

Models for Content Distribution: Evaluating the Performance of Content-Centric Networks

Fabio Martignon

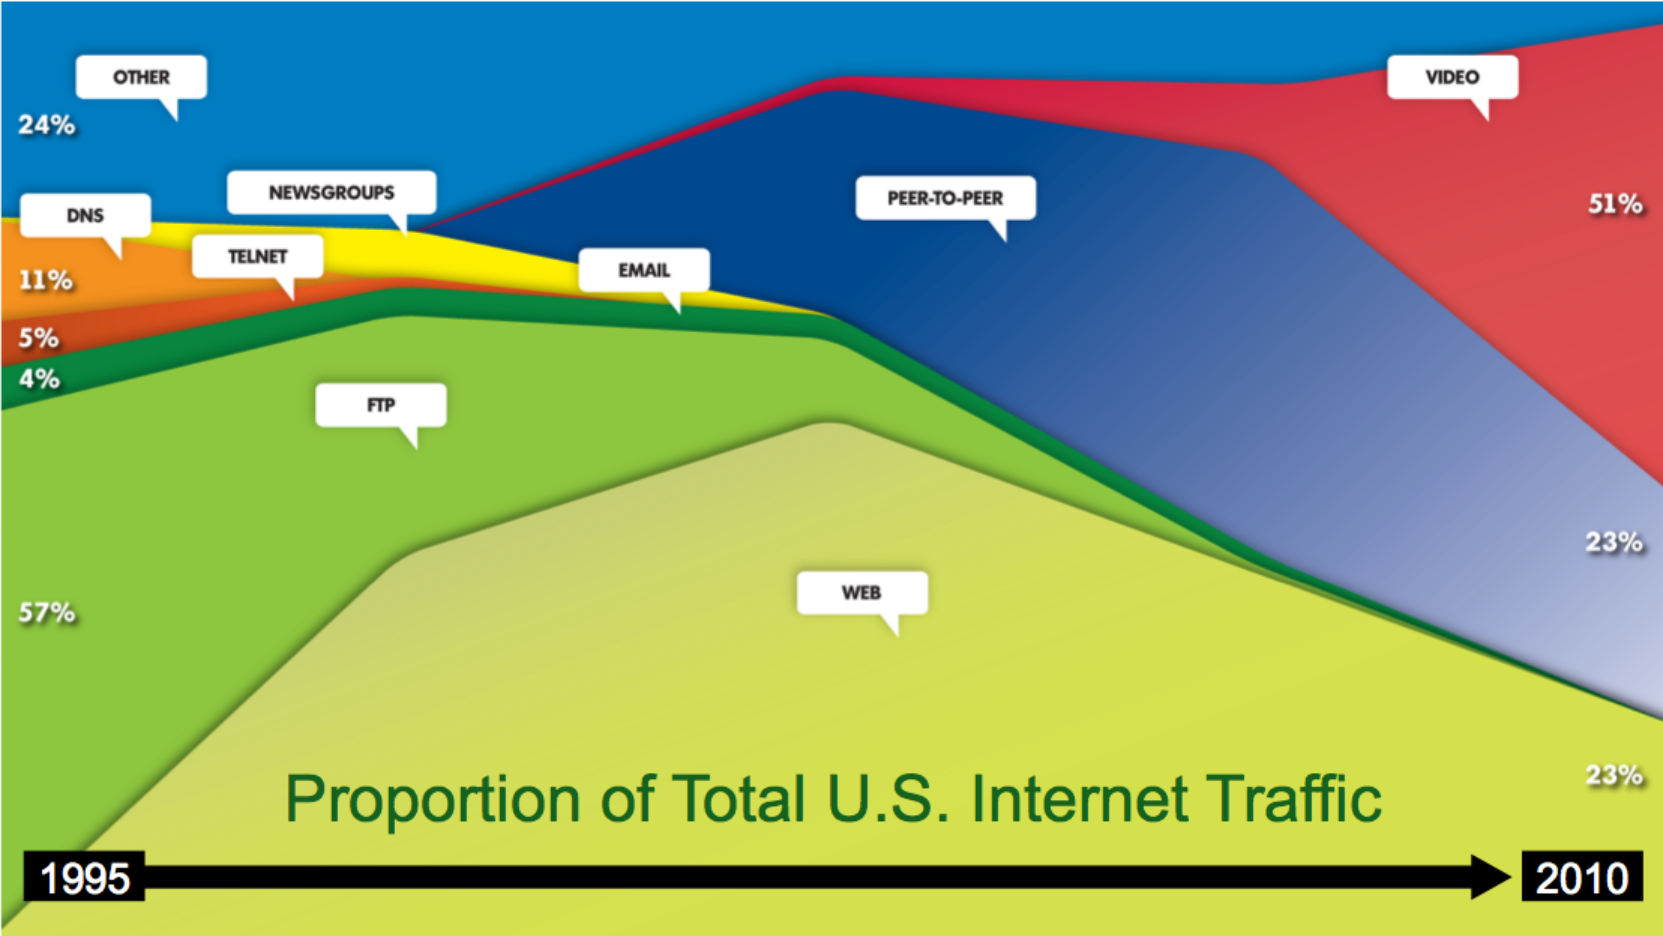
Outline of the Presentation

1. The Past / The Present / The Future
 - Traffic Trends for Internet
 - Network Models & Protocols
2. Content Distribution in Internet
 - CDN and ICN
 - Named-Data Networking (NDN/CCN)
3. Evaluating the Content-Distribution performance
 - Performance models for CCN / CDN
 - Addressing the content popularity evolution
4. Conclusion

The Past / The Present / The Future

Traffic Trends for Internet

Traffic Trends: The Past




Limelight Networks – Investor Overview – Nov. 2012

Traffic Trends: The Present

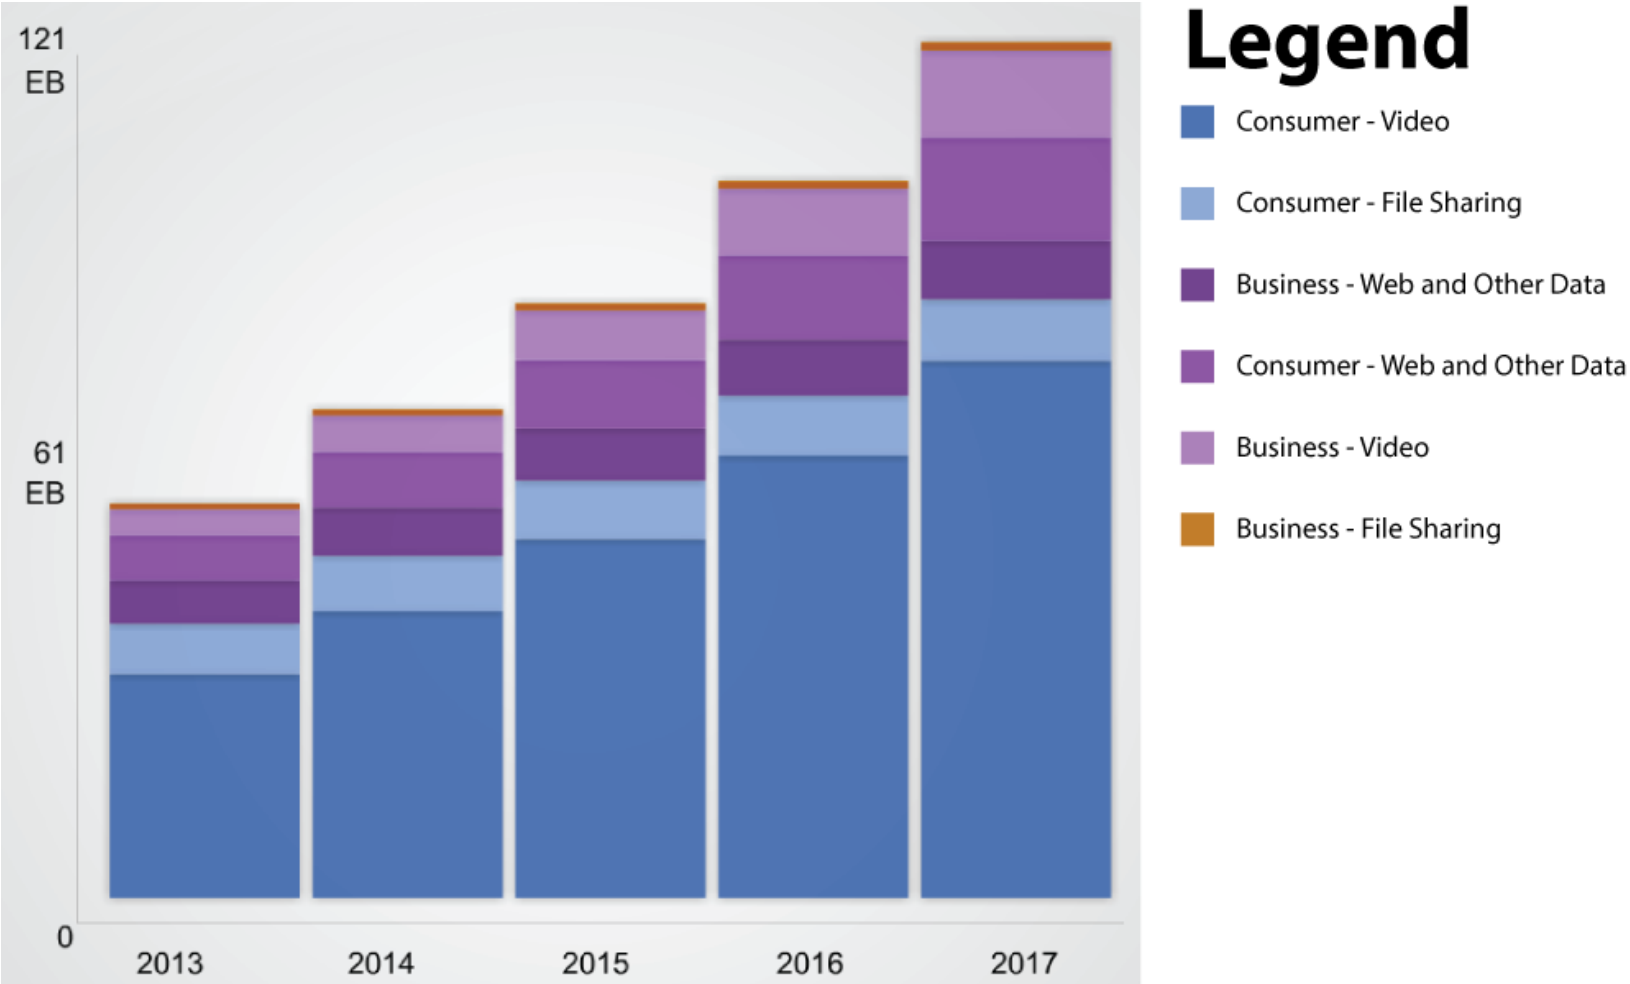
Netflix and YouTube Are America's Biggest Traffic Hogs

Share of peak period downstream traffic in North America, by application*



* September 2013. Fixed access only.

Traffic Trends: The Future



Cisco VNI report, May 2013

The Past / The Present / The Future


Network Models & Protocols

Network Models: The Past

- Telephony model
 - Host-to-Host communication
 - Circuit switching
 - The focus is on ***the path*** connecting the two endpoints (the wire)

- Internet Protocol
 - Host-to-Host communication
 - Packet switching
 - The focus is **on the two end-points**

- IP was originally designed in the late '60s to exploit expensive computational resources available remotely



Network Models: The Present

- The nature of today's TCP/IP Internet is the one of a **communication network**

- However, users mostly exploit Internet as a **distribution infrastructure:**
 - They want to retrieve the content
 - Does not really matter which server provides the data

- To fill the gap of this misalignment, content distribution is performed using **overlay** approaches:
 - CDN
 - P2P



Network Models: The Future

- Proposals for the Future Internet:
 - Information-Centric Networks (ICNs)
 - Change the Network layer protocols to turn the Internet into a *content distribution architecture*
 - They jointly enforce new requirements:
 1. *Scalable content distribution*
 2. *Mobility support*
 3. *Network security*

Content Distribution in Internet

Content Delivery Networks (CDNs)

Information Centric Networks (ICNs)

Content Delivery Networks

- Today's TCP/IP Internet relies on Content Delivery Networks to accommodate current users' needs

- A Content Delivery Network is:
 - A distributed system;
 - Composed of many "*surrogate servers*" (or also "*replica servers*")
 - Used to offer a *transparent content distribution* service
 - Run by *a single owner* (in a **centralized** manner)

- Being centralized means that the CDN owner can optimize:
 - ***Surrogate server placement***
 - ***Replica object placement***
 - ***Request routing***

Content Delivery Networks

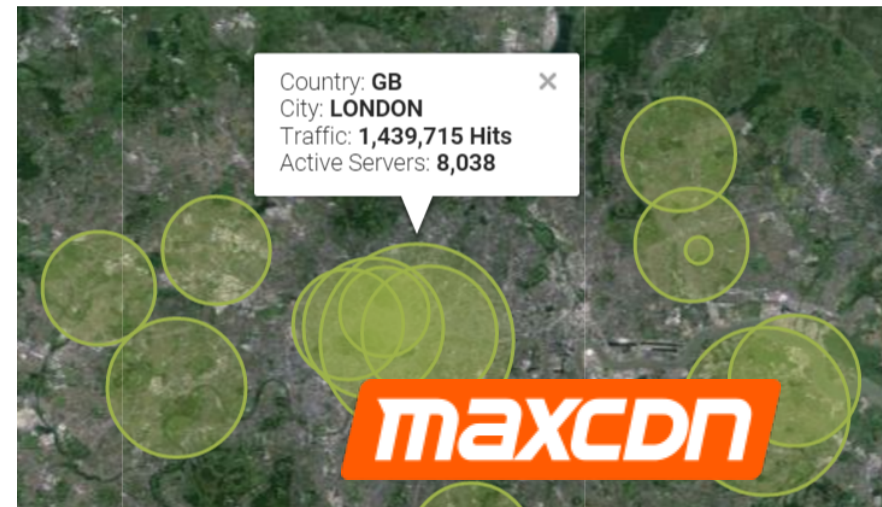
Content Delivery Networks (CDNs) make it possible to accommodate users' traffic requests in today's TCP/IP Internet

Content Delivery Networks



Content Delivery Networks

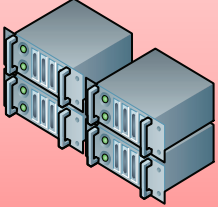
- ❑ Akamai serves 15-30% world's web traffic
- ❑ On average handles
 - 25 million req **per SEC**
 - ~ 2 trillion req per day
- ❑ Comprises more than 200'000 servers
 - Connected to 1400+ networks



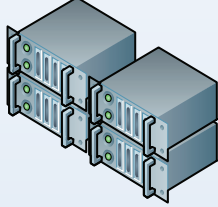

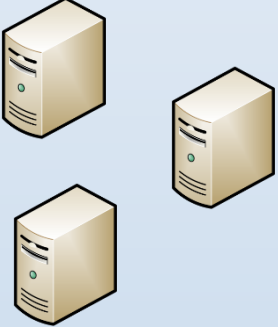
CDN: Architecture

ISP

Local DNS (LDNS)

A rack of four blue server units representing the ISP's Local DNS (LDNS).


Authoritative DNS
for akamaihd.net

A rack of four blue server units representing the authoritative DNS for akamaihd.net.
The Akamai logo, featuring a blue wave icon and the word "Akamai" in orange.
Three tan server units representing replica servers.


Replica servers

**Content
Provider**

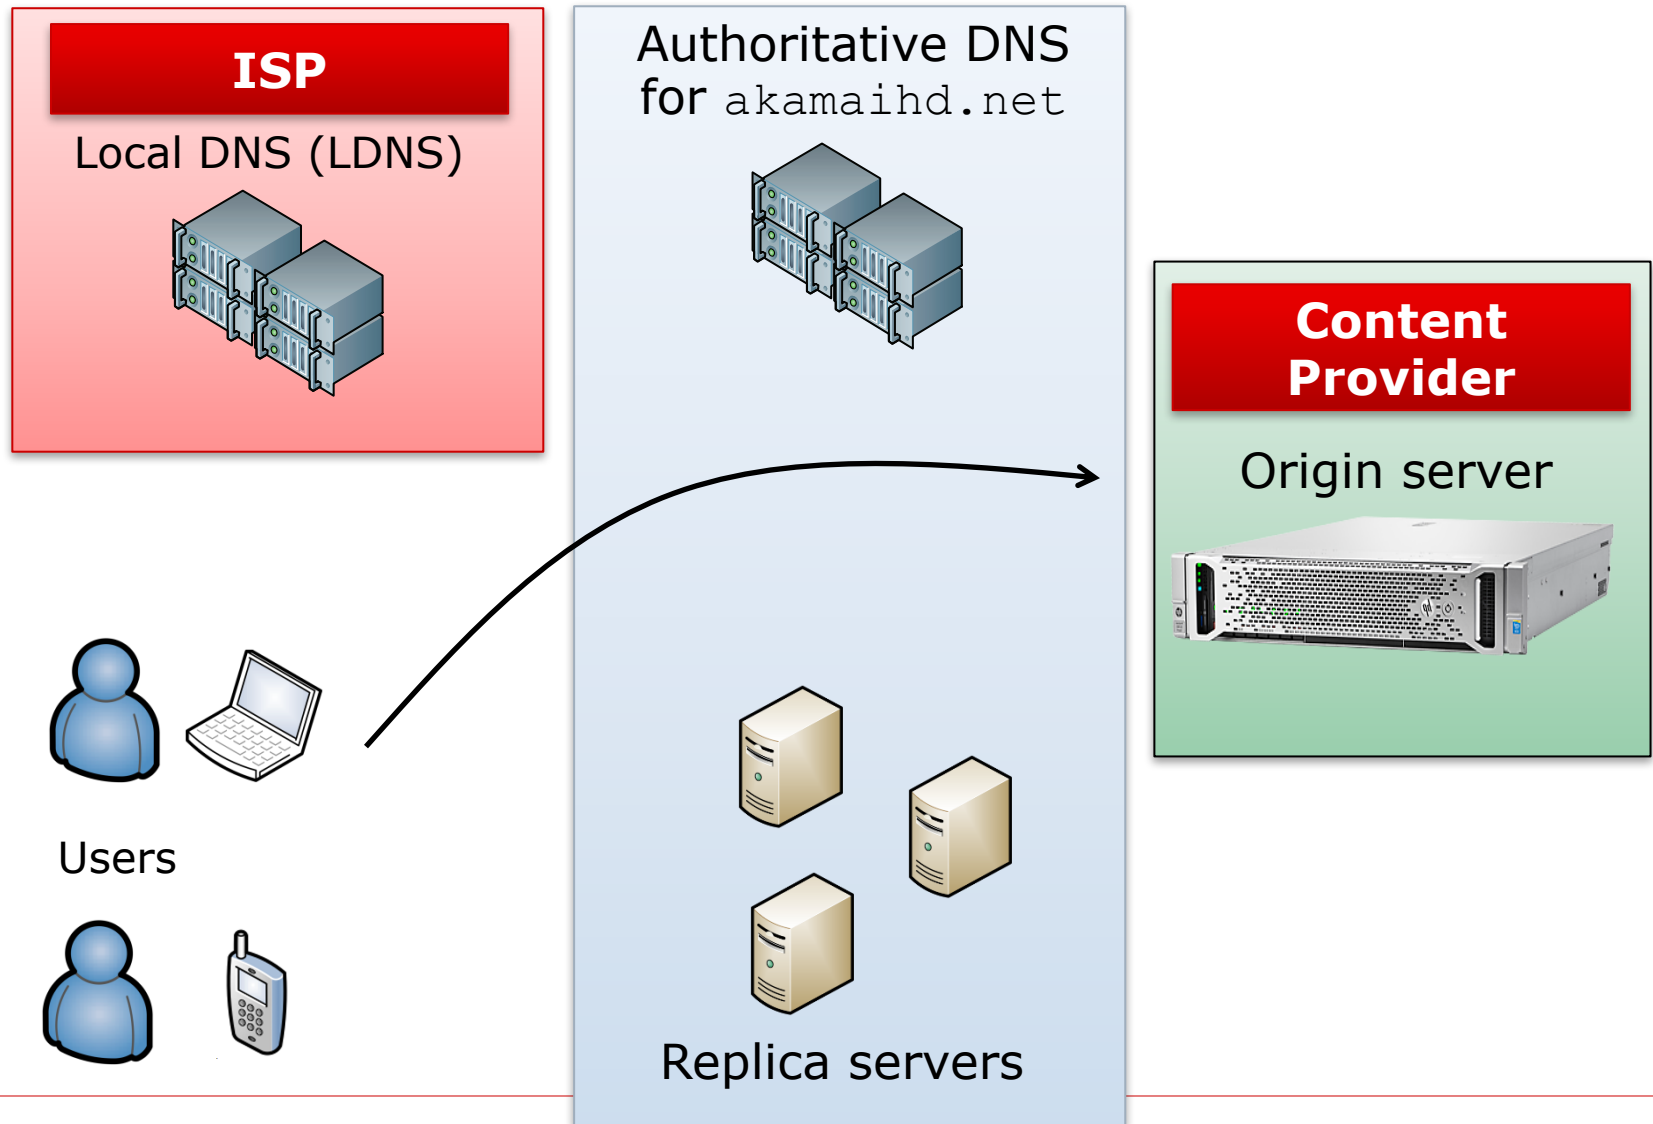
Origin server

A large white server rack representing the origin server.

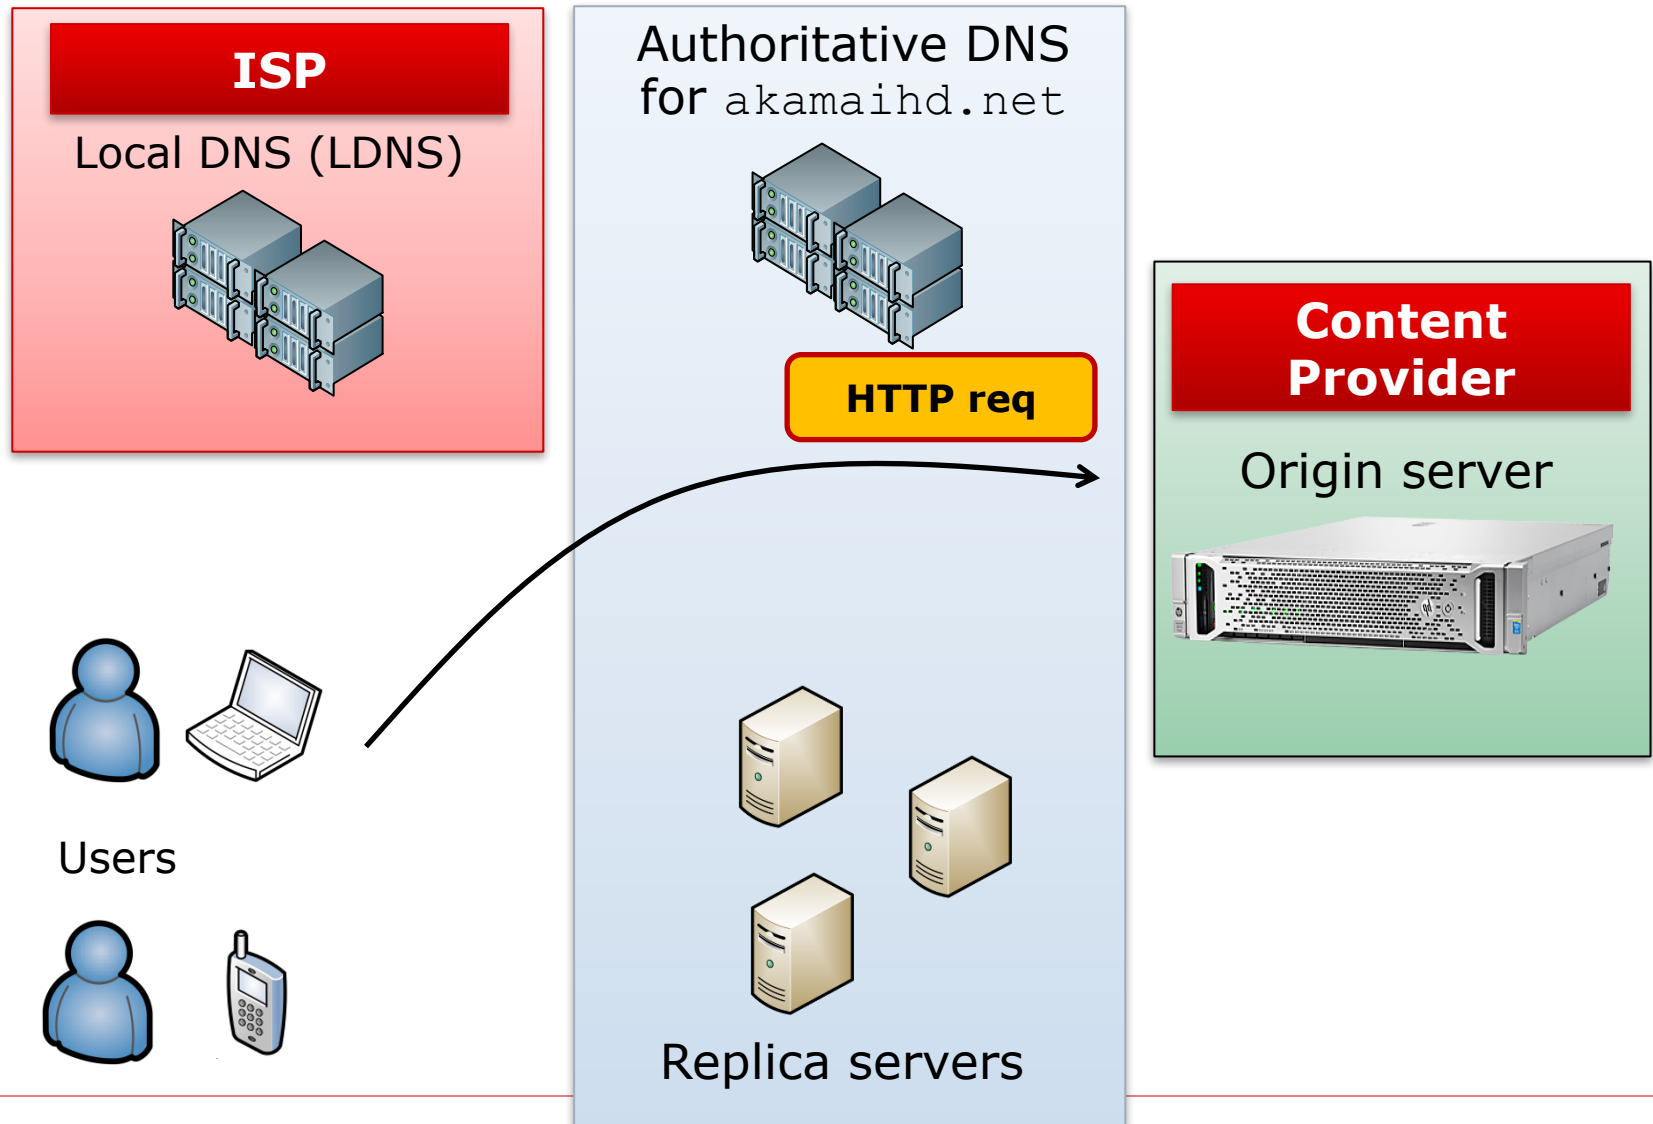
Users



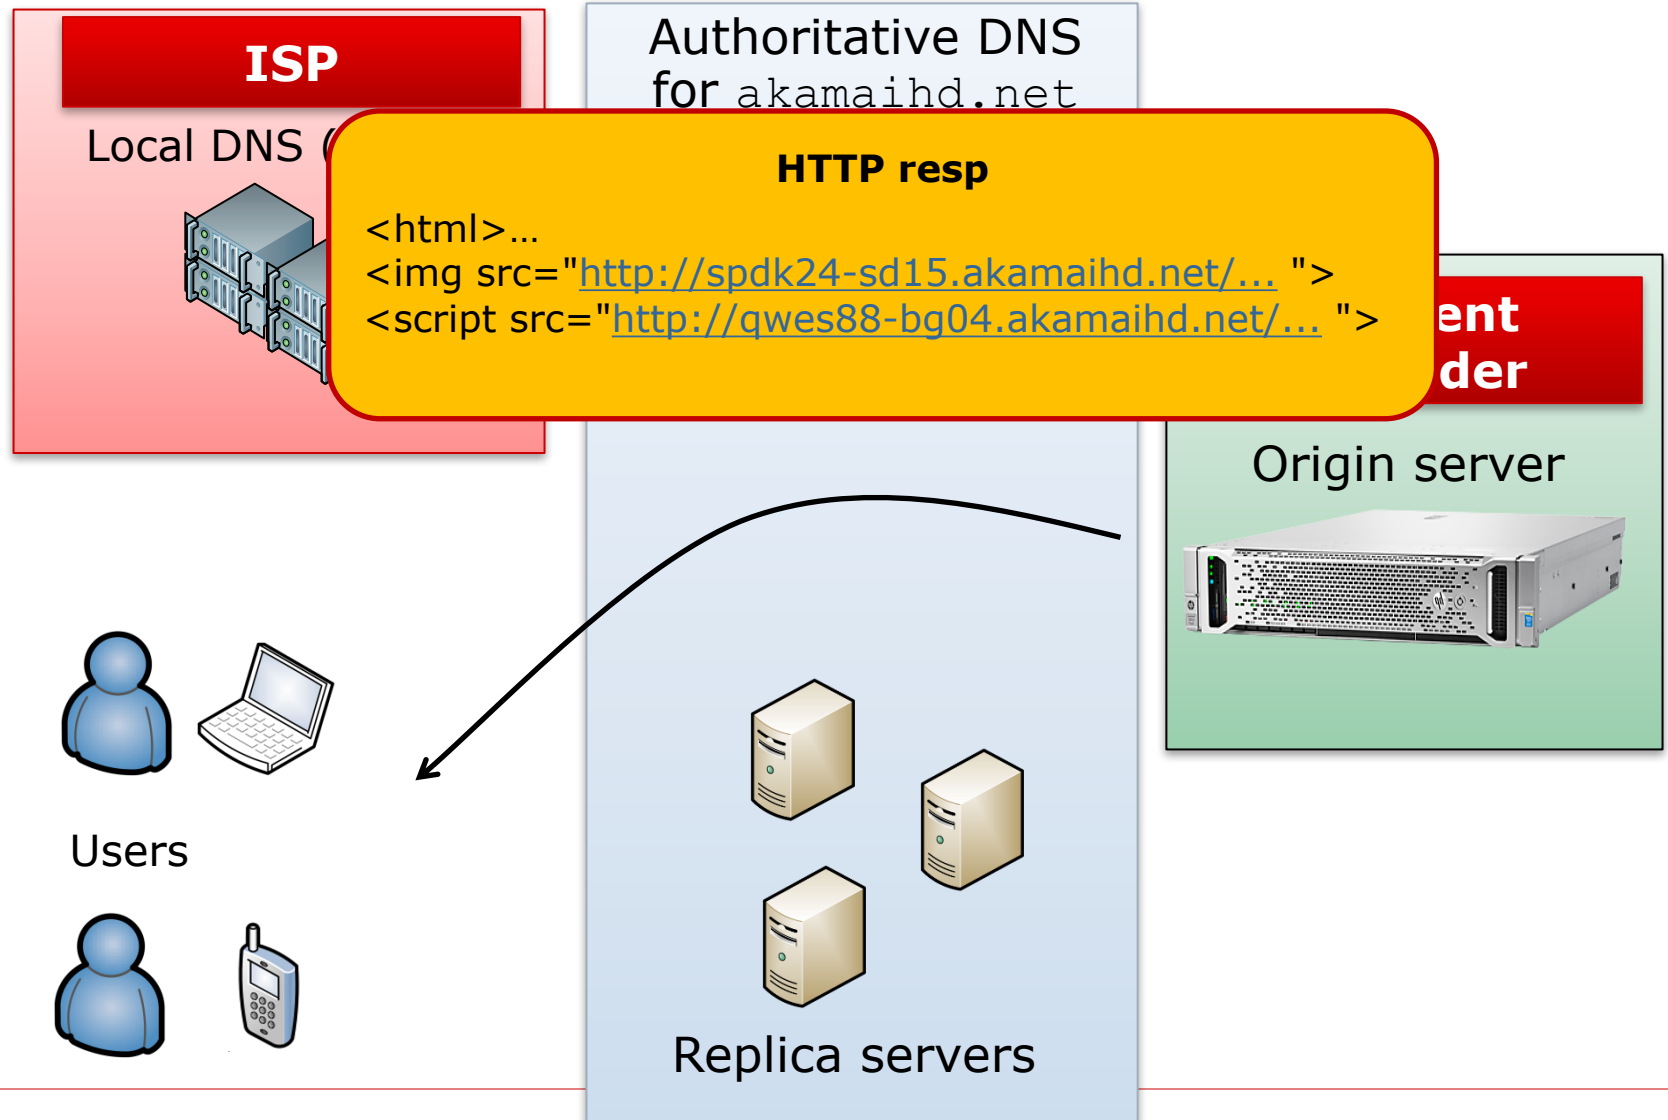
CDN: Architecture



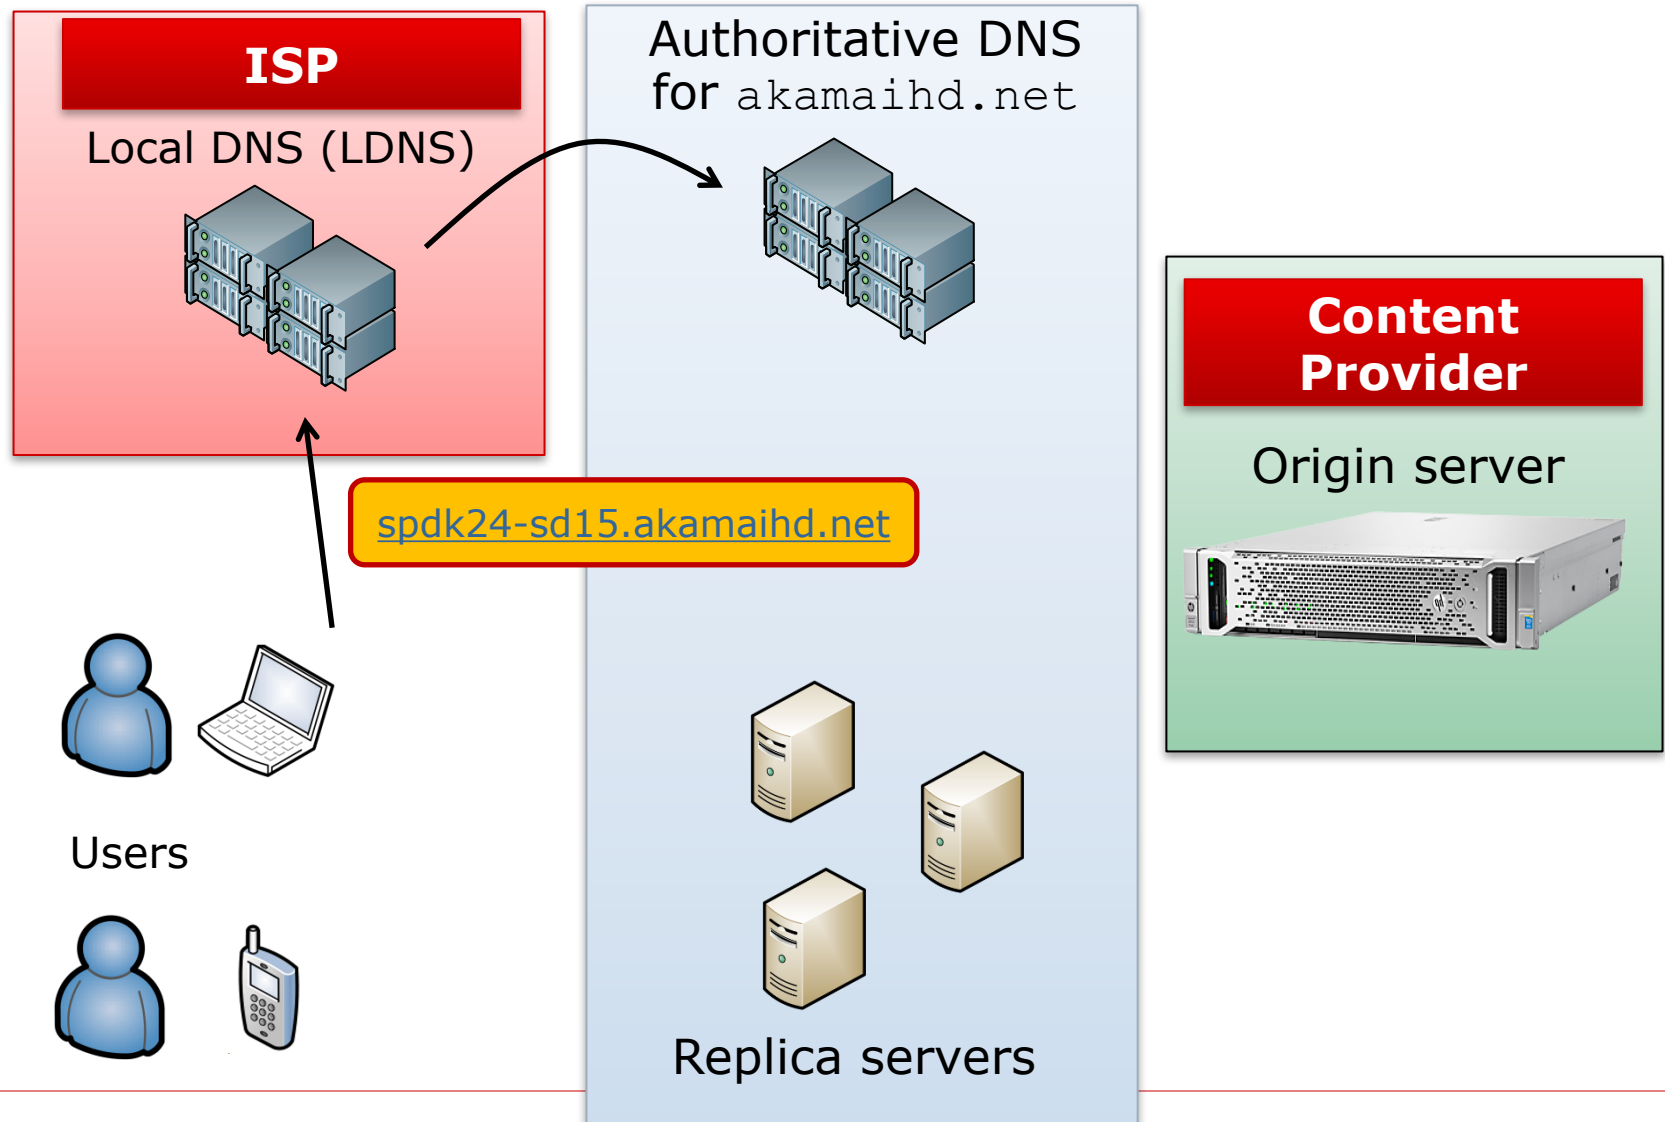
CDN: Architecture



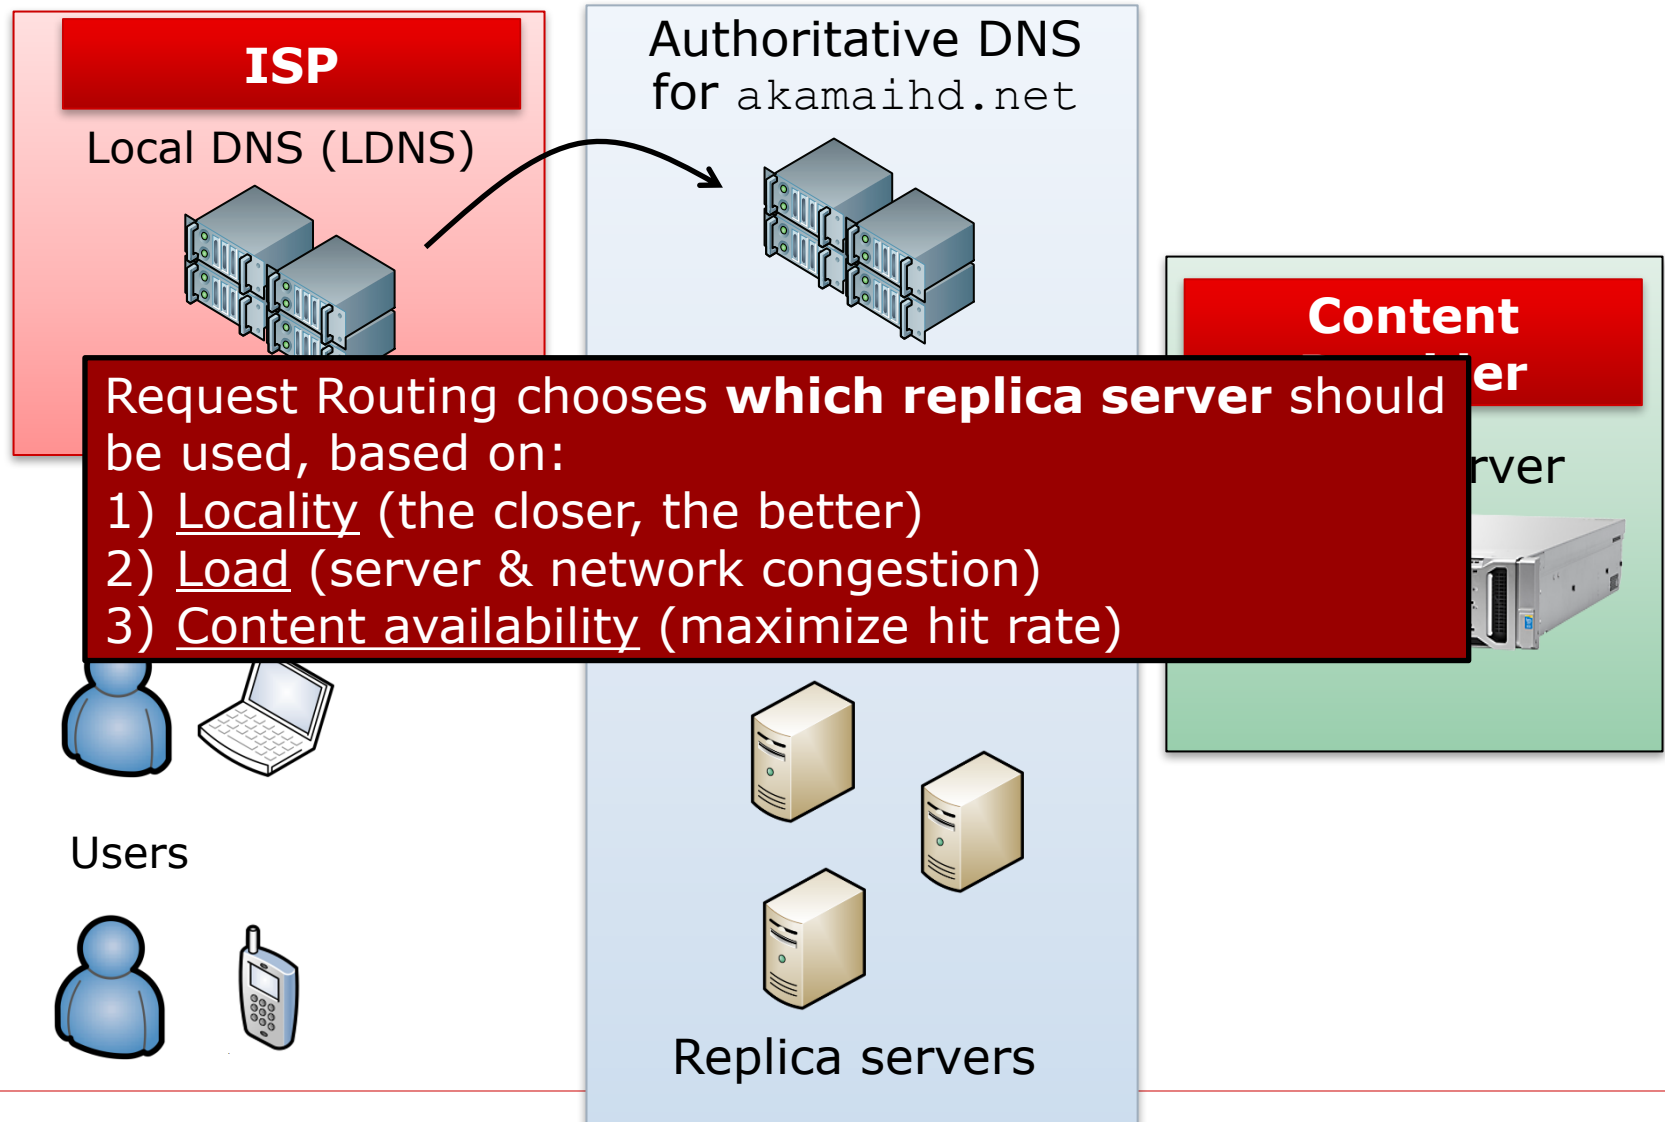
CDN: Architecture



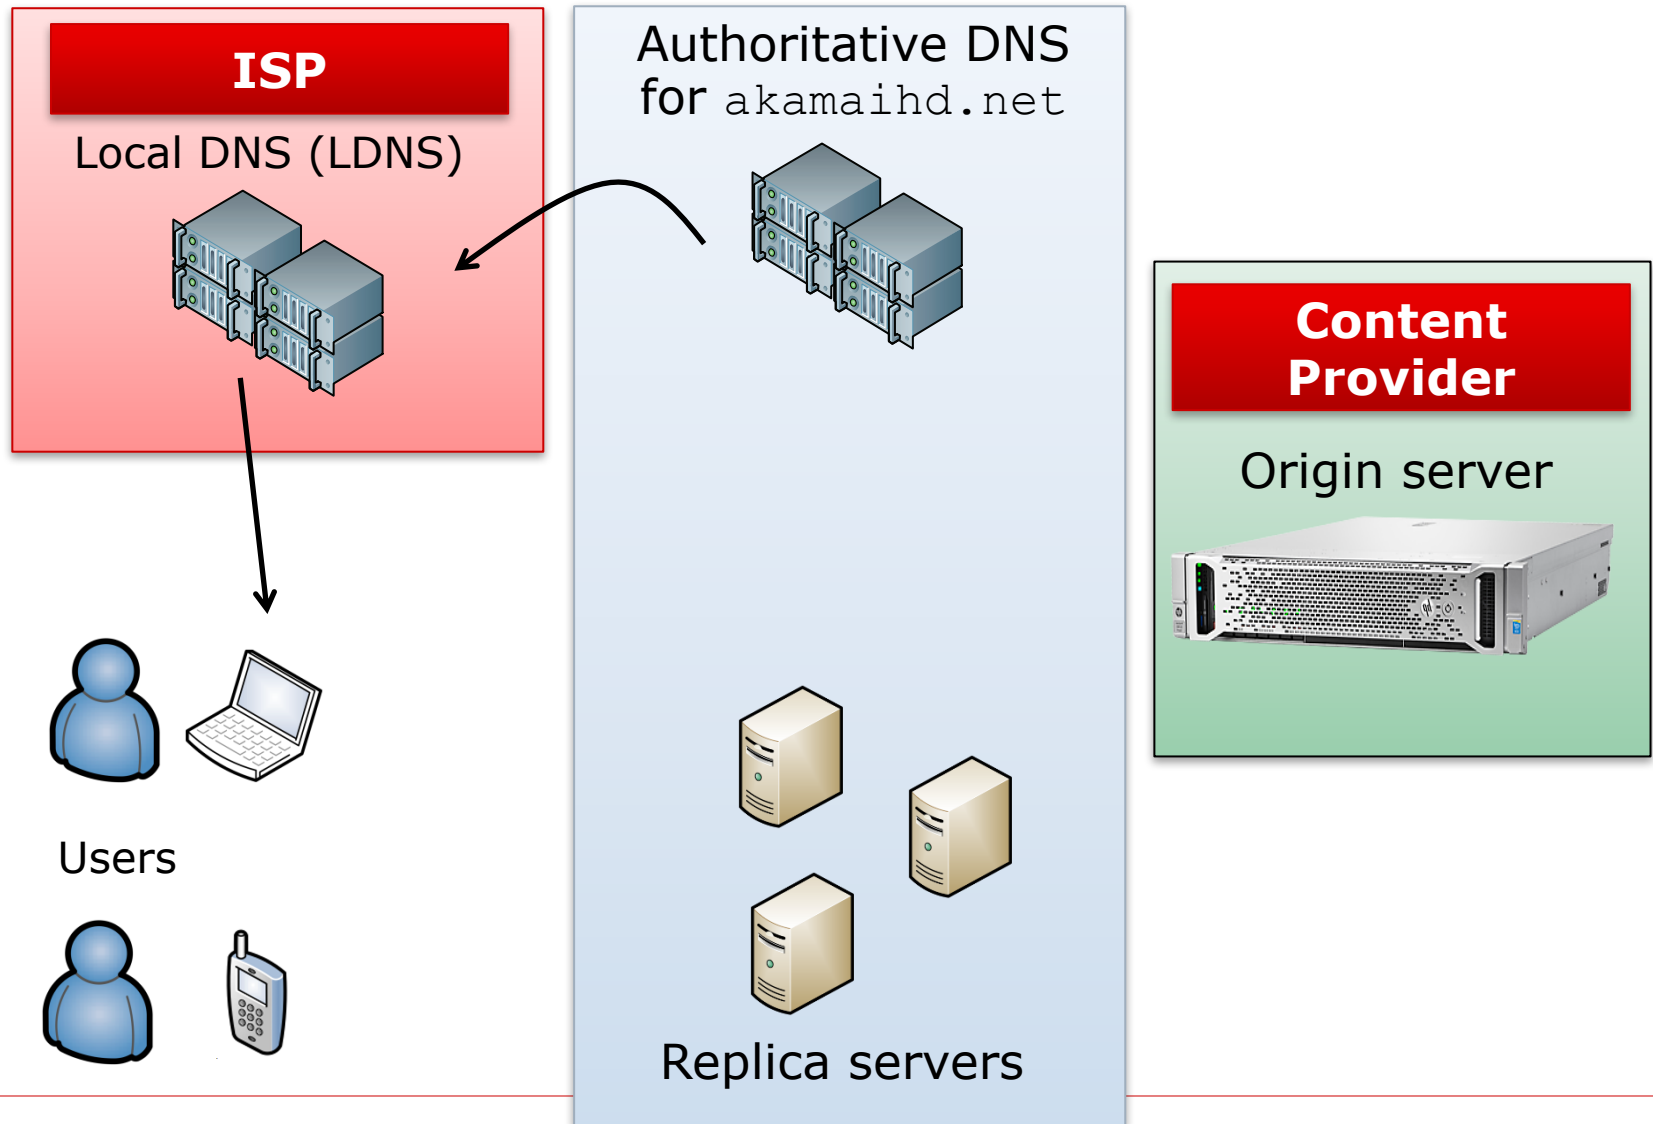
CDN: Architecture



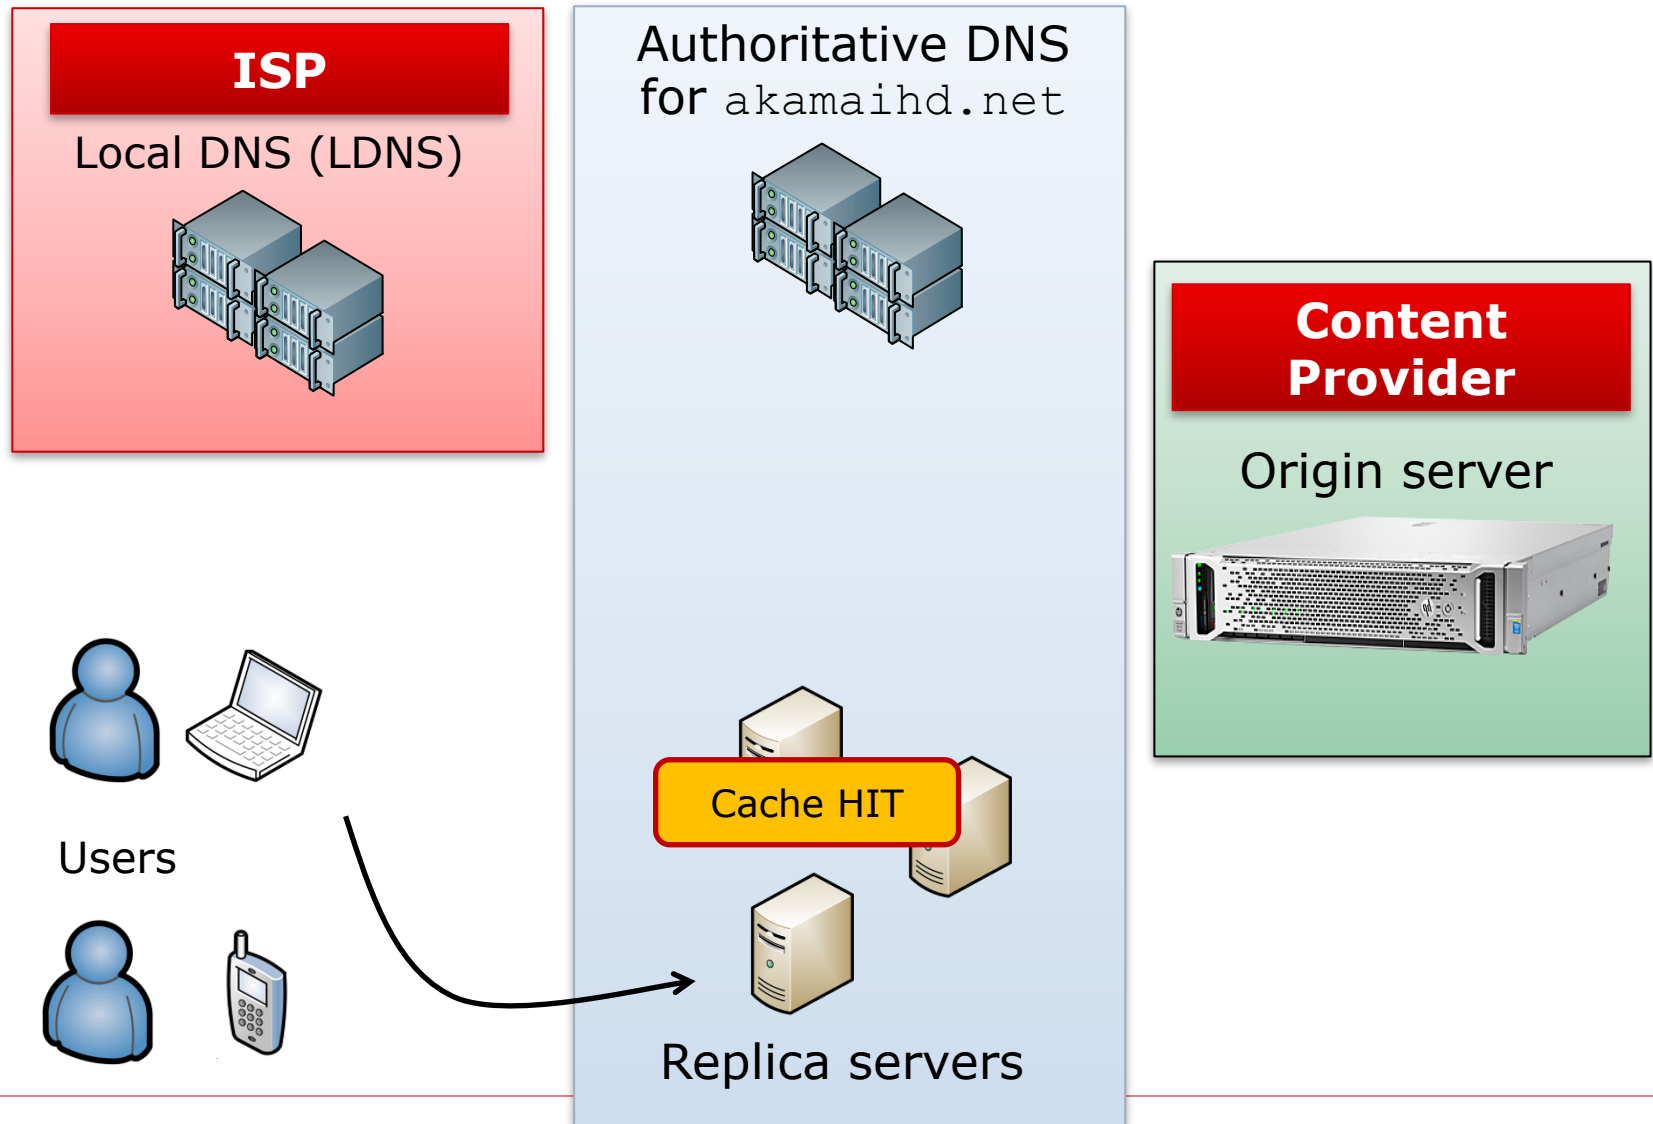
CDN: Architecture



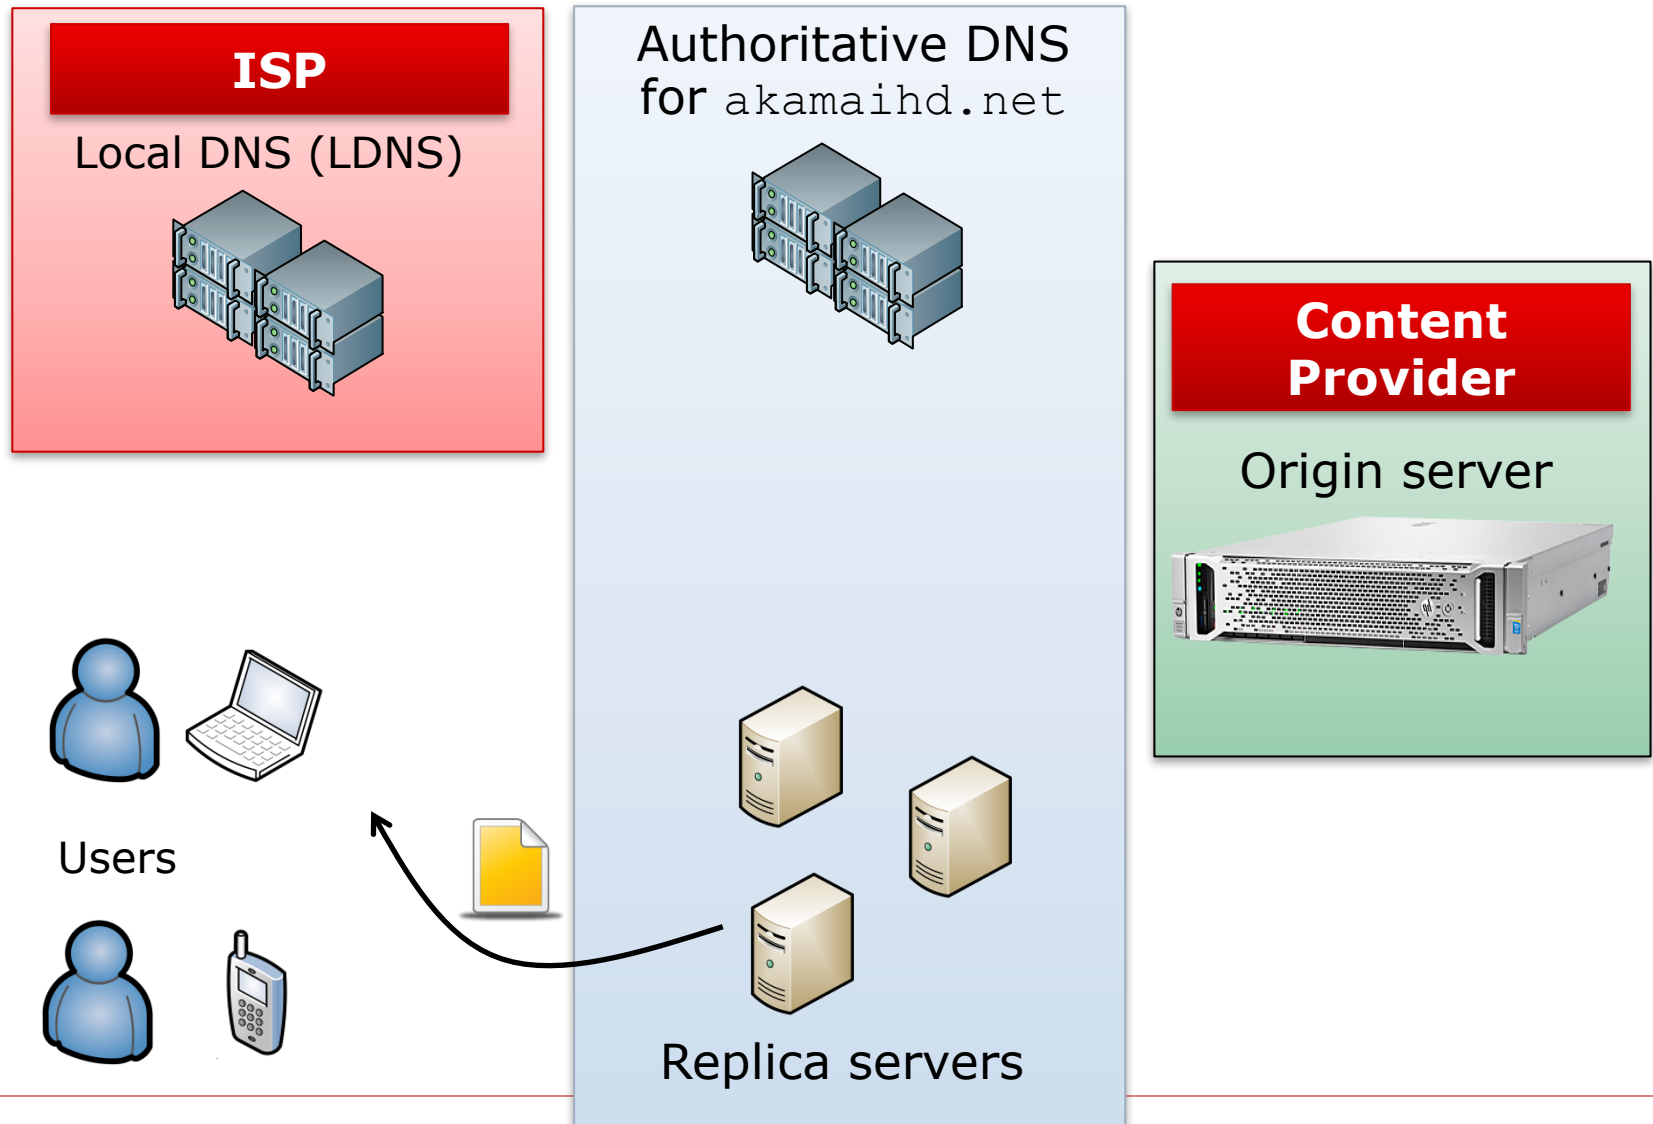
CDN: Architecture



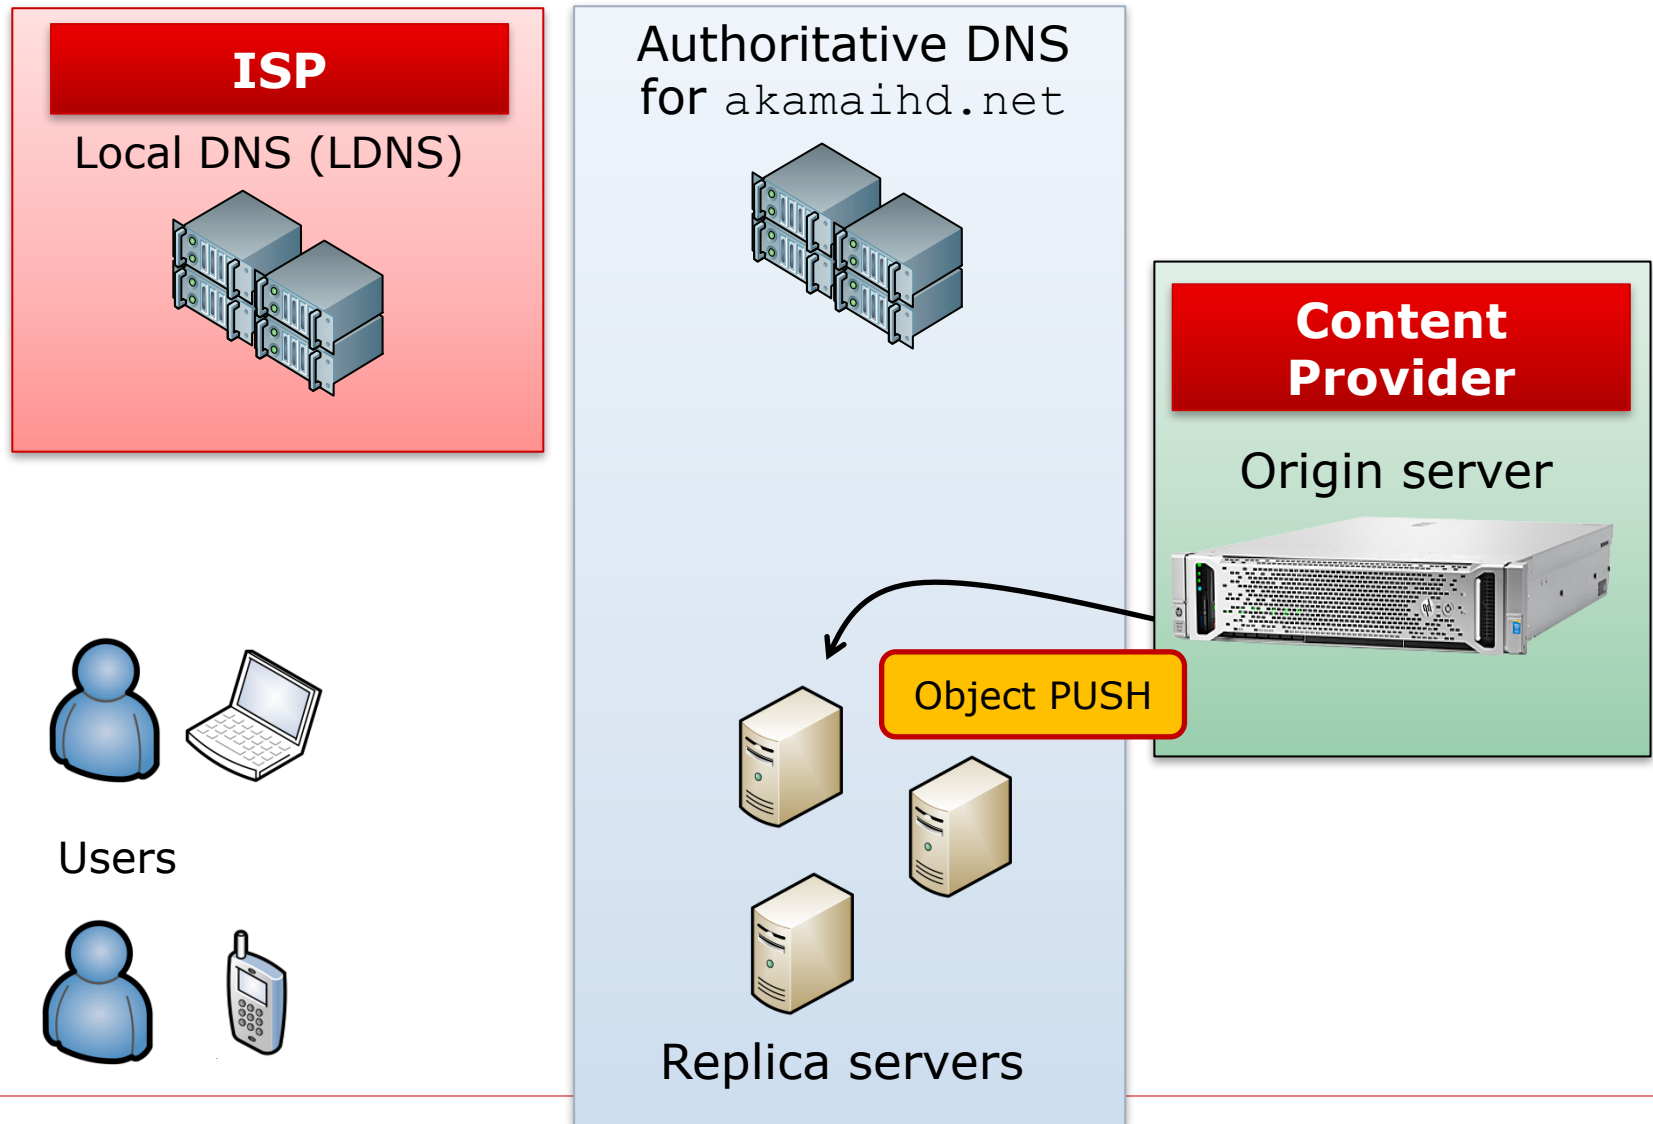
CDN: Architecture



CDN: Architecture



CDN: Architecture



CDN Challenges

1. Request Routing mechanisms are a **key component**:
 - DNS-based / URL Rewriting / HTTP redirection
2. Managing a CDN is **very expensive!**
 - Huge expenses for replica servers deployment
 - It is also expensive to manage the “*Mapping*” infrastructure
3. Security is *non-trivial*:
 - Use Akamai shared certificate
 - Or let them securely transfer the content under your domain

Information-Centric Networks

- Information-Centric Networks (ICNs)
 - A common label for many Future Internet designs
 - U.S. Funded projects (DONA, CCN/NDN, MobilityFirst)
 - EU Funded projects (Connect, PSIRP/PURSUIT, 4WARD/SAIL)

- Like CDN, ICN designs propose to:
 - “Flatten” the Internet
 - Replicate content to the “edge” of the network, closer to the location of end-users

- Unlike CDN:
 - ICNs are run by many operators (in a distributed way)
 - They are usually built as *clean-slate* network layer protocols (there are some exceptions, e.g.: DONA)

Information-Centric Networks

- All ICN proposals share a common set of principles:
 1. A new, content oriented **addressing namespace**
 - *From location-based* (IP address of the machines)
 - *To location-independent content names*
 2. **Caching** functionalities added to network nodes
 - Wires move the data **in space**
 - Memories move the data **in time**
 3. The same content can potentially be retrieved from **many locations** (*even untrusted third parties*)
 - **Security** cannot be enforced by securing the communication channel, rather, it must *be built right inside the content* itself

- Among all ICN proposals we focus our attention on NDN/CCN
 - Open source code base: www.ccnx.org
 - This is the proposal that received most of the attention from the scientific community

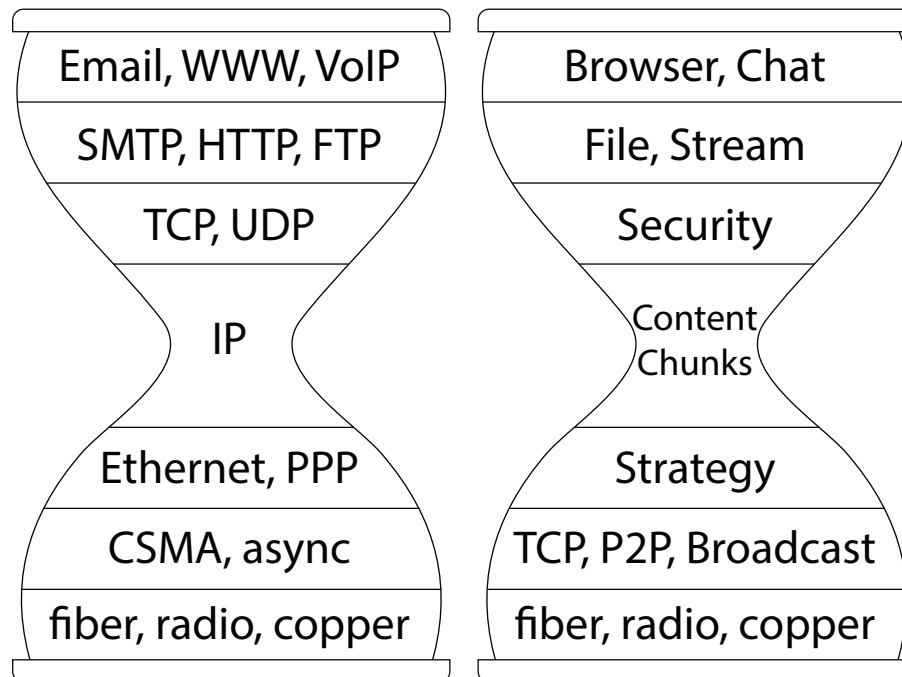
Content Distribution in Internet

Named Data Networking
Content-Centric Networking

(NDN/CCN)


Content-Centric Networking

- The Content-Centric Networking (CCN) design from PARC was originally described in a [Google tech talk](#) by Van Jacobson
- The Named Data Networking (NDN) project, funded by the US Future Internet Architecture program, is further developing the pioneering work accomplished in CCN
- **Design principle** : Create a new thin waist for the hourglass of the Internet

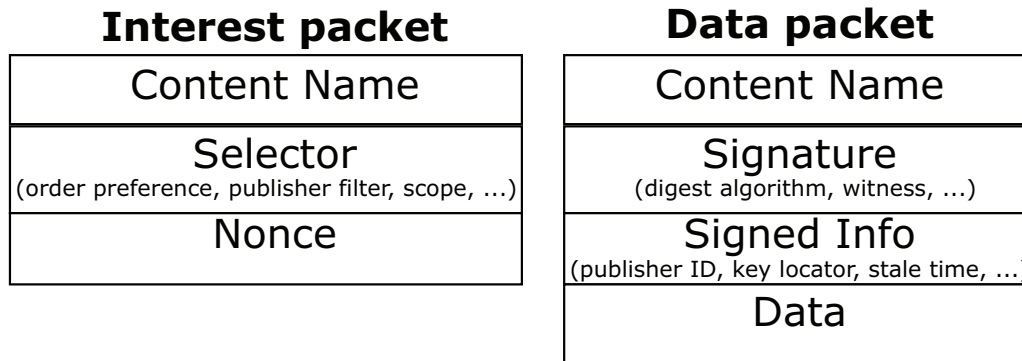


CCN: Naming & Packets


- Names are hierarchical structures, opaque to the applications. Each name is a list of variable-length components



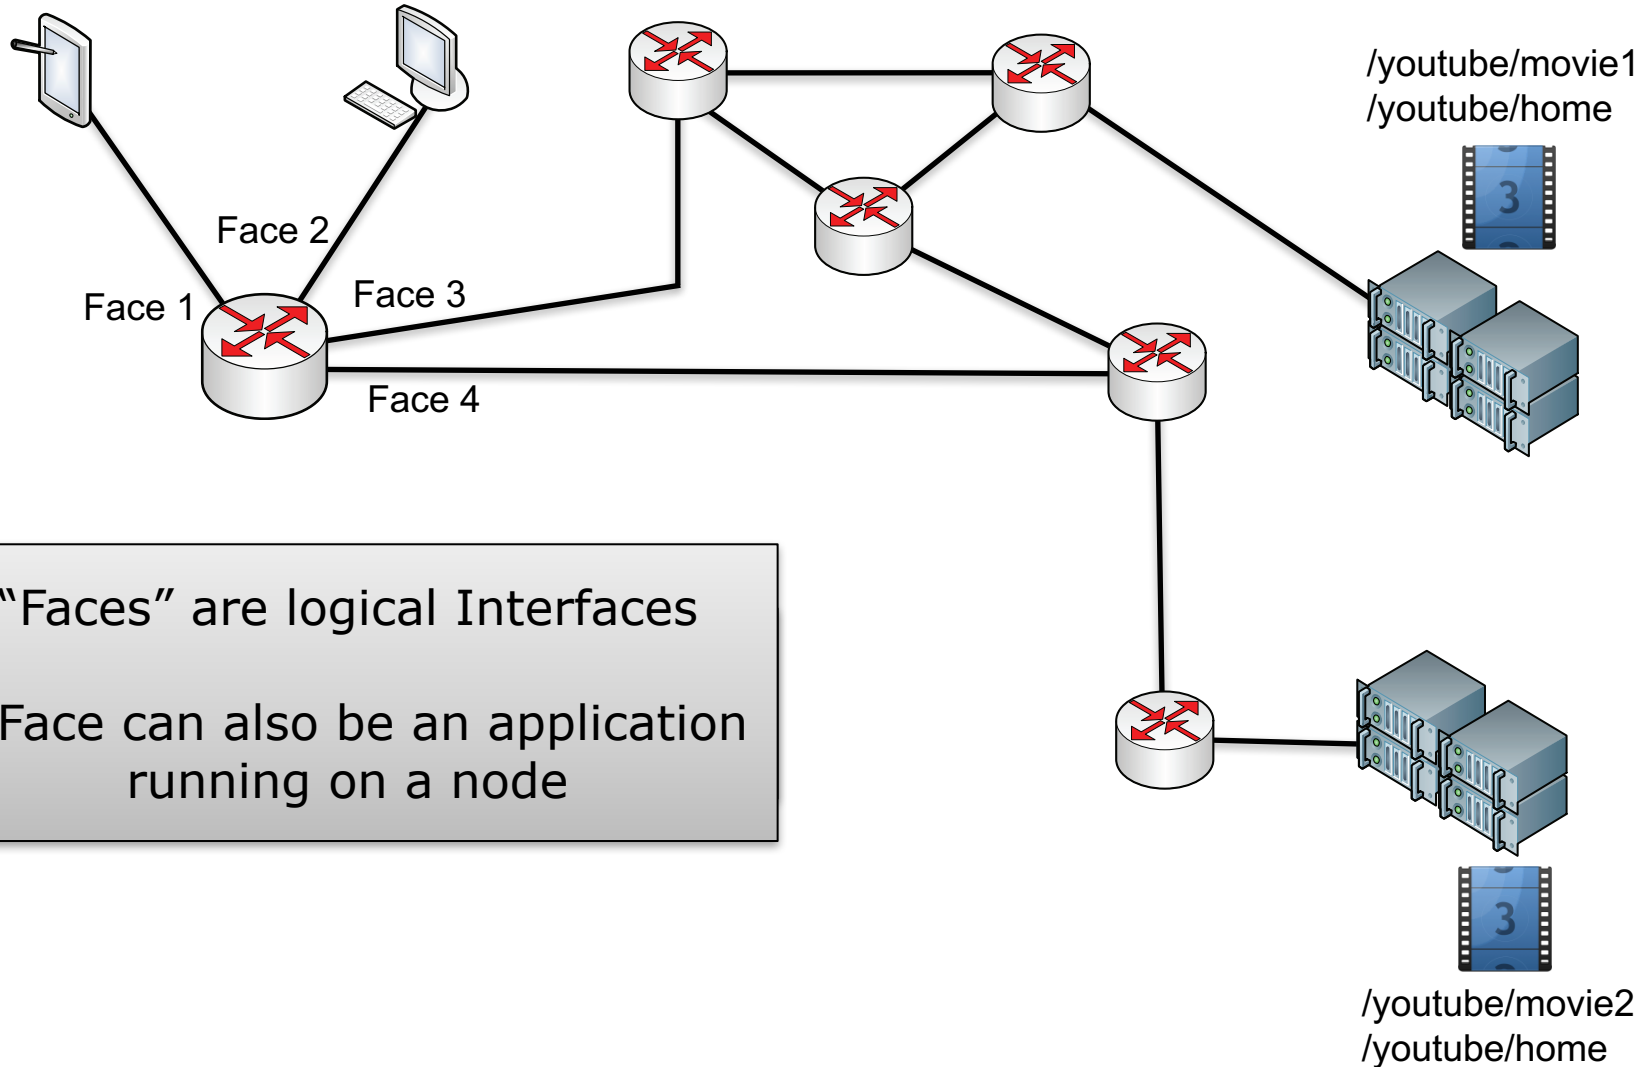
- Two type of packets are envisioned:
 - Interest Packet
 - Data Packet



CCN: The Node Model

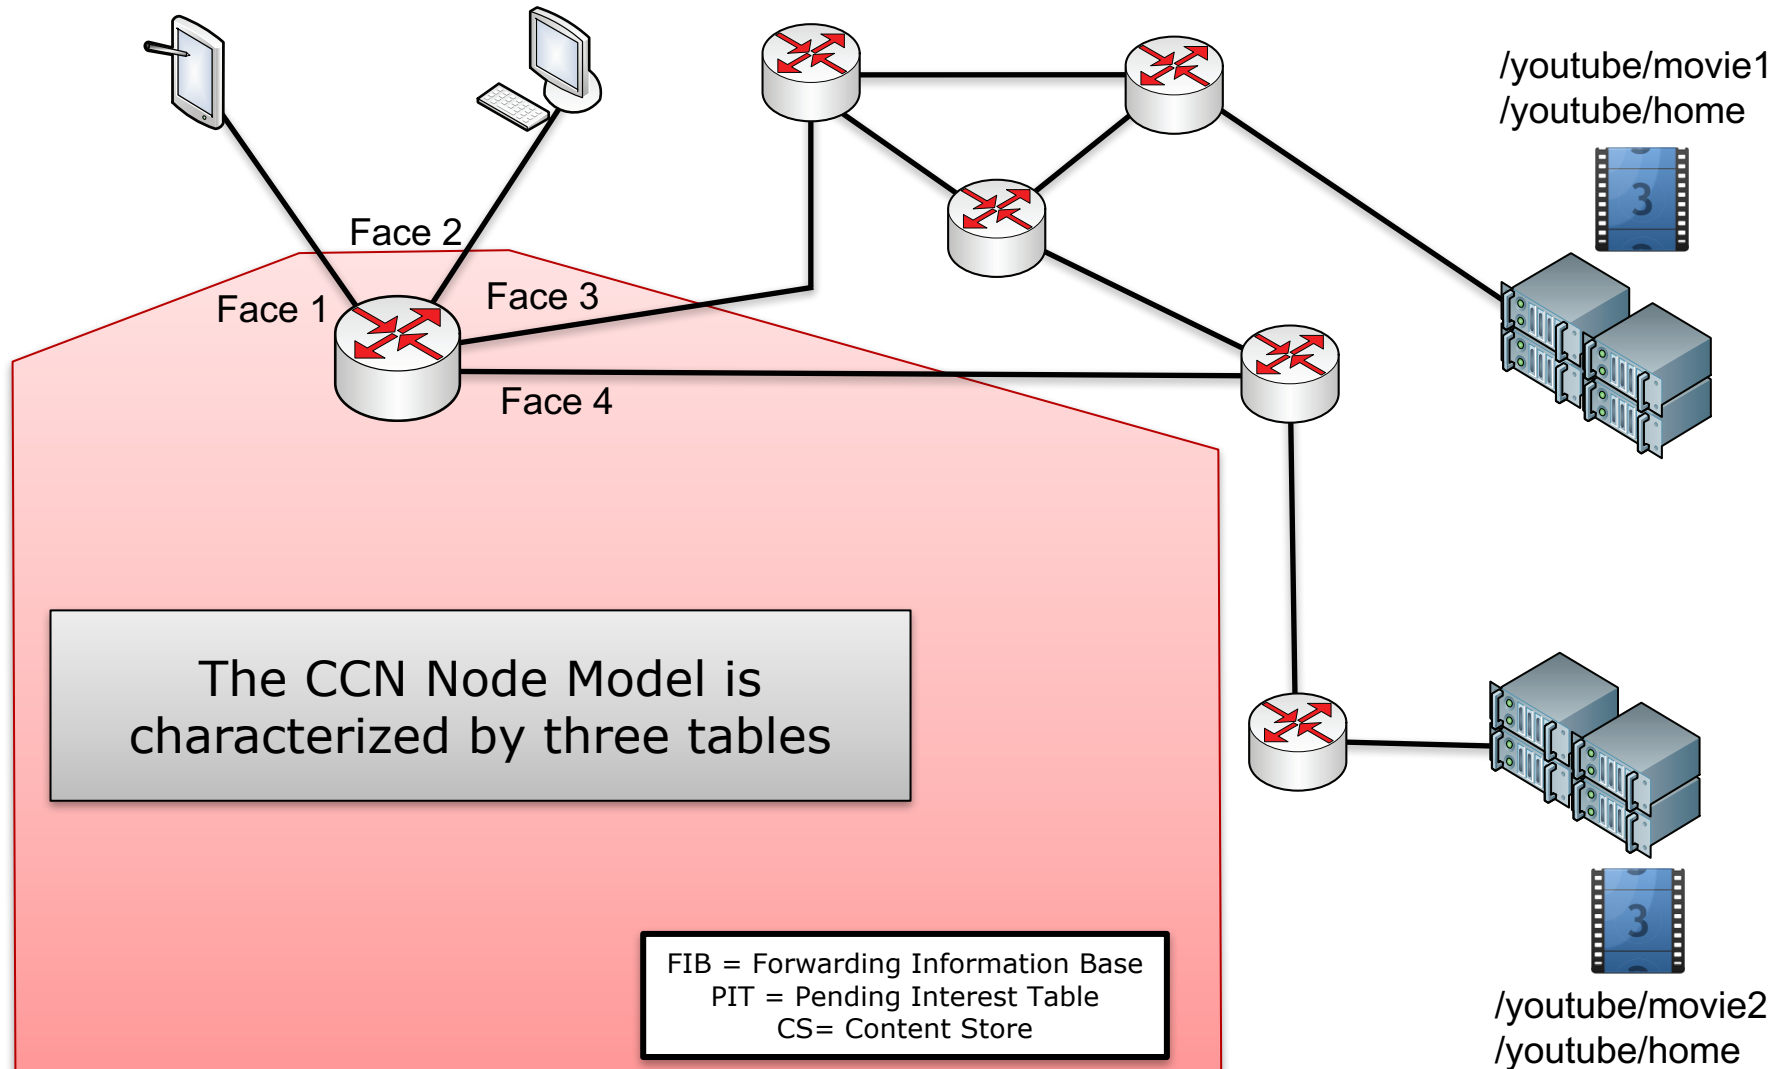


CCN: The Node Model

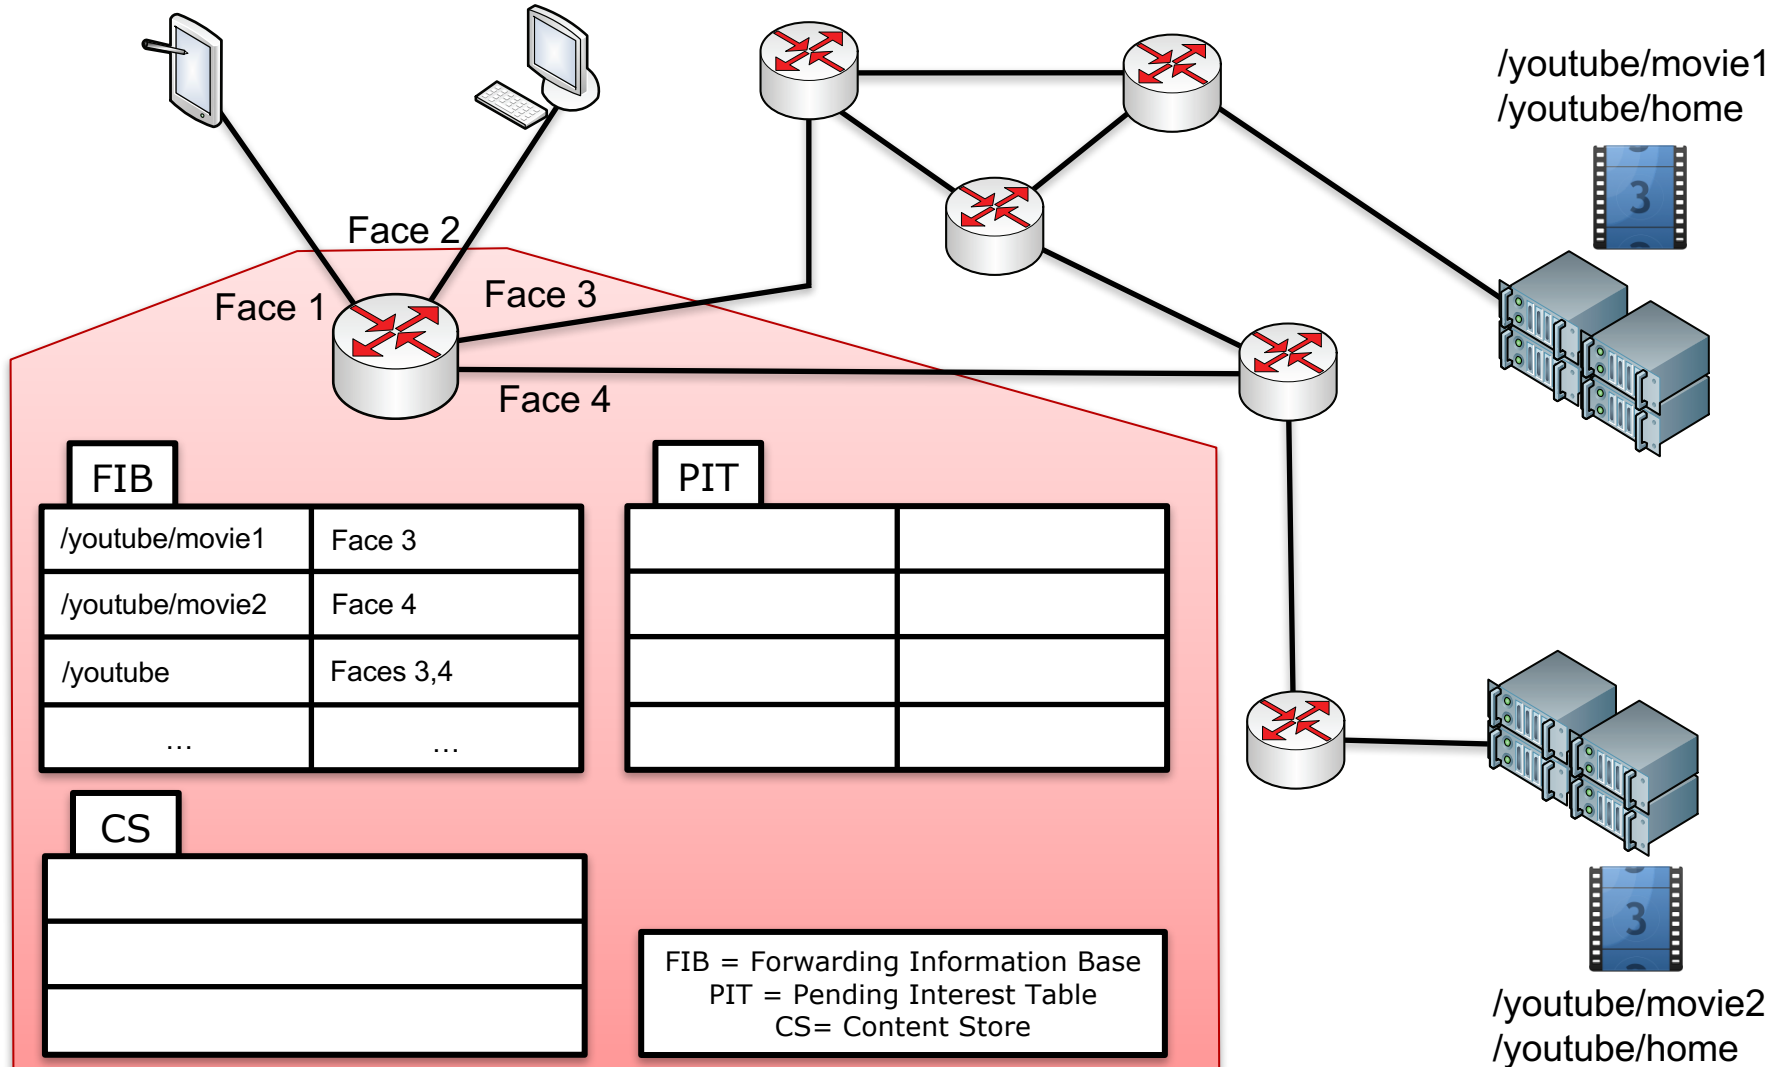


“Faces” are logical Interfaces
A Face can also be an application running on a node

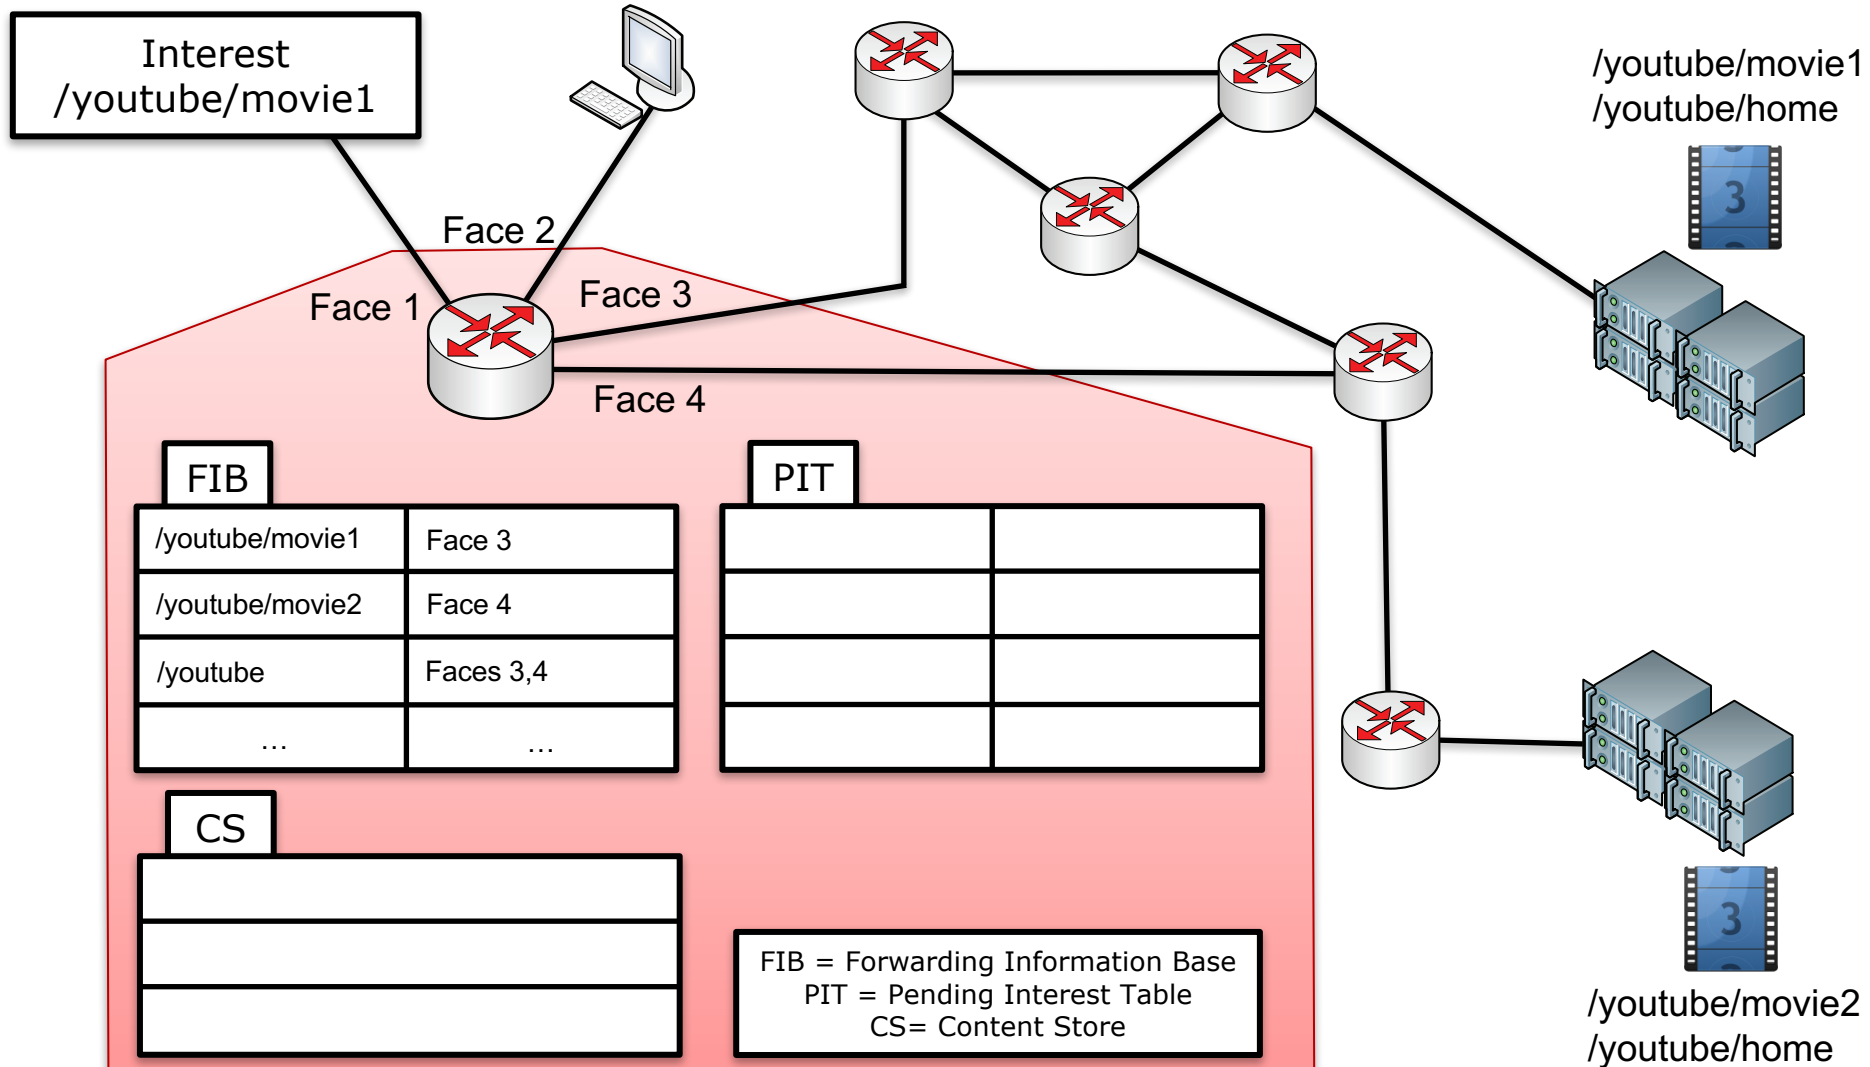
CCN: The Node Model



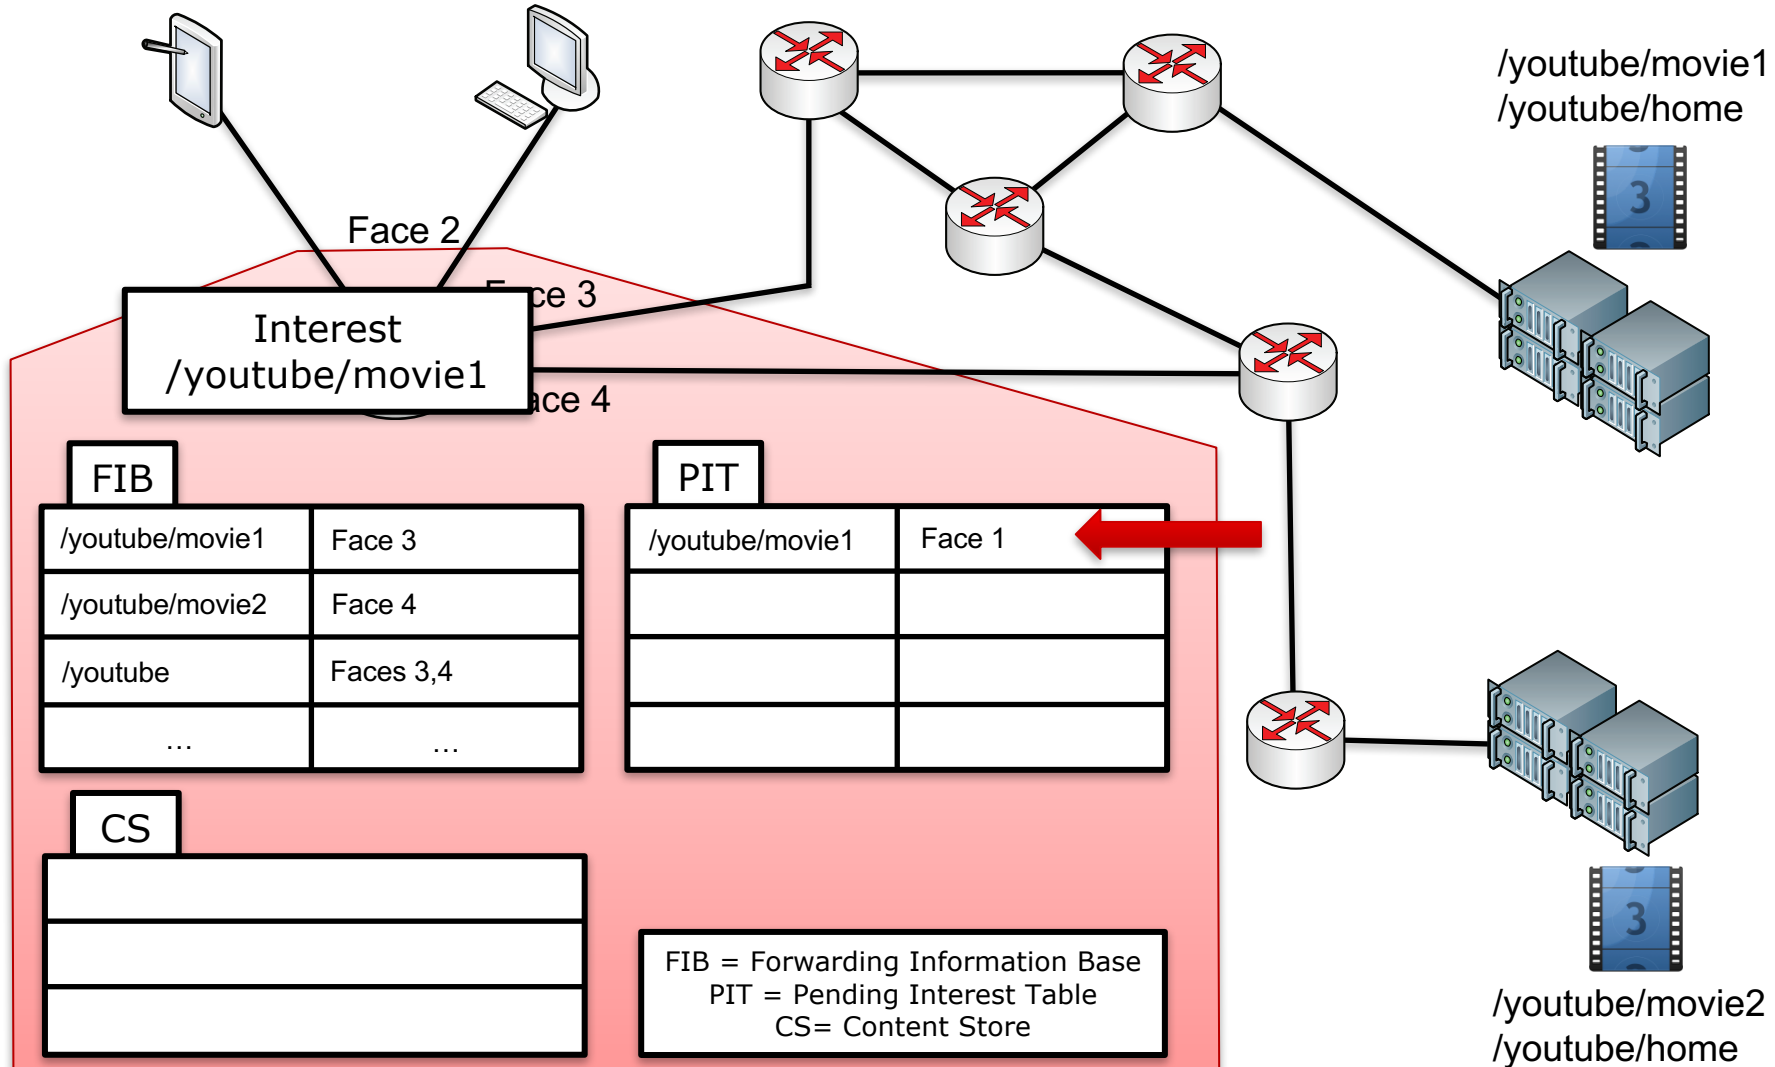
CCN: The Node Model



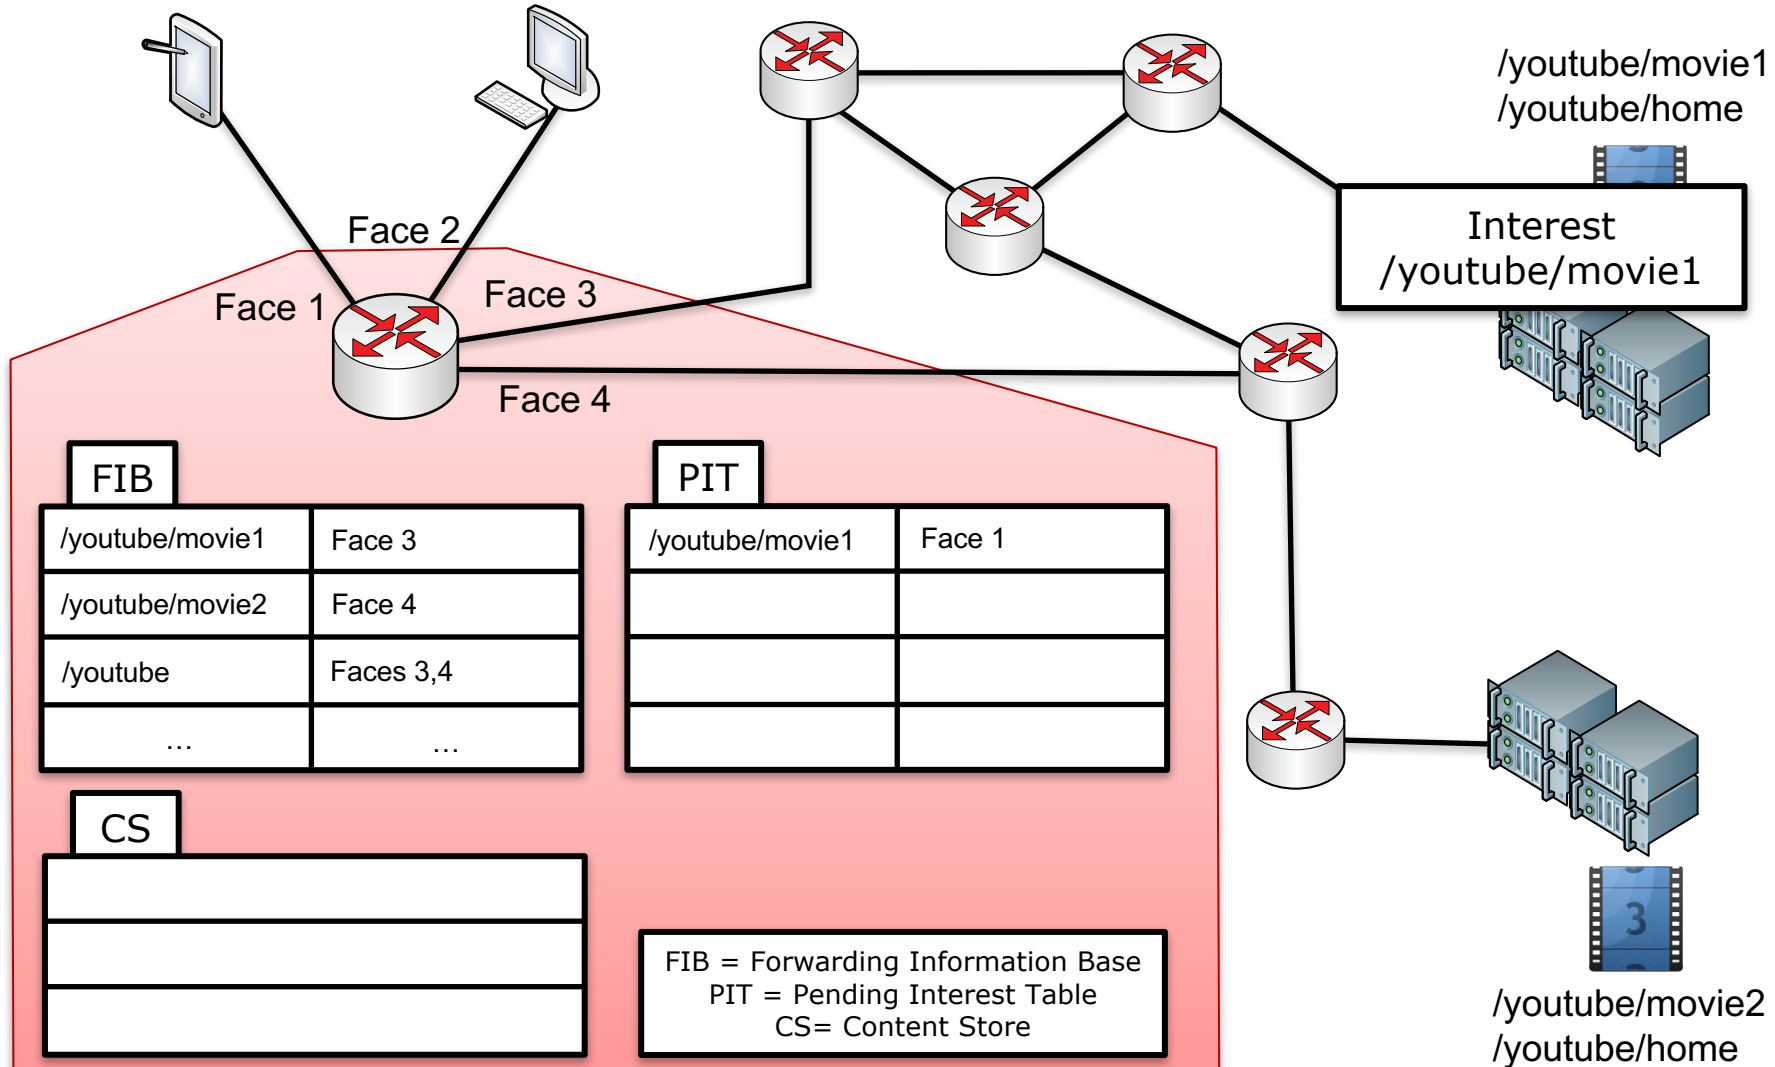
CCN: The Node Model



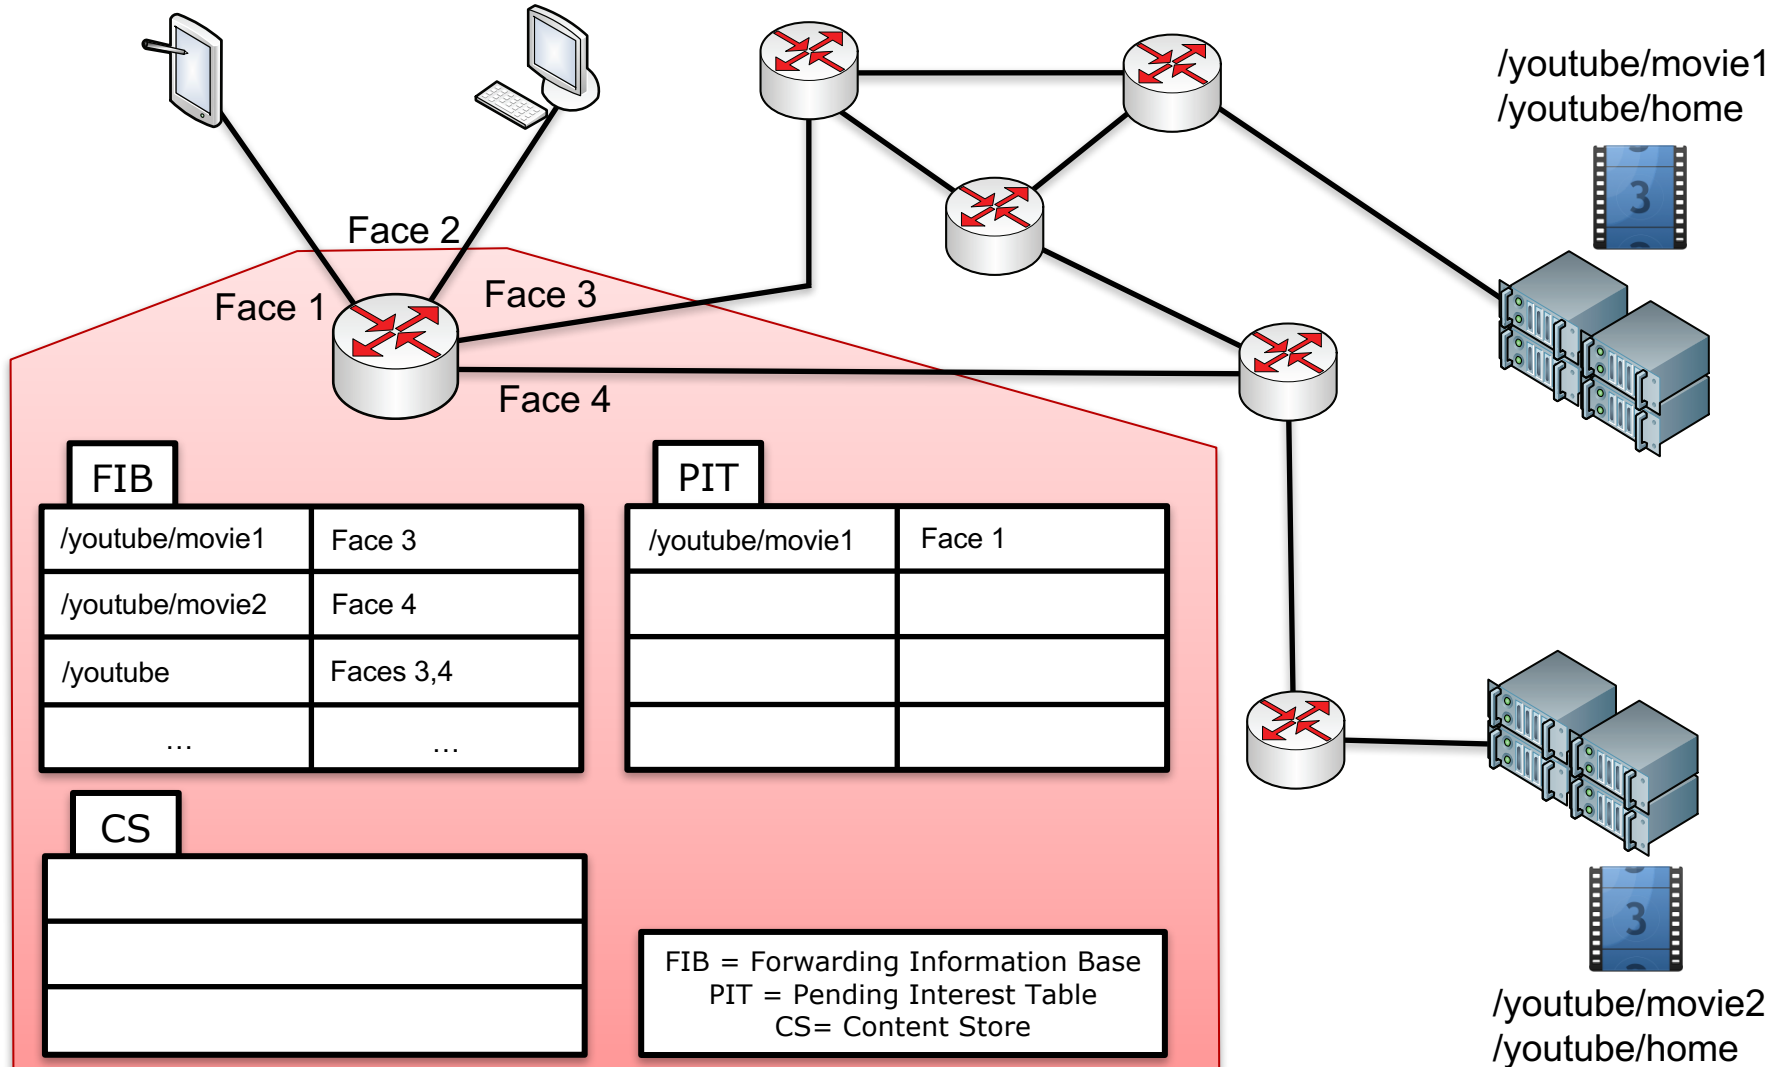
CCN: The Node Model



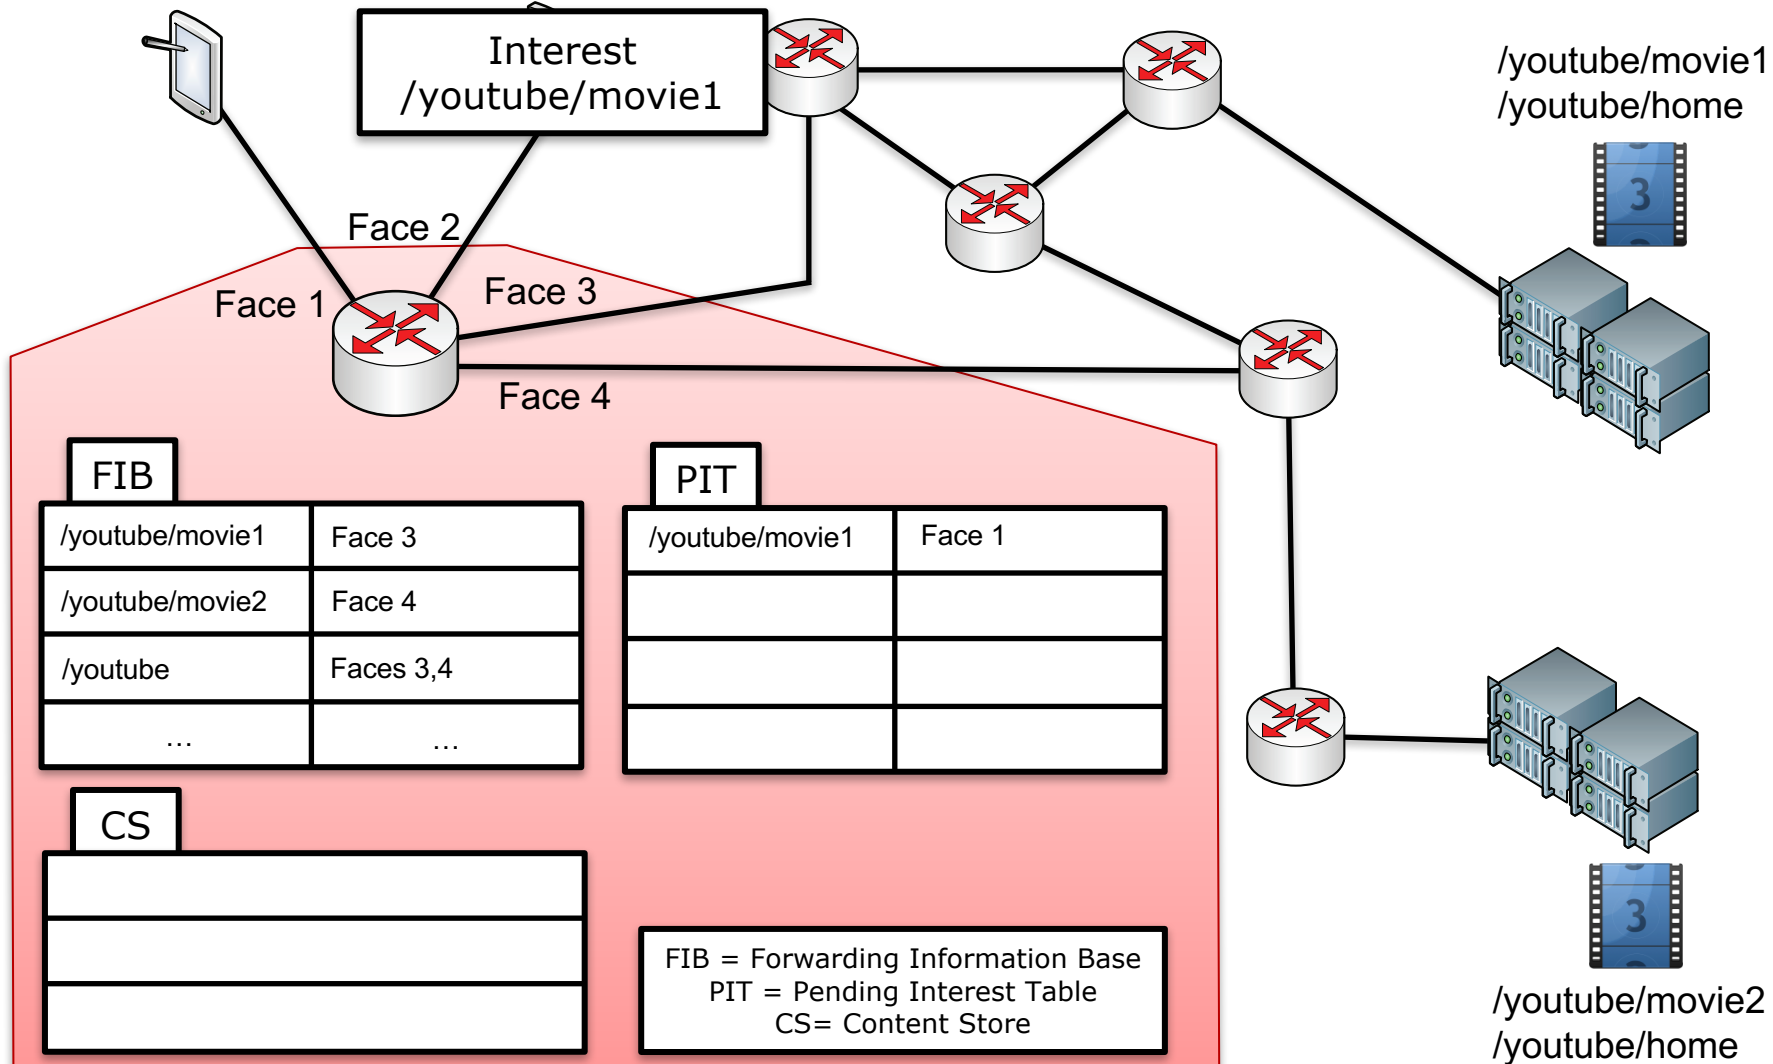
CCN: The Node Model



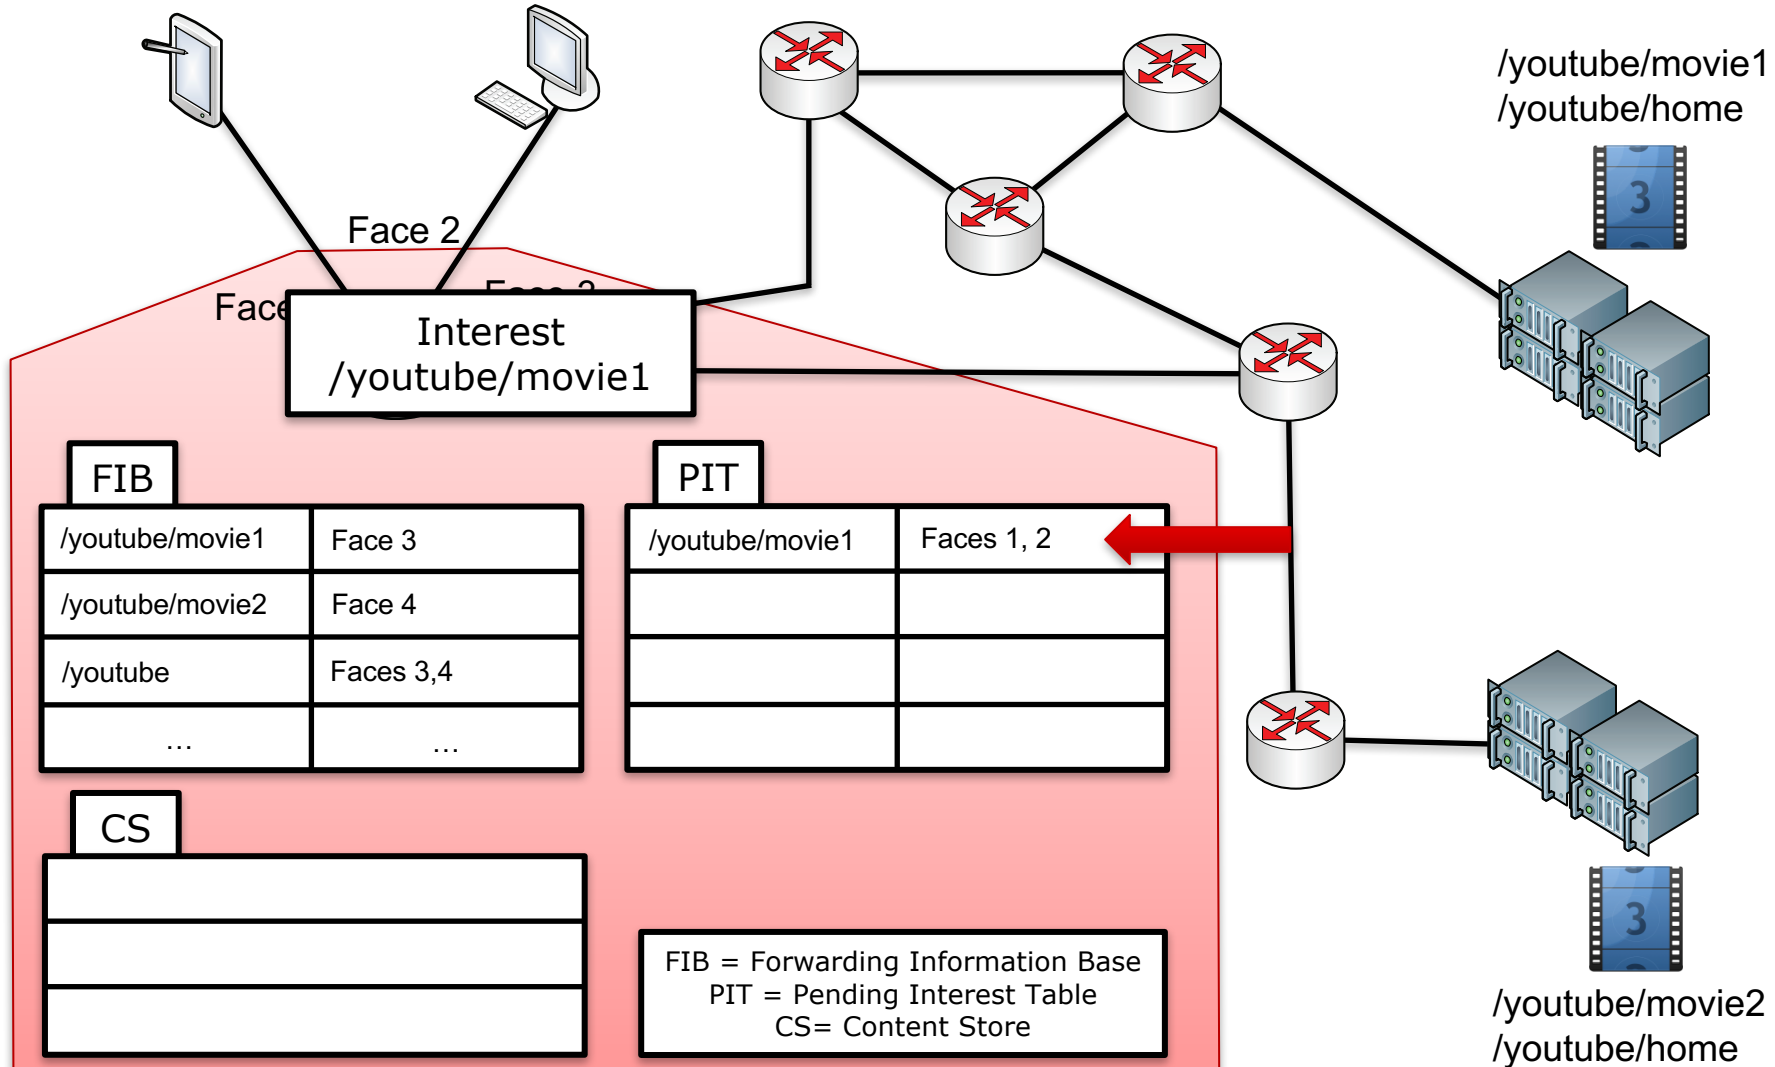
CCN: The Node Model



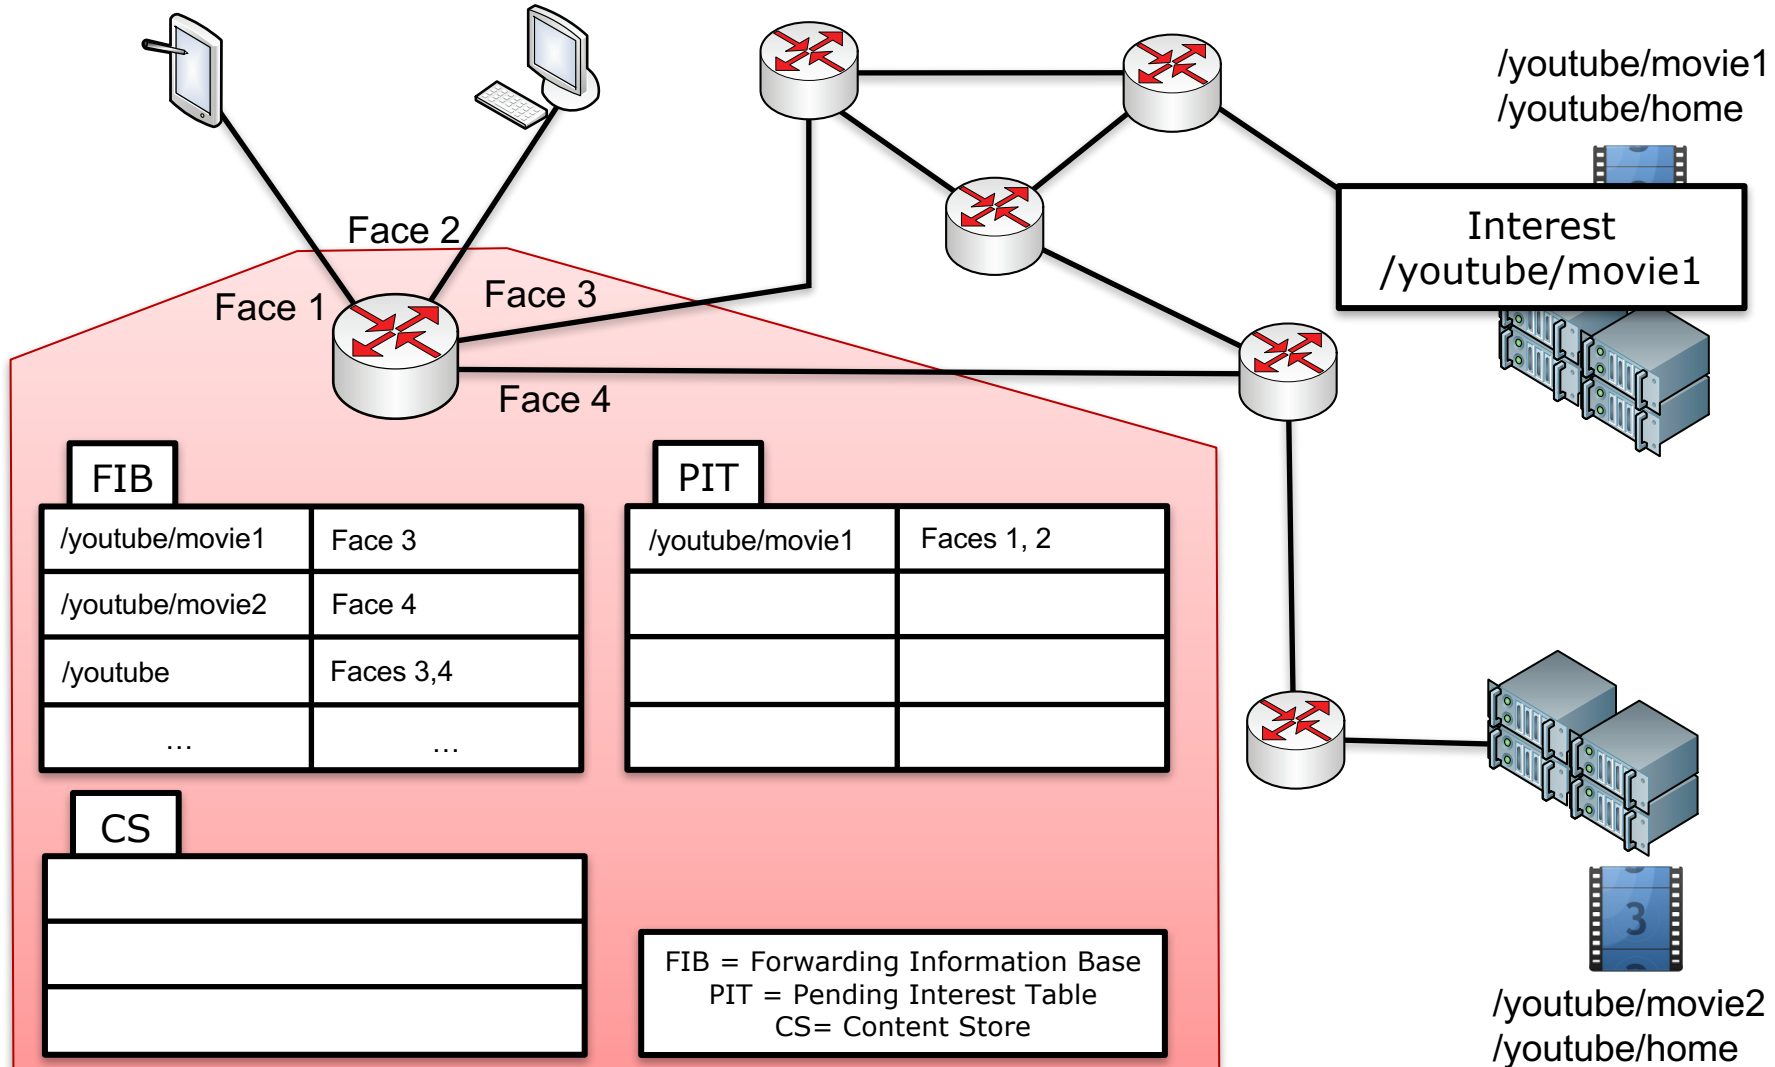
CCN: The Node Model



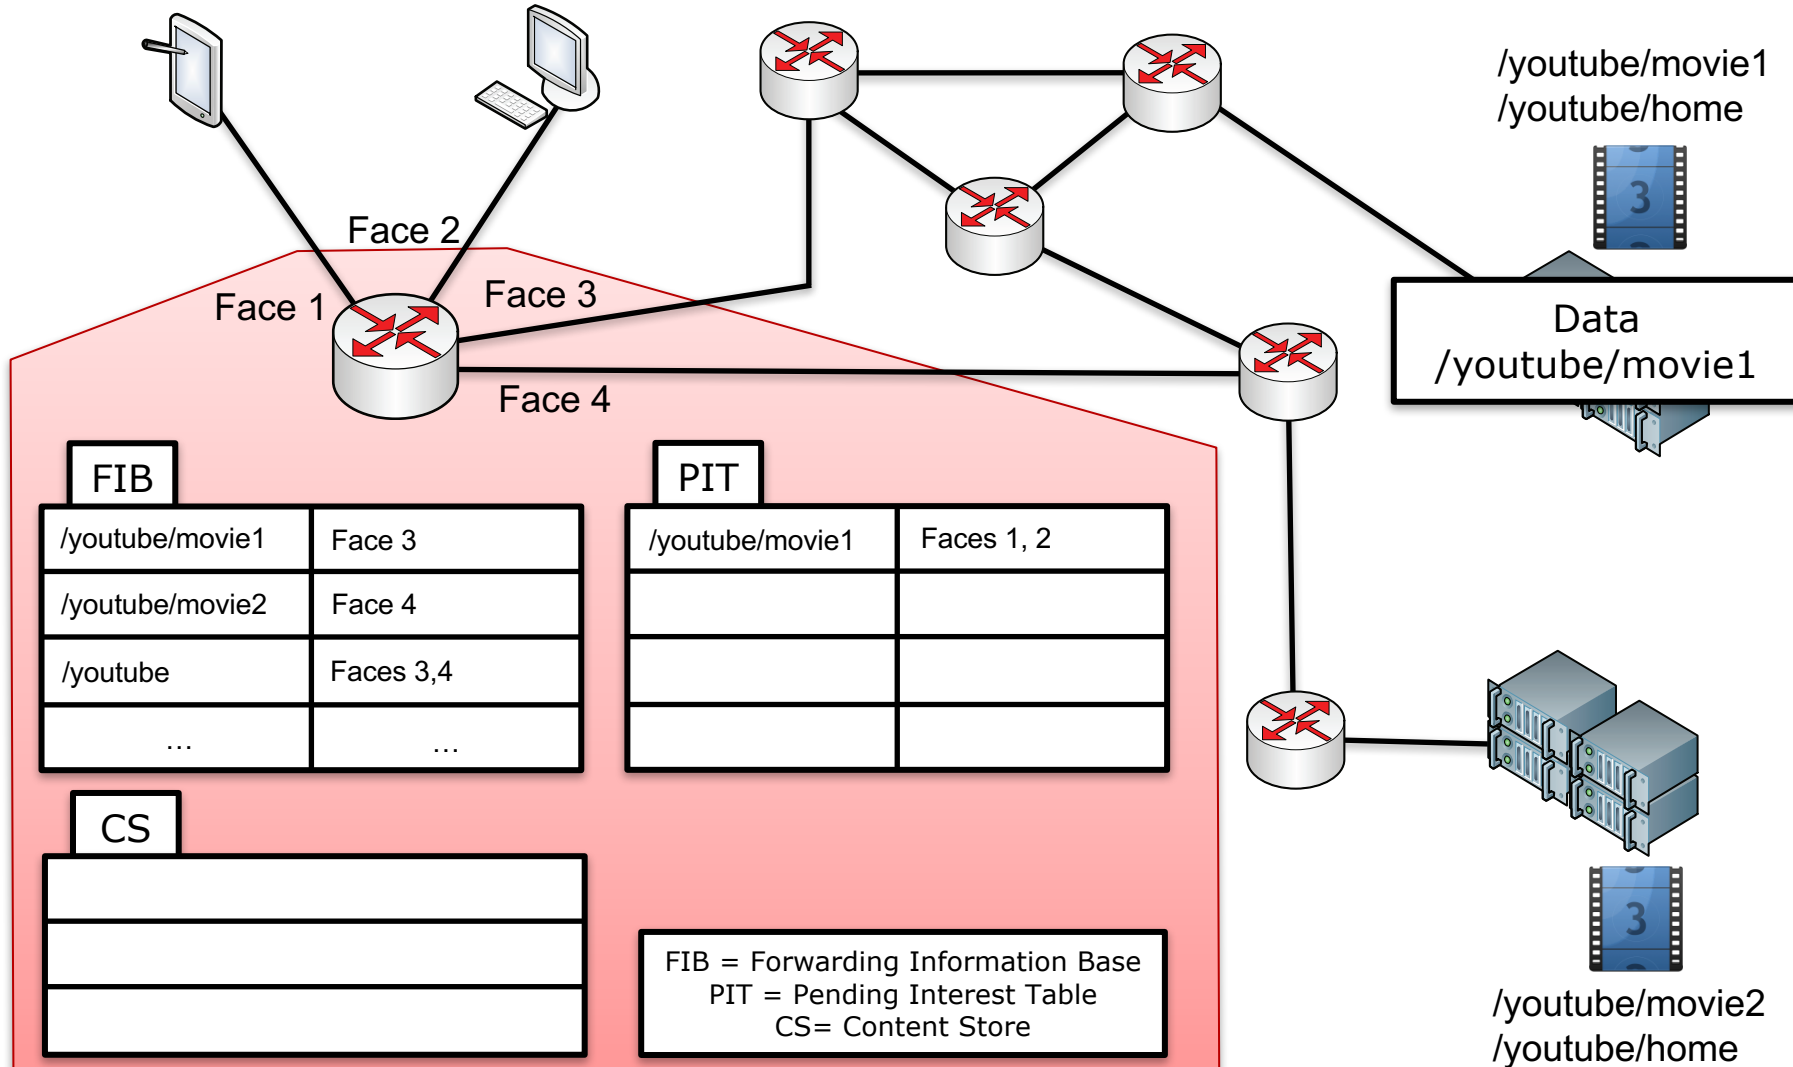
CCN: The Node Model



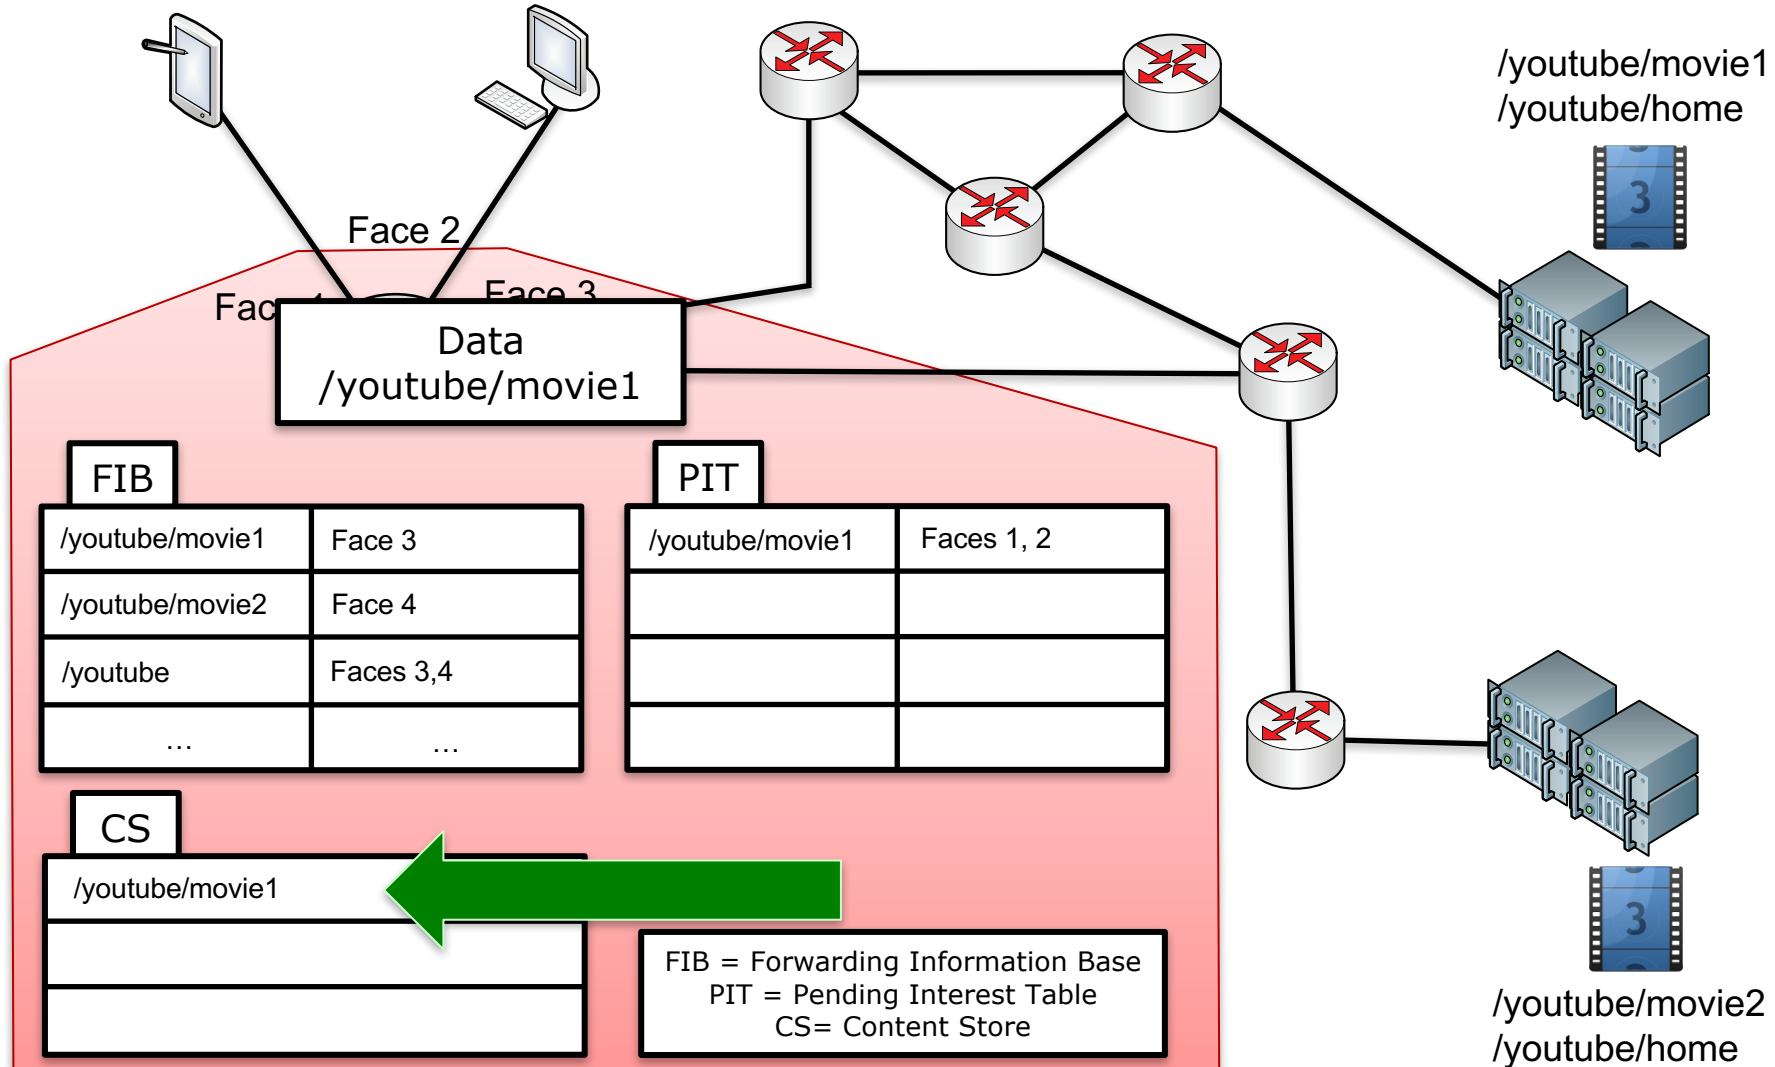
CCN: The Node Model



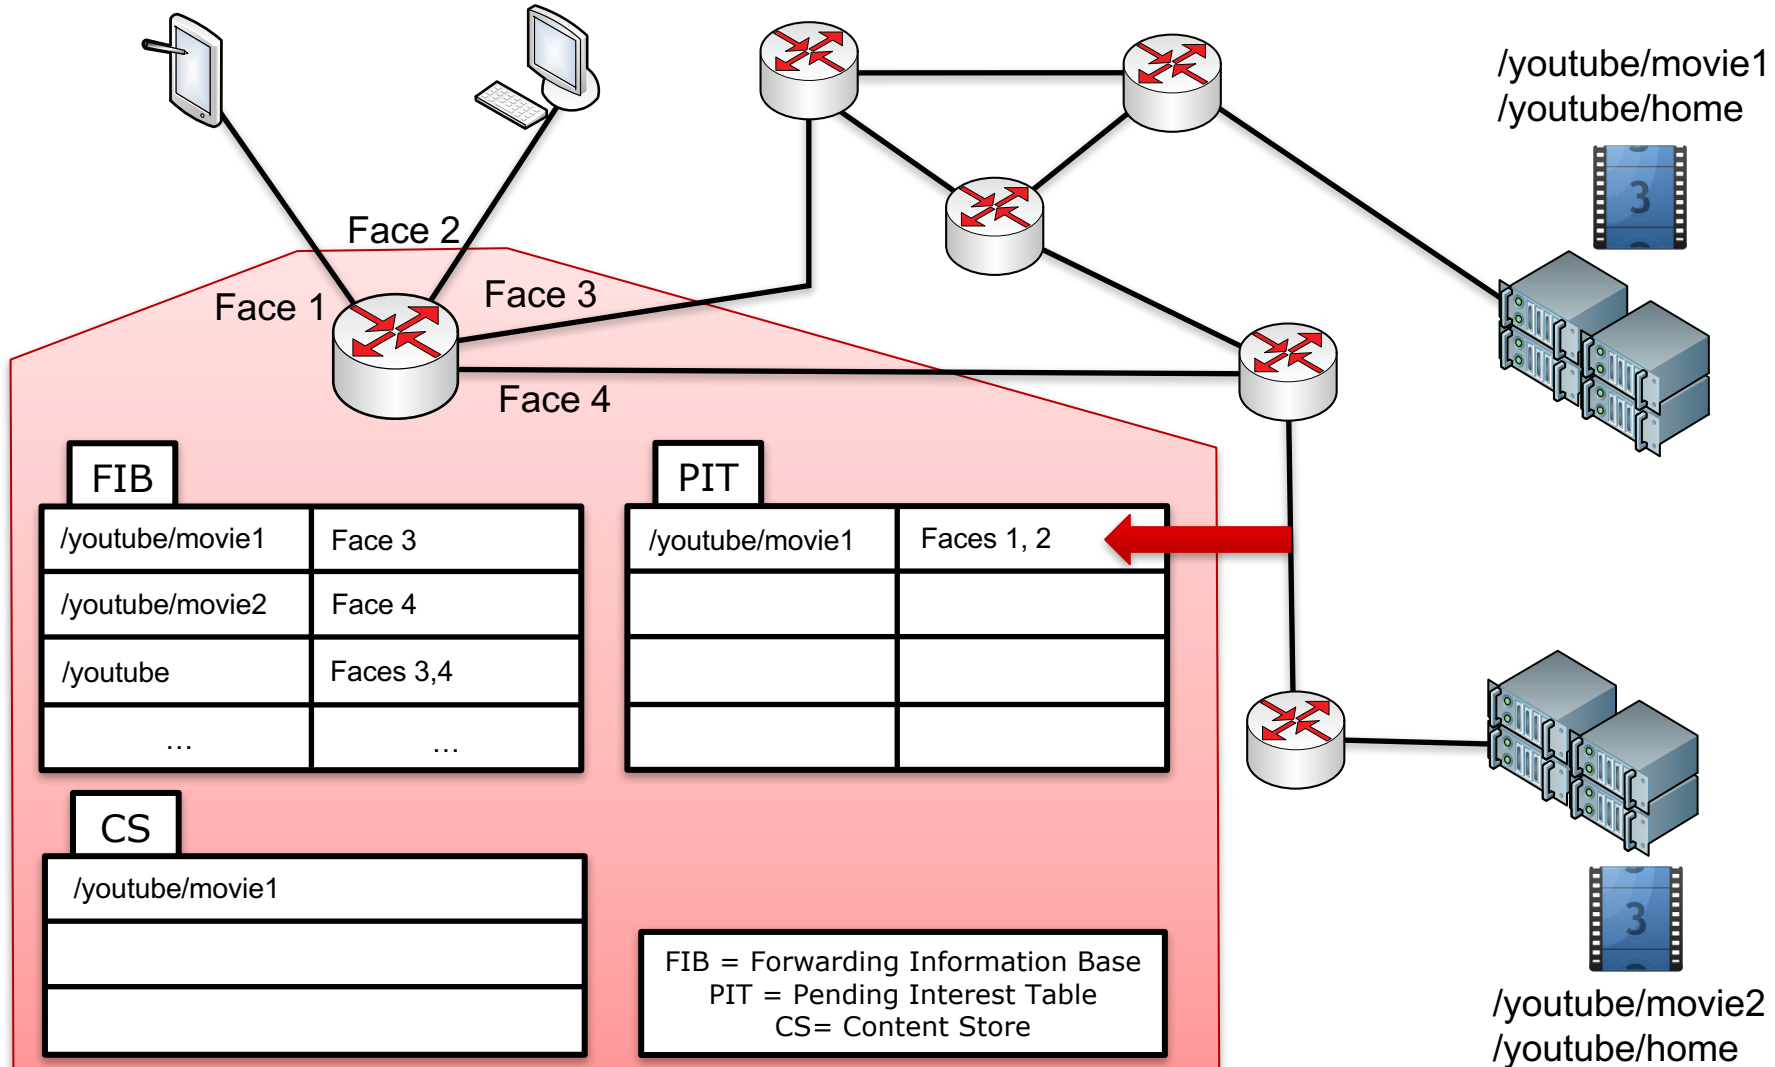
CCN: The Node Model



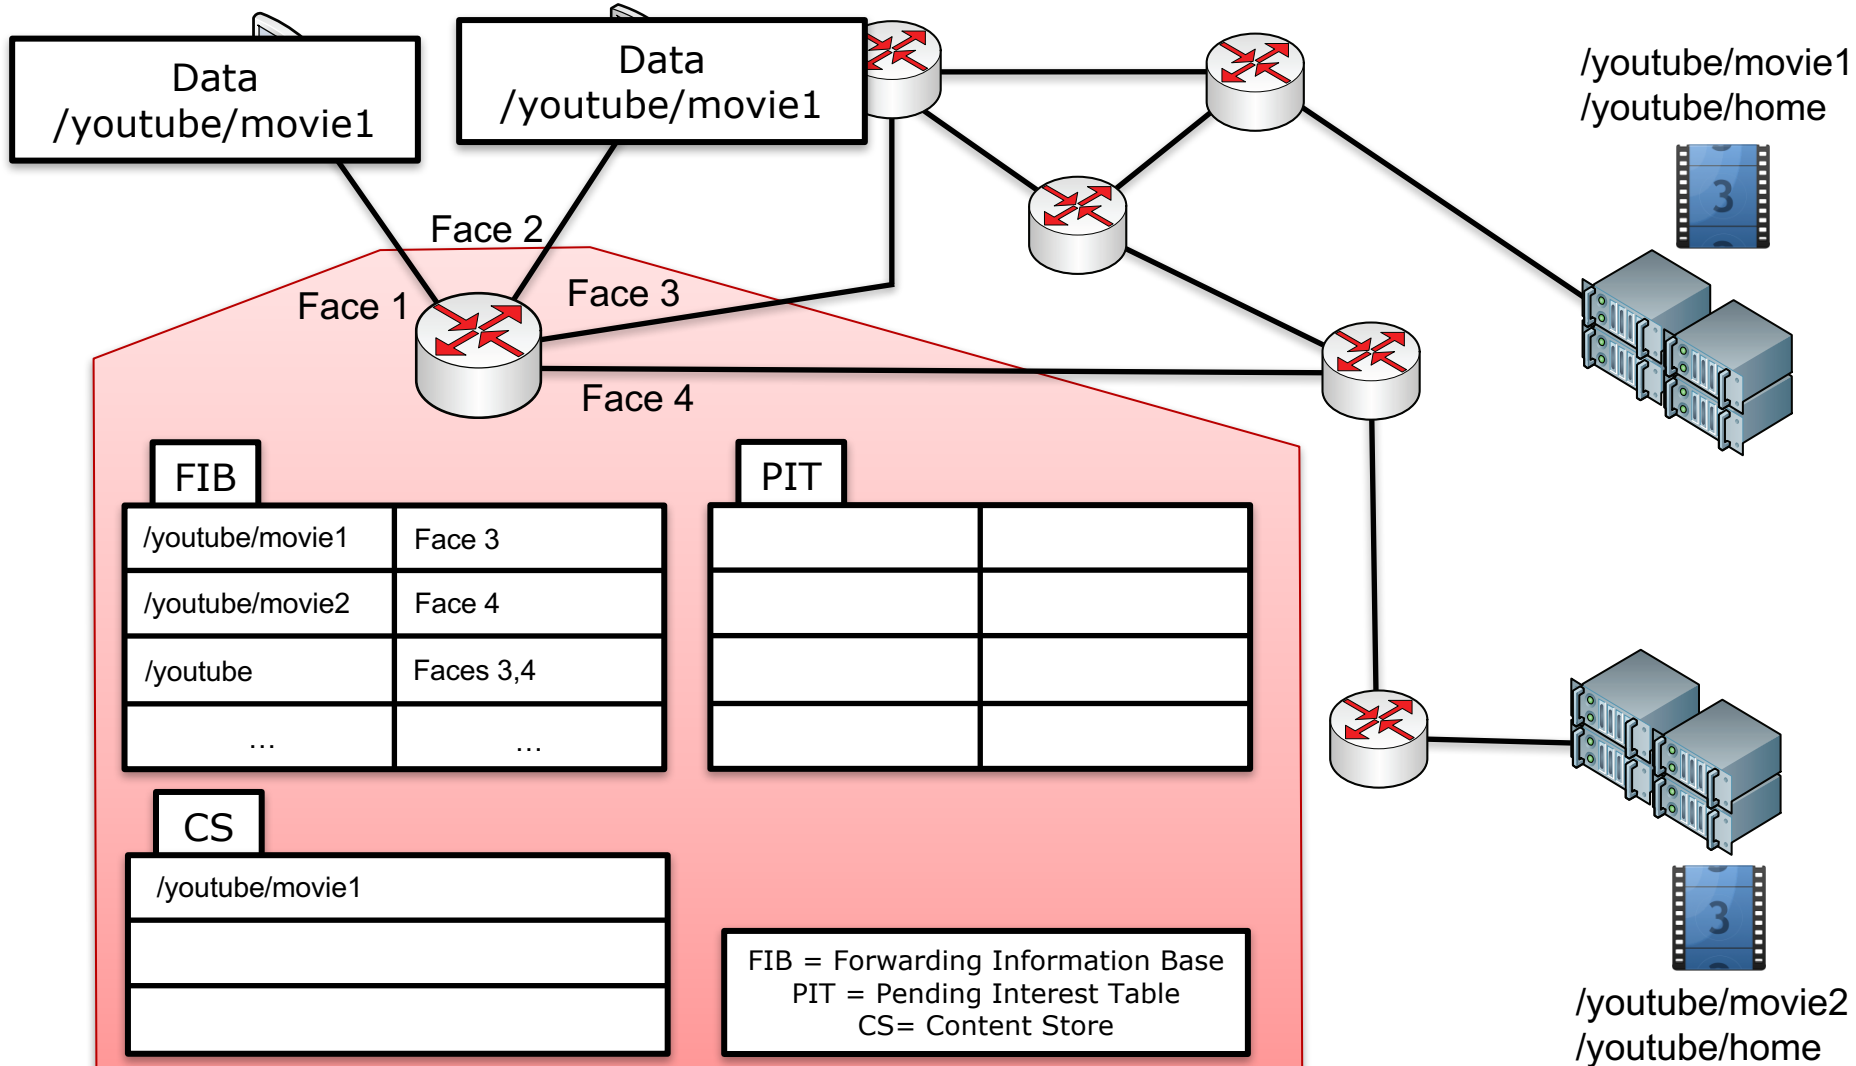
CCN: The Node Model



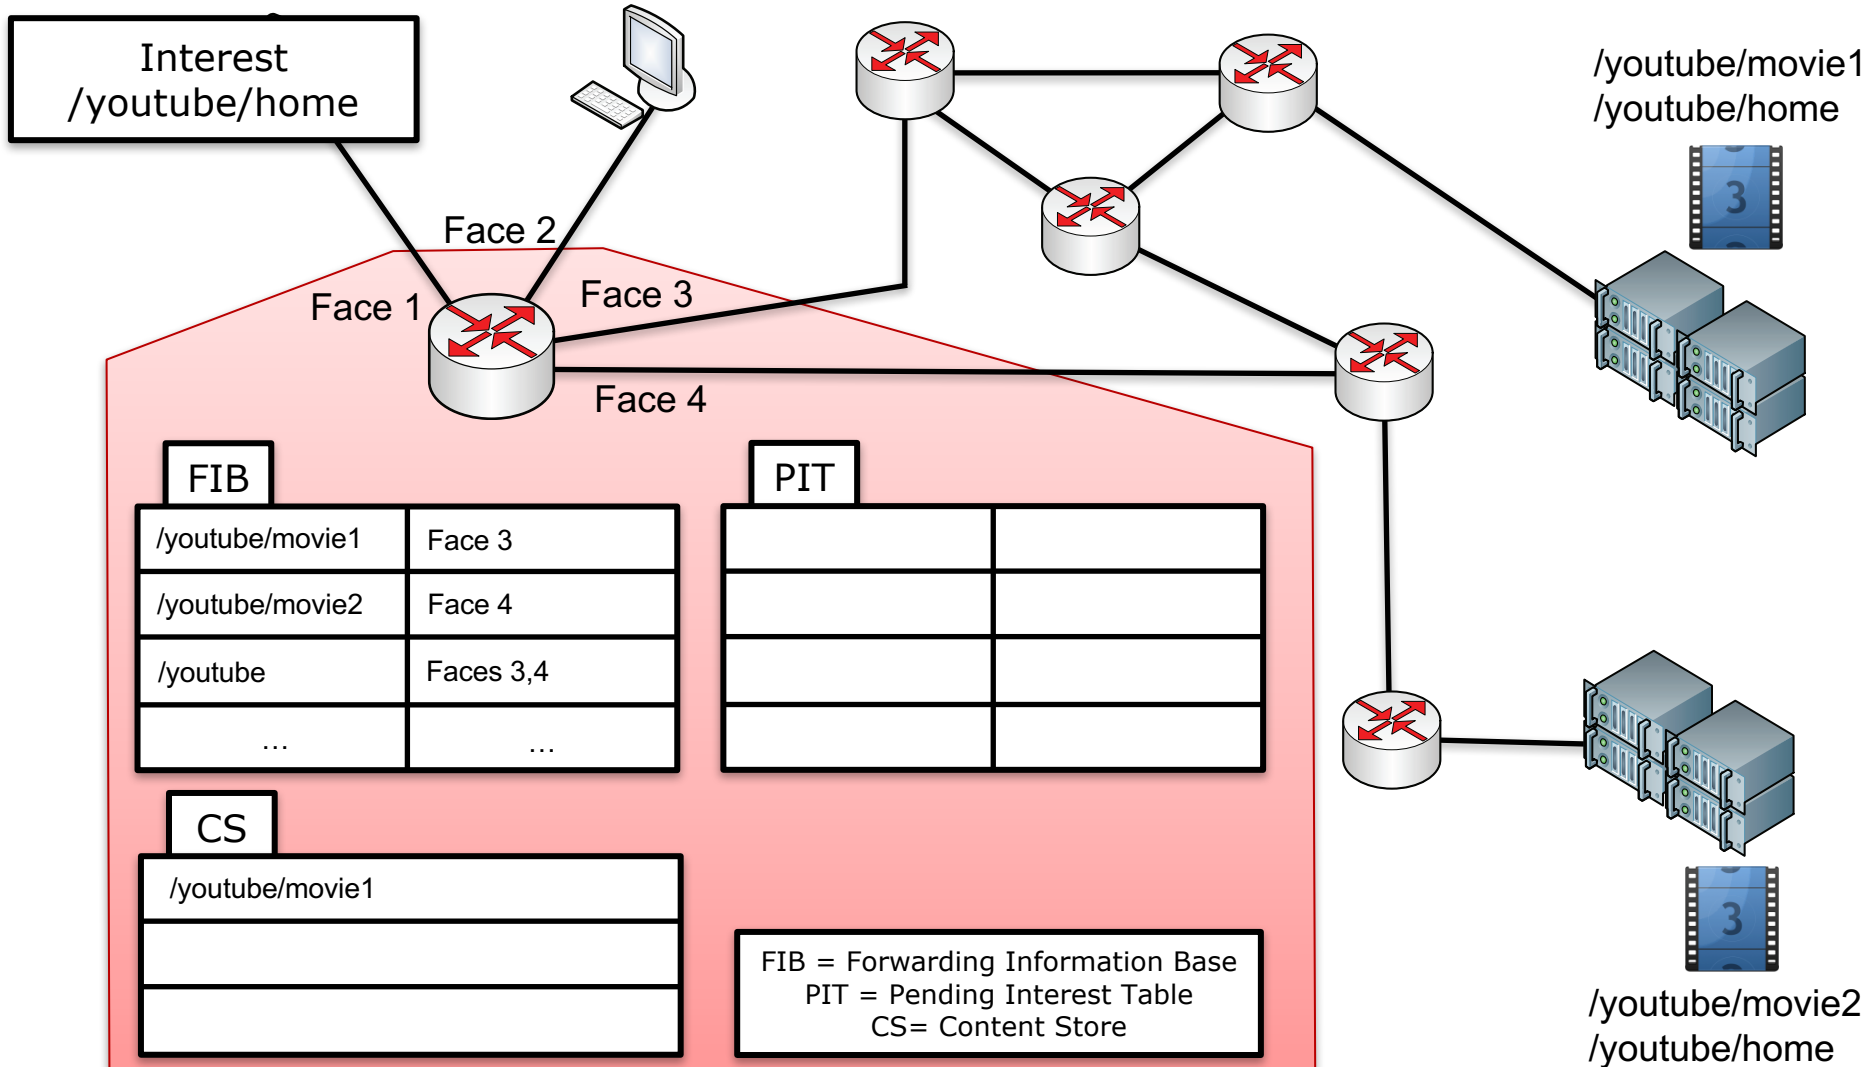
CCN: The Node Model



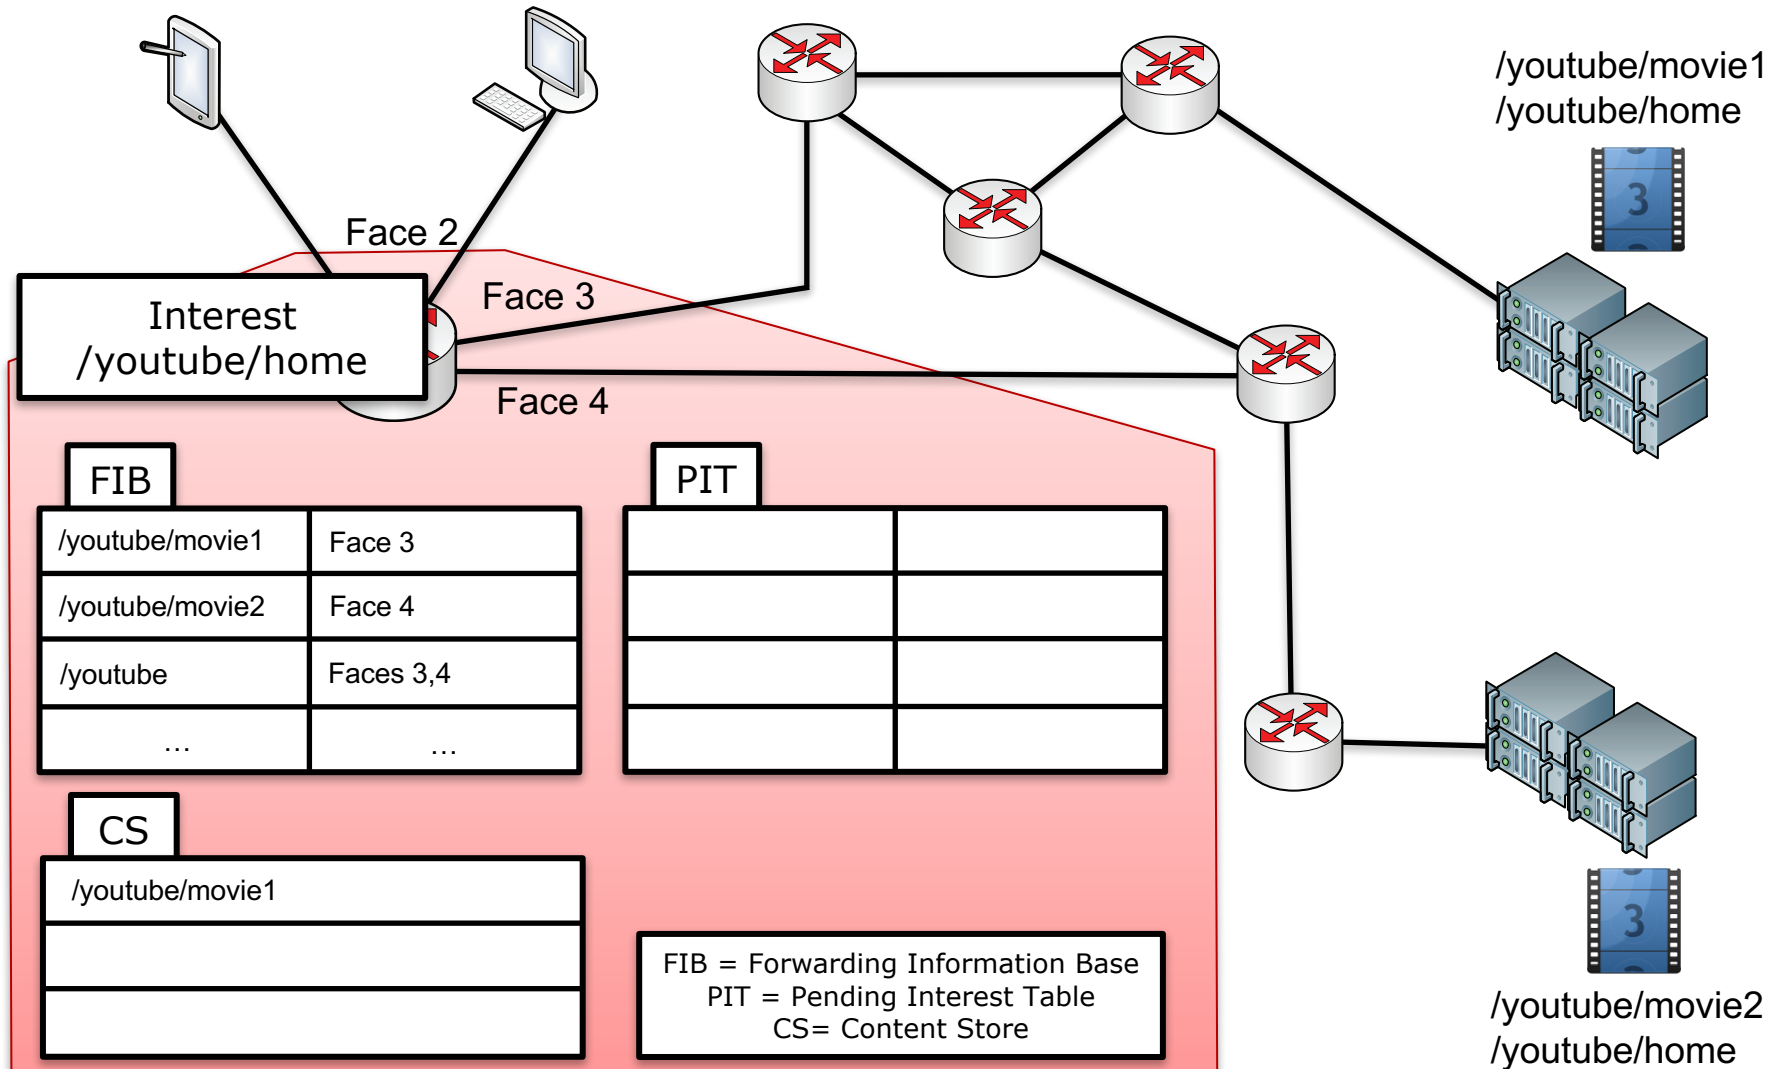
CCN: The Node Model



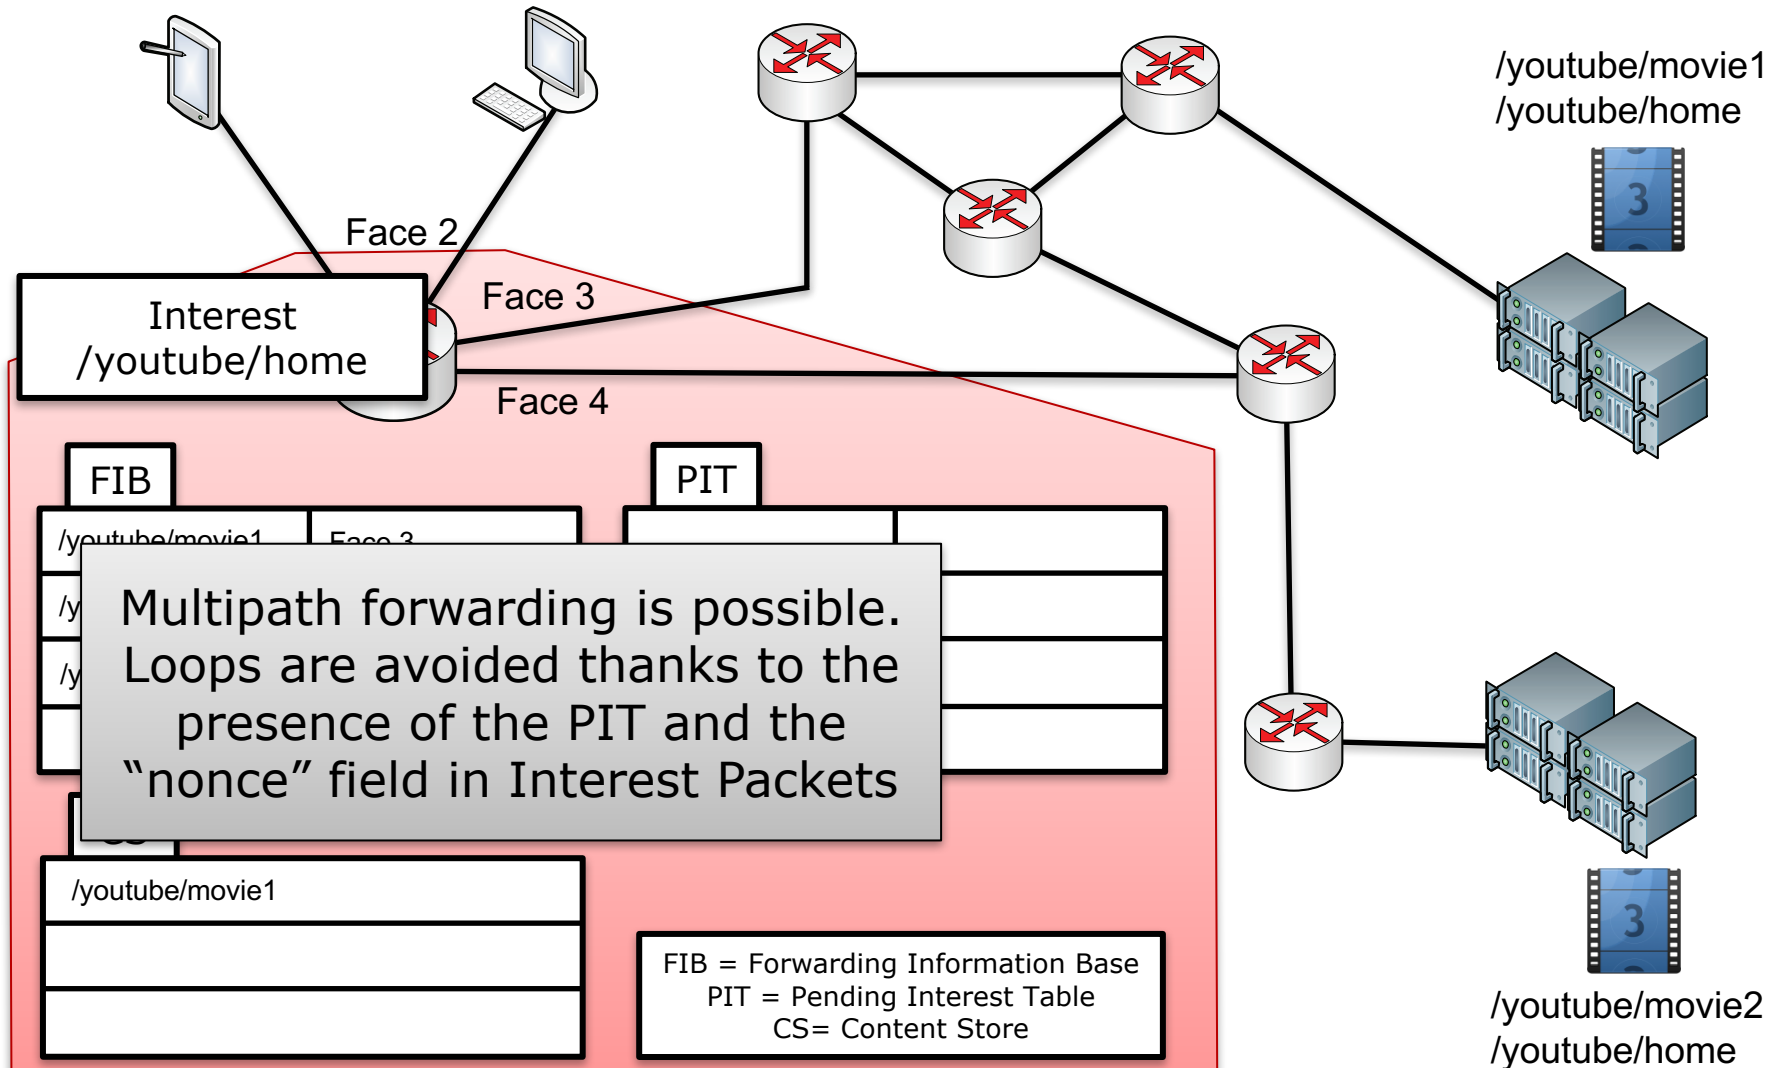
CCN: The Node Model



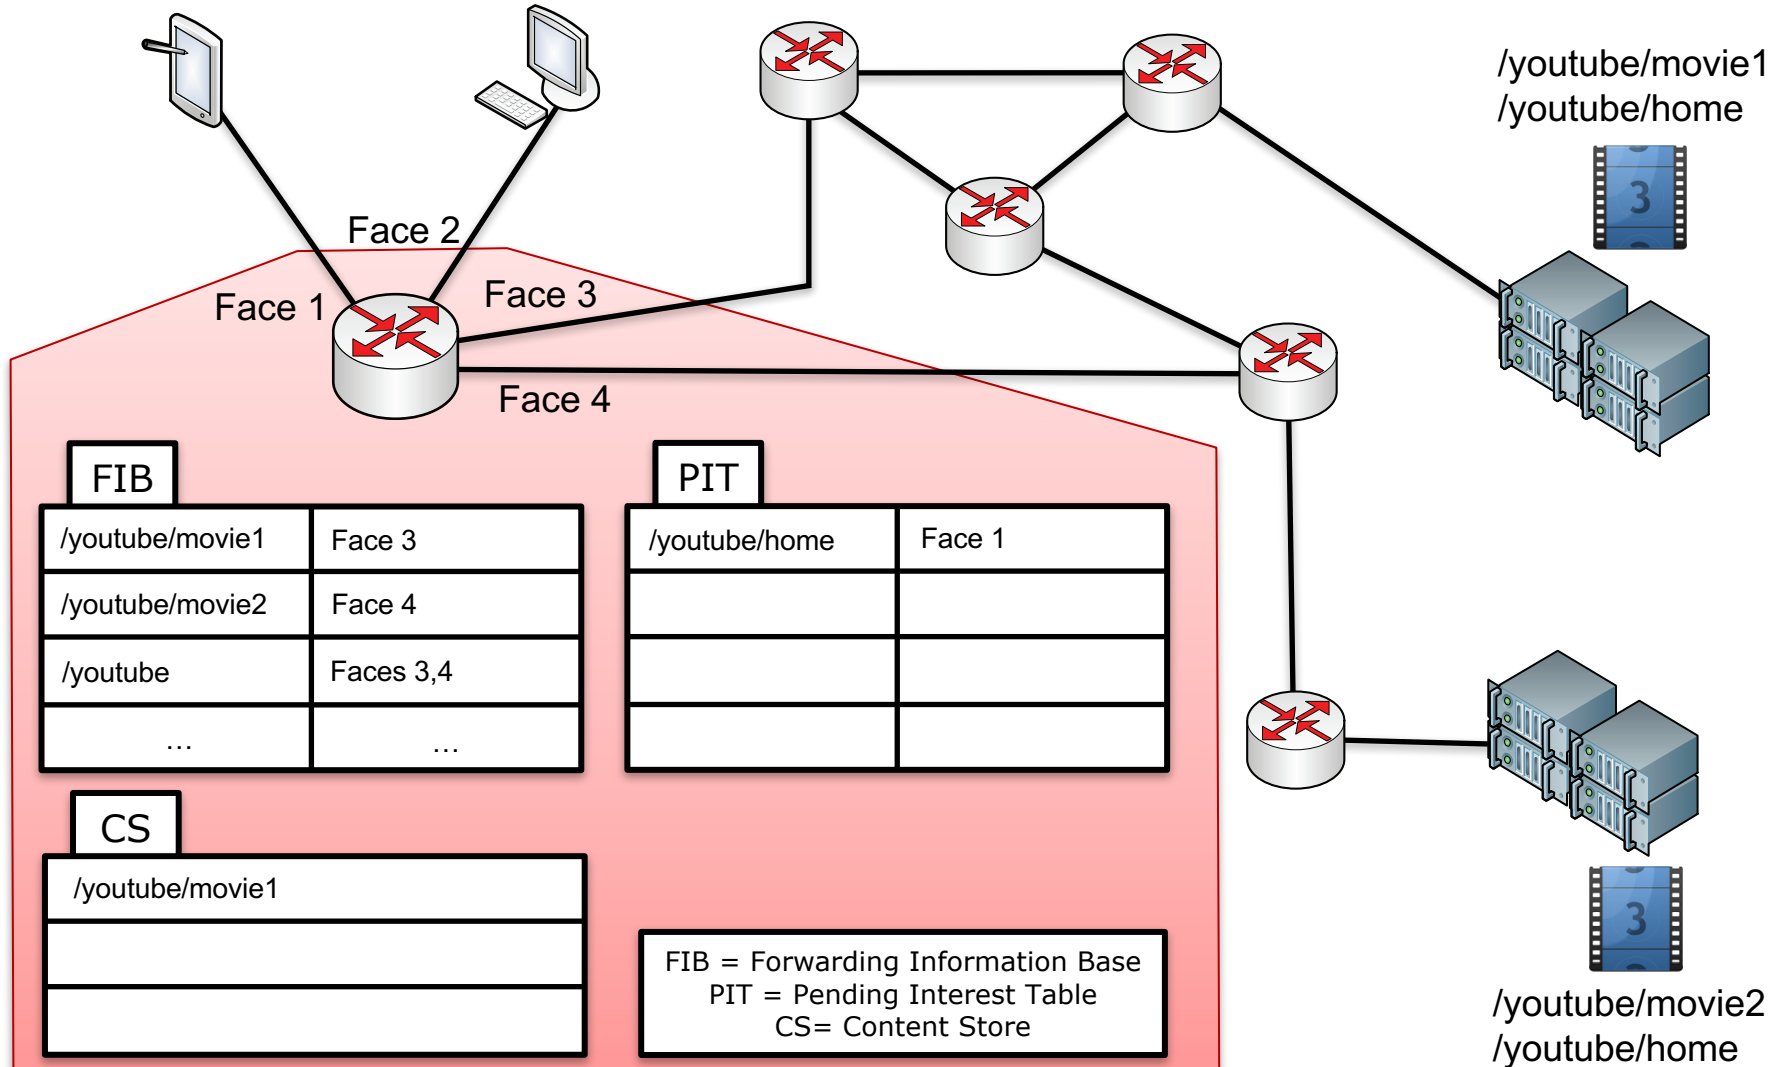
CCN: The Node Model



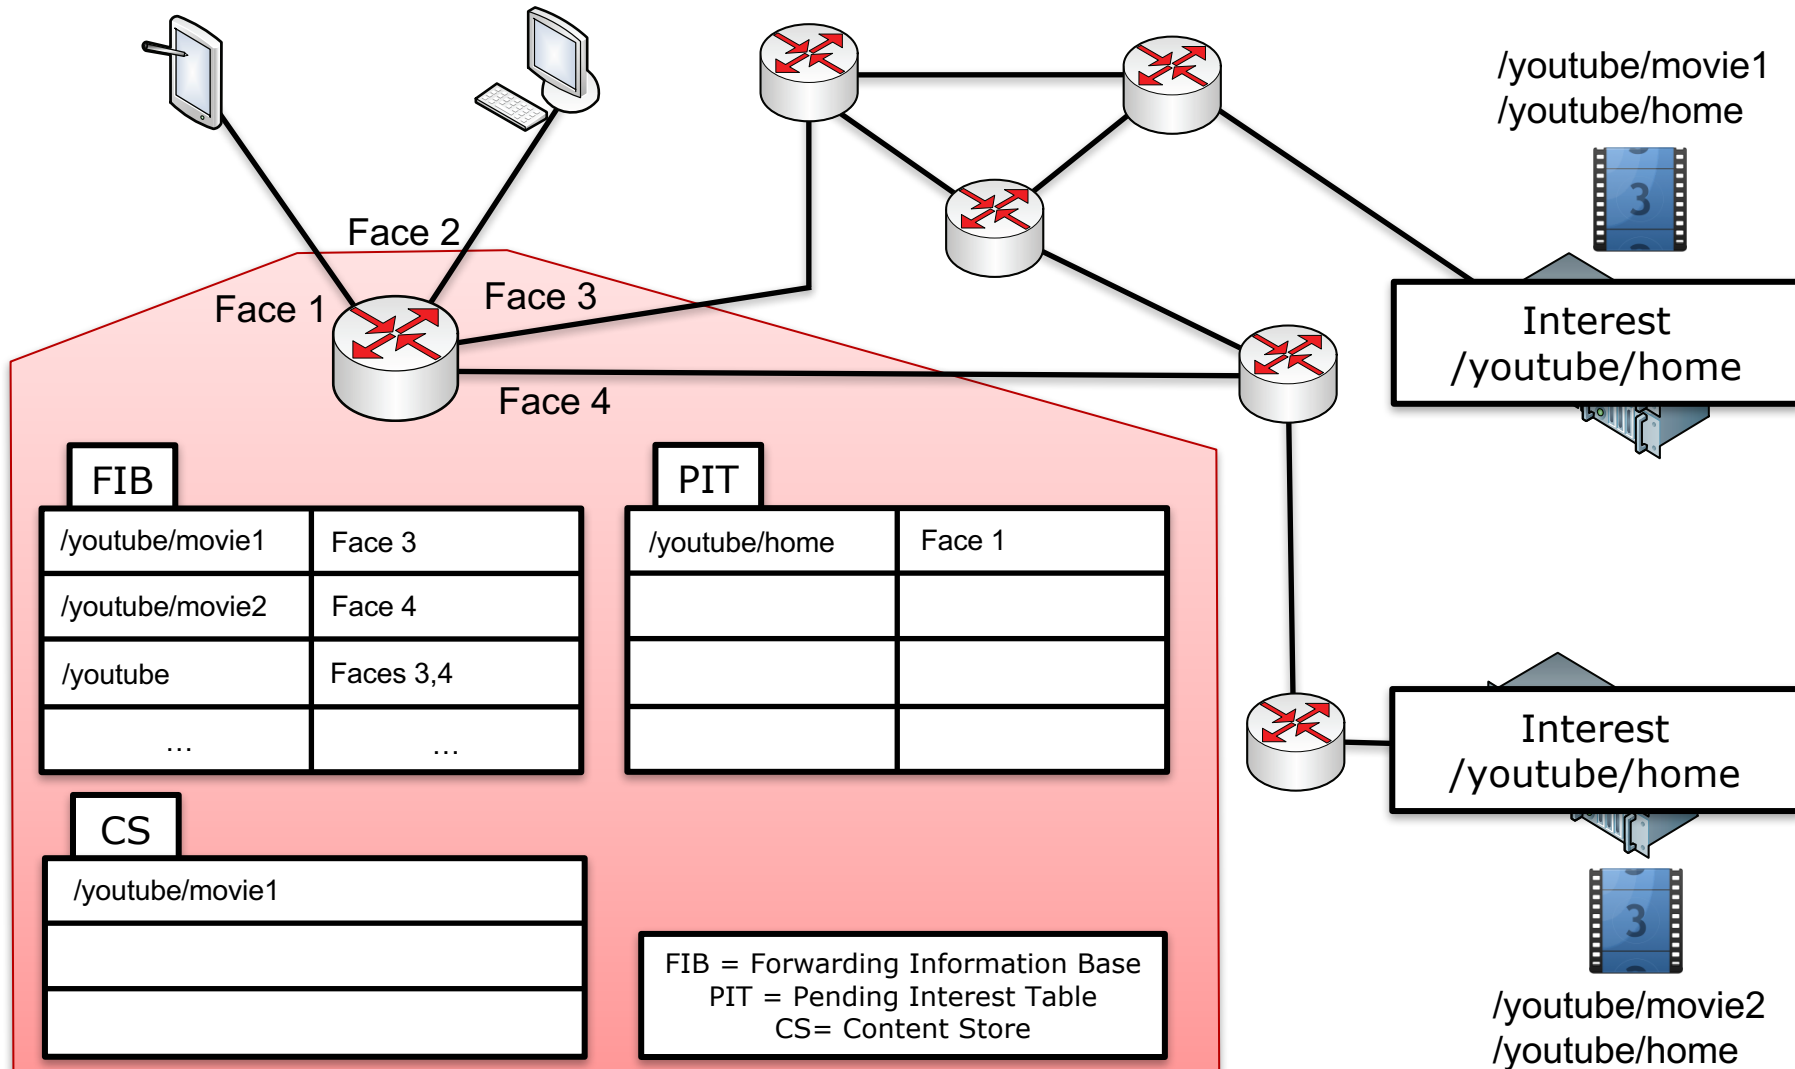
CCN: The Node Model



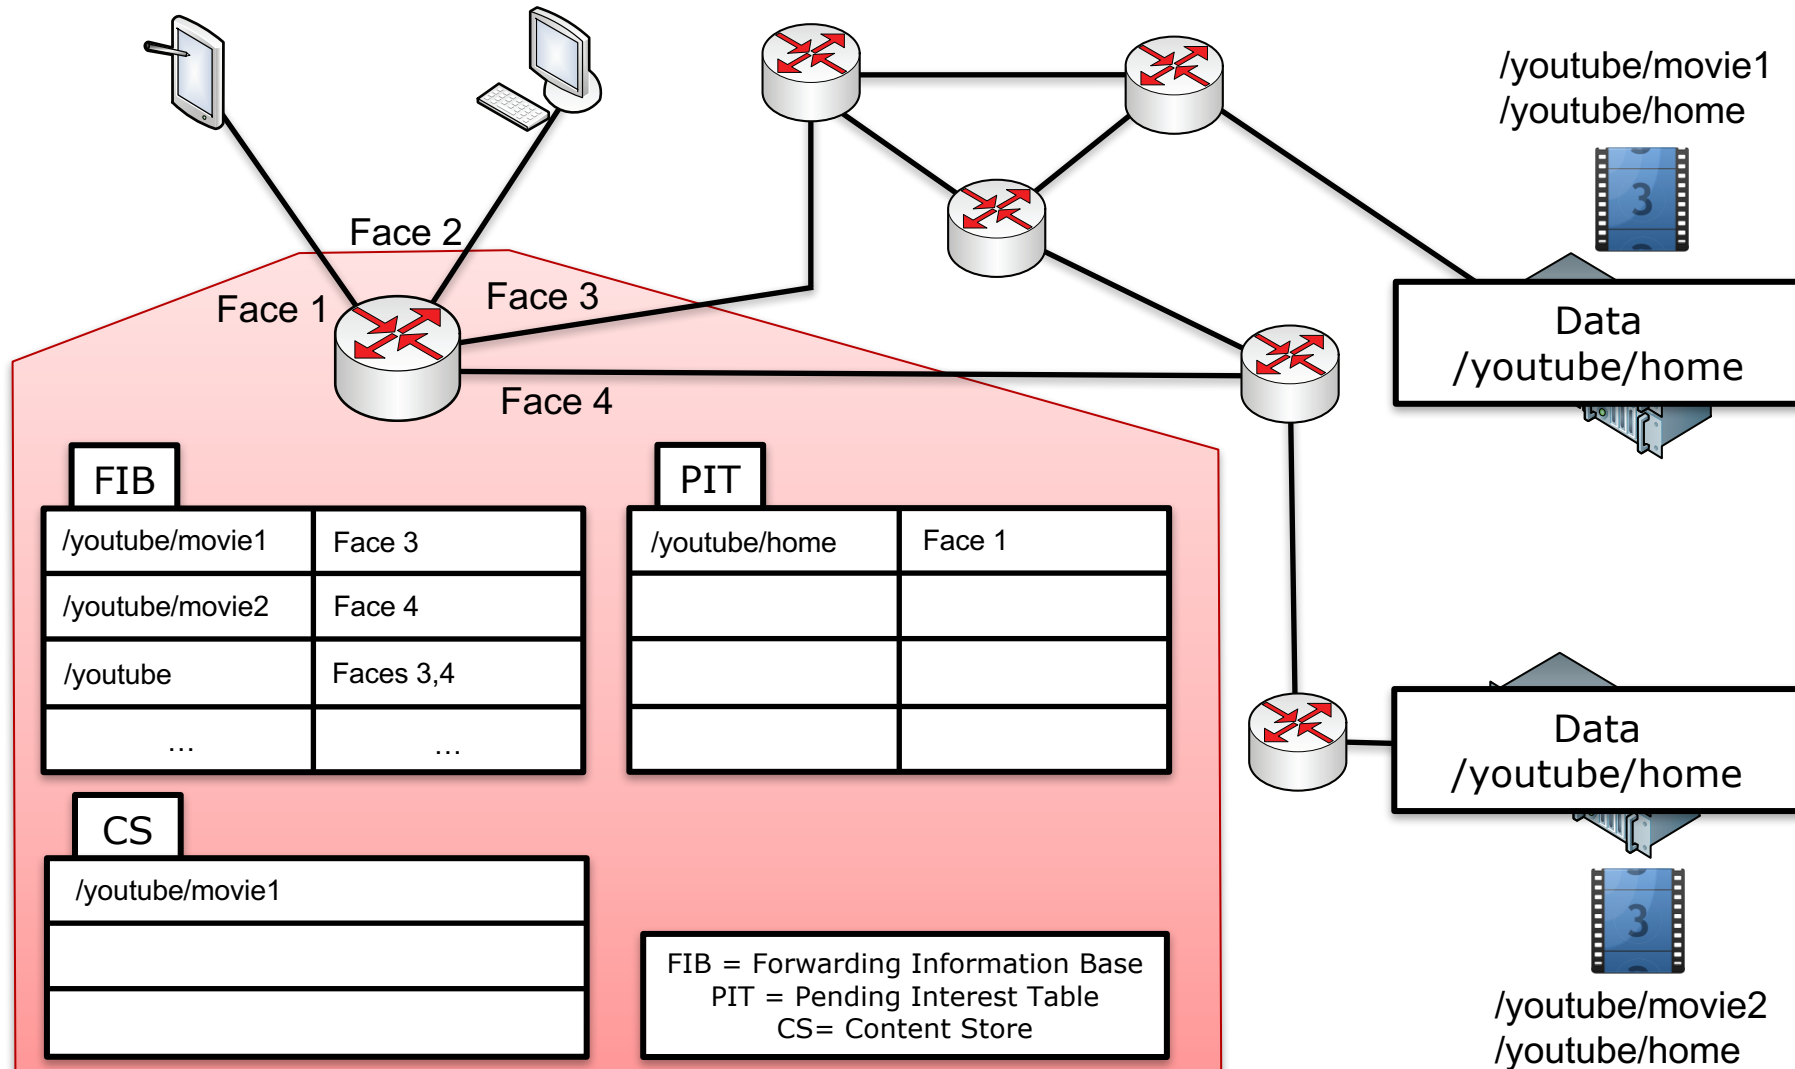
CCN: The Node Model



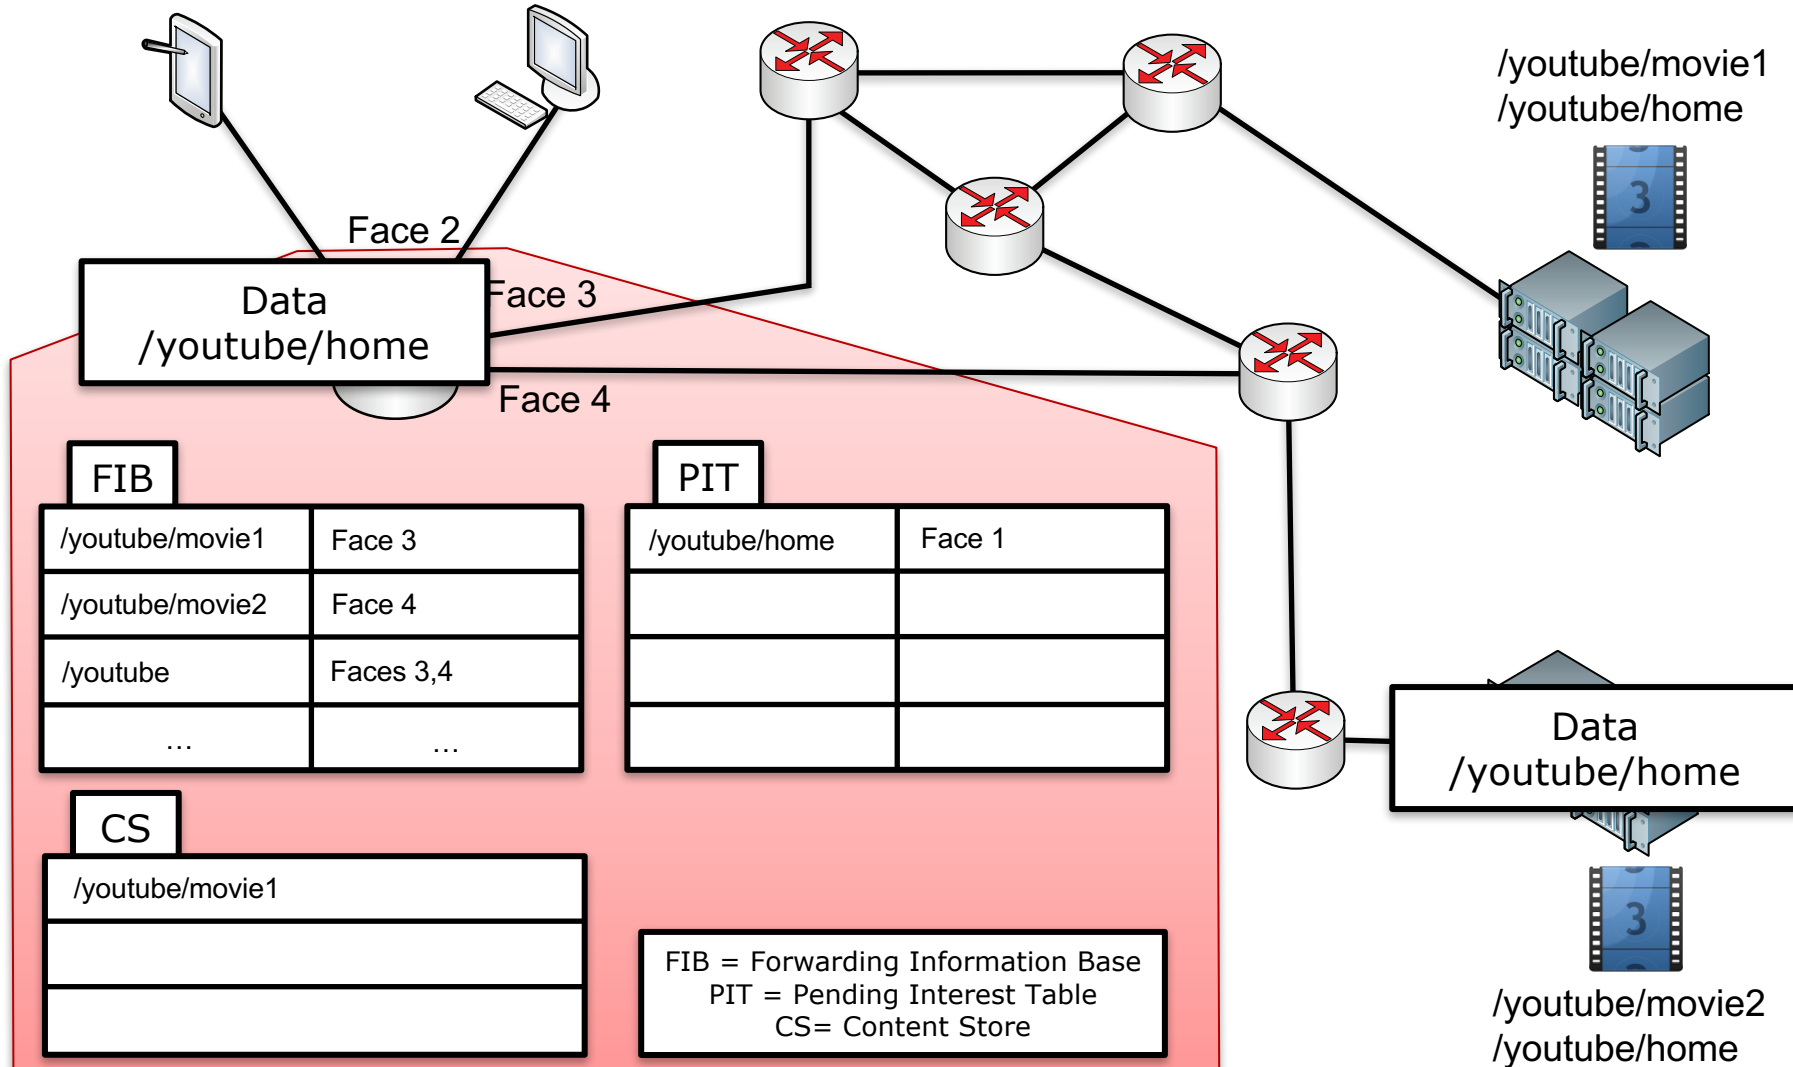
CCN: The Node Model



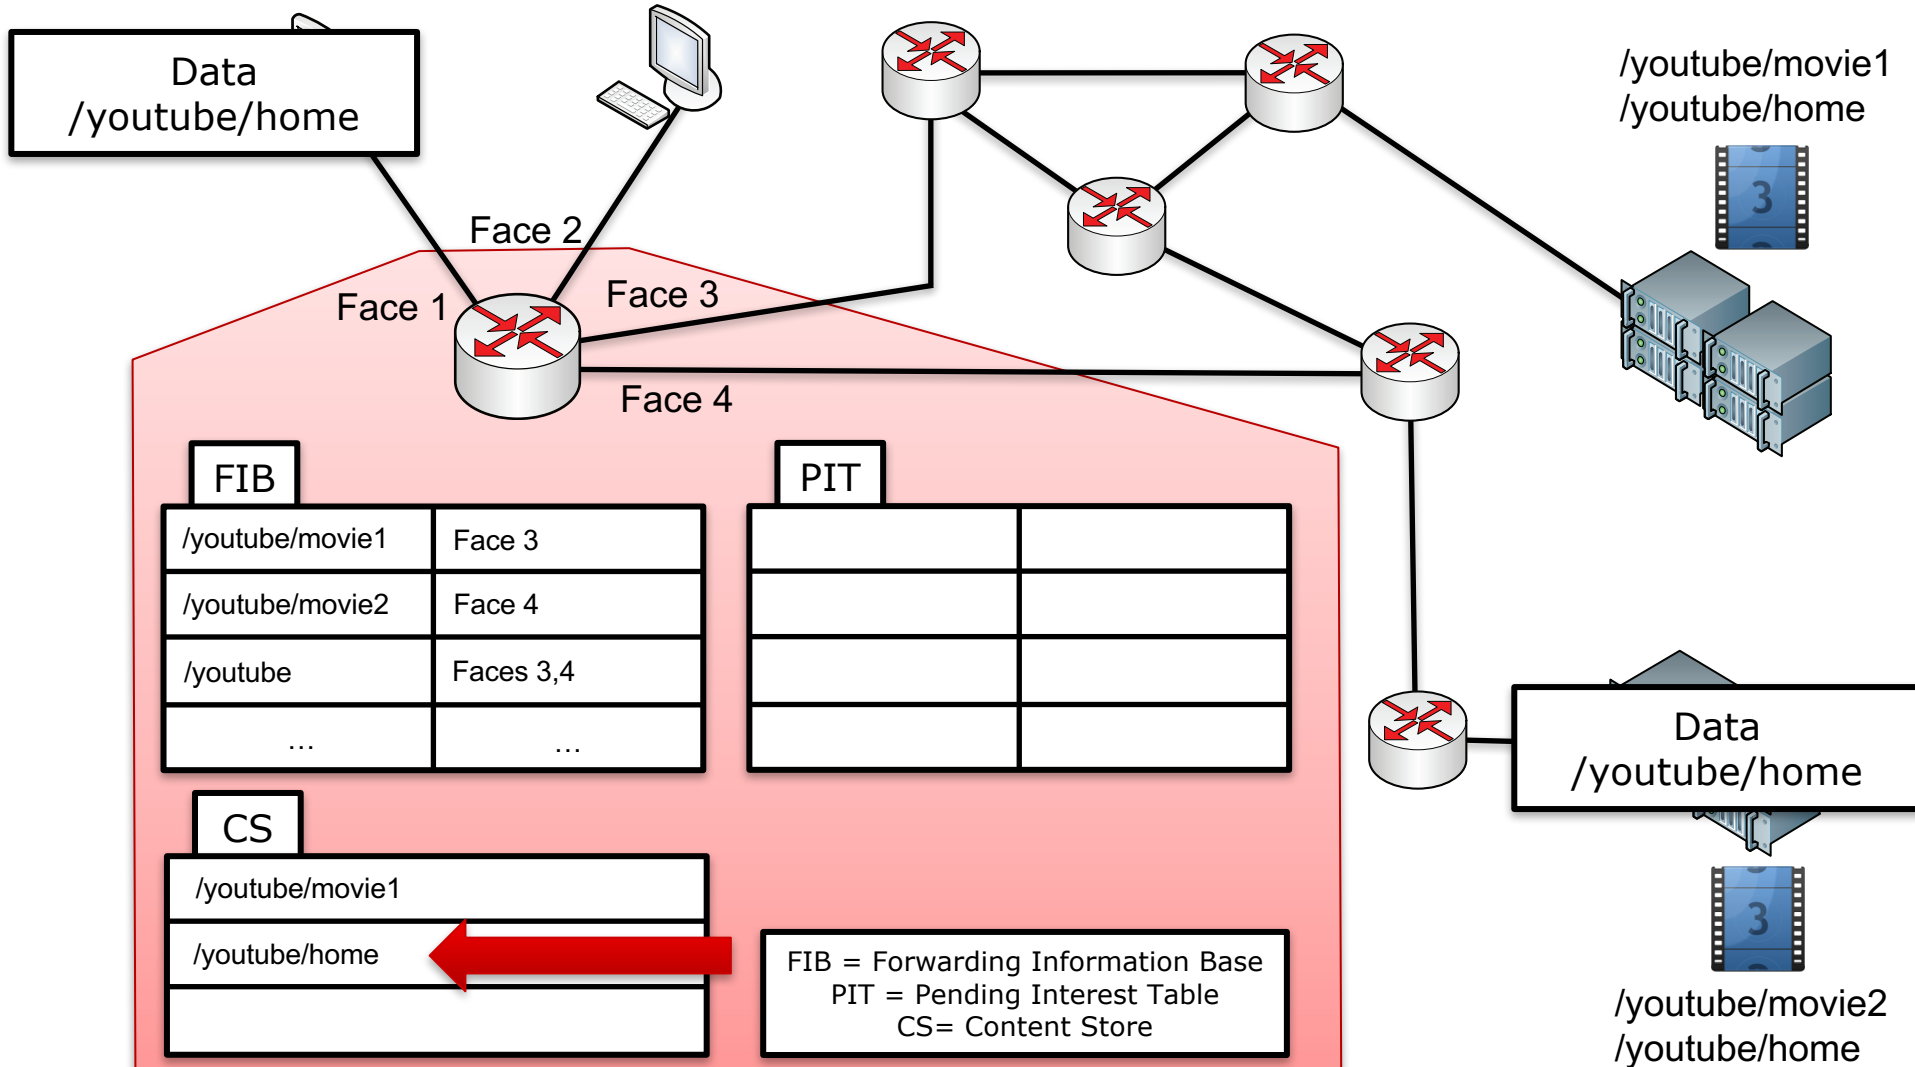
CCN: The Node Model



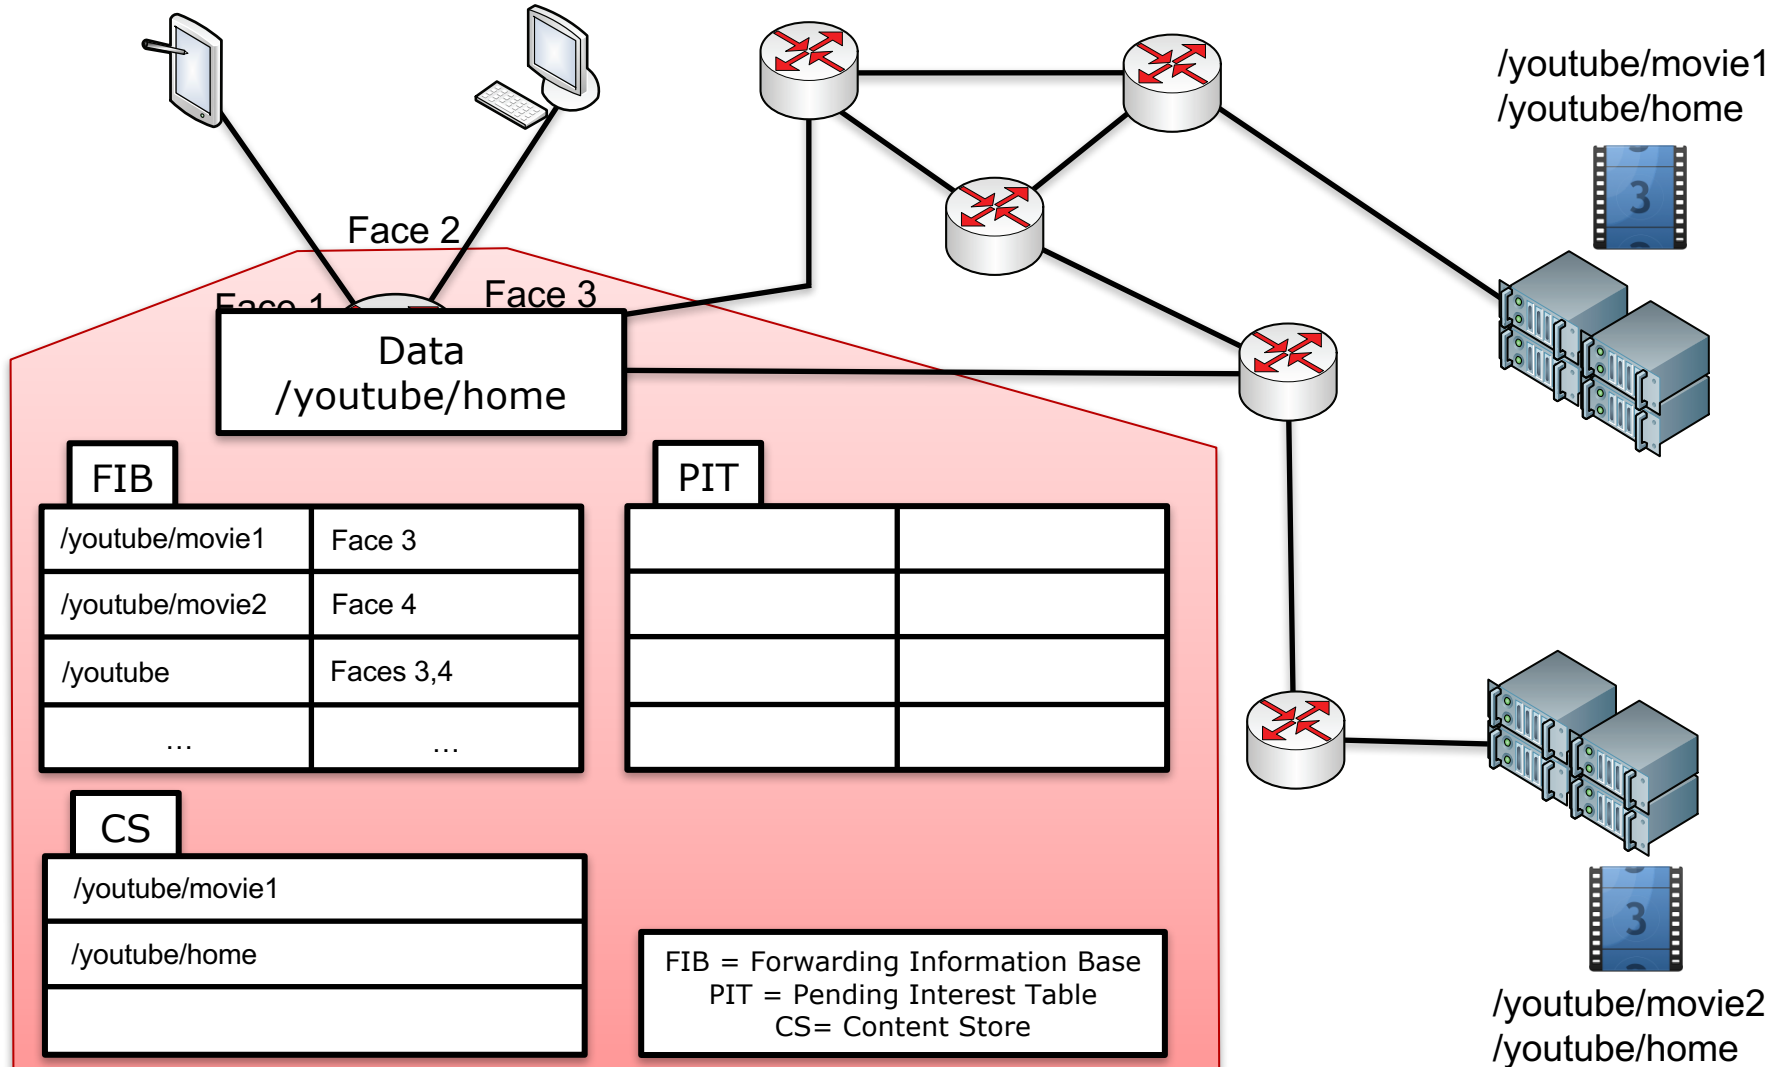
CCN: The Node Model



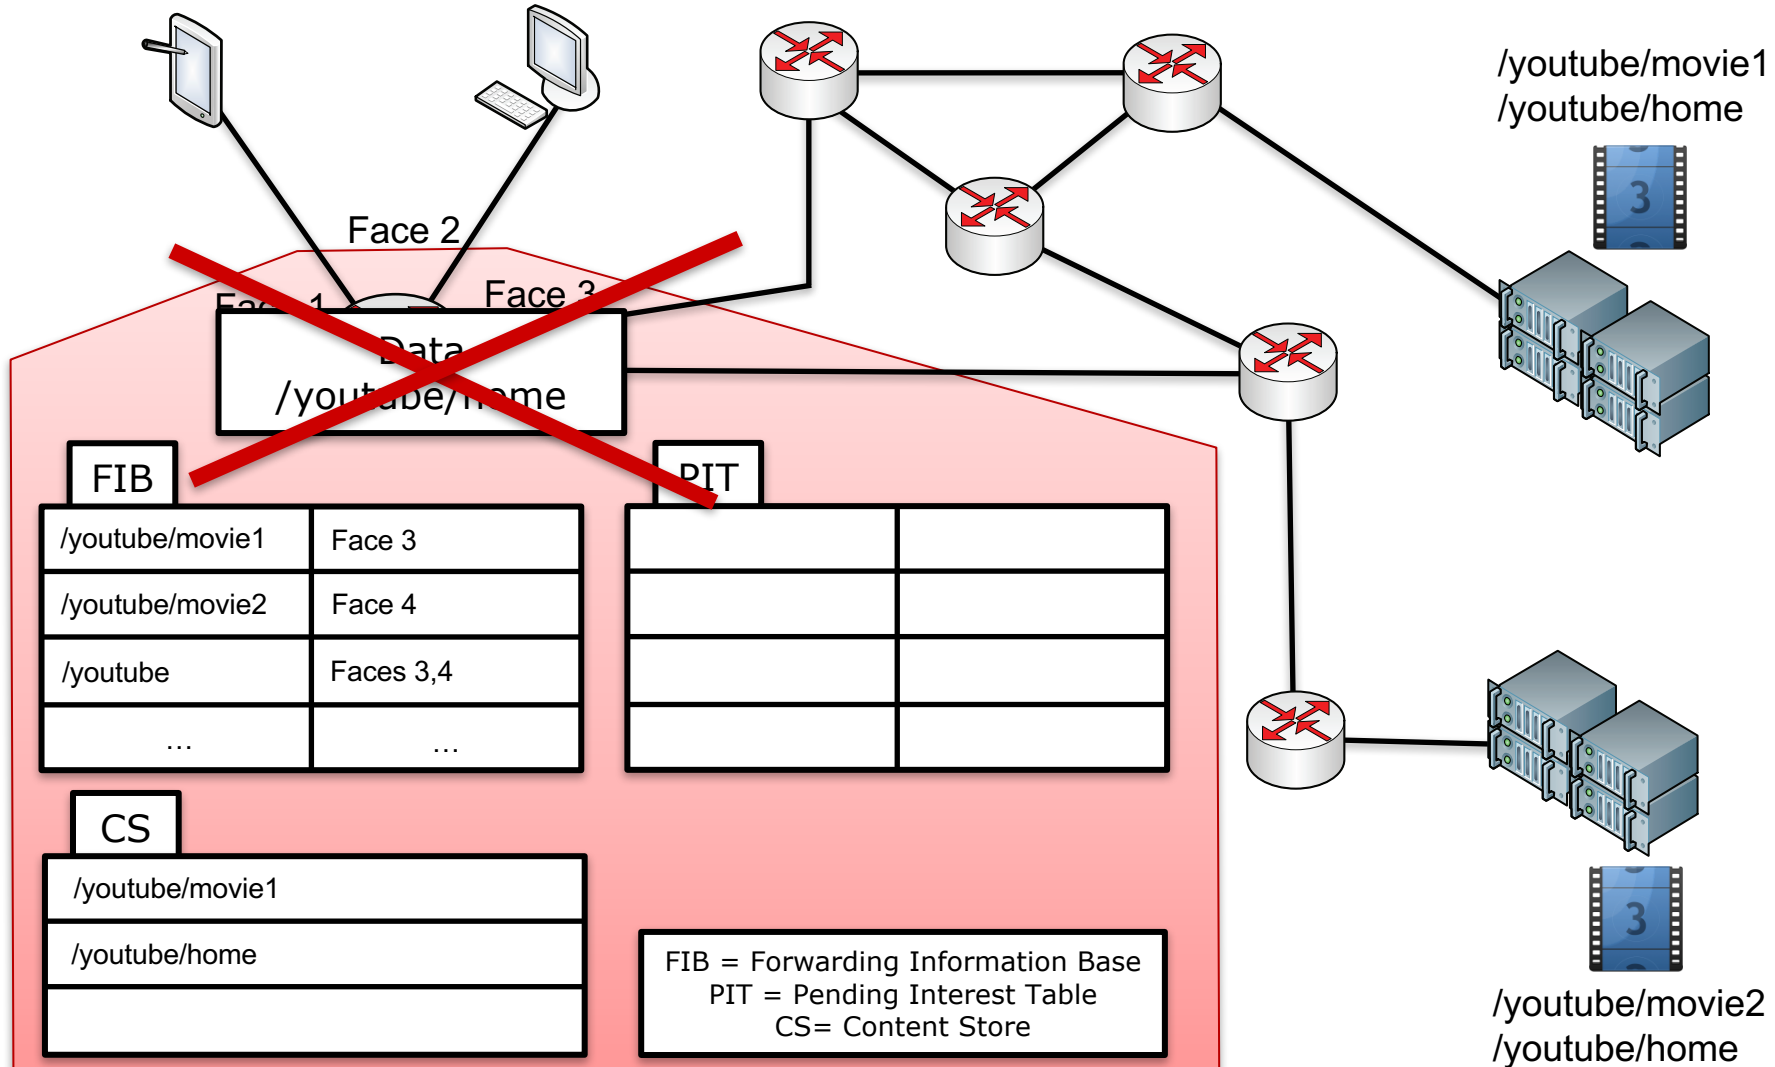
CCN: The Node Model



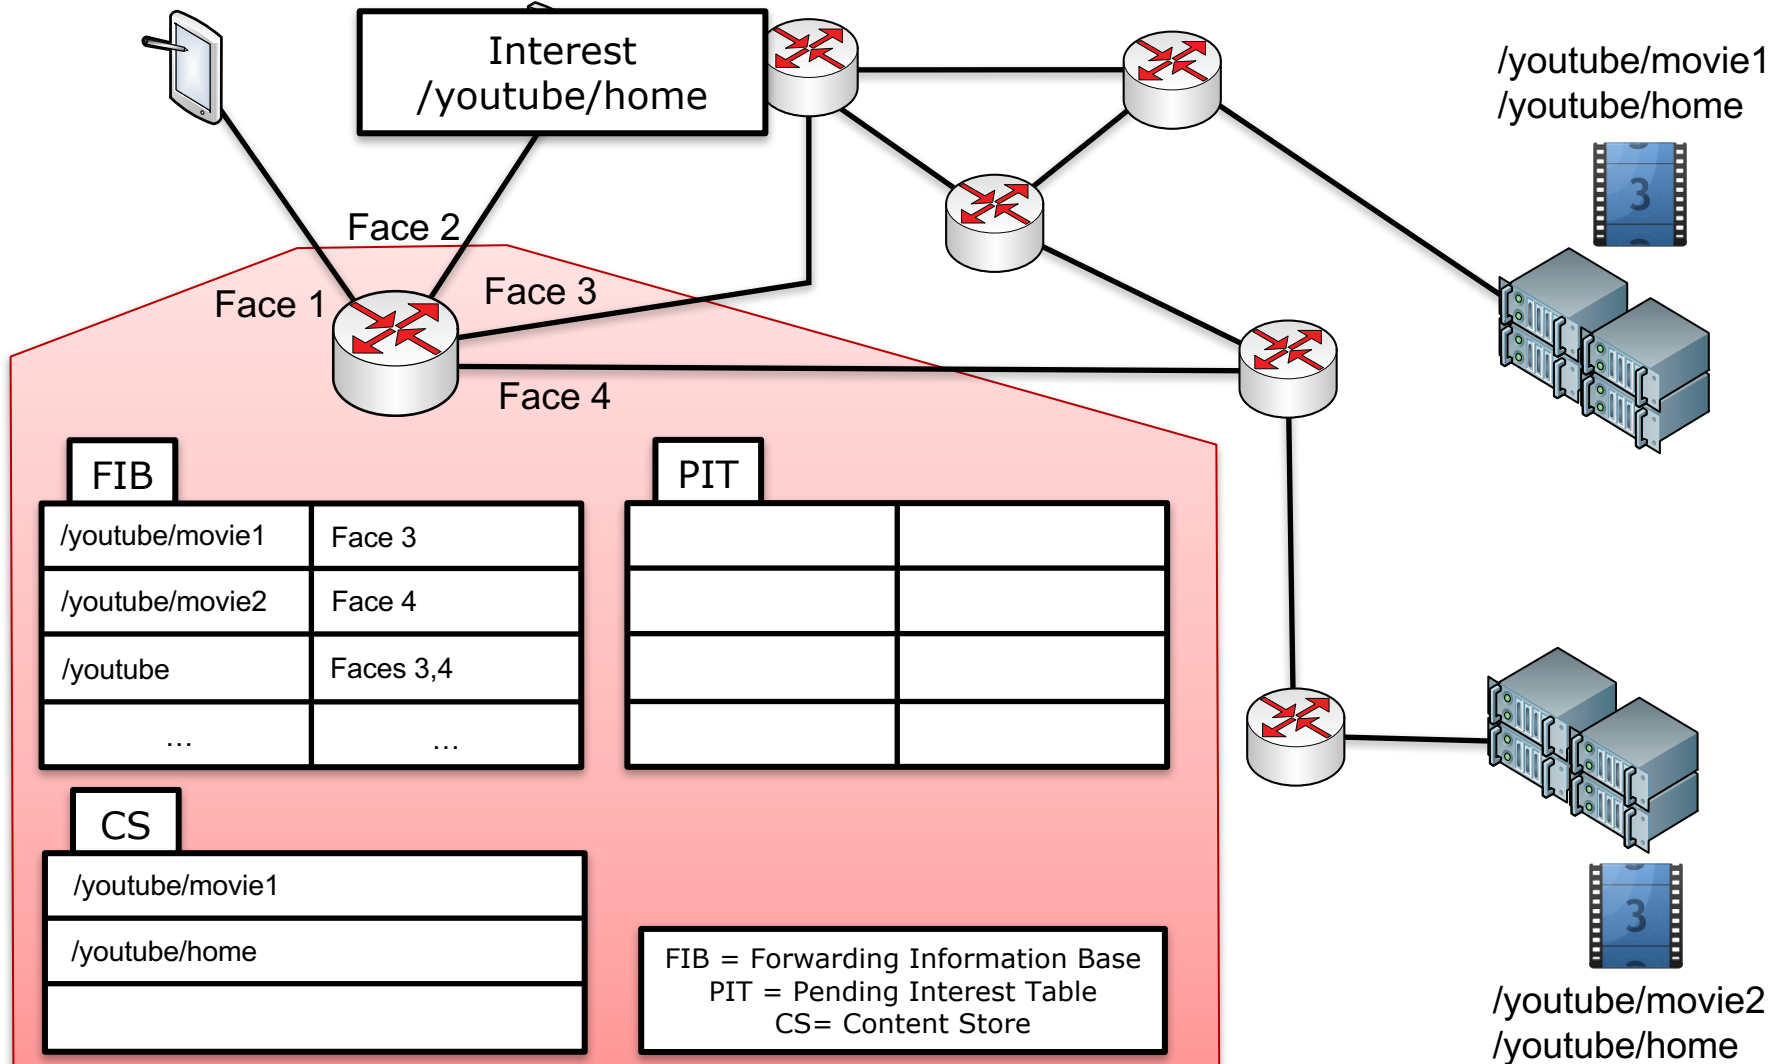
CCN: The Node Model



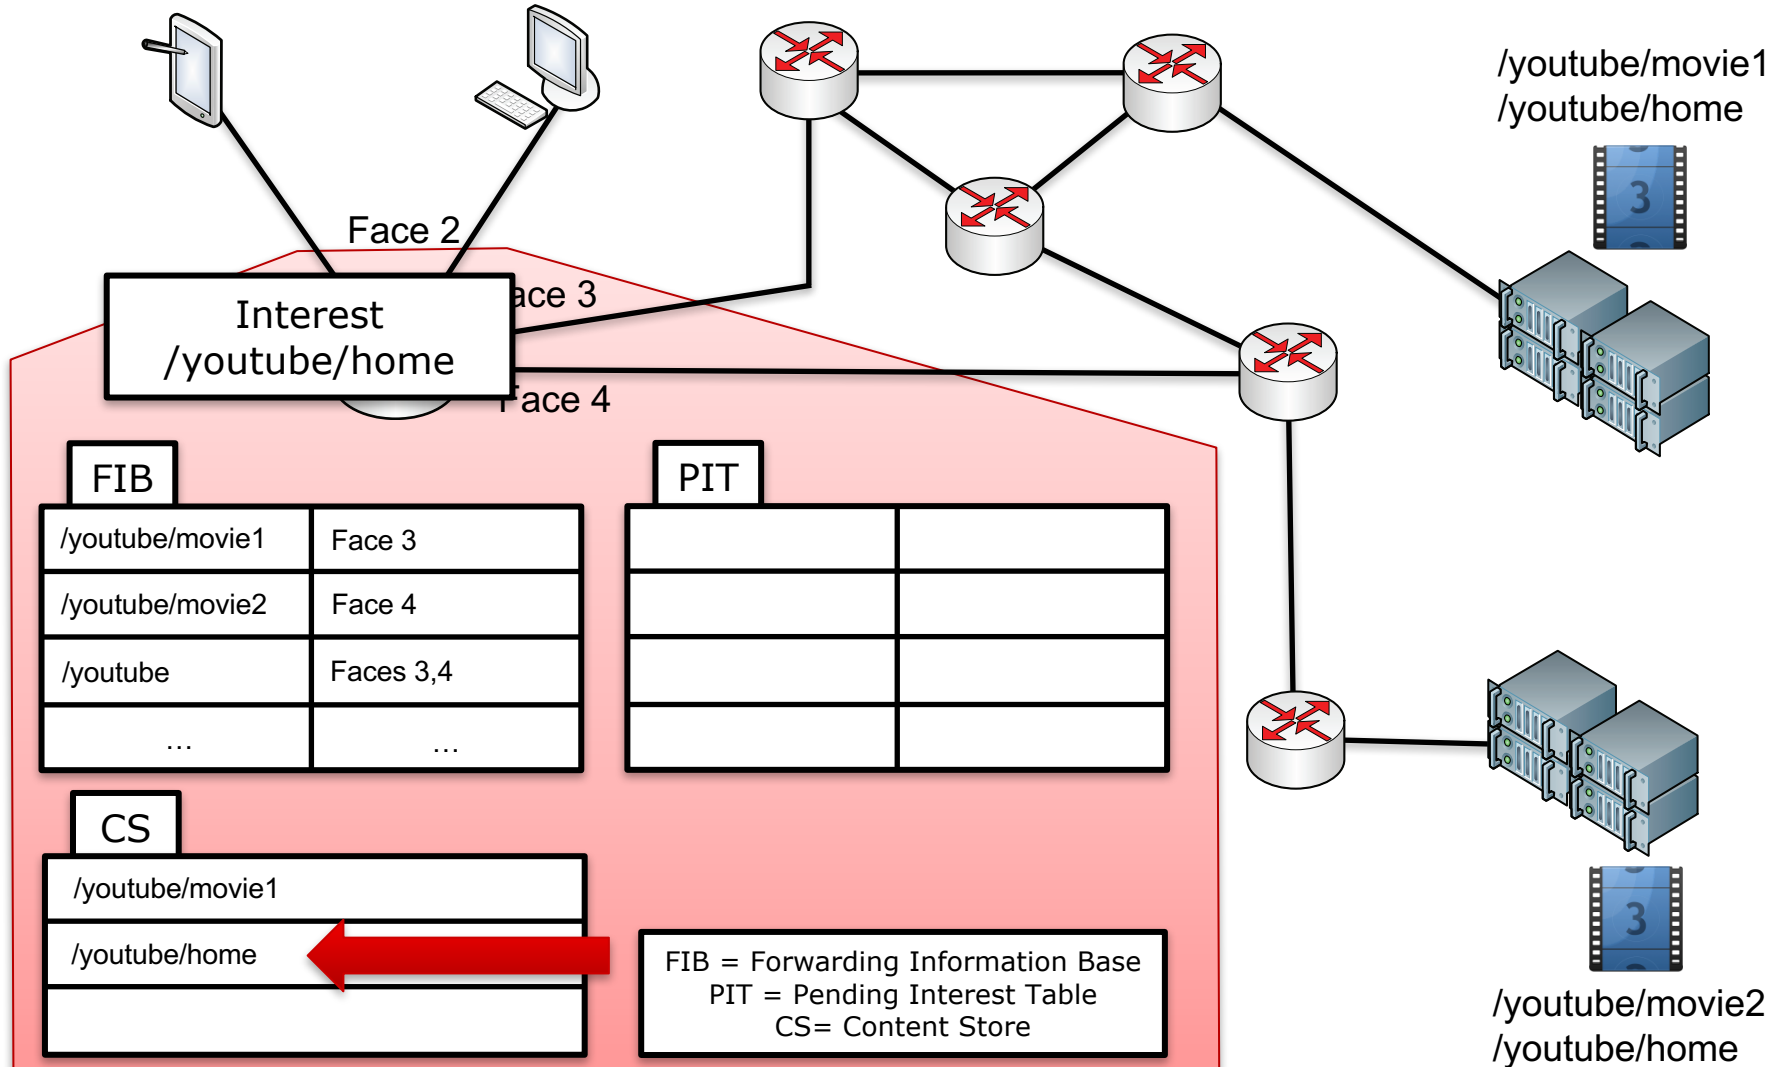
CCN: The Node Model



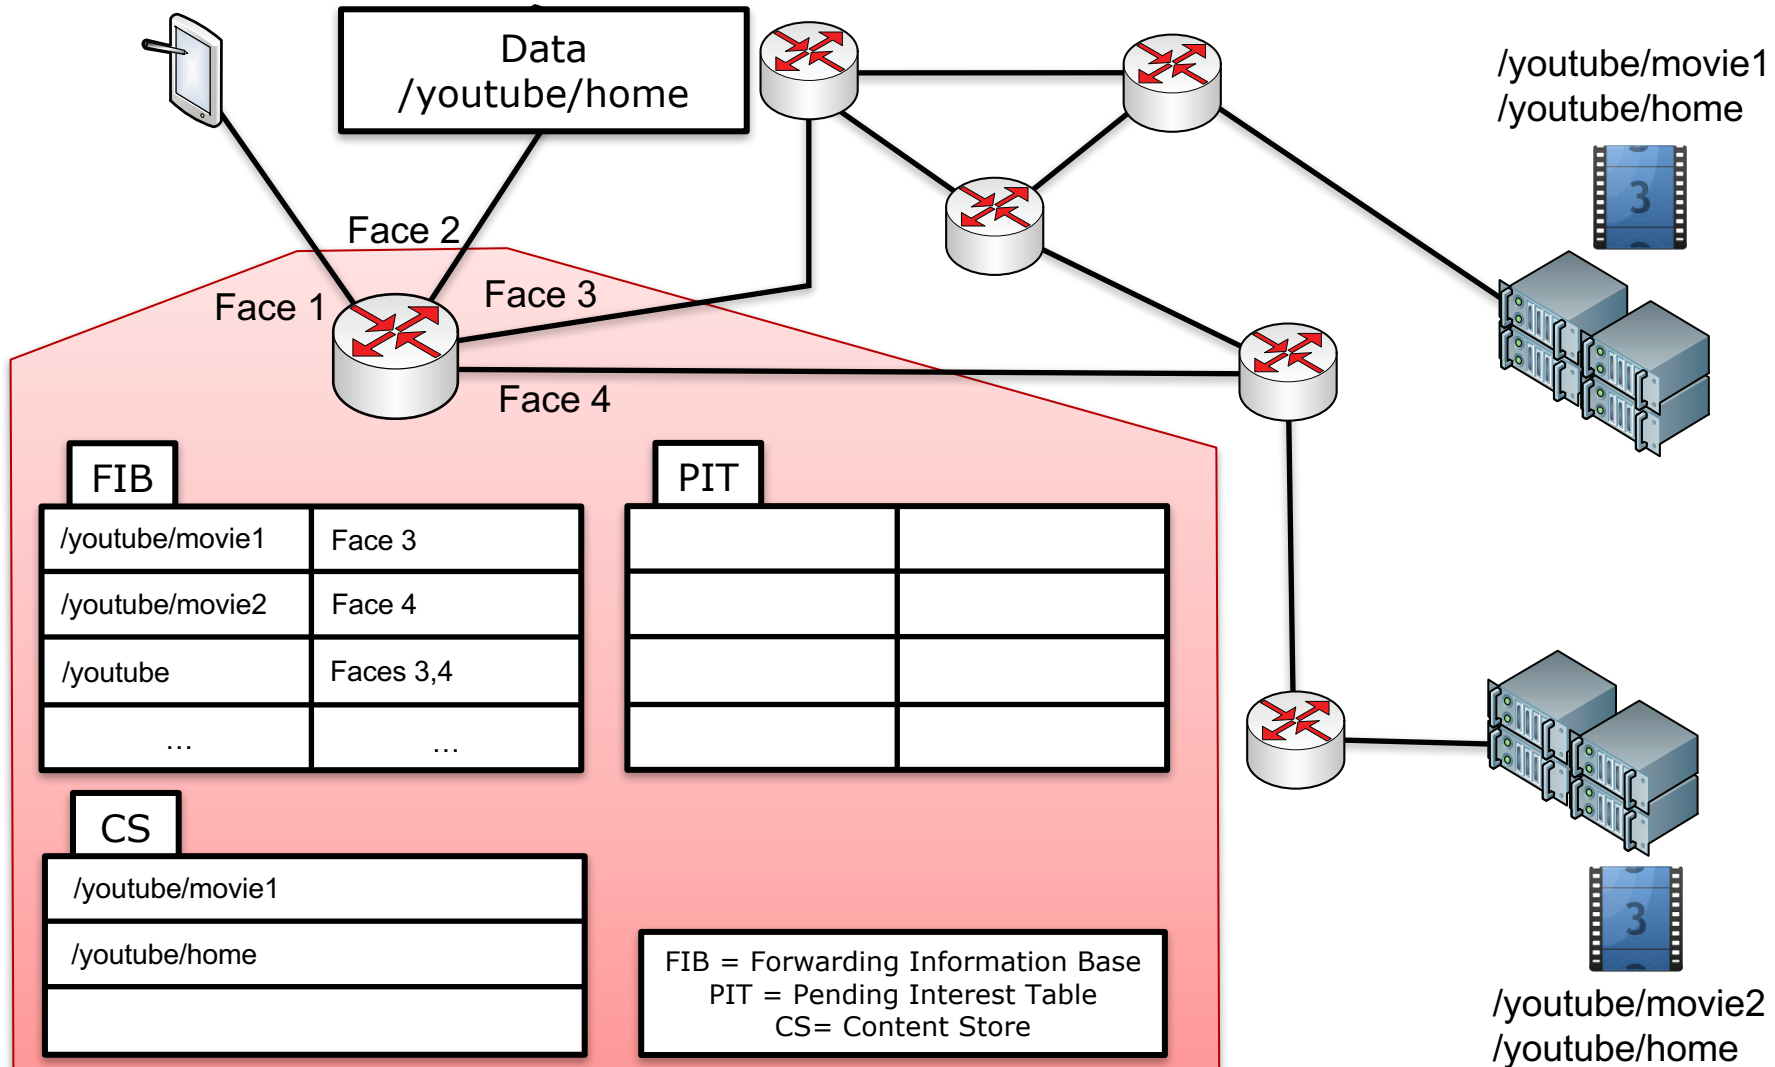
CCN: The Node Model



CCN: The Node Model



CCN: The Node Model



Security in NDN/CCN

Basic principles and techniques

Security in NDN/CCN

- Secure communications in NDN/CCN can be implemented by enforcing three properties to the *consumers*:
 - **Data Integrity**: Is a complete and uncorrupted copy of what the producer sent?
 - **Provenance**: Is the producer one the consumer is willing to trust to supply this content?
 - **Relevance**: Is the content an answer to the question the consumer asked?
- While in Today's TCP/IP Internet we instead secure the communication channel between the two endpoints

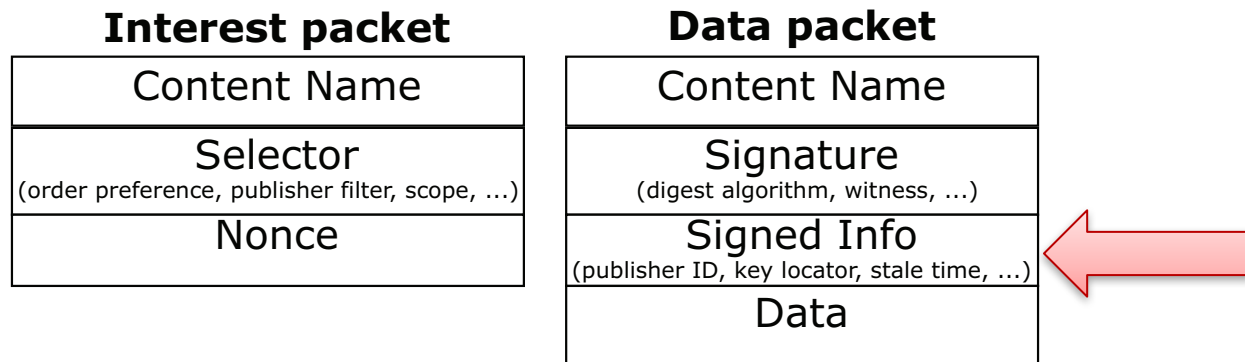
NDN: Security

- In NDN *each packet* is signed by the legitimate producer
 - The key-locator is specified in the packet

- Each node can serve the content, but the client can verify:
 - **Validity/Data integrity**: the content integrity has not been altered
 - **Provenance**: the content has been published by the legitimate producer
 - **Relevance**: the content satisfies the request originally sent by the consumer

Security in NDN/CCN

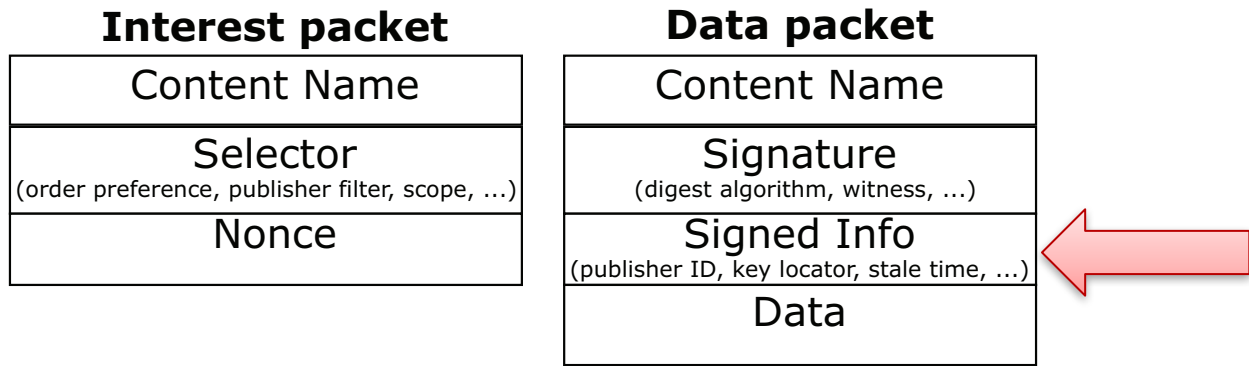
- Only Data packets are signed
- The signature creates a link between:
 - The content Name = C
 - The data = D
 - The private key PK of the producer
 - $\langle C, D, \text{Sign}_{\text{PK}}(C,D) \rangle$



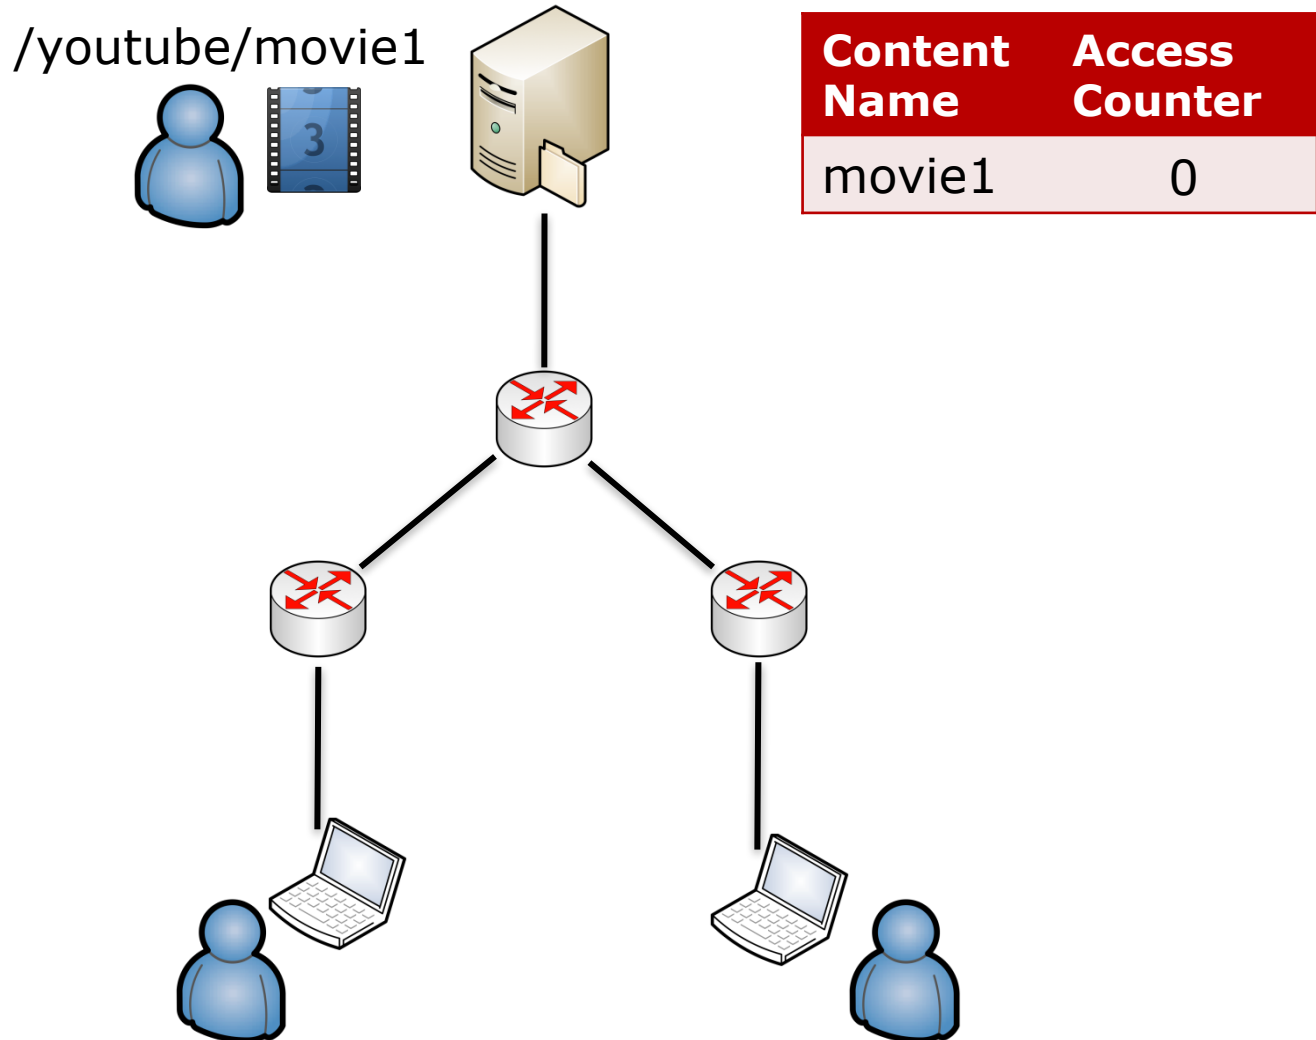
Security in NDN/CCN

- Only Data packets are signed
- The signature creates a link between:
 - The content Name = C

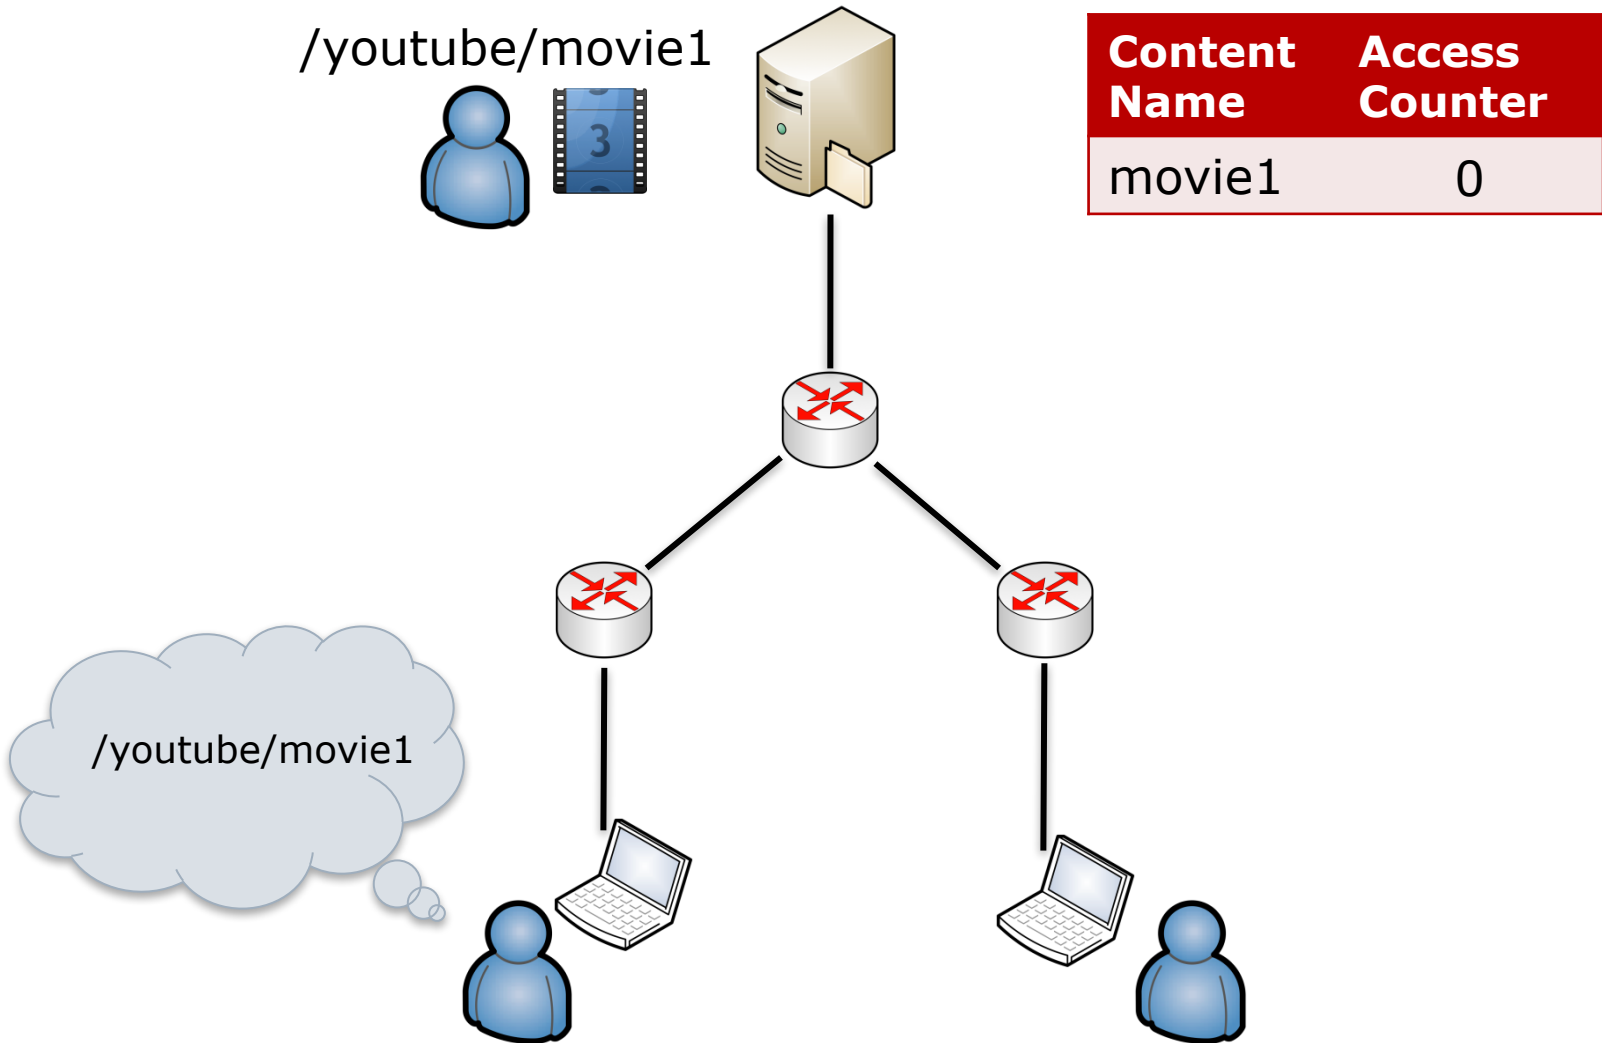
■ The location of the public key can be provided in the Key-Locator field



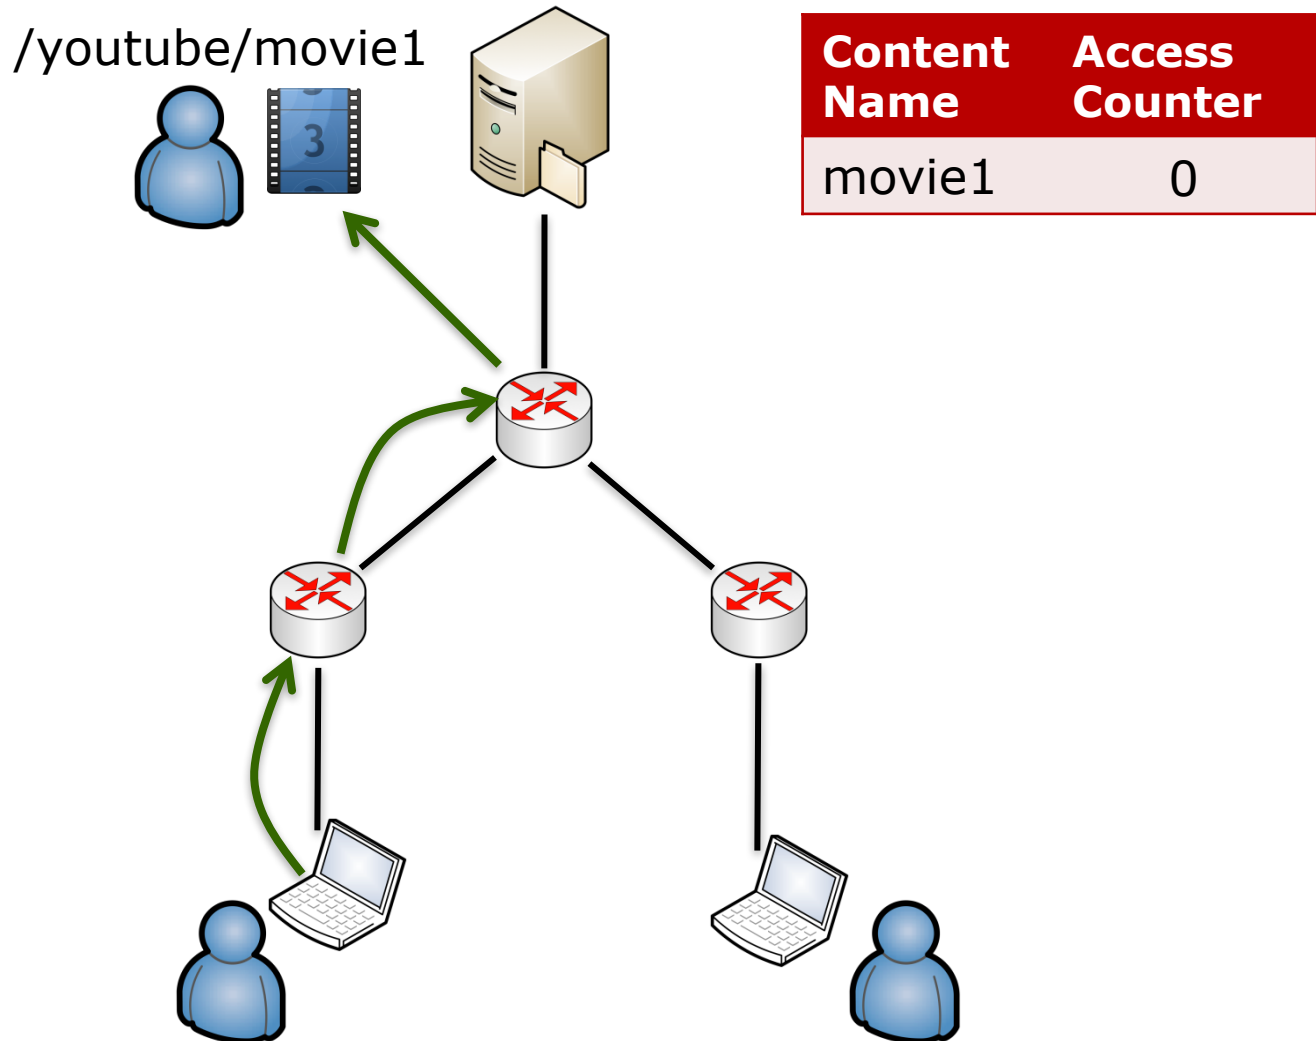
Security Issues in ICN



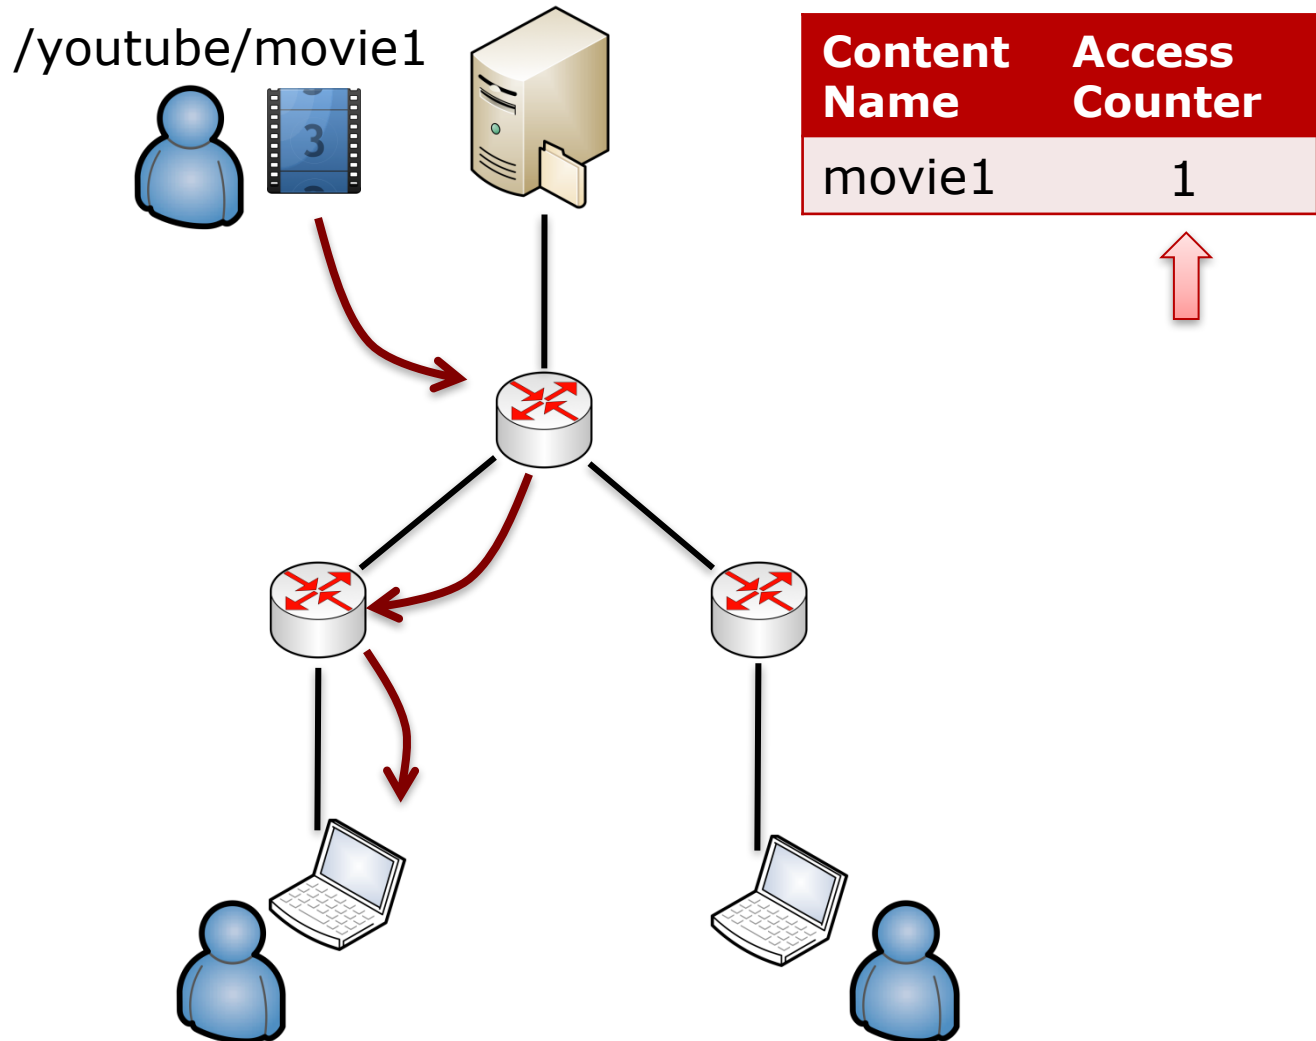
Security Issues in ICN



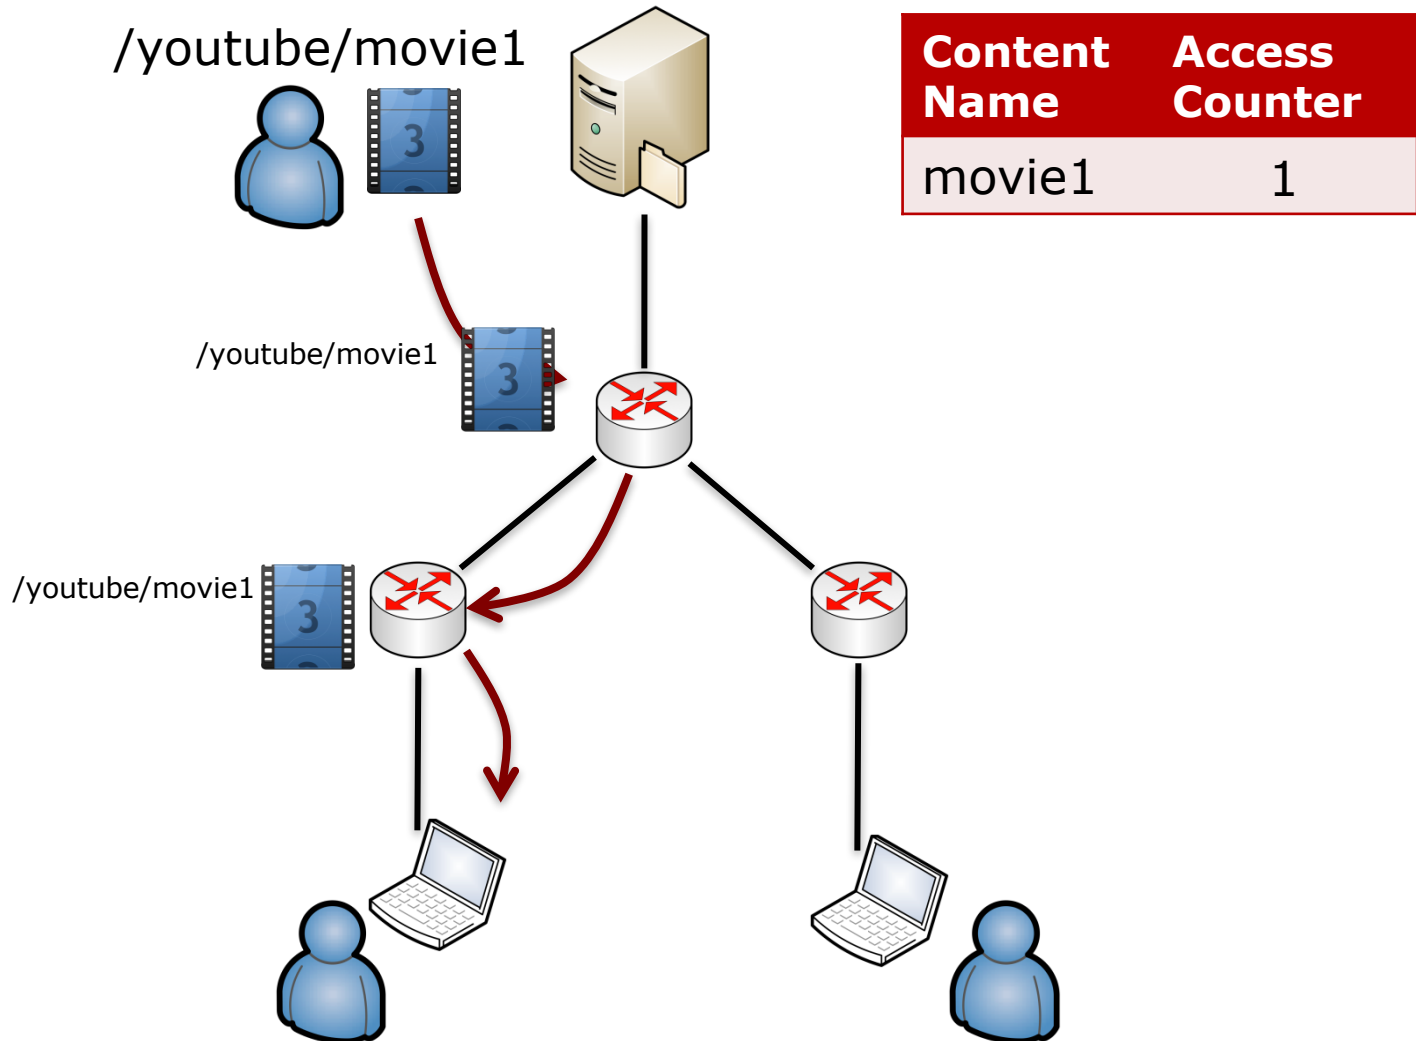
Security Issues in ICN



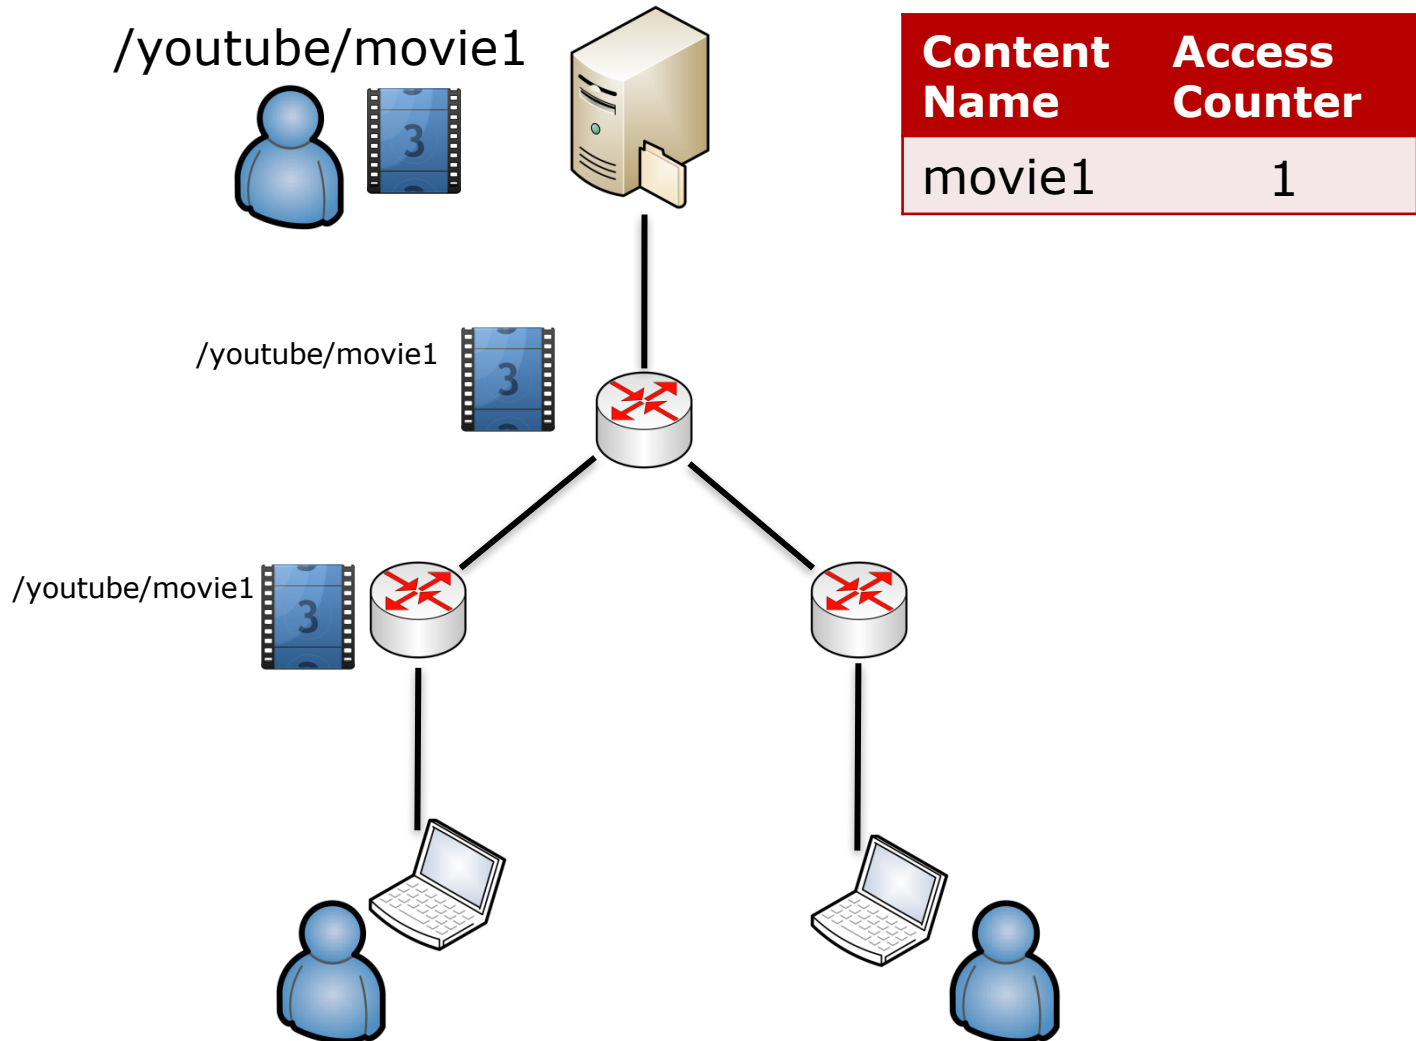
Security Issues in ICN



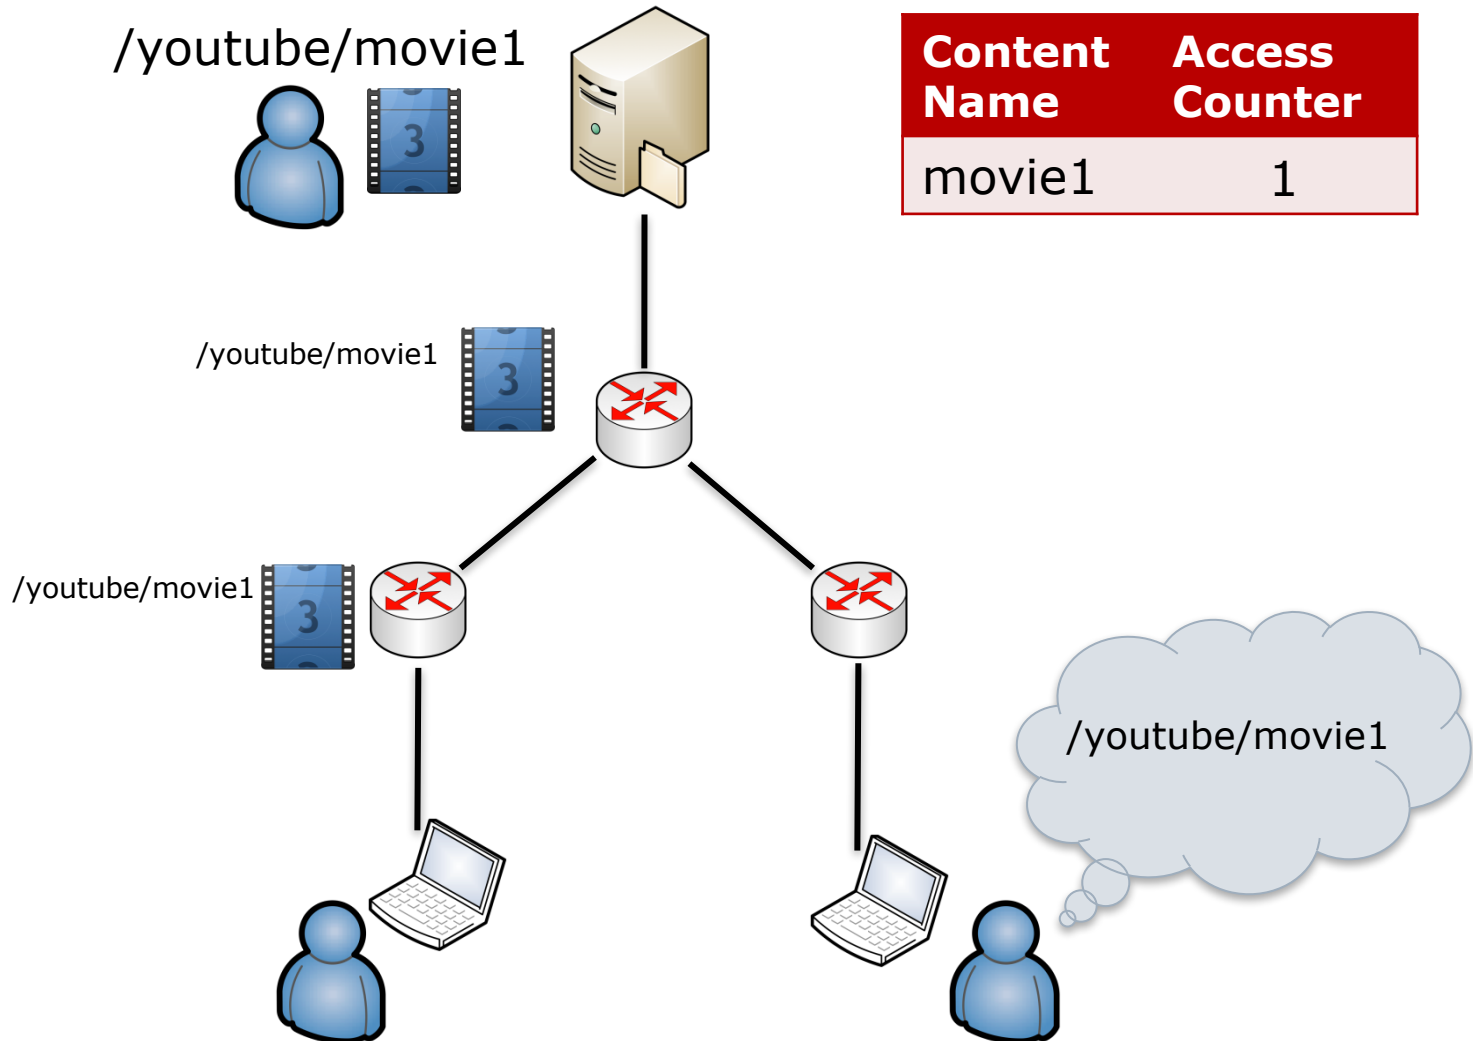
Security Issues in ICN



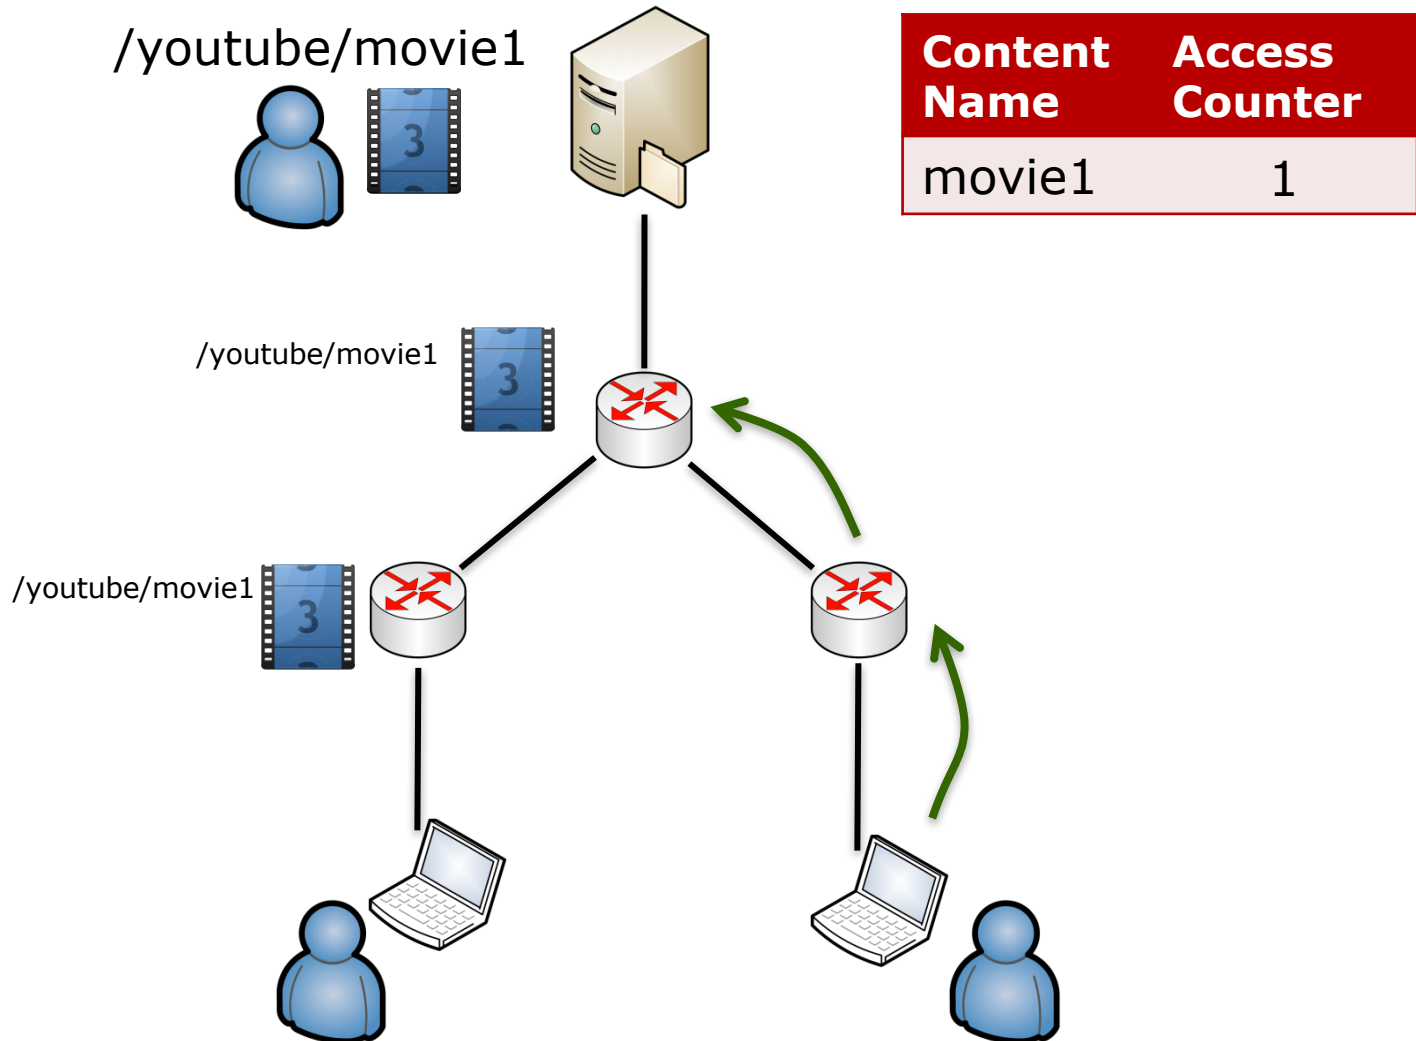
Security Issues in ICN



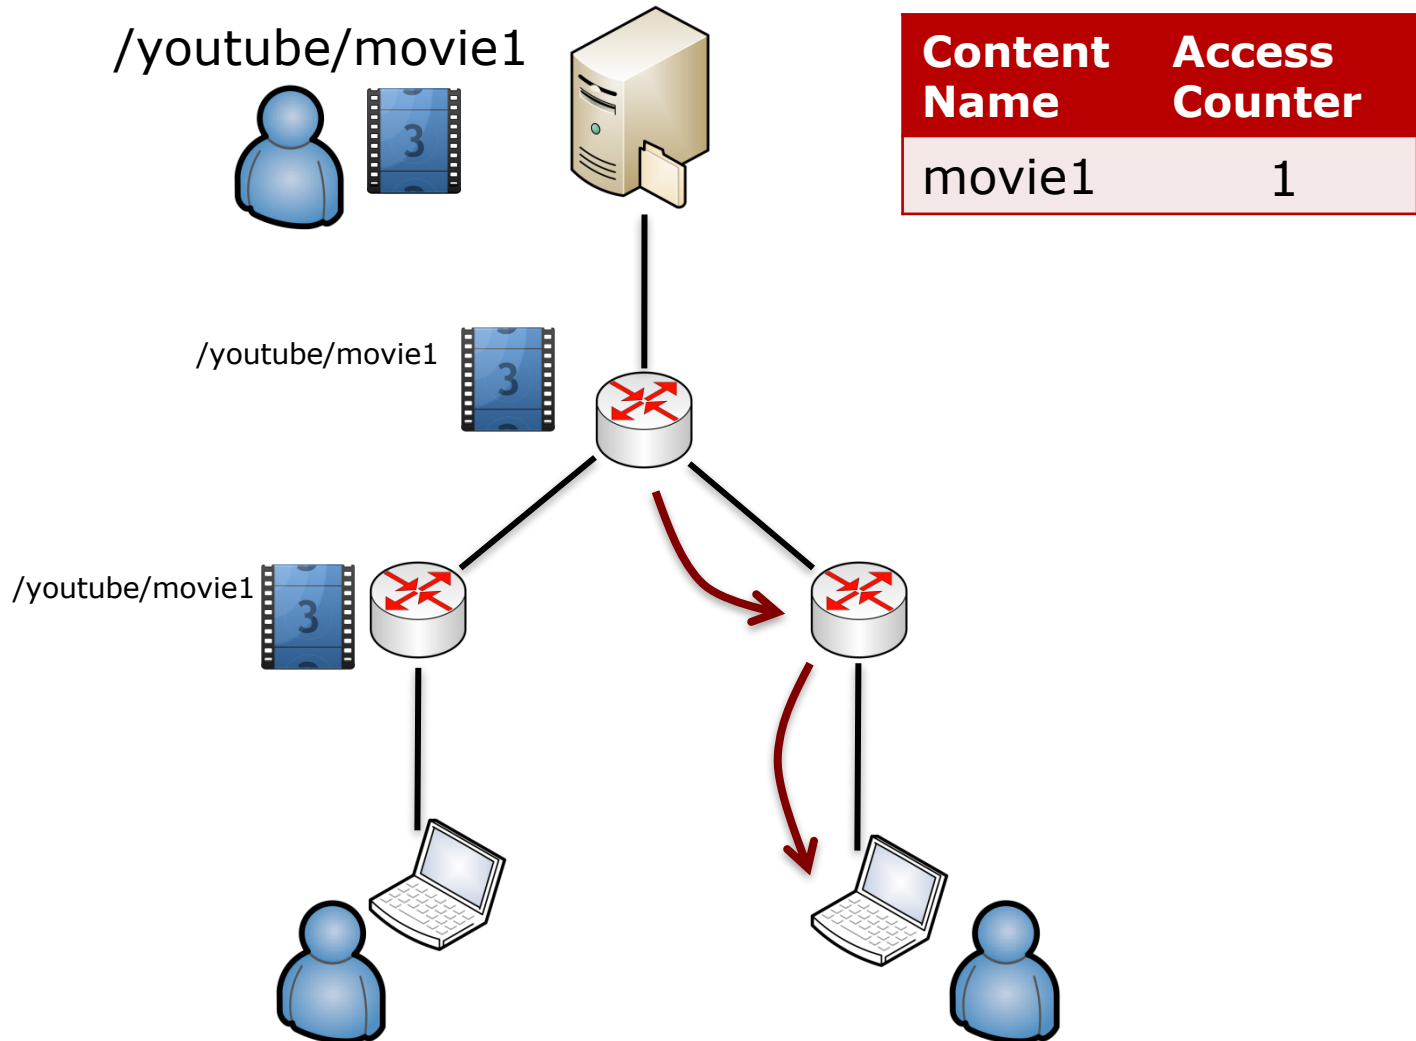
Security Issues in ICN



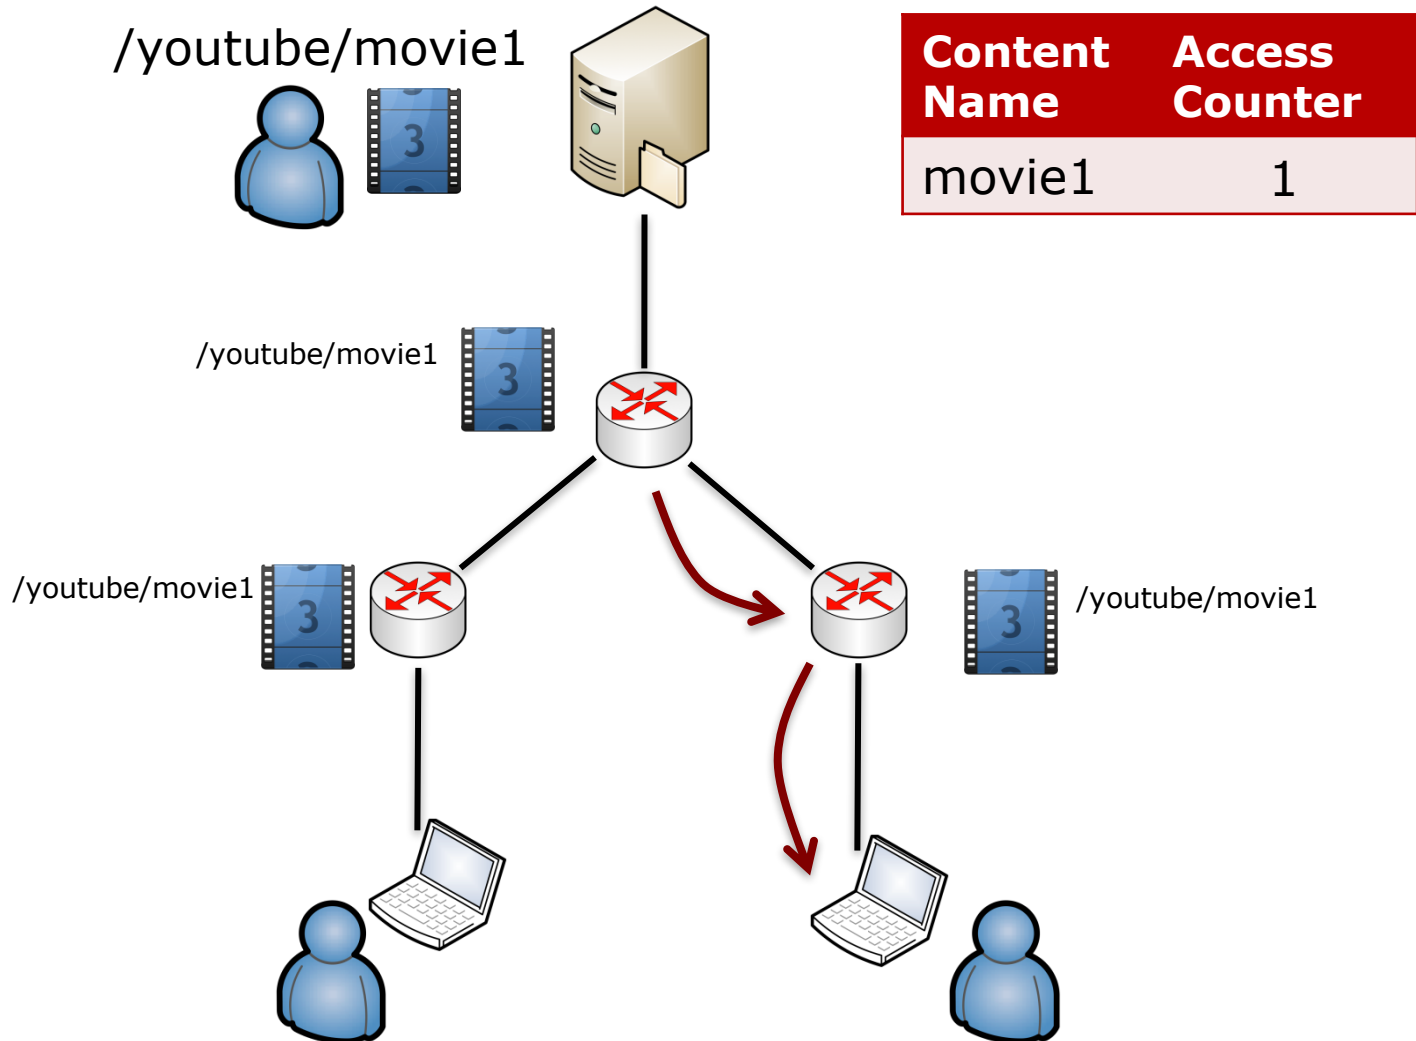
Security Issues in ICN



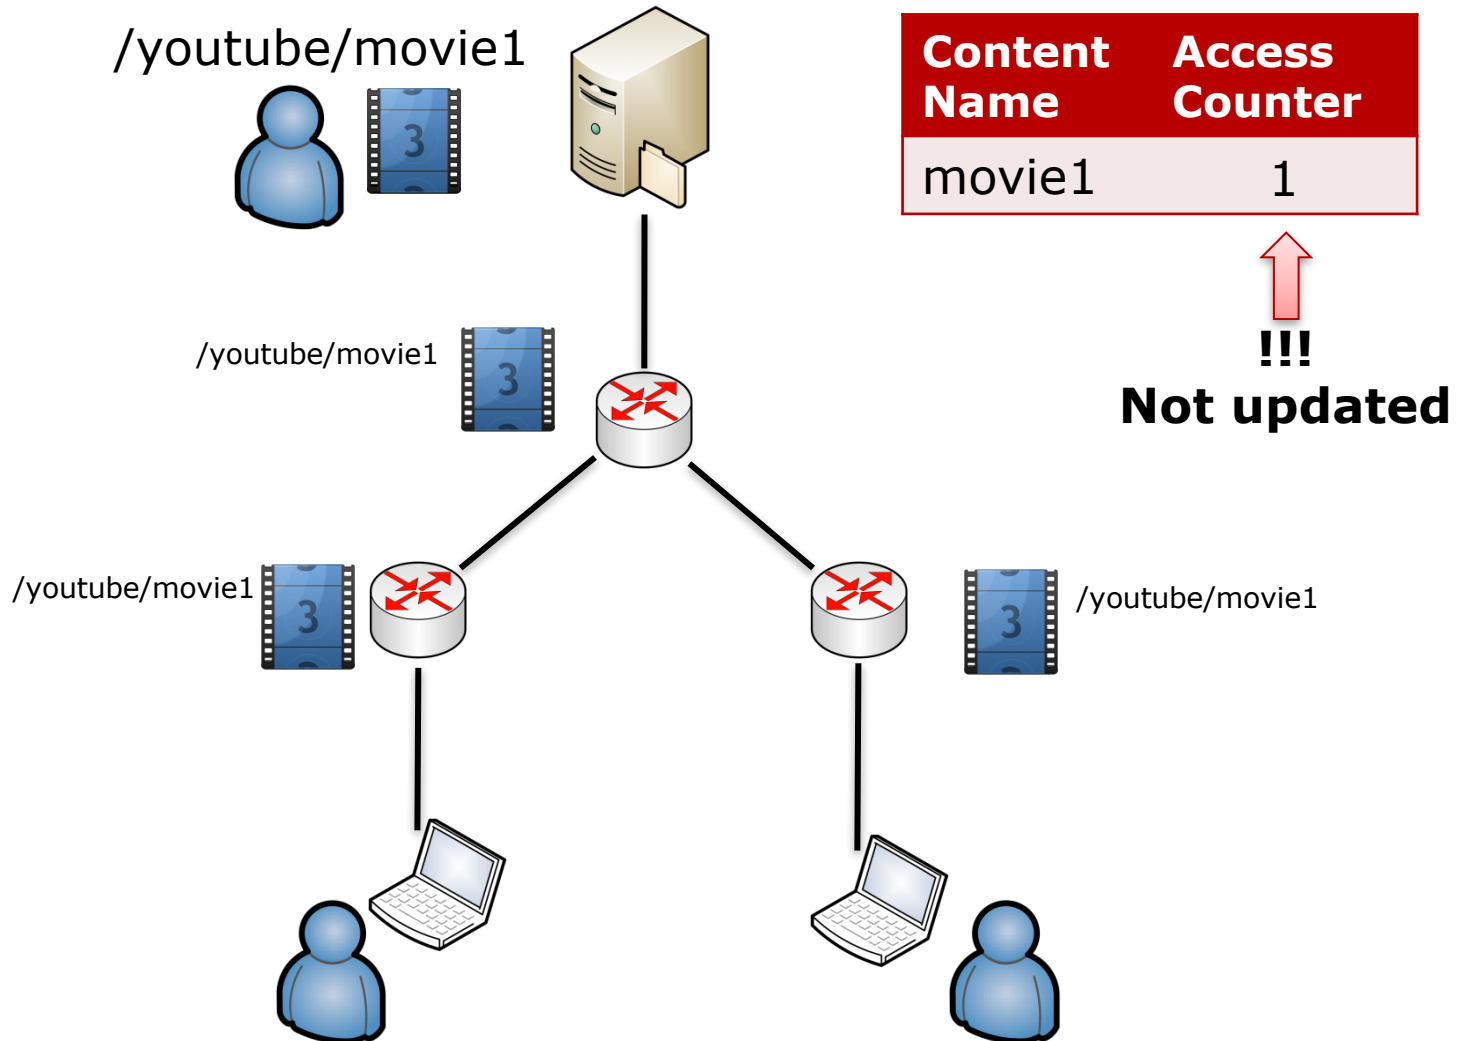
Security Issues in ICN



Security Issues in ICN

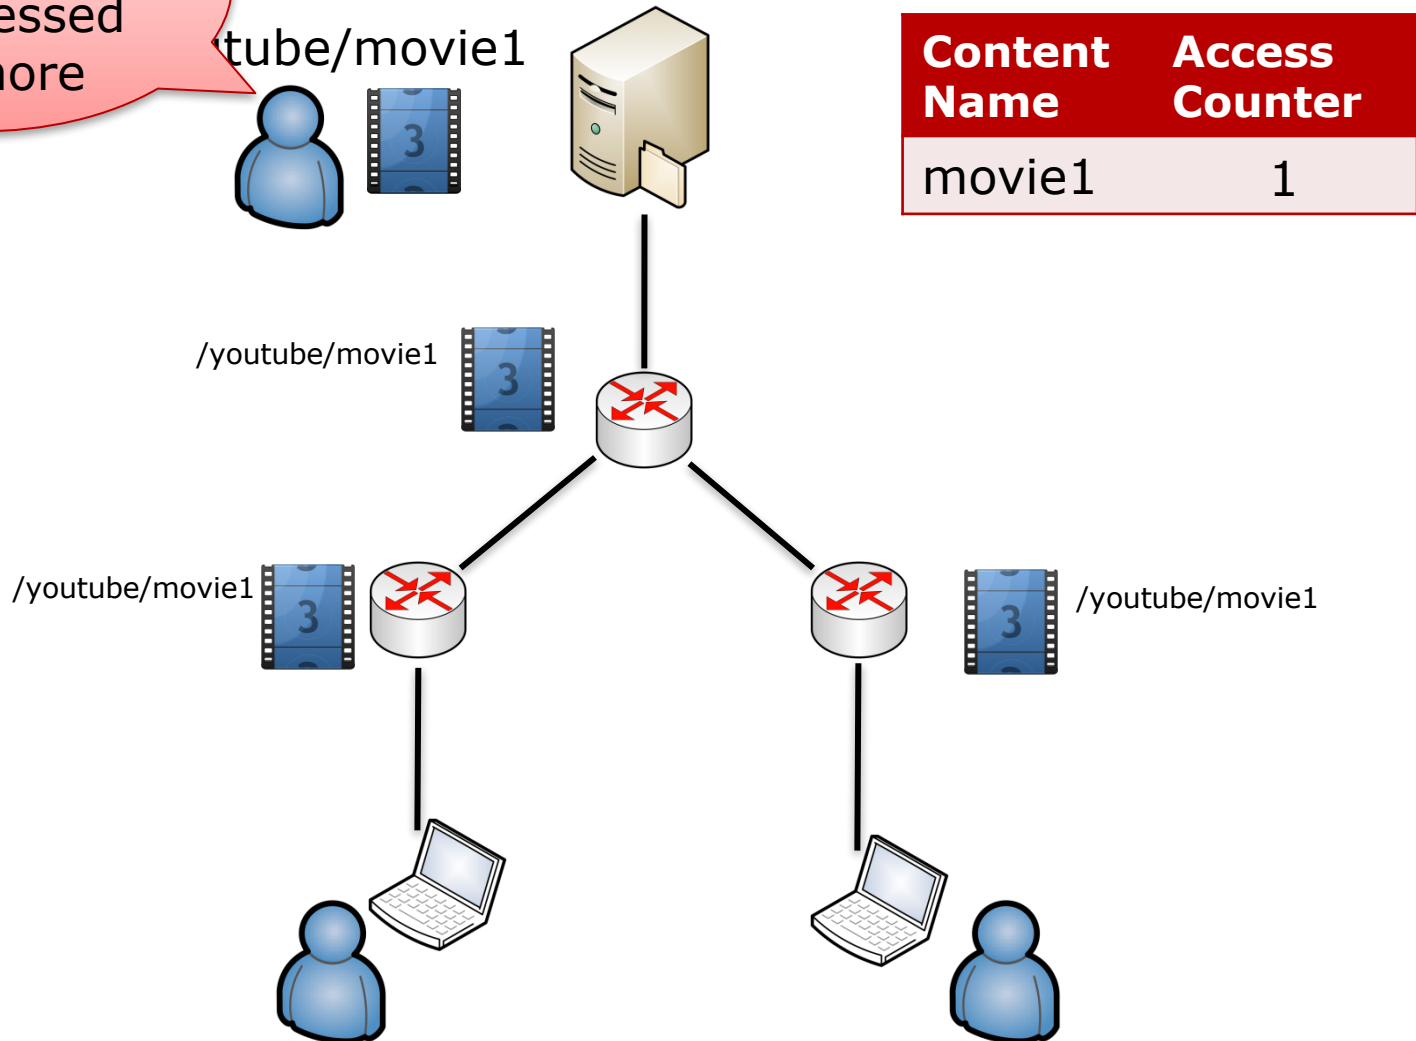


Security Issues in ICN



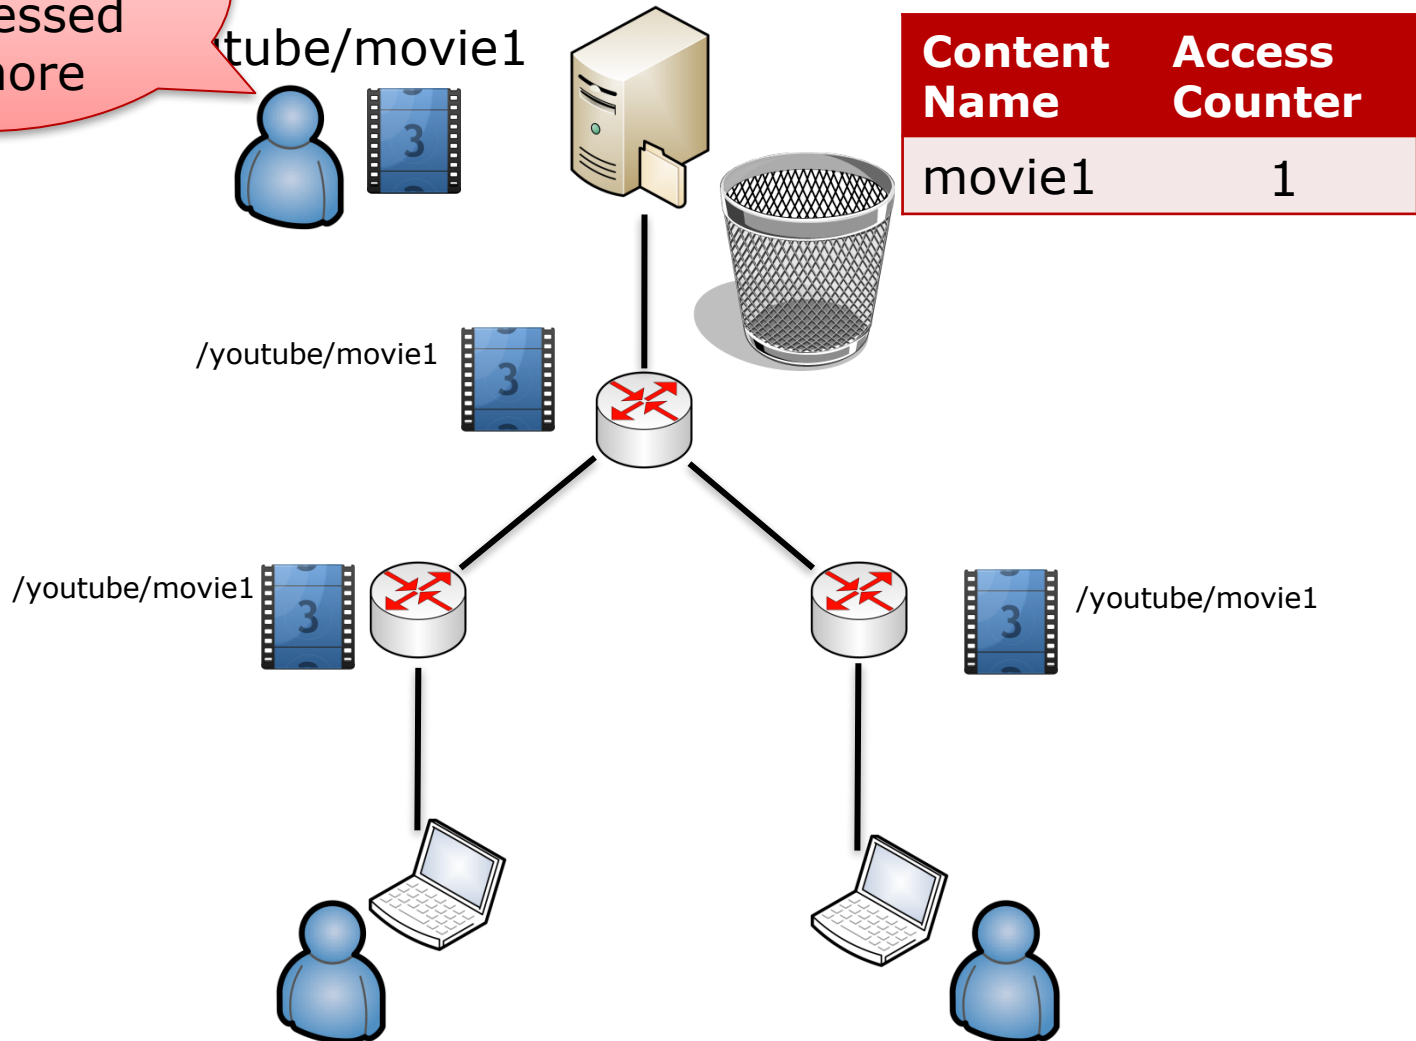
Security Issues in ICN

"movie1"
should not
be accessed
anymore



Security Issues in ICN



"movie1"
should not
be accessed
anymore



Security Issues in ICN

"movie1"
should not
be accessed
anymore


/youtube/movie1




Content Name	Access Counter
--------------	----------------

movie1	1
--------	---

/youtube/movie1



/youtube/movie1

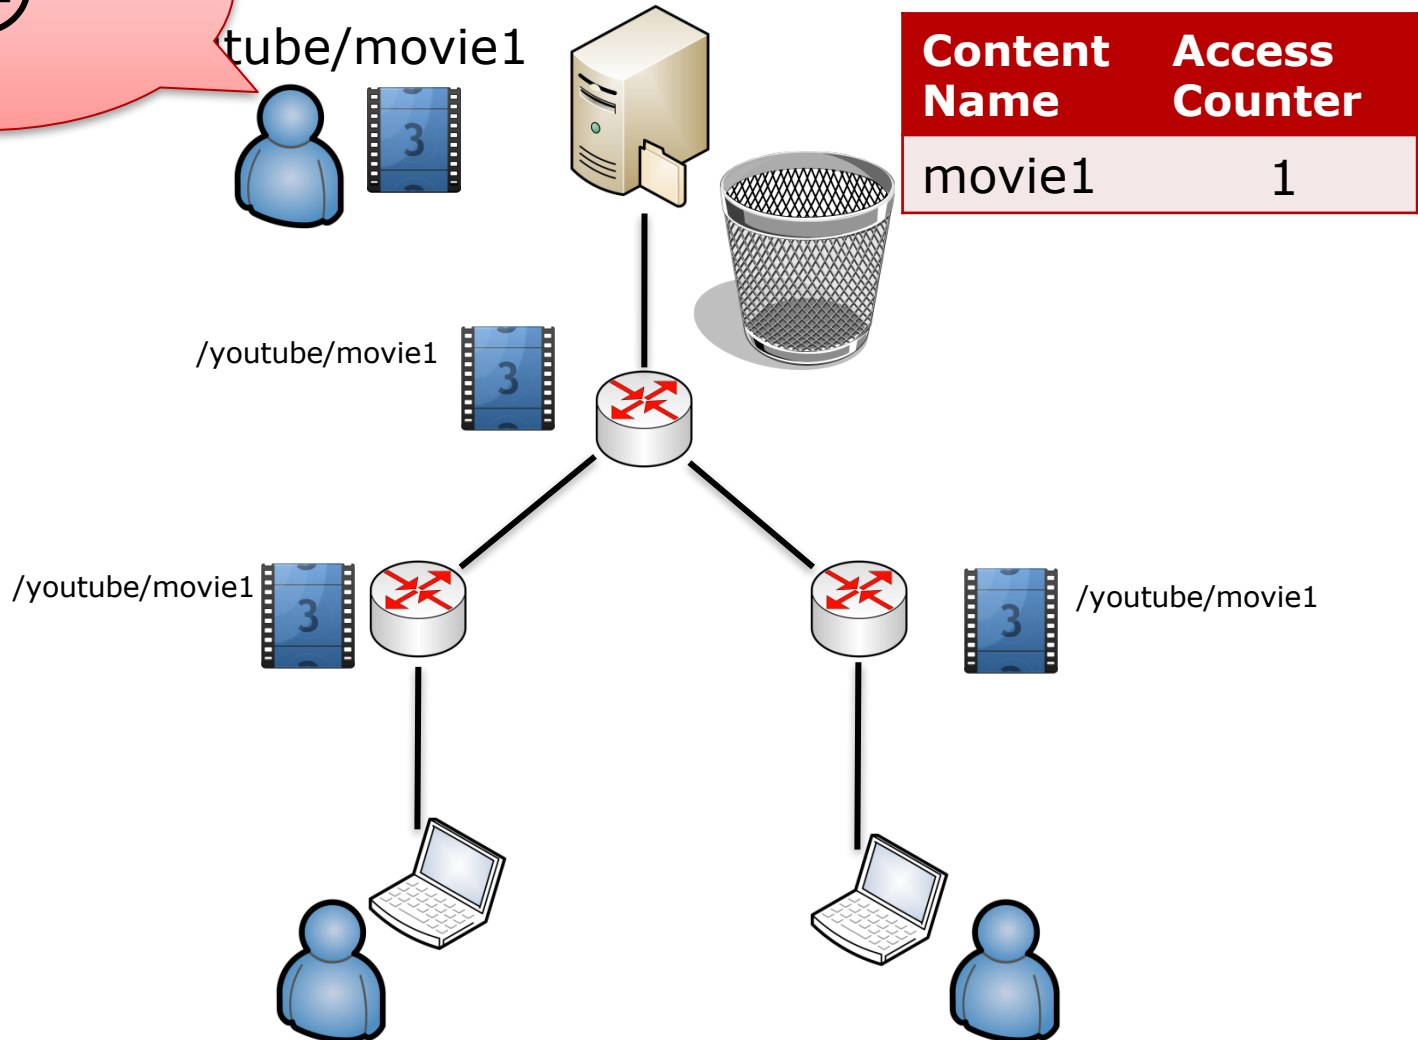


/youtube/movie1



/youtube/movie1

Security Issues in ICN



Security Issues in ICN

- Current CCN proposals do not provide means to enforce:
 - ***Confidential*** communications
 - ***Trackable*** content access
 - And do not support access ***policy evolution***

- *ConfTrack-CCN*:
 - The *first layer* of encryption enforces confidentiality
 - The *second layer* enforces trackability
 - *Key derivation*: to support policy evolution

 - + *ConfTrack-CCN* is cache-friendly

Models for Content Distribution in Internet

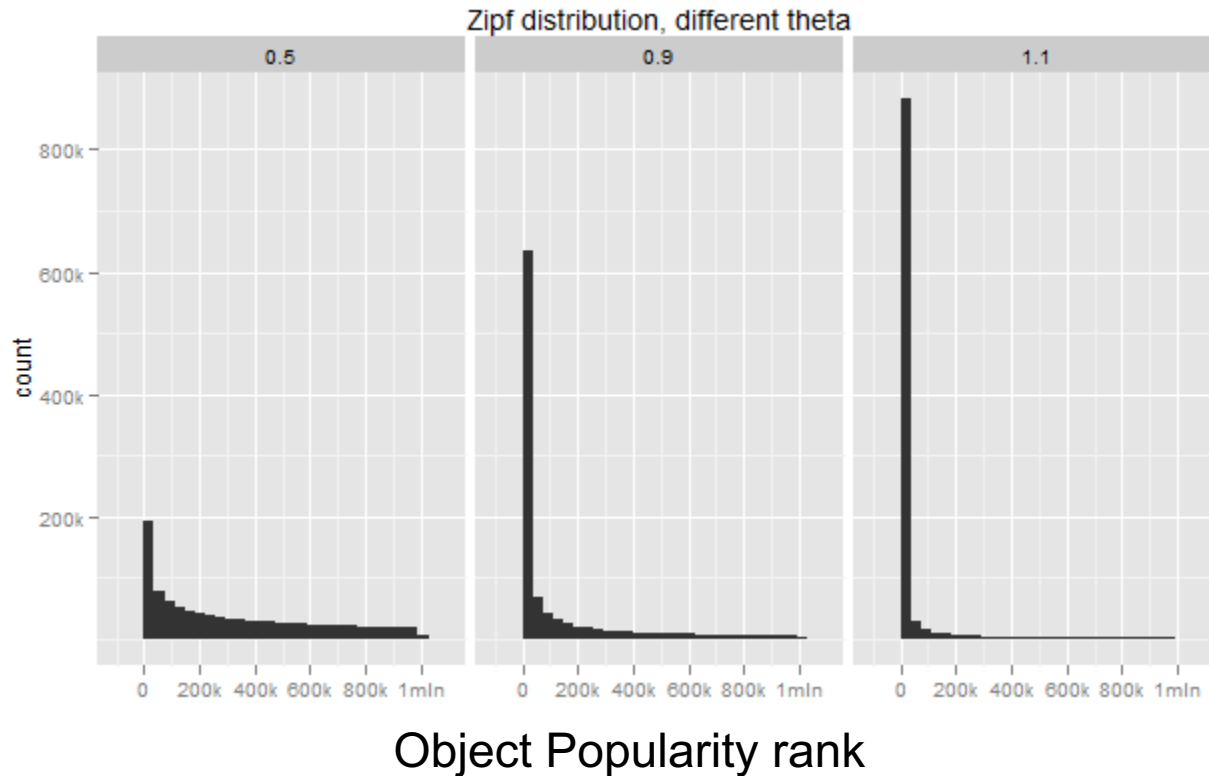
Fabio Martignon

Popularity Models

- Content popularity is the parameter that mostly affects the final results
- Common assumptions made in the literature
 - Requests follow the Independent Reference Model (IRM)
 - Object Popularity is Zipf

$$P(X = n) = \frac{1/n^\alpha}{\sum_{i=1}^N (1/i^\alpha)},$$

Zipf Distribution



Surrogate Server Placement

The K-Median Model

The K-Median Model for CDN

- We tackle the surrogate server placement problem

- We denote:
 - The set of surrogate servers, routers and consumers with: \mathcal{D} , \mathcal{R} , \mathcal{C} , respectively
 - For each router $j \in \mathcal{R}$, the binary variable k_j is set to 1 if we connect a CDN node to that router
 - We denote with $c_{i,j}$ the number of hops on the shortest path between consumer $i \in \mathcal{C}$ and router $j \in \mathcal{R}$
 - We assign each consumer to 1 CDN node. We set $y_{i,j}$ to 1 if consumer $i \in \mathcal{C}$ is assigned to the CDN deployed at router $j \in \mathcal{R}$

The K-Median Problem

- The K-Median Problem can be formulated as follows:

$$\min \sum_{\substack{\forall i \in \mathcal{C} \\ \forall j \in \mathcal{R}}} y_{i,j} \cdot c_{i,j}$$

subject to:

$$x_j \geq y_{i,j} \qquad \forall (i, j) \in \mathcal{C} \times \mathcal{R}$$

$$\sum_{j \in \mathcal{R}} y_{i,j} = 1 \qquad \forall i \in \mathcal{C}$$

$$\sum_{j \in \mathcal{R}} x_j = |\mathcal{D}|$$

$$x_j \in \{0, 1\} \qquad \forall j \in \mathcal{R}$$

$$y_{i,j} \in \{0, 1\} \qquad \forall (i, j) \in \mathcal{C} \times \mathcal{R}$$