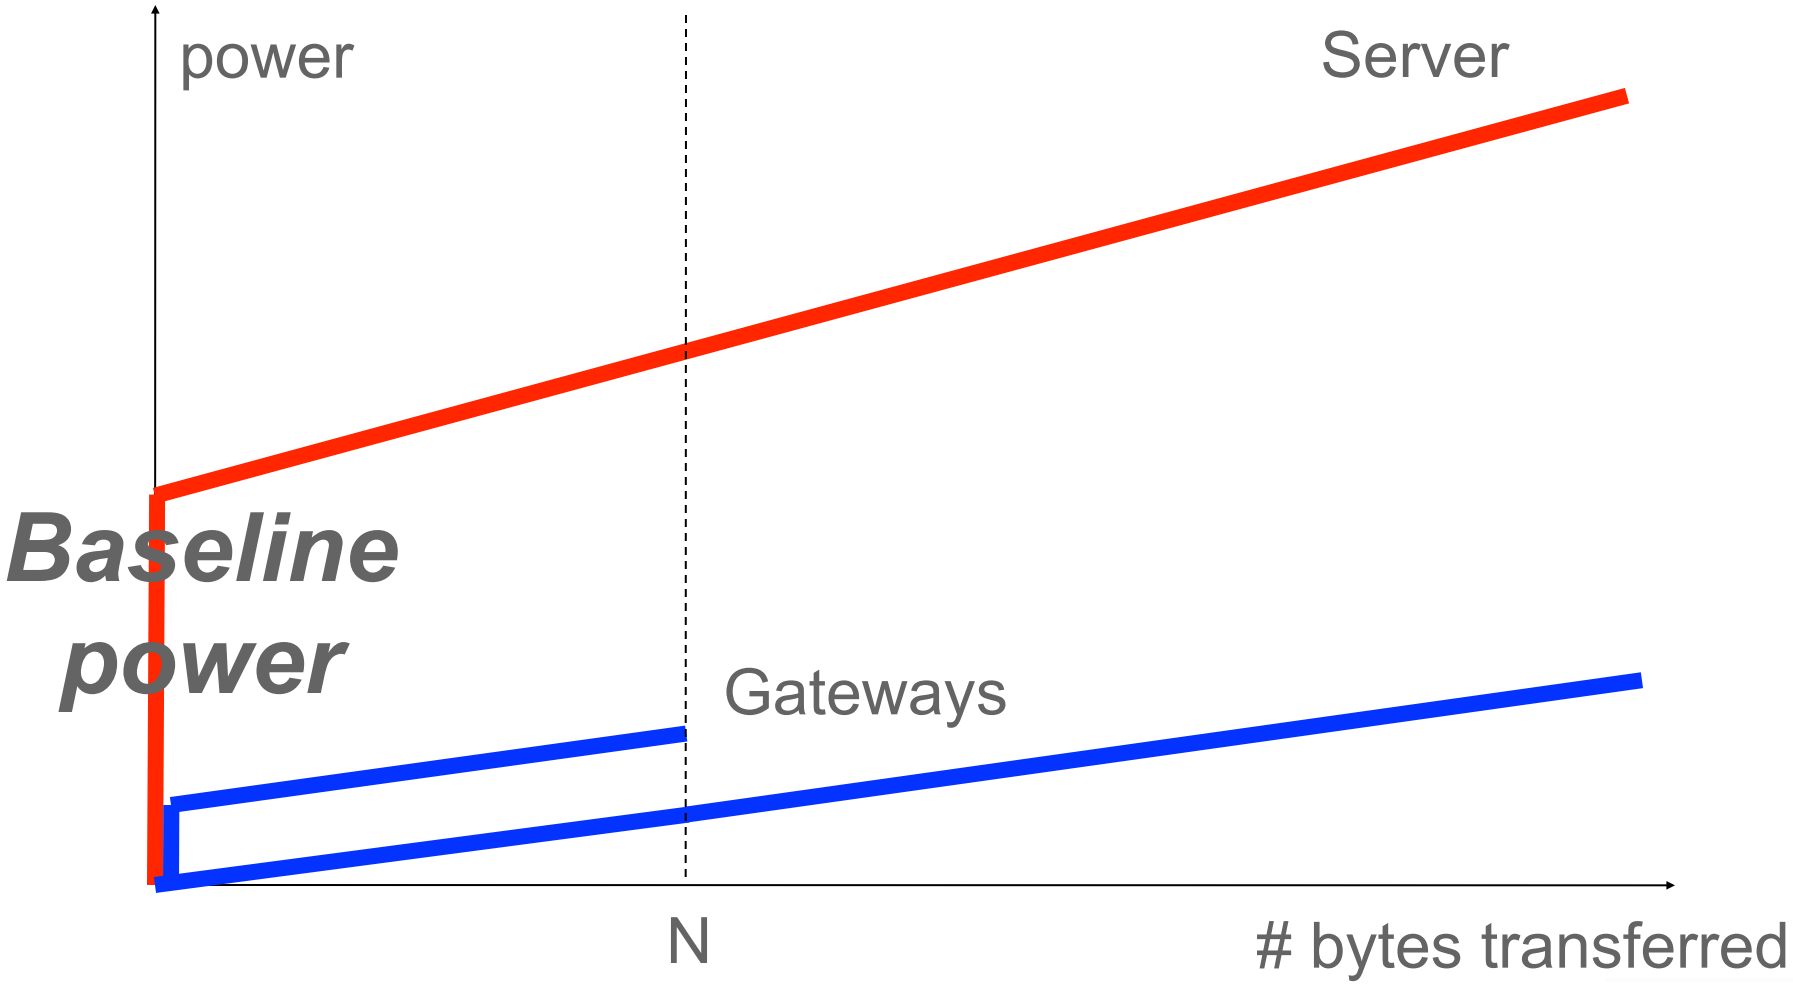

Trace driven simulations

- **Traces from**
  - Netflix, IPTV (Telefonica), YouTube
- **Content popularity from Netflix**
- **Topologies and workload from Telefonica**
- **Power numbers from Thomson's gateway and IPTV servers**
- **Popularity aware placement**

Simulation parameters

Gateway Storage	100MB-10GB
Gateway Upstream	0.1-2Mbps
Content characteristics	from data set
Users	10k-30k
Content window	10s-120s
Replicas for warm content	1 (20s windows)
Simulation duration	1 day - 86400 s
Router energy/bit	150 10^{-9}
Server energy/bit	40 10^{-9}
Gateway energy/bit	18 10^{-9}
Power Usage Effectiveness (PUE)	1.7
Home electricity cost factor	1.1
Hops to server	4
Hops to peer	2

Total energy use (YouTube)

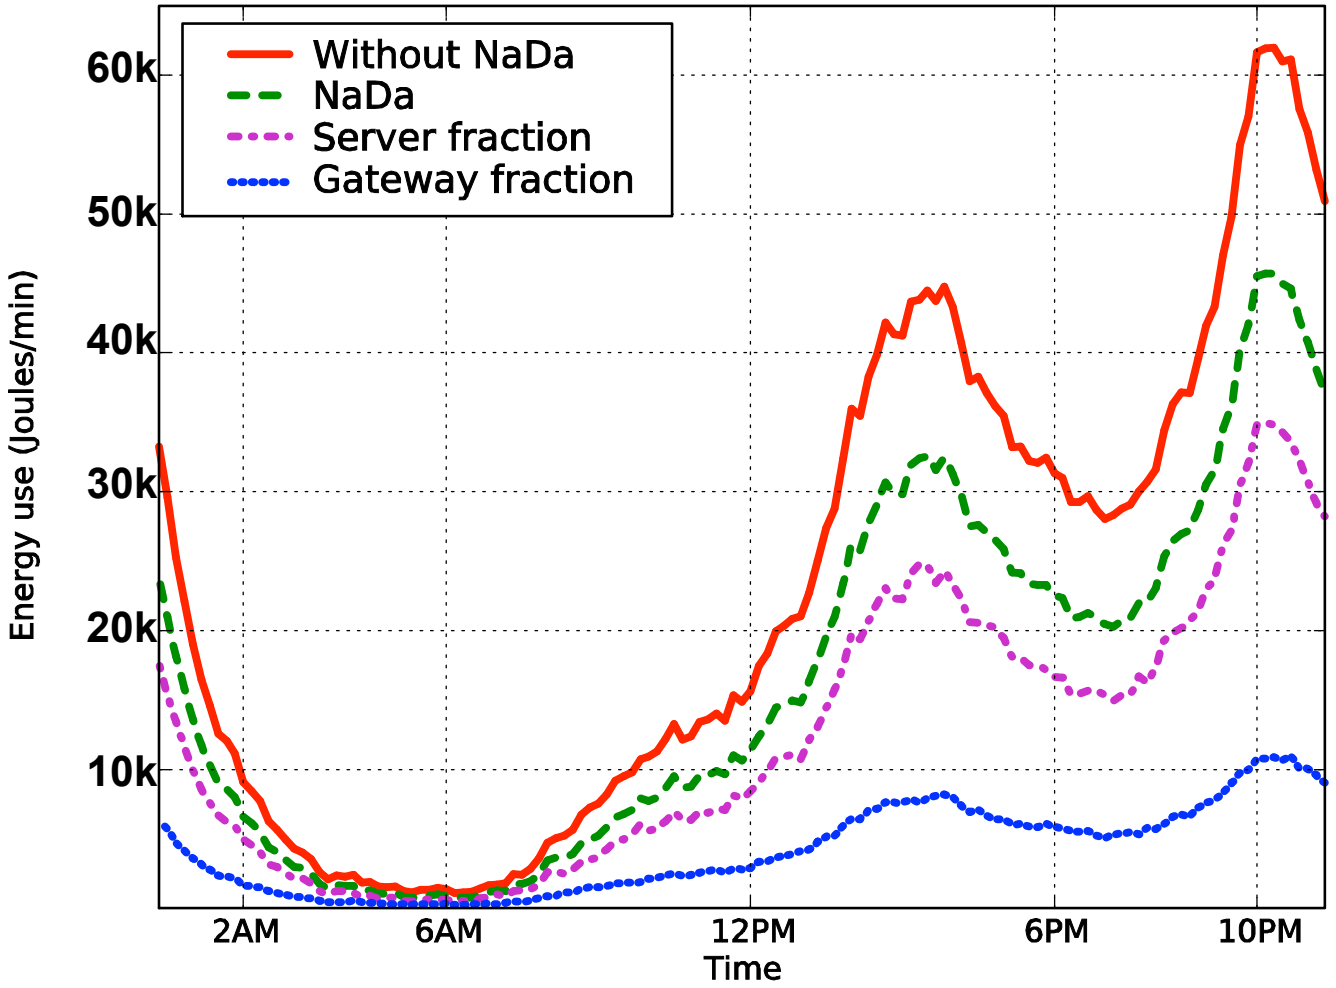

THOMSON
images & beyond

Gateway storage

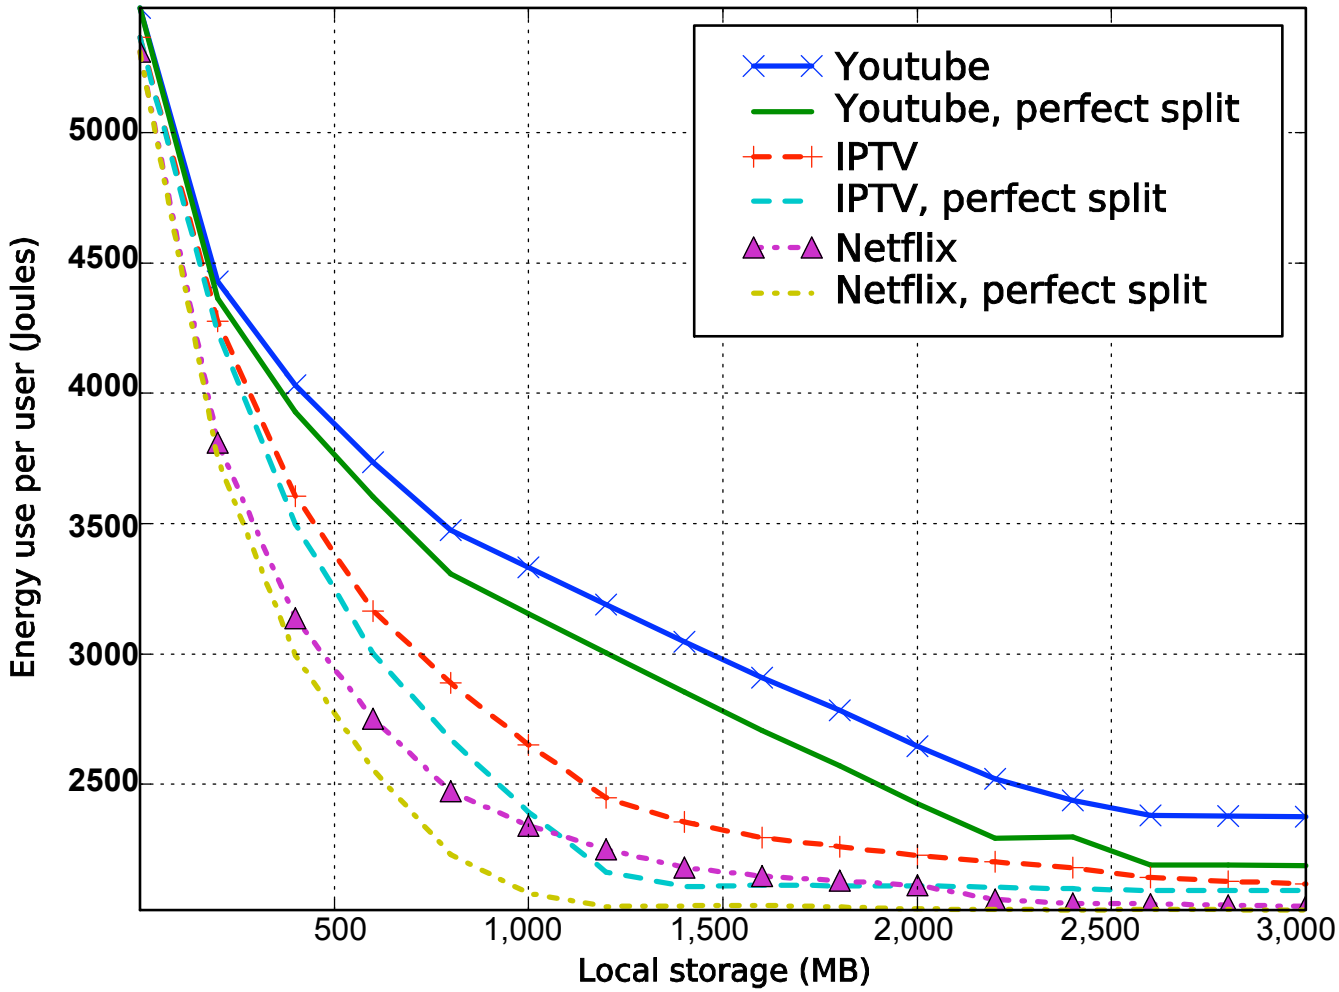

THOMSON
images & beyond

Number of users

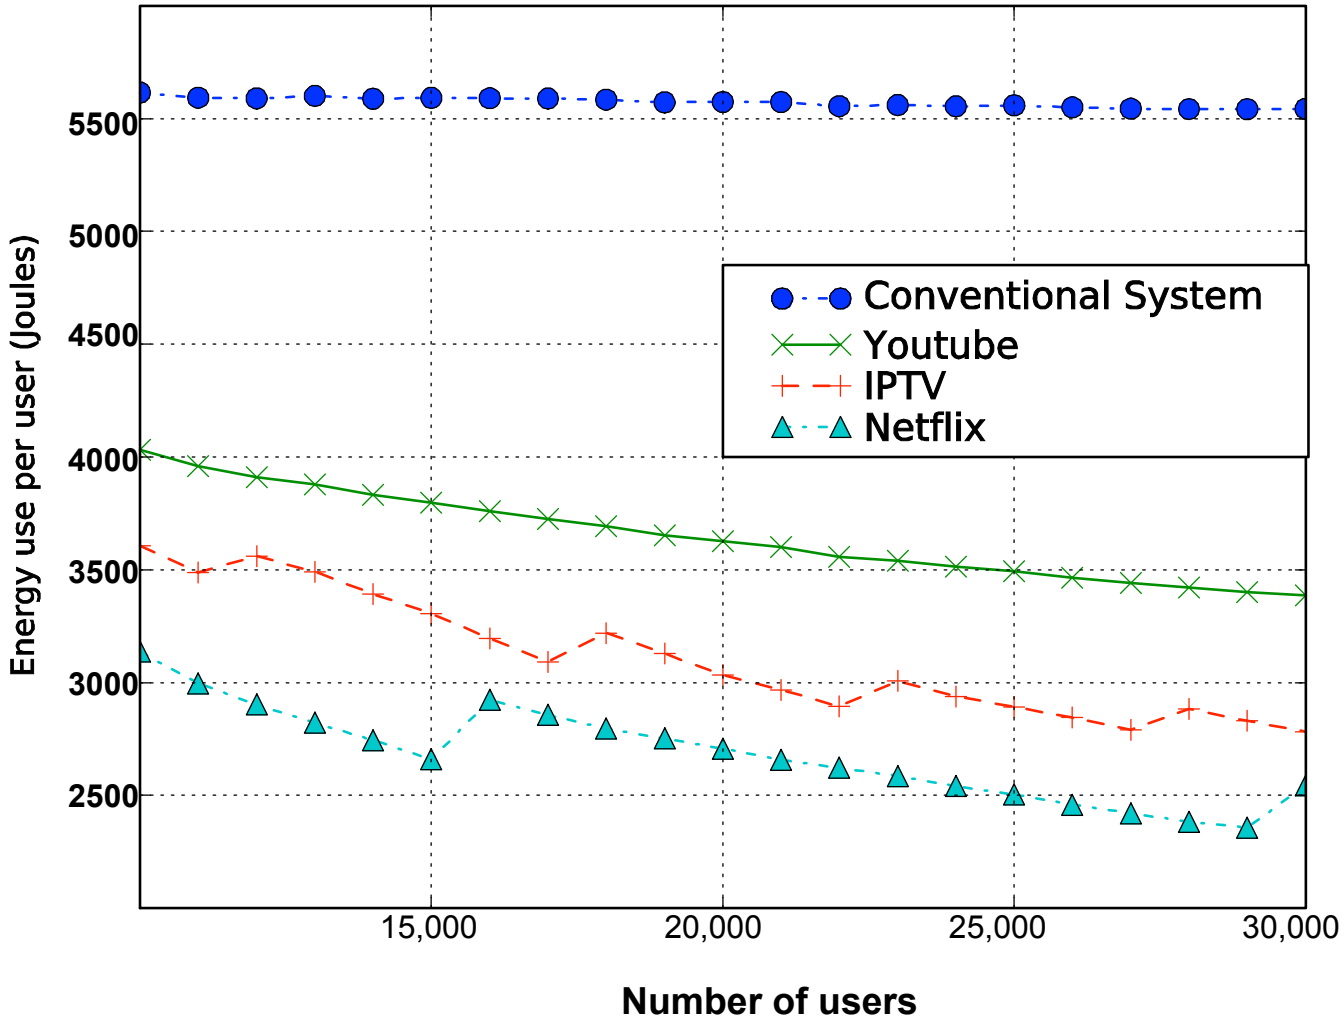

Upstream bandwidth

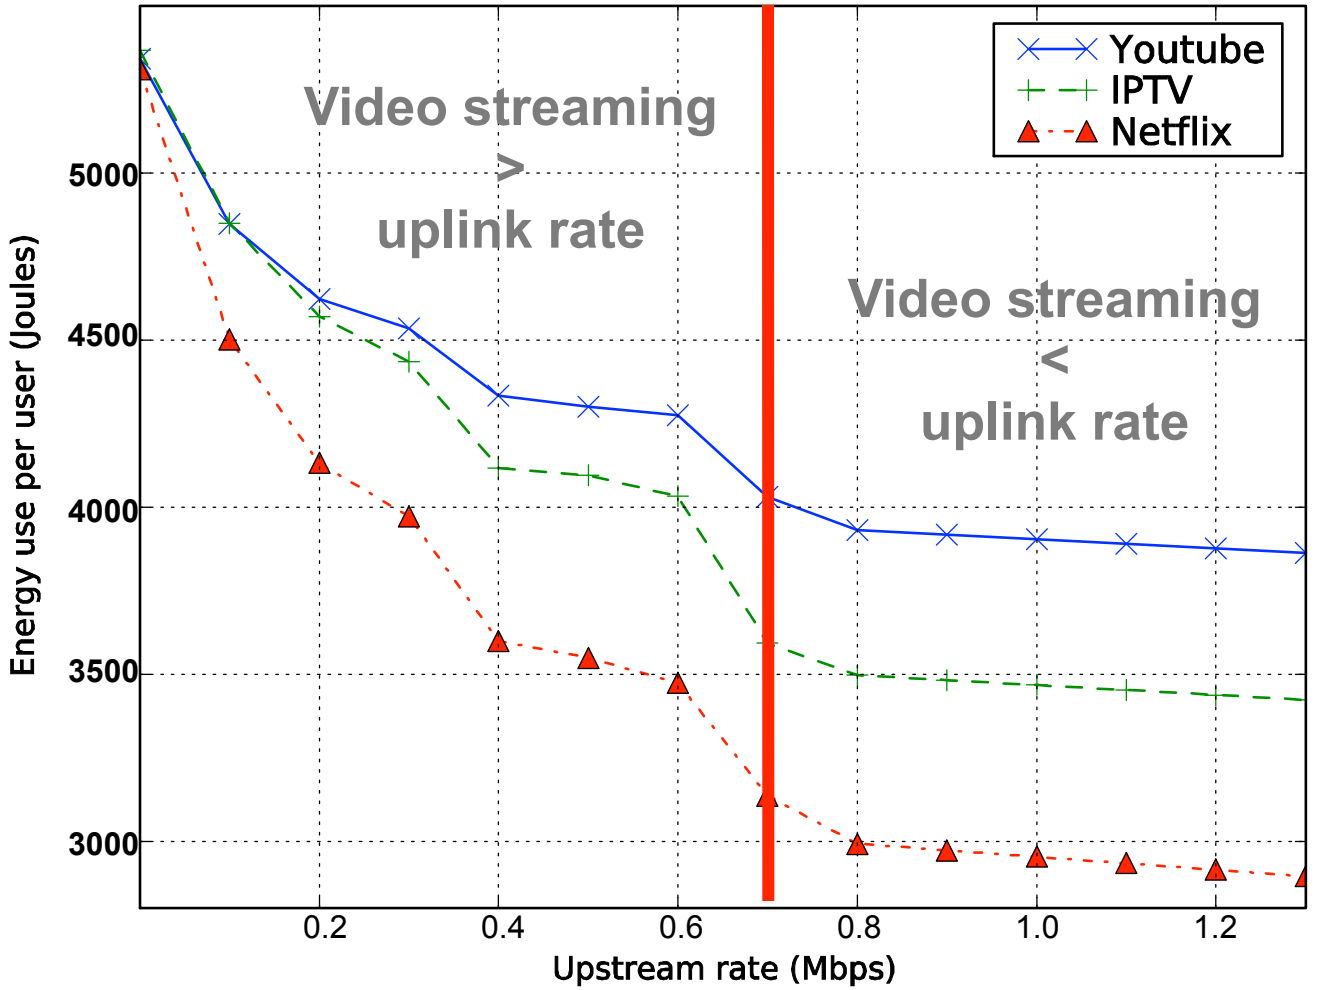

THOMSON
images & beyond

Conclusions

- Free-riding on existing infrastructure can significantly reduce load on conventional servers
- Simulations demonstrated energy savings ranging from 20% to 60% versus data centers
- Gateways can accomplish this with only modest resources (a few GB of storage, limited upload)

Questions?

- Potential QoS issues moving control from content providers (YouTube) to ISPs
- Effects of consumer line overprovisioning
- Security of content serving from home gateways