



# Architectural Development and Routing Design in Named Data Networking

Lan Wang

University of Memphis

Oct. 28, 2015

[www.named-data.net](http://www.named-data.net)

# Outline

- Named Data Networking (NDN) concepts
- NDN development overview
- Potential solution to scale NDN routing


# The NDN Project <http://www.named-data.net/>

- Part of the NSF Future Internet Architecture FIA initiative
- Goal: design the next generation Internet Architecture
- NSF funding for NDN: 2010-13 and 2014-16, ~\$15M



# Today's Internet Traffic

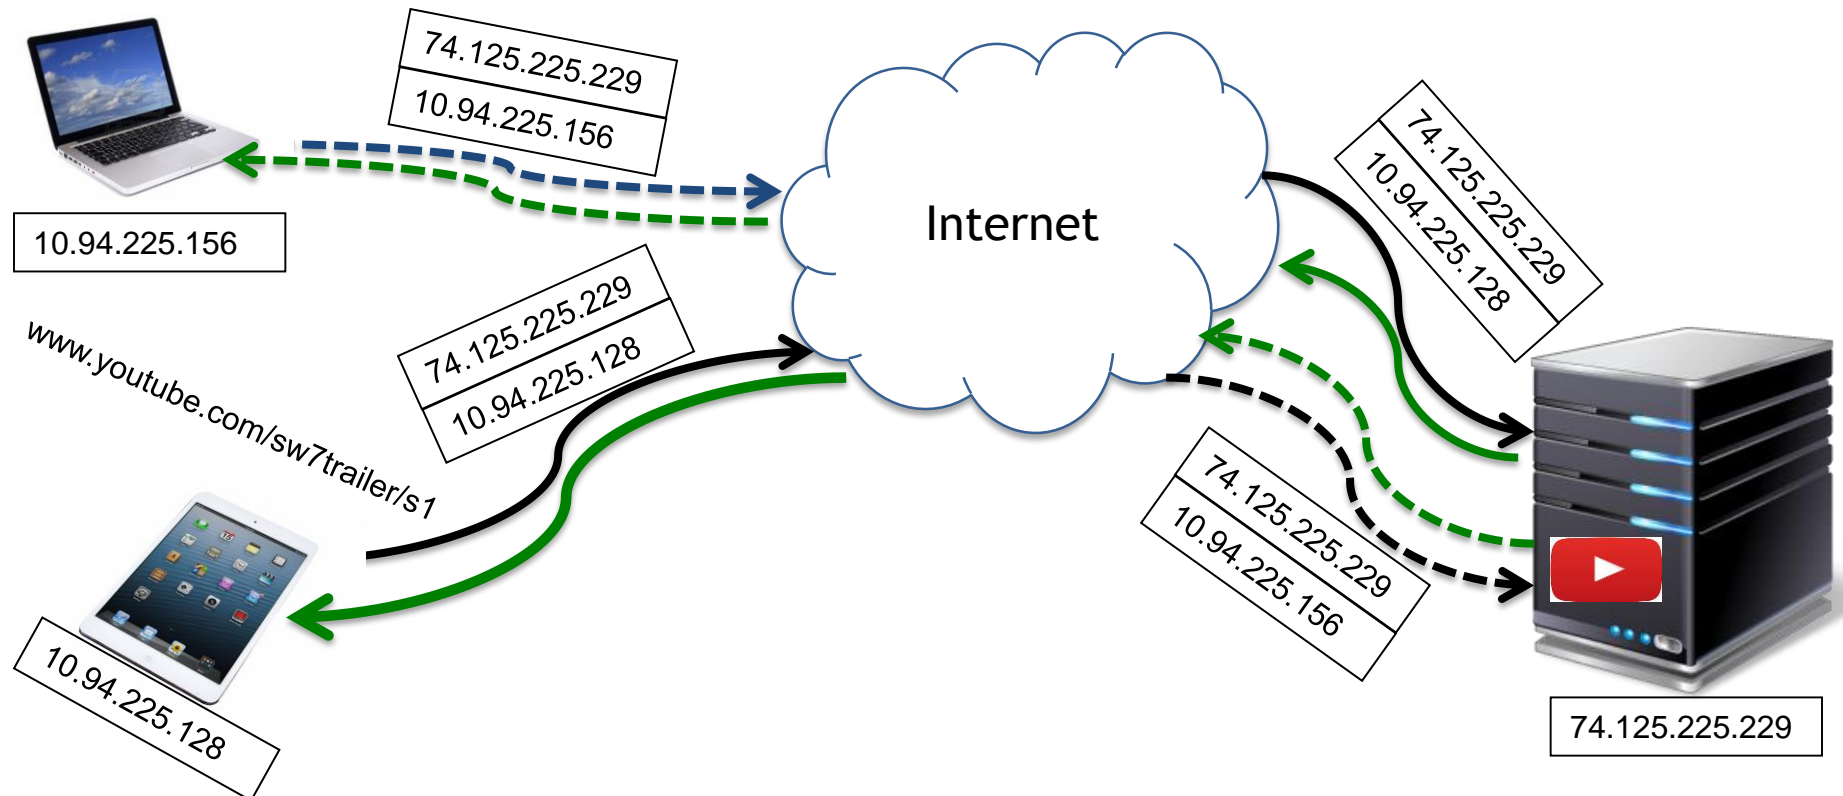
Communication is “Content” driven.



# Internet Protocol (IP)

But underlying communication is “*destination*” driven.

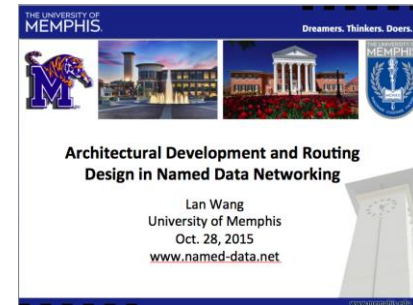
[www.youtube.com/sw7trailer/s1](http://www.youtube.com/sw7trailer/s1)



# A Simpler Way

1. **Name the data, not the hosts!**
2. tell the network what you want..
3. let the network find it for you

</memphis.edu/lanwang/talks/OleMiss15.pdf>




</room/thermostat/1/status>



# A Simpler Way

1. **Name the data, not the hosts!**
2. tell the network what you want..
3. let the network find it for you

`/memphis.edu/lanwang/talks/OleMiss15.pdf`



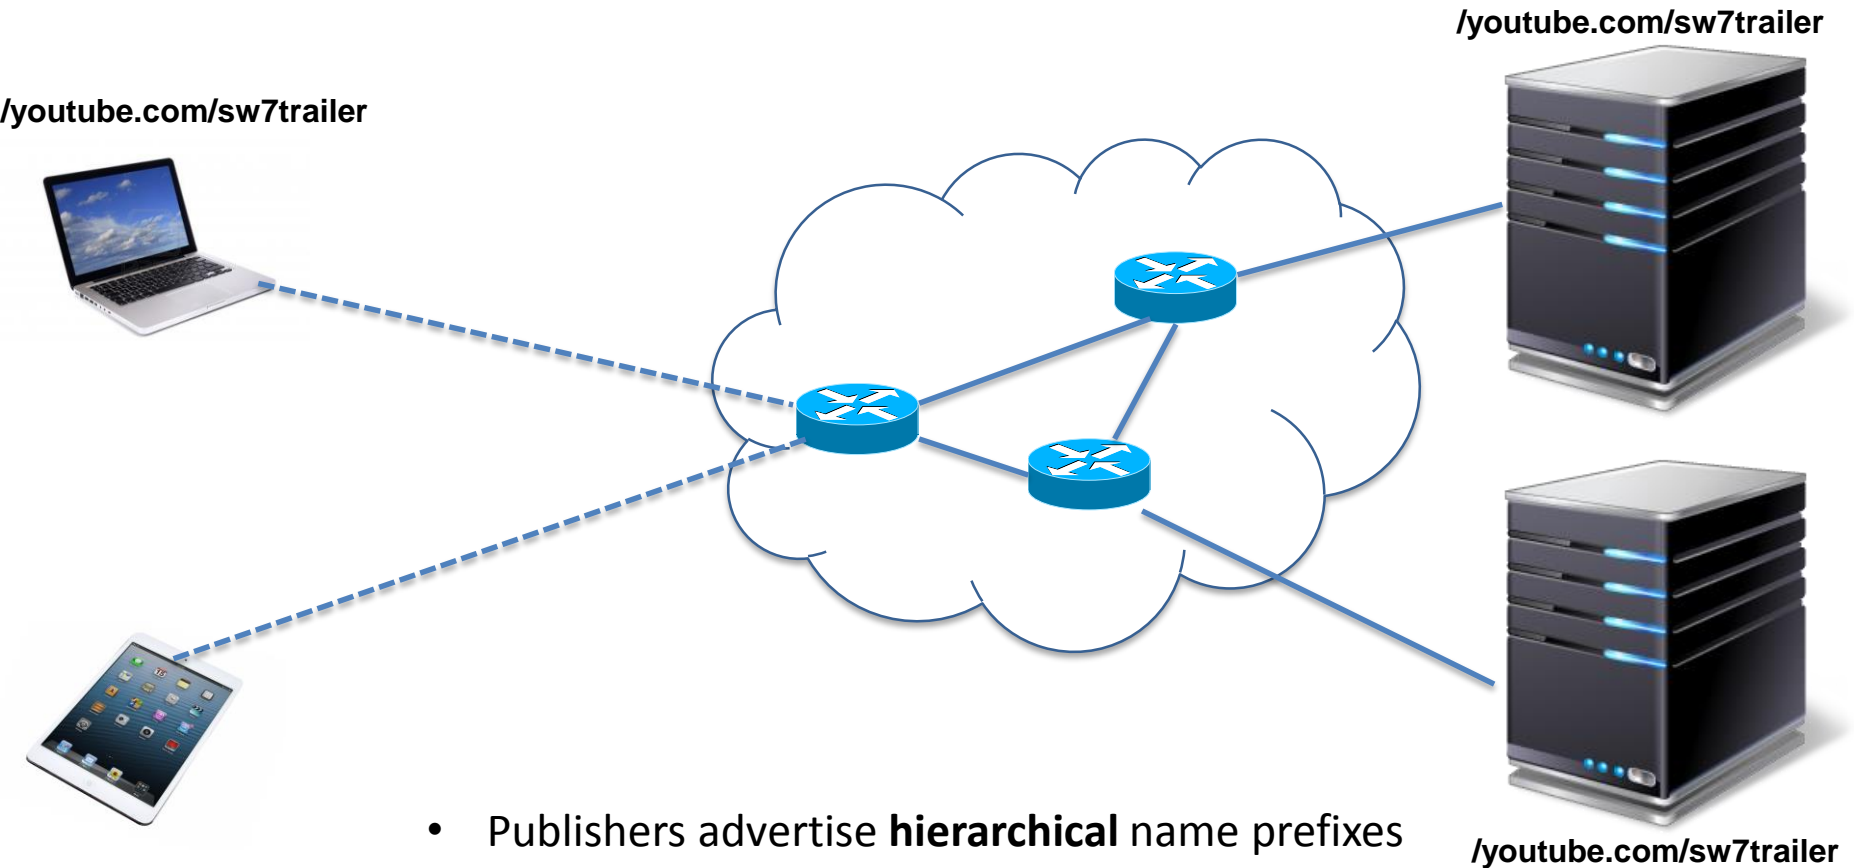
`/room/thermostat/1/status`



Note: this is not Google Search!



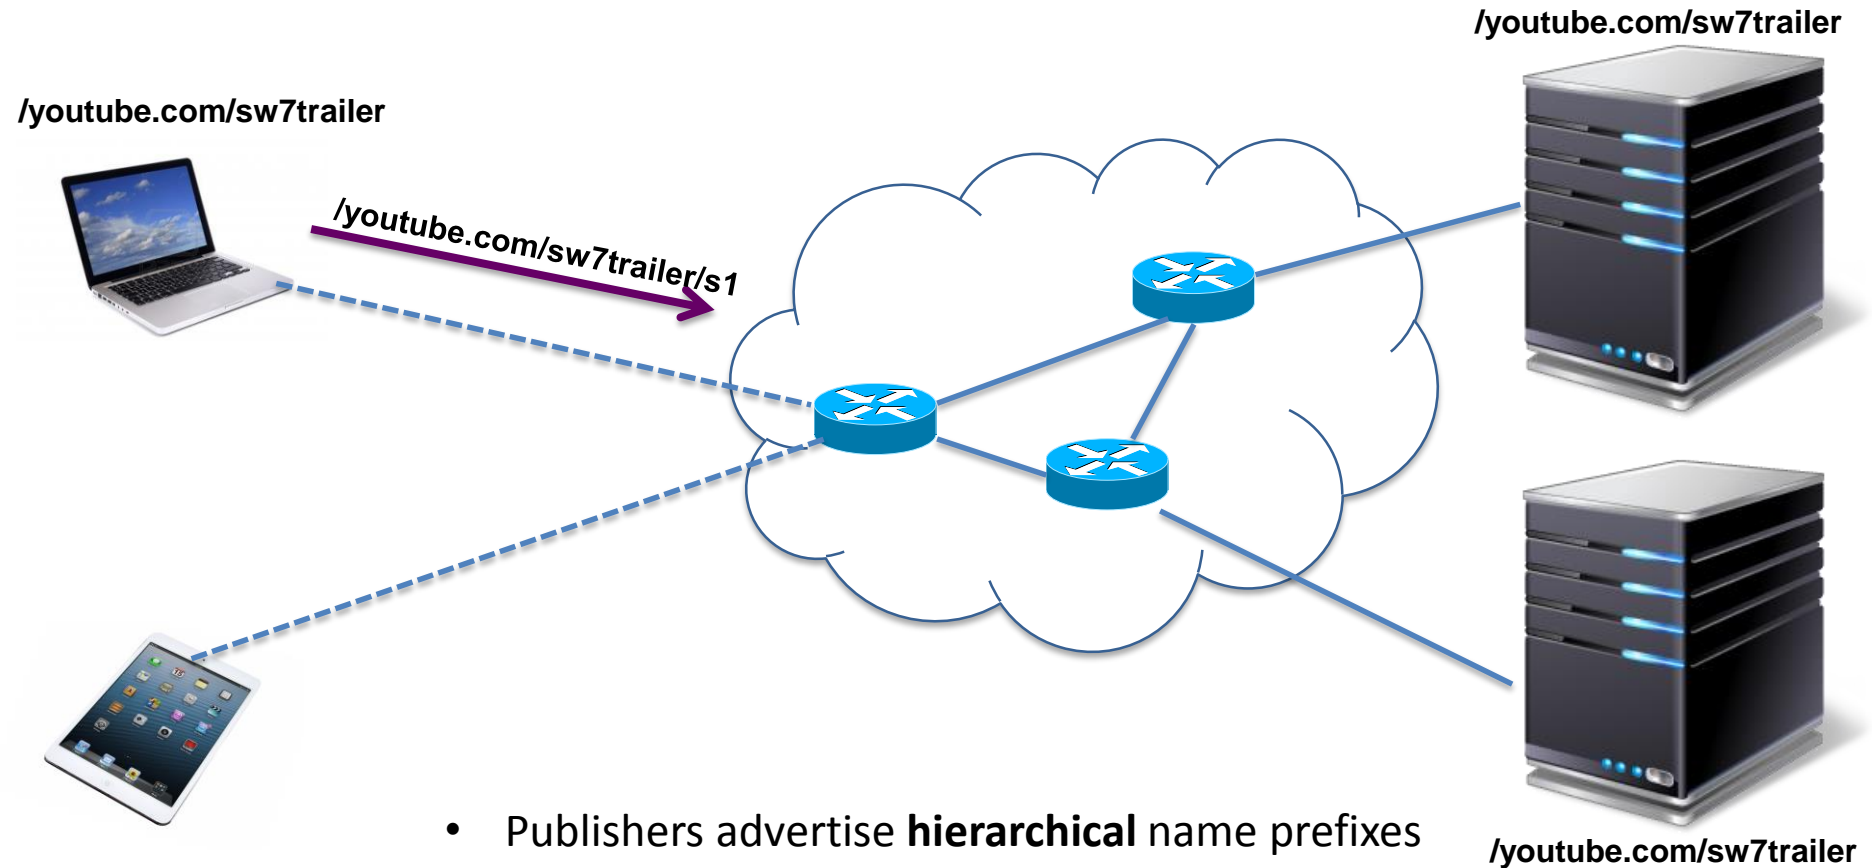
# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

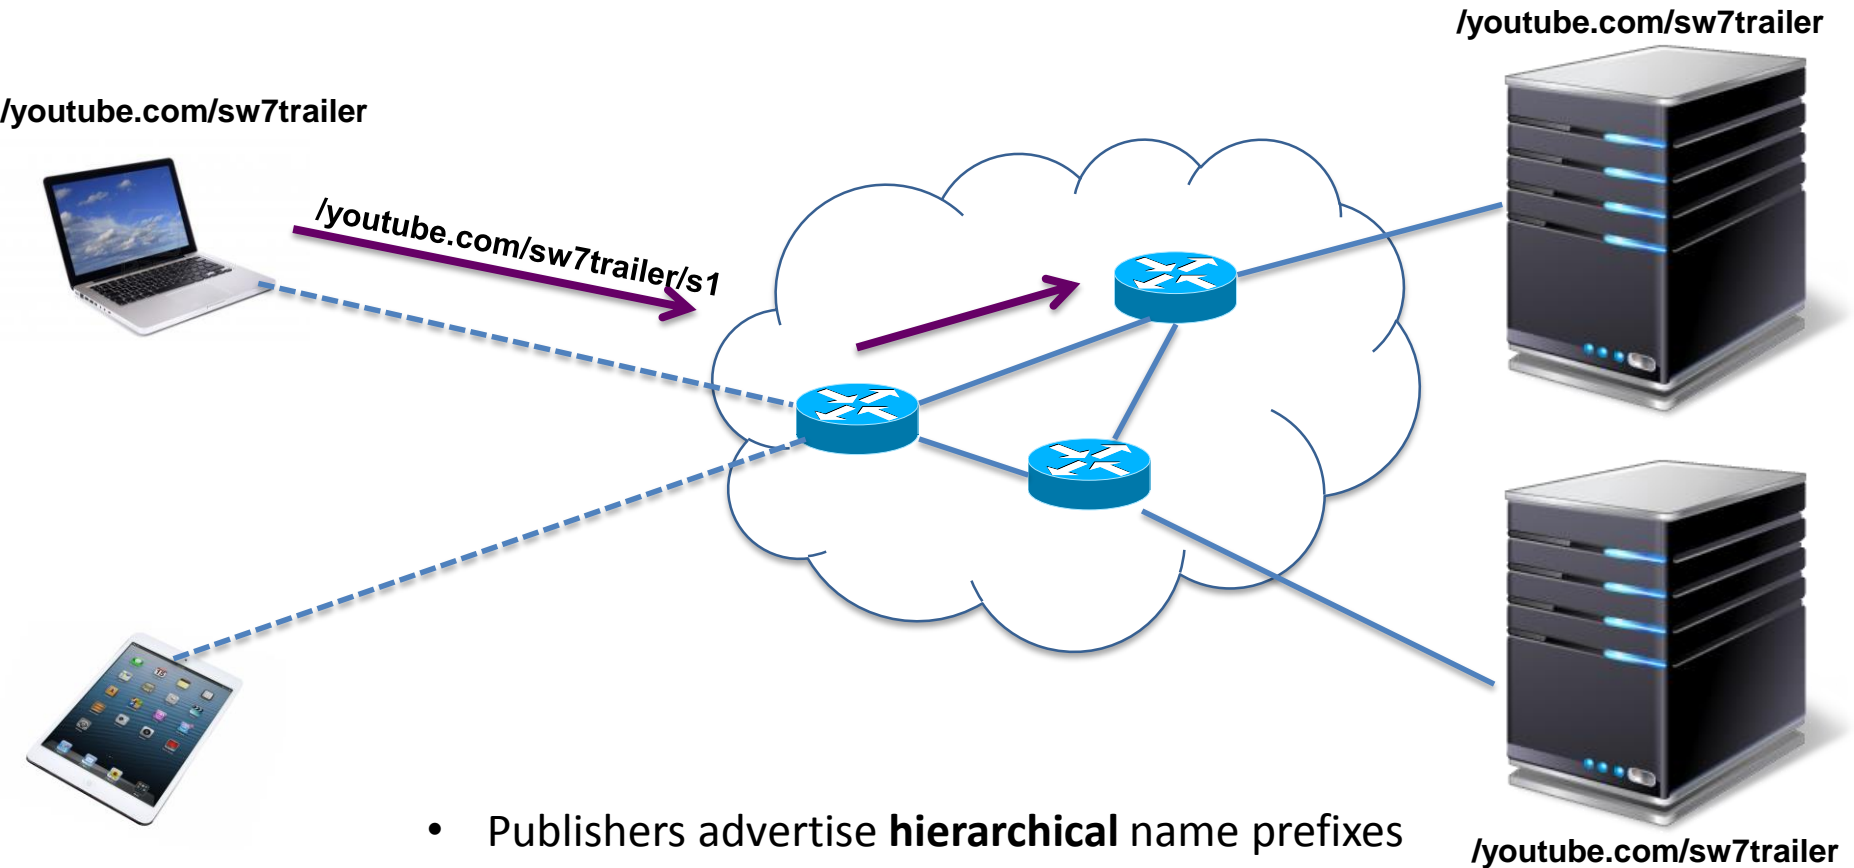


# Named Data Networking (NDN)



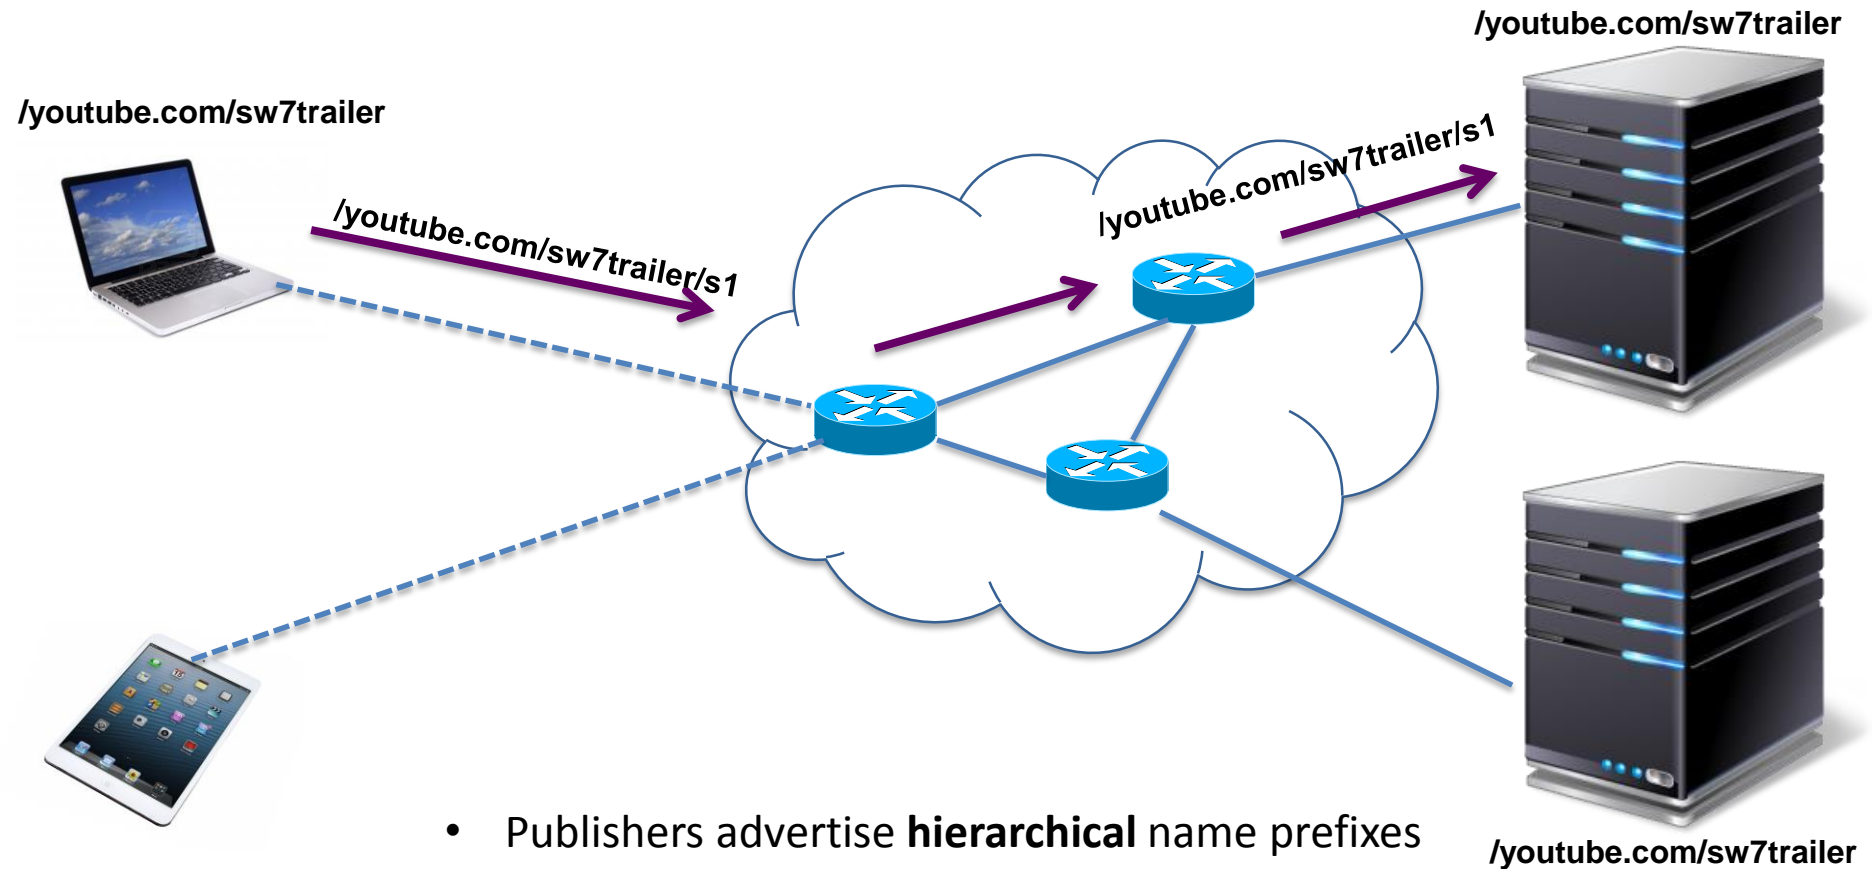
- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network


# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network




# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network




# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network


# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network




# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network


# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

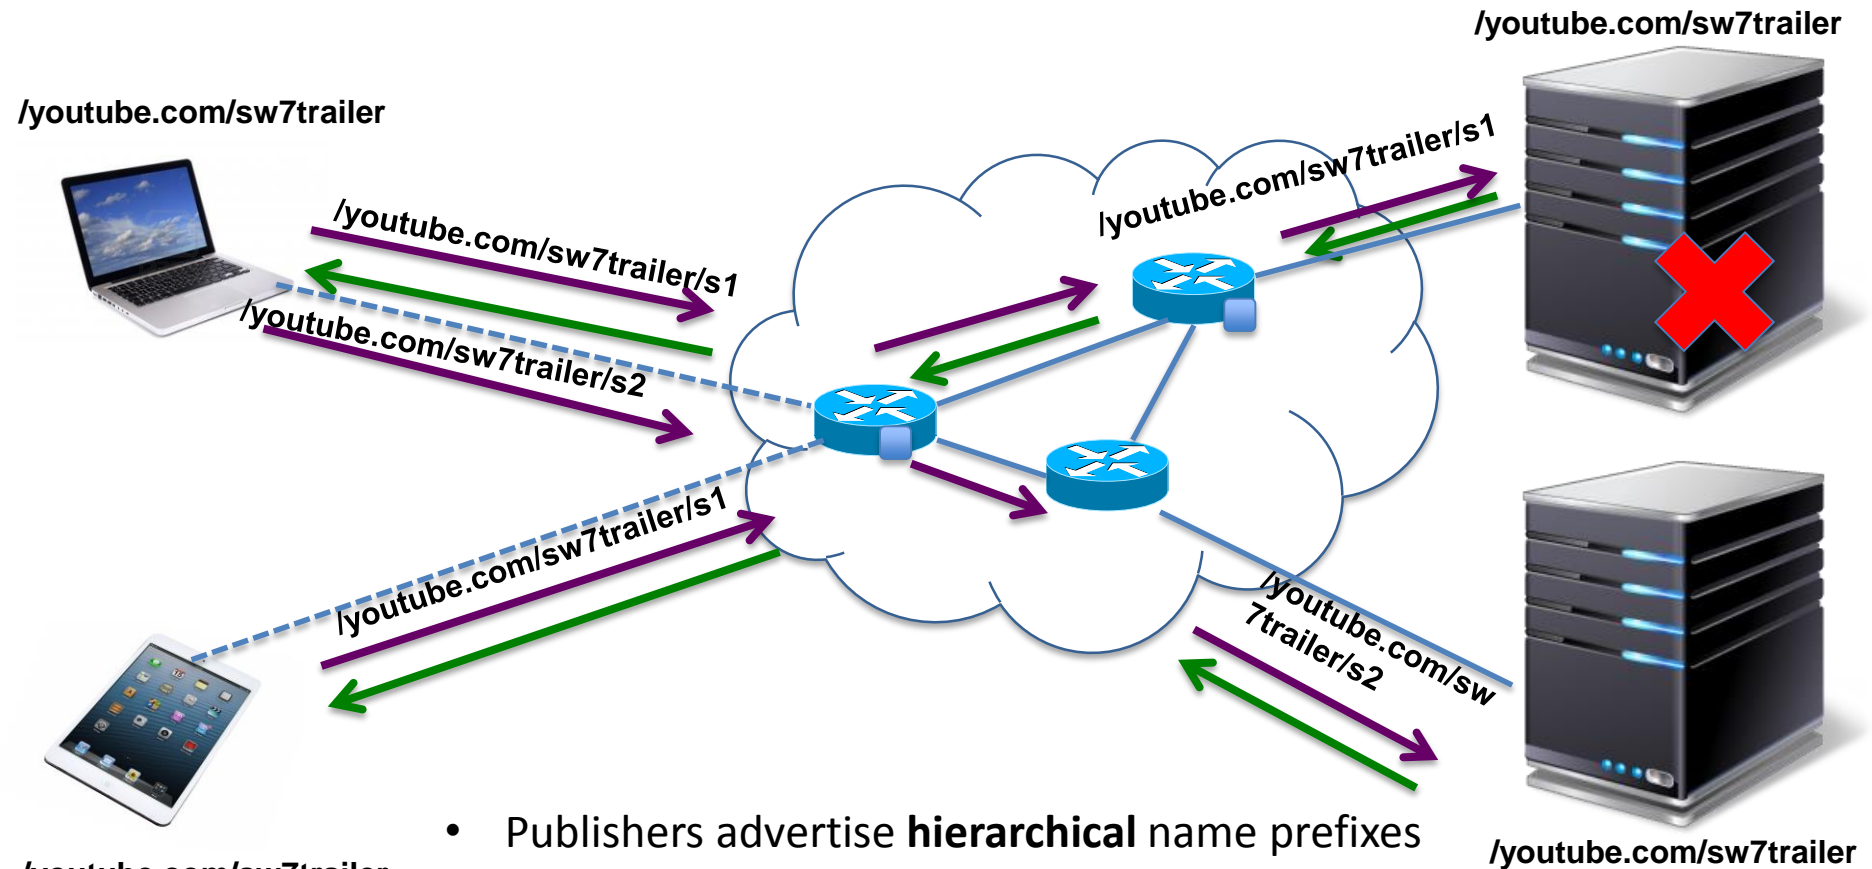


# Named Data Networking (NDN)



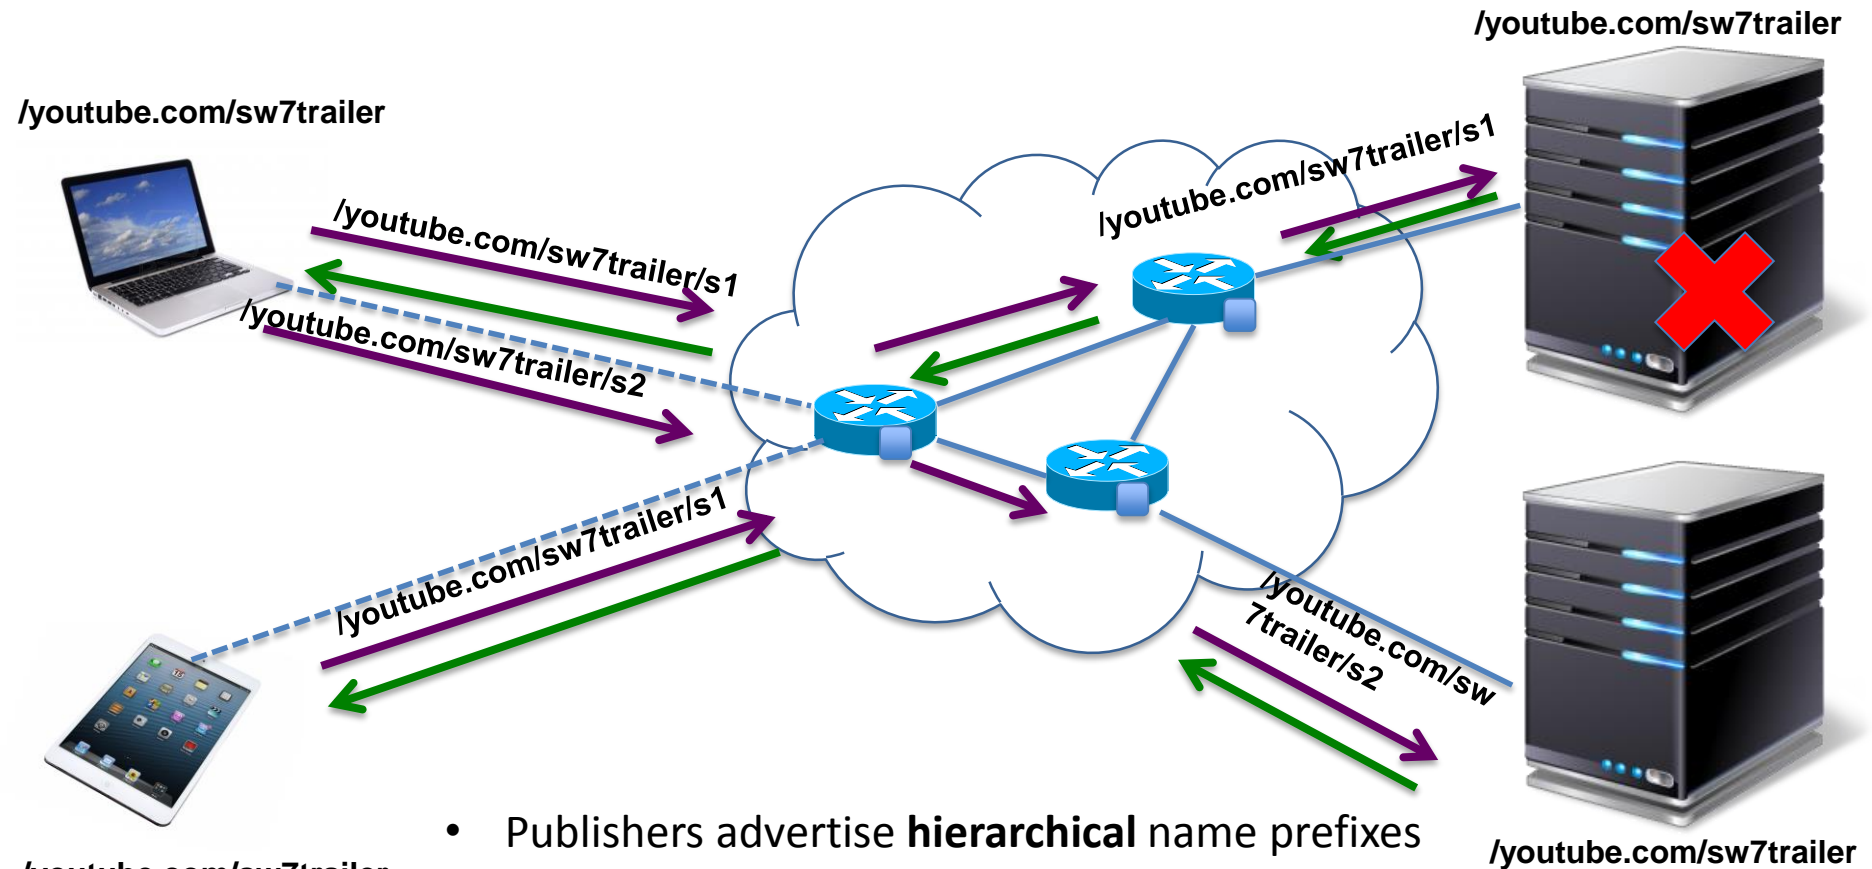
- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)



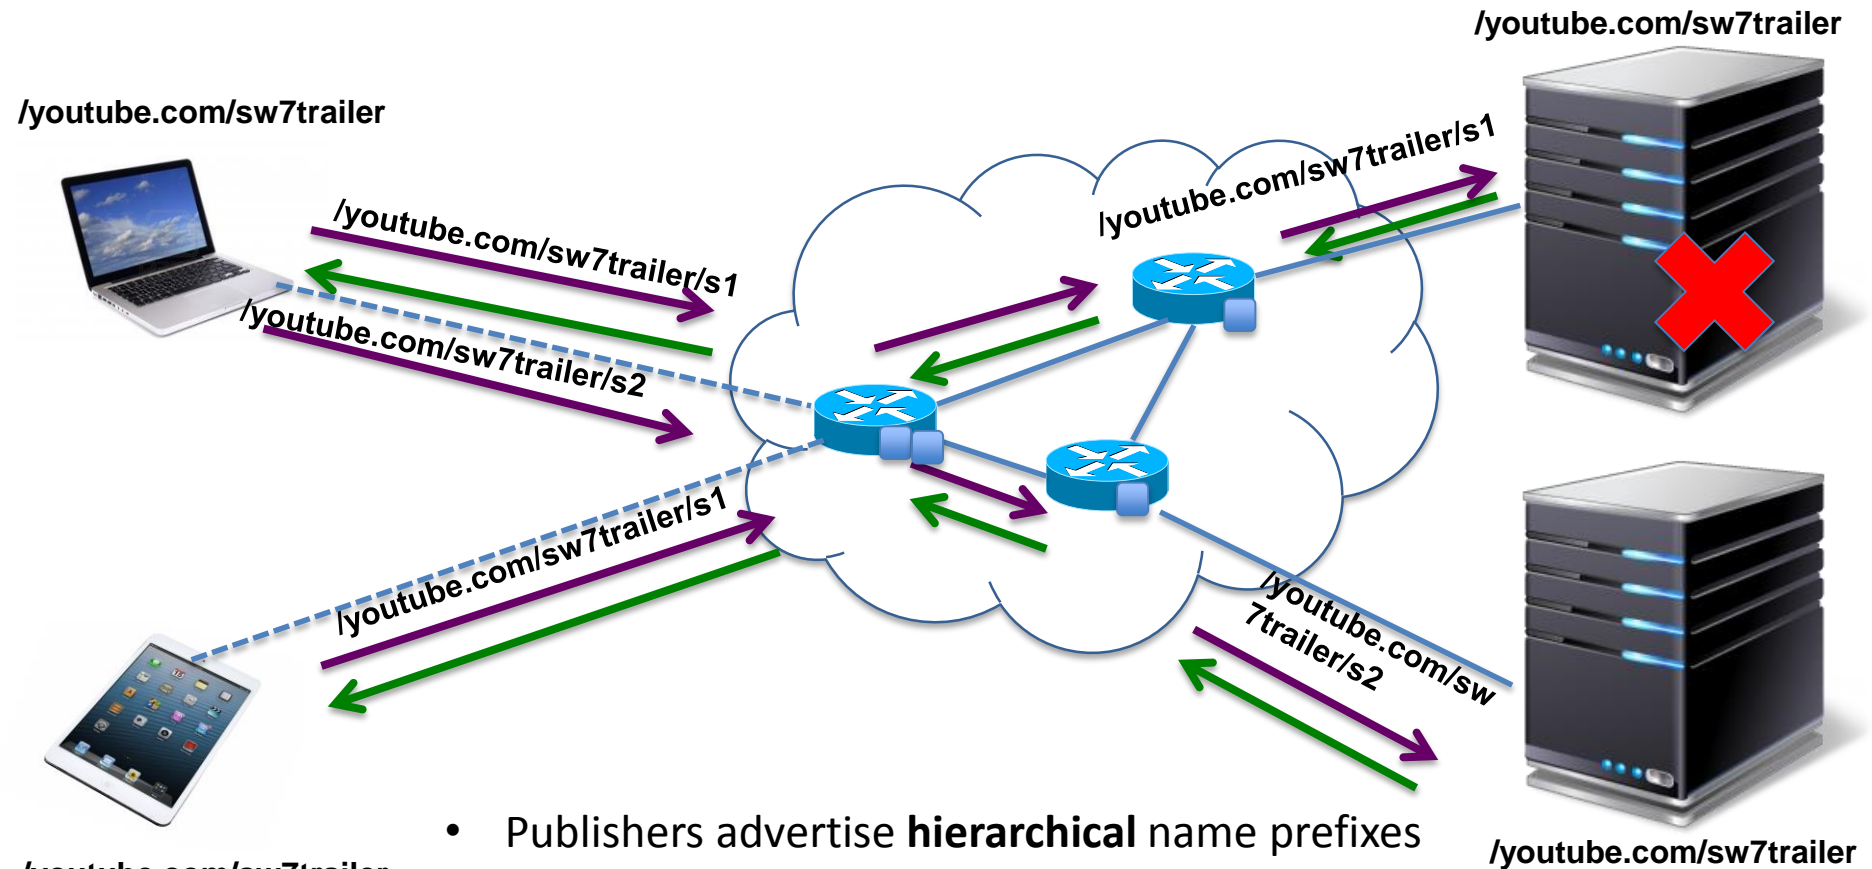
- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

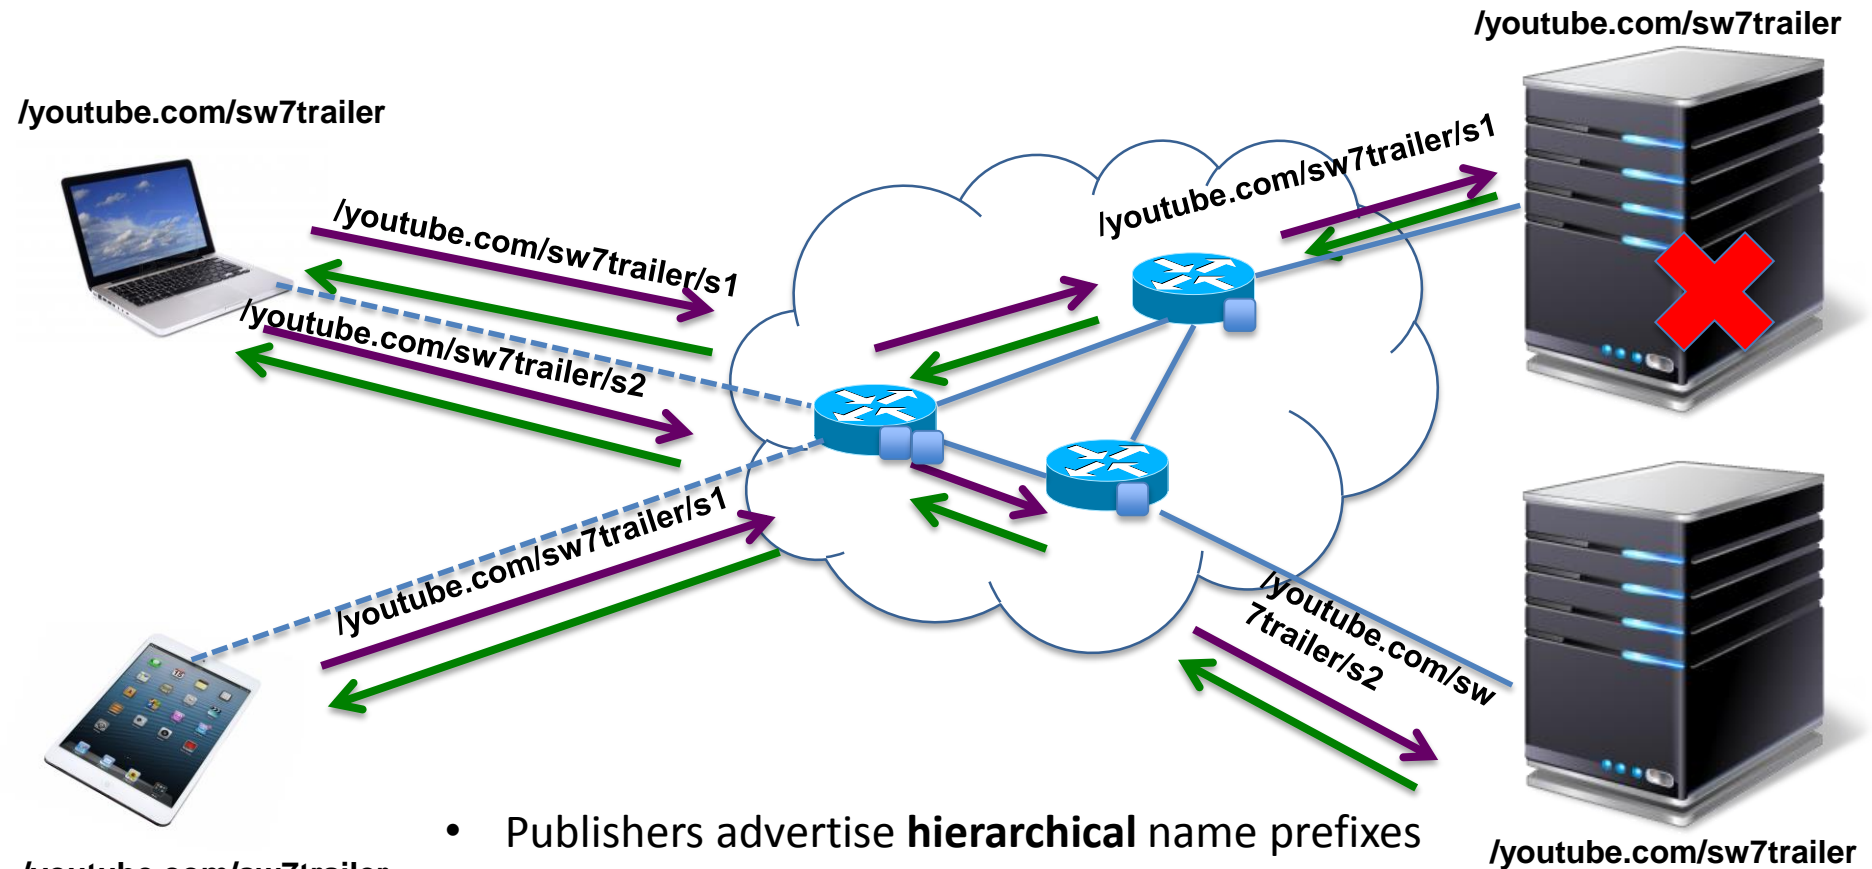
# Named Data Networking (NDN)



- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network




# Named Data Networking (NDN)




- Publishers advertise **hierarchical** name prefixes
- Users send **Interests** to published prefix
- “Breadcrumbs” direct **data** back to the user
- Data is **cached** into the network

# Today's Data distribution







# Today's Data distribution




# Today's Data distribution




# Today's Data distribution



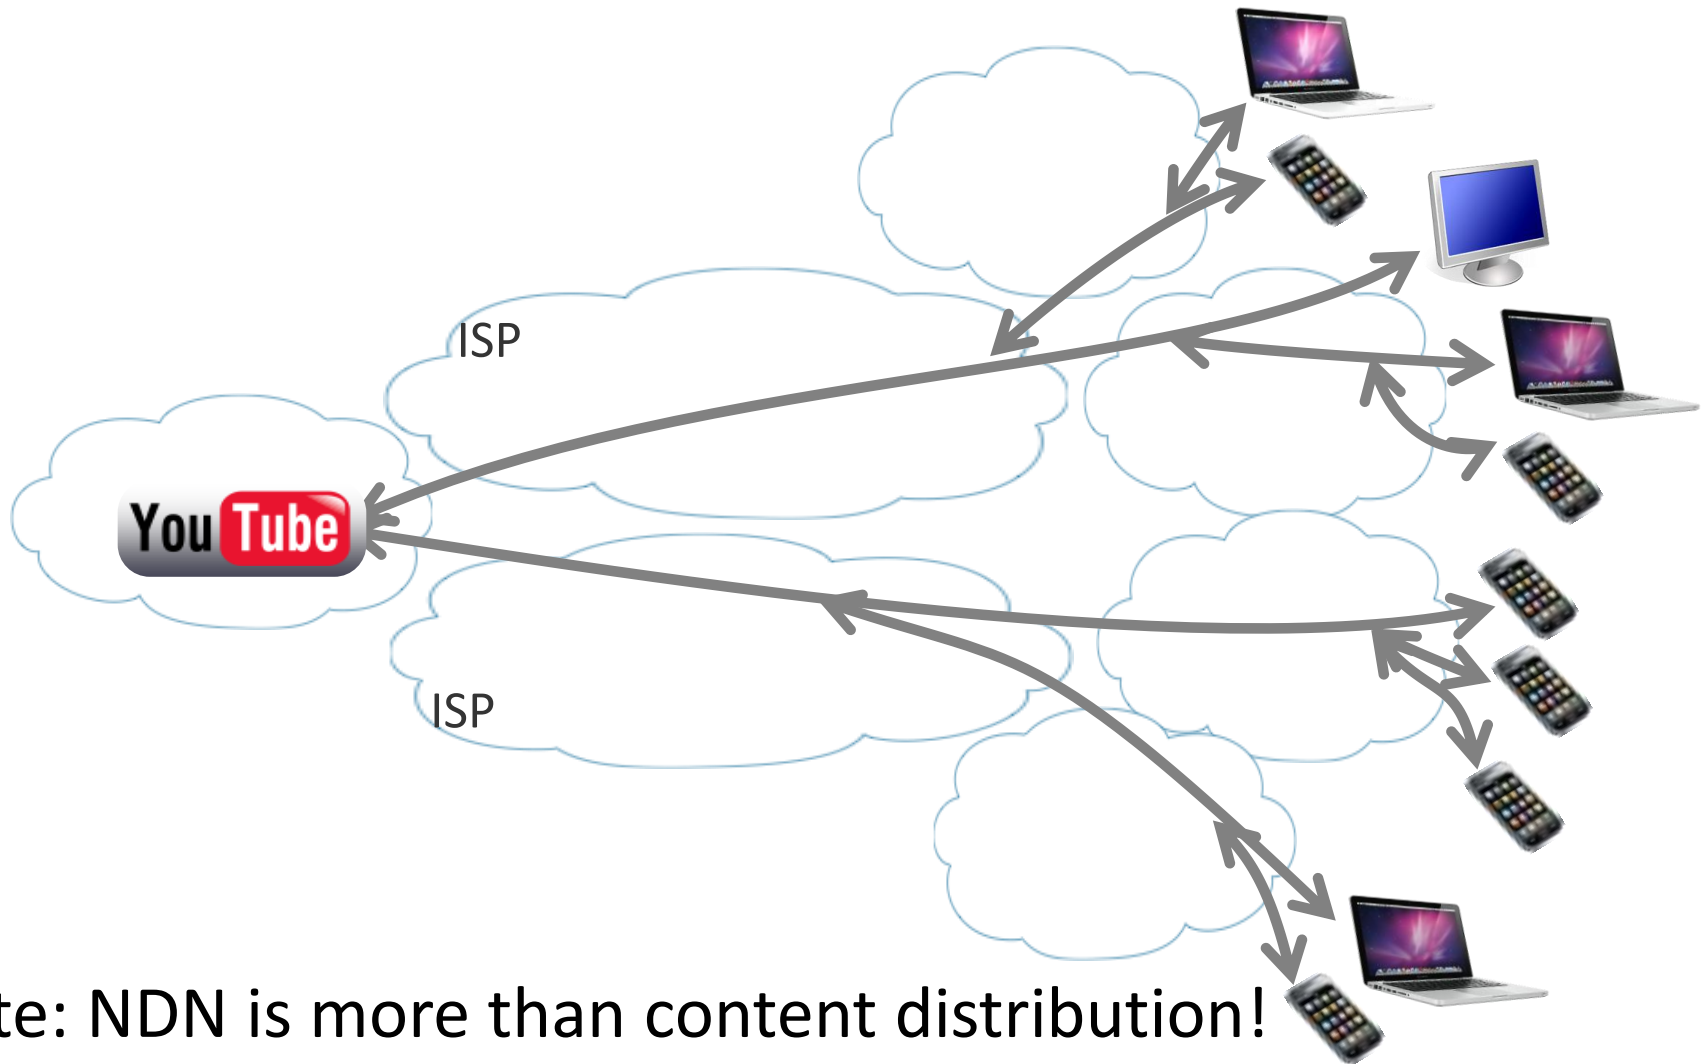
# Today's Data distribution



# Data distribution via NDN



# Data distribution via NDN

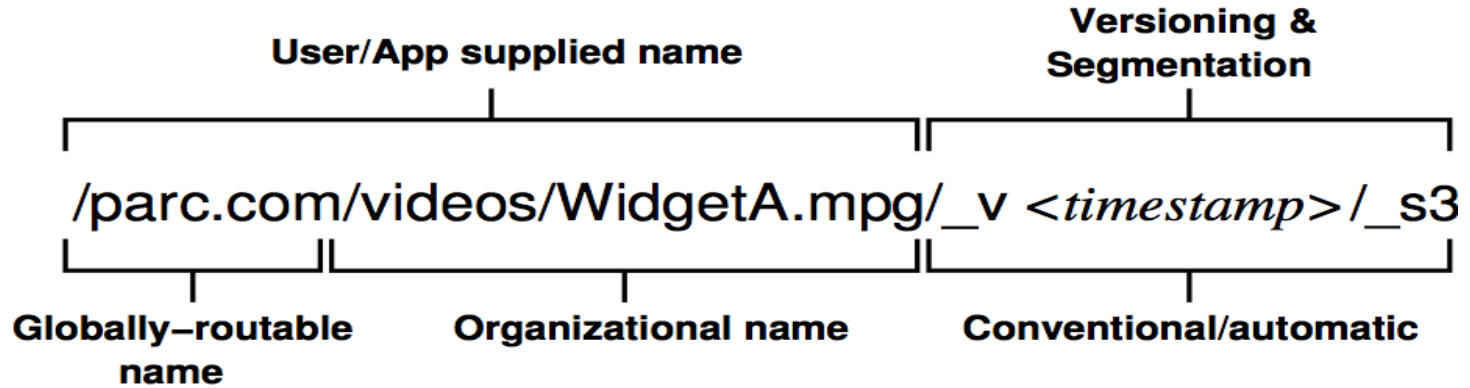


Note: NDN is more than content distribution!

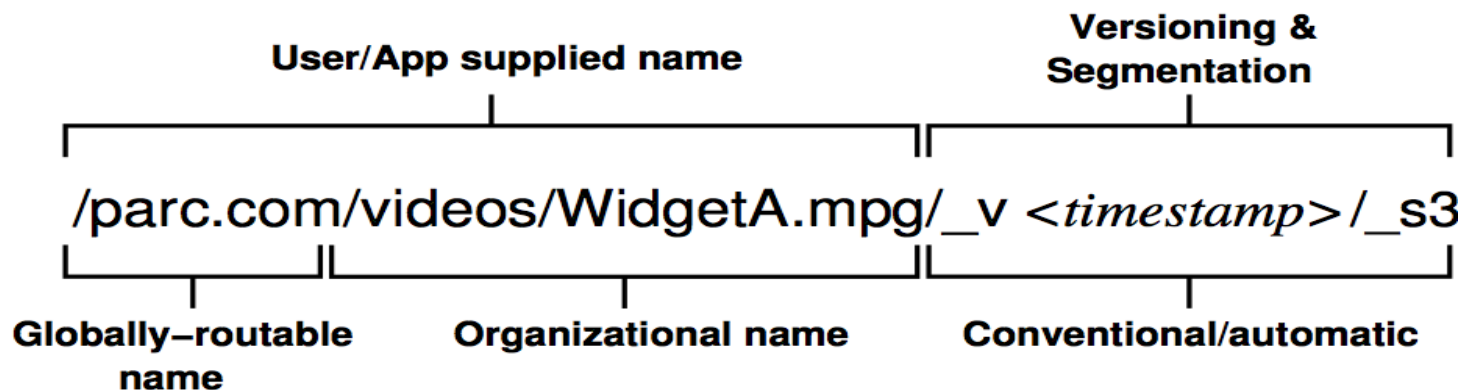
# NDN: Retrieving Named Data




# NDN: Retrieving Named Data



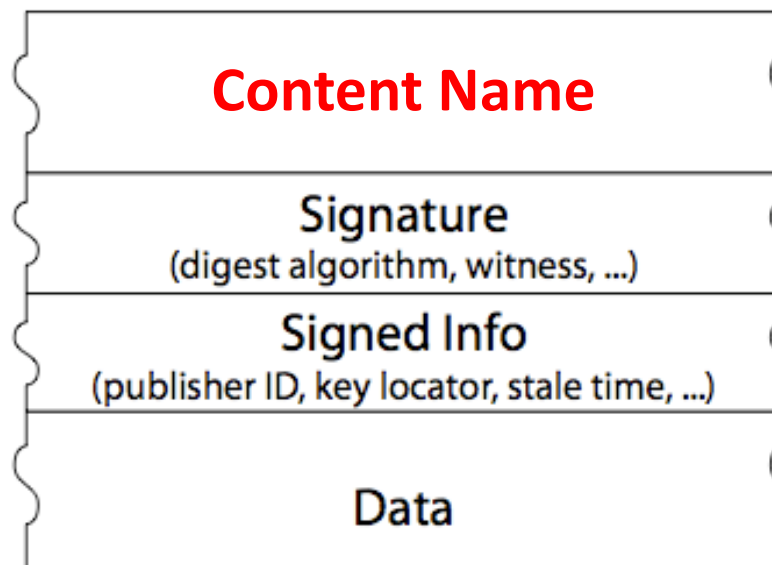
# NDN: Retrieving Named Data



## Interest packet




## Data packet



# Named Data is Easy to Secure

- In the Internet you secure your path..
- ..but the server may still be hacked!
- In NDN you **sign** the data with a **digital signature**..
- ..so the users know when they get bad data!



# Other NDN Benefits

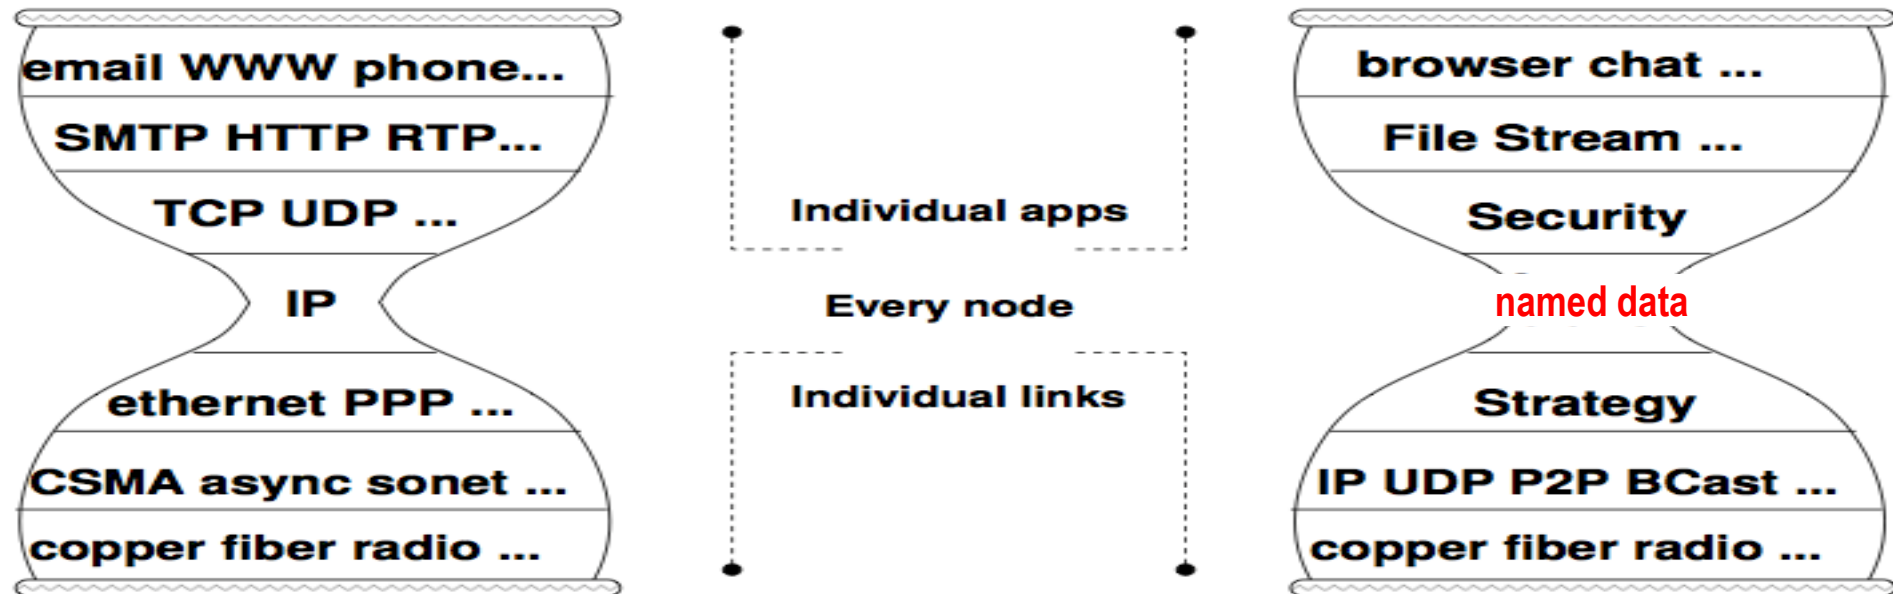
- Built-in multicast delivery
- Loop-free: enabling multipath forwarding
- Feedback loop, at every hop
  - Be able to detect packet delivery problem.
- Flow balance
  - Interest/data one-to-one match
- DoS mitigation

# NDN ARCHITECTURAL DEVELOPMENT [1]

- use applications to drive design and implementation
- test and deploy on operational testbed
- conduct real-world demos

[1] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, kc claffy, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang, Named Data Networking, in *ACM SIGCOMM CCR*, July 2014 (also *NDN Technical Report 0019*)

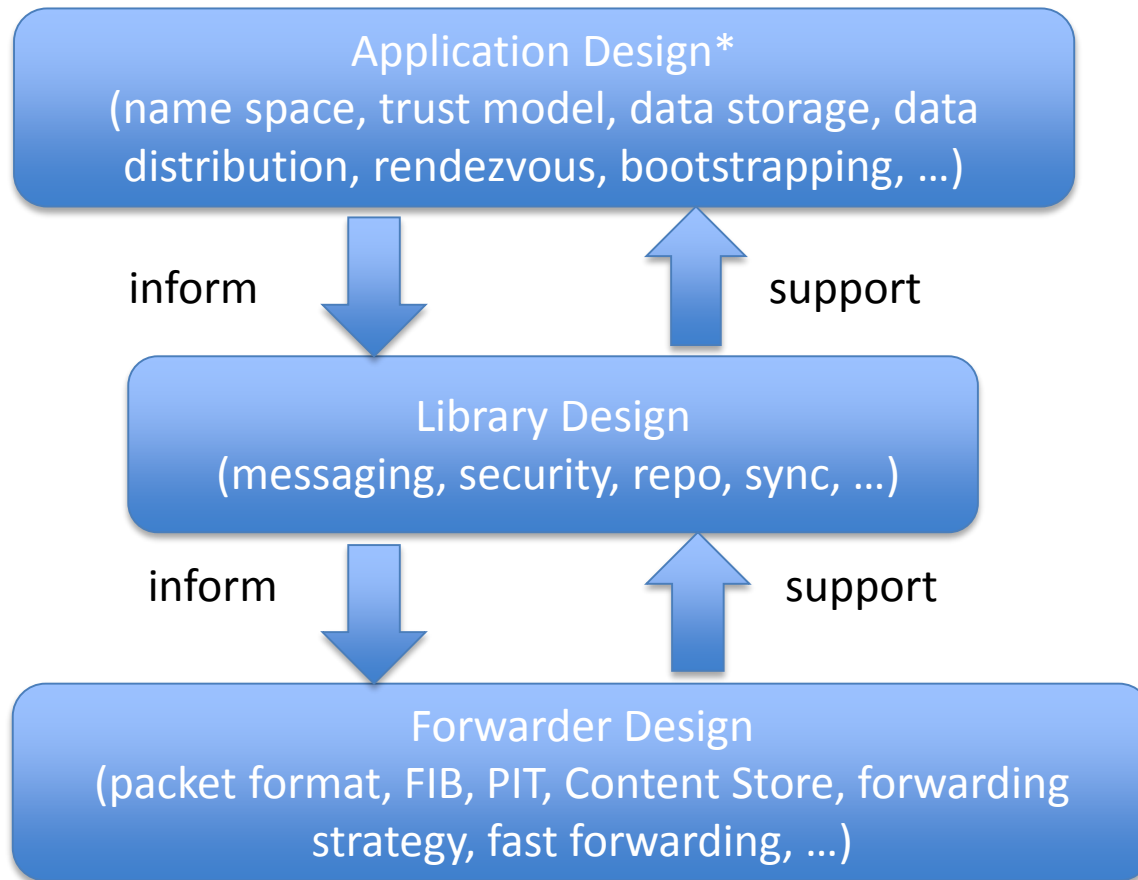
# Preserving the Hourglass Shape



- ◆ The Narrow waist: common interface, the network layer
  - ◆ IP -- address format, IP packet forwarding
  - ◆ NDN – data name format, NDN Interest/Data forwarding



# Development Approach



# NDN Software

# NDN Software

- Multimedia applications
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing

# NDN Software

- Multimedia applications
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing
- IoT applications
  - EBAMS: Building automation and management [5, 6]
  - Vehicular networking [7]
  - NDNFit: exercise/fitness tracking

# NDN Software

- Multimedia applications
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing
- IoT applications
  - EBAMS: Building automation and management [5, 6]
  - Vehicular networking [7]
  - NDNFit: exercise/fitness tracking
- Big data applications
  - Climate science [14, 15]
  - High energy physics [15]

# NDN Software

- Multimedia applications
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing
- IoT applications
  - EBAMS: Building automation and management [5, 6]
  - Vehicular networking [7]
  - NDNFit: exercise/fitness tracking
- Big data applications
  - Climate science [14, 15]
  - High energy physics [15]
- Libraries: NDN-CCL, NDN-CXX, pyNDN2, ndn-js, ChronoSync [9], NDN repo



# NDN Software

- Multimedia applications
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing
- IoT applications
  - EBAMS: Building automation and management [5, 6]
  - Vehicular networking [7]
  - NDNFit: exercise/fitness tracking
- Big data applications
  - Climate science [14, 15]
  - High energy physics [15]
- Libraries: NDN-CCL, NDN-CXX, pyNDN2, ndn-js, ChronoSync [9], NDN repo
- Forwarding and routing
  - NDN Forwarding Daemon (NFD)
  - Named-data Link State Routing [8]

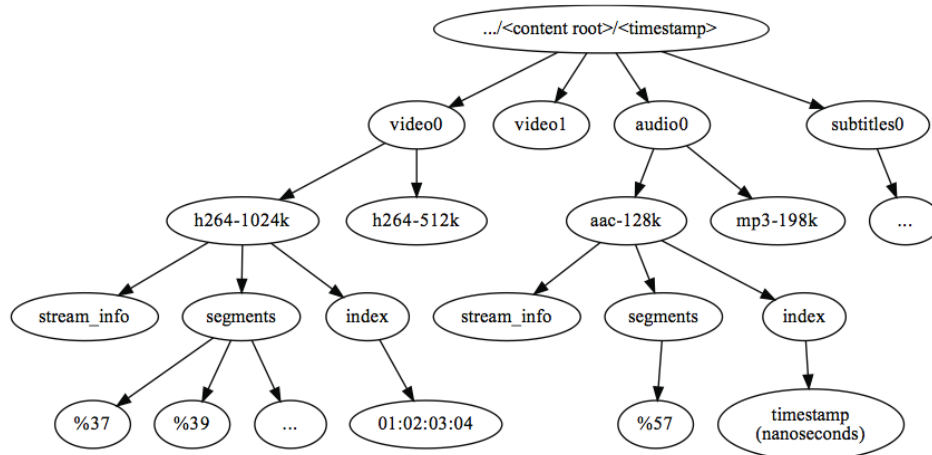
# NDN Software

- **Multimedia applications**
  - NDNVideo [2]: video broadcasting
  - ChronoChat [3]: multiparty chat
  - NDNrtc and NDNCon: realtime conferencing
  - ChronoShare [4]: file sharing
- **IoT applications**
  - EBAMS: Building automation and management [5, 6]
  - Vehicular networking [7]
  - NDNFit: exercise/fitness tracking
- **Big data applications**
  - Climate science [14, 15]
  - High energy physics [15]
- **Libraries: NDN-CCL, NDN-CXX, pyNDN2, ndn-js, ChronoSync [9], NDN repo**
- **Forwarding and routing**
  - NDN Forwarding Daemon (NFD)
  - Named-data Link State Routing [8]


All code is open source at <https://github.com/named-data/>.

# NDNVideo [2]

- Live and pre-recorded streaming to multiple consumers.
- No session semantics => scalability.  
Tested for ~1000 clients from 1 src
- First Interest sent can randomly access a keyframe at any timecode
- Leverages caching.

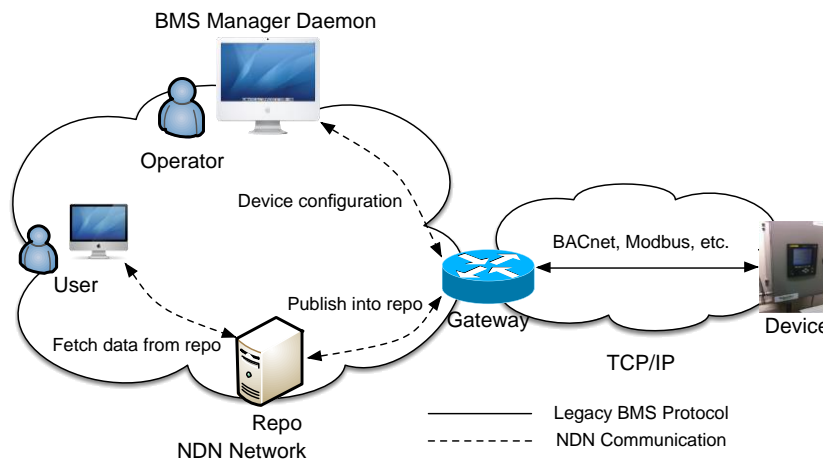


[2] D. Kulinski, J. Burke, and L. Zhang. "Video Streaming over Named Data Networking," *IEEE COMSOC MMTC E-Letter*, 2013.

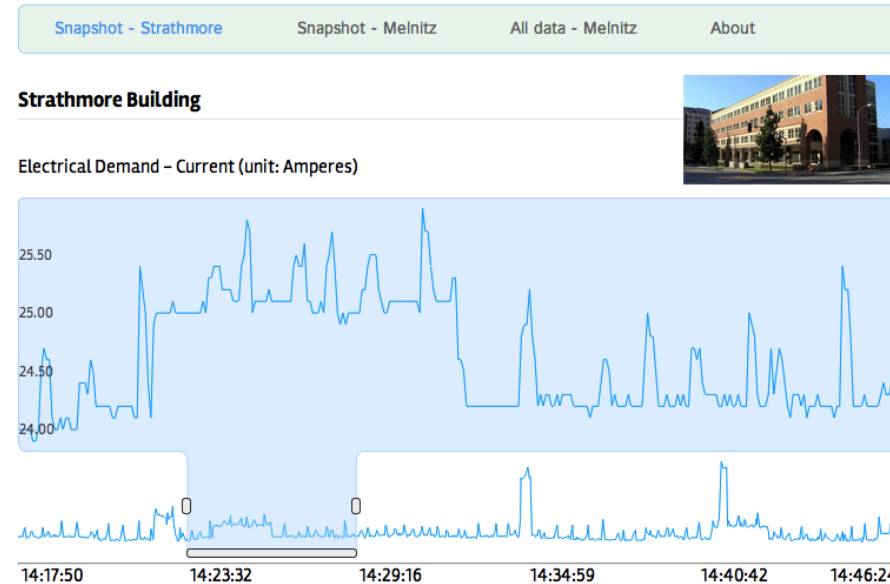


# Building Automation and Management [5,6]

- Improve application development process, management, interoperability and security.
- Practical work so far: NDN interfaces to BacNET and Modbus sensing, authenticated lighting control.
- Partner: UCLA Facilities Management.



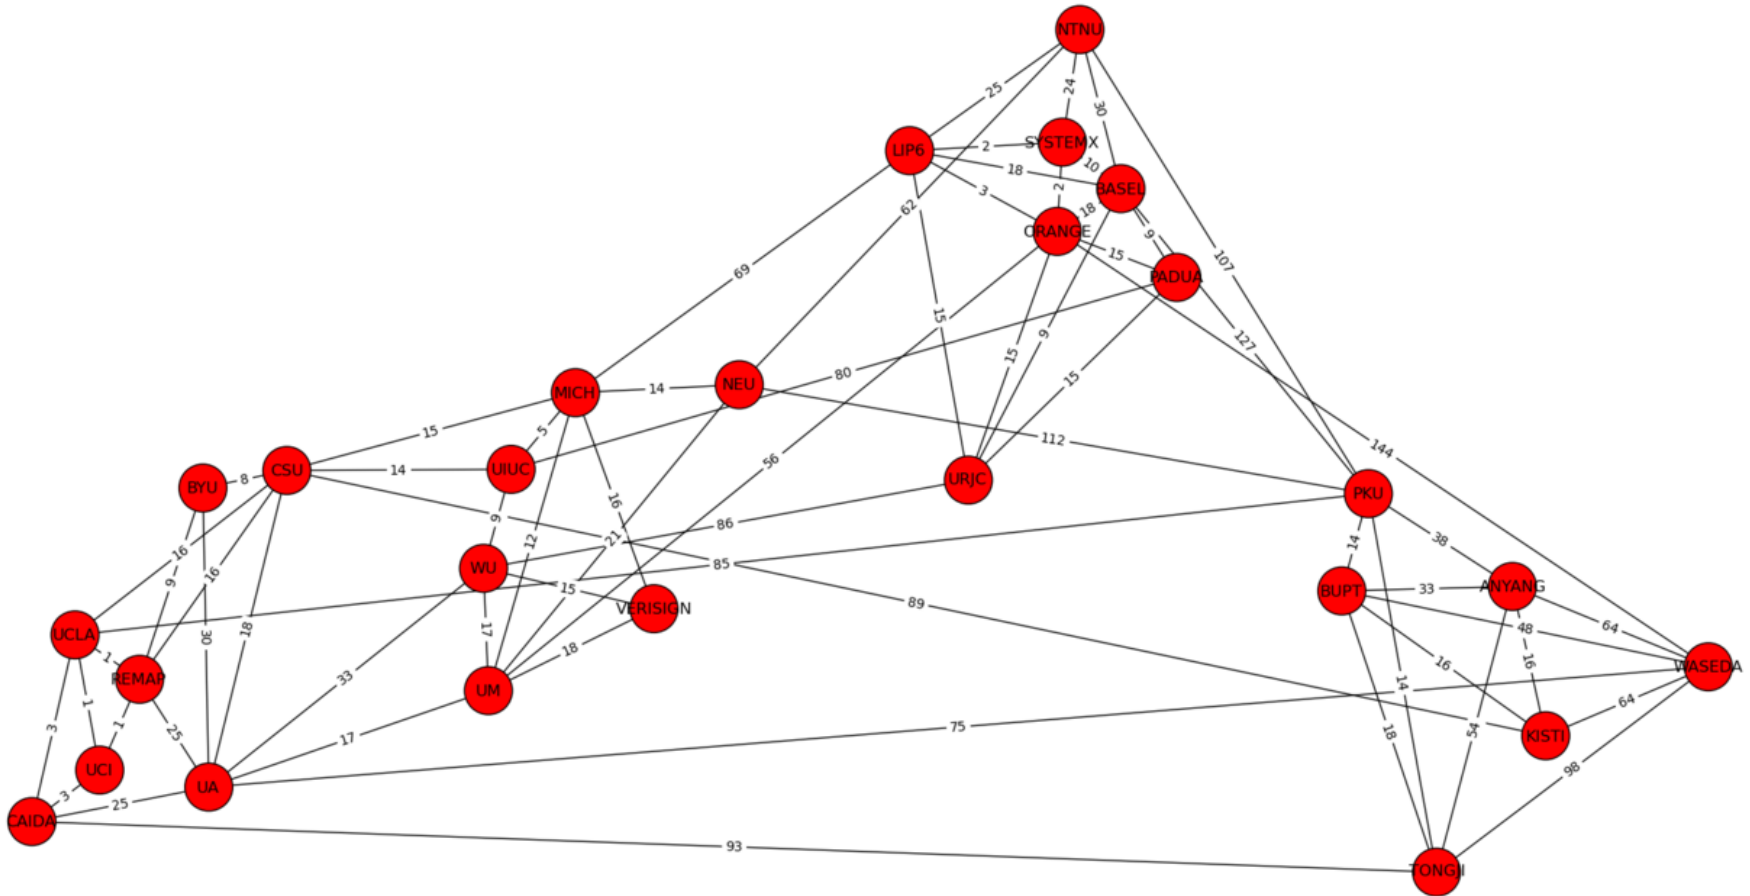
## UCLA NDN Building Monitoring Testbed



[5] Burke et al. Securing instrumented environments over Content-Centric Networking: the case of lighting control. In *IEEE INFOCOM NOMEN Workshop*, Apr. 2013.

[6] Shang et al., "Securing Building Management Systems Using Named Data Networking," *IEEE Network*, May/June 2014.

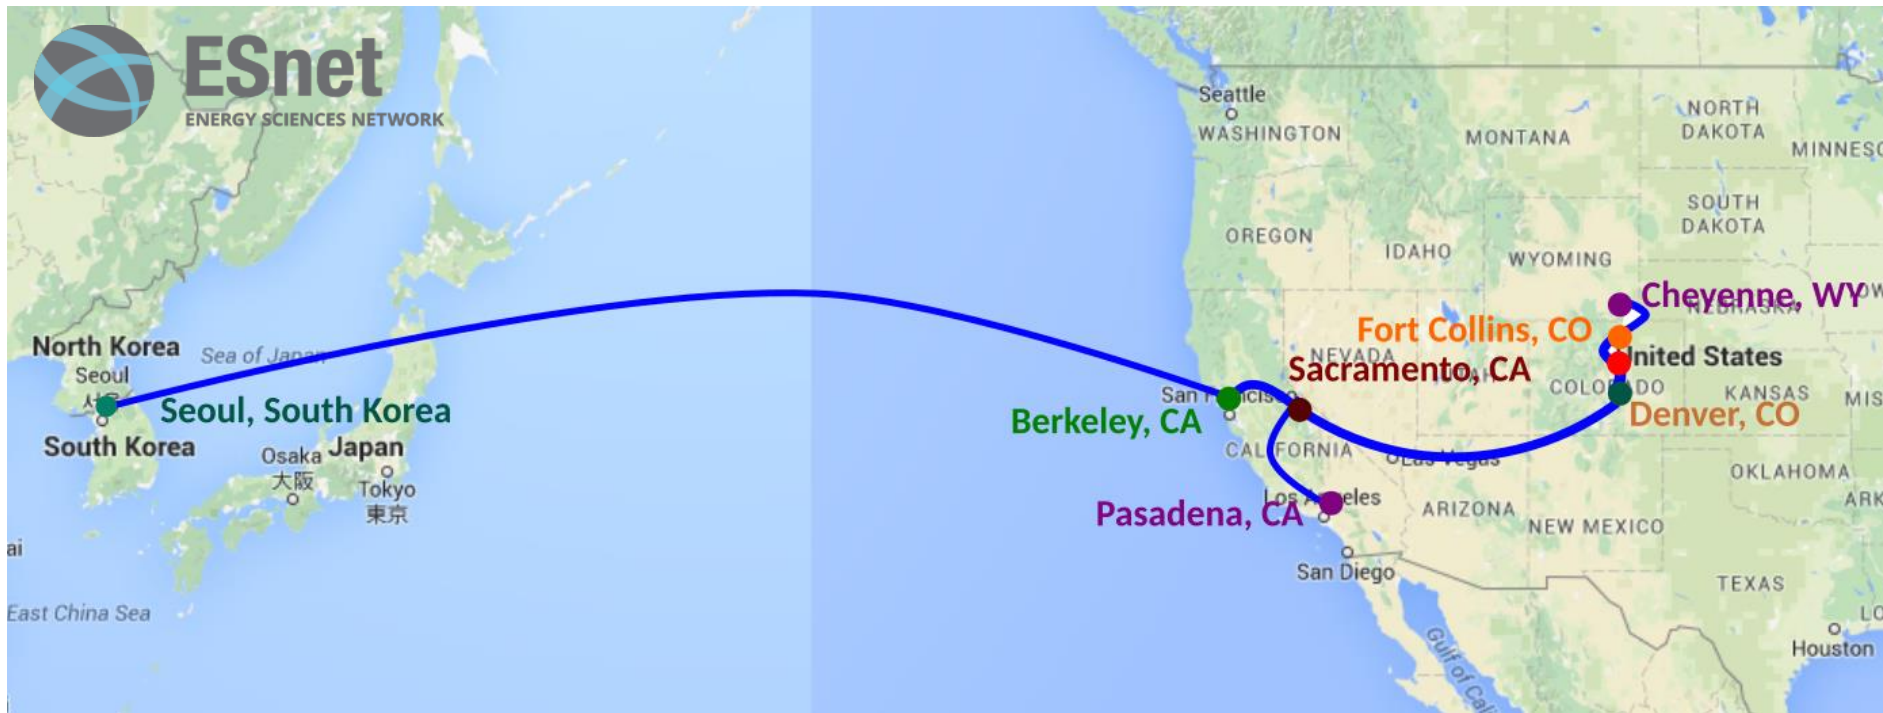
# NDN Testbed



26 nodes, 66 links, more info at <http://named-data.net/ndn-testbed/>

Contact us if interested in joining the testbed.

# Science NDN Testbed



- NSF CC-NIE campus infrastructure award
  - 10G testbed (courtesy of ESnet, UCAR, and CSU Research LAN)
- Currently ~50TB of CMIP5, ~20TB of HEP data




# ROUTING SCALABILITY


- NDN name space is unbounded, though hierarchical
  - $10^8$  DNS names today
- Traditional routing schemes won't scale.
  - Forwarding table (FIB) size
  - Number of routing updates

# Hyperbolic Routing

Hyperbolic coordinates (radius, angle) encode network geometry.



Name Prefix	Next hops
/ndn/D	{A, cost=10} {B, cost=30}


/ndn/D  
  
**D**

Greedy forwarding:


- Find the neighbor closest to the name prefix
- Forward the packet to that neighbor

# Hyperbolic Routing

Hyperbolic coordinates (radius, angle) encode network geometry.



Name Prefix	Next hops
/ndn/D	{A, cost=10} {B, cost=30}


/ndn/D  
  
**D**

Greedy forwarding:

- Find the neighbor closest to the name prefix
- Forward the packet to that neighbor

# Hyperbolic Routing

Hyperbolic coordinates (radius, angle) encode network geometry.




Name Prefix	Next hops
/ndn/D	{A, cost=10} {B, cost=30}

Greedy forwarding:

- Find the neighbor closest to the name prefix
- Forward the packet to that neighbor

# Why Hyperbolic Routing (HR)?

- In the ideal case, no FIB is needed
- Low communication cost  
Few routing updates, as coordinates rarely change
- Drawbacks?
  - Suboptimal paths
  - Local minima
  - Does not react to network dynamics
- **How to mitigate these drawbacks?**  
NDN can try multiple paths and find best path (forwarding strategy)



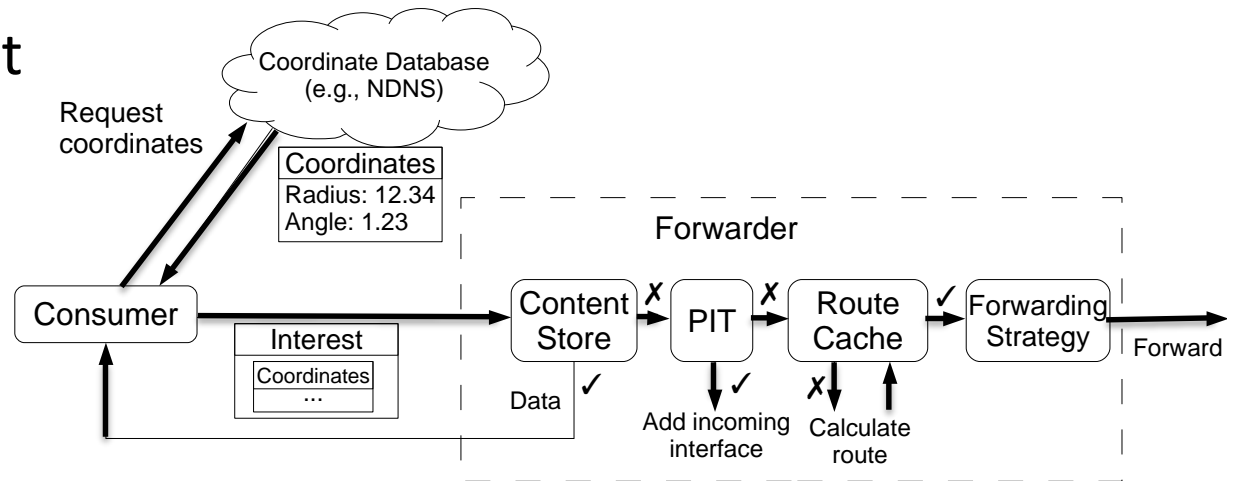
# Forwarding Strategy

- Use Hyperbolic Routing's ranking as a hint, but probe alternative routes periodically
- Adaptive Smoothed RTT-Based Forwarding
  - Best SRTT-Based Forwarding
  - Probabilistic SRTT-Based Probing



# HR Deployment in NDN


- Consumer can fetch coordinates from a distributed database (e.g. NDNS)
- Interest carries name and coordinates
- Forwarder picks next hop using neighbors' distances to coordinates
- **Note:** Name is first matched against CS, so still Data centric



# Evaluation Goals


- We know HR has no FIB and updates, but:
  - Under HR, can forwarding strategy find optimal paths during failures and recoveries?
  - Is performance similar to link-state routing implemented by Named Data Link-State Routing (NLSR)?
  - Is probing overhead less than routing update overhead?
  - Does overhead scale as topology size increases?

# Delay Stretch



Hyperbolic routing/ASF's delay stretch (over Link State Routing) has median close to 1 and 95th-percentile below 2.

# Loss Rate



# Message Overhead

# Nodes	LS Overhead	HR Overhead
22	2.2 pps	0.28 pps
41	7.8 pps	0.28 pps
58	17.5 pps	0.36 pps
78	39.4 pps	0.47 pps

Hyperbolic routing has much lower overhead than link-state routing.

# Conclusion

- NDN is a *data-centric* architecture that names data, not hosts, at the *network* layer
  - Efficient large-scale data distribution and multi-party data sharing
  - Built-in security: signature for provenance, name-based trust model
  - Mobility: data name does not change with location.
- Many playgrounds and opportunities for research

# Research Opportunities

- Big-data and small-data applications
  - hierarchical naming still challenging, not fully expressive, but compromise between expressiveness, scalability, security
- Routing – traditional and coordinate-based (hyperbolic)
- Forwarding strategies
- Caching as an economic problem
- Security, trust models, privacy, encryption, advertising models
- Android, WRT implementations, IoT (light bulbs), Arduino devices
- .. and much more



# References

1. L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, kc claffy, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang, Named Data Networking, to appear in *ACM SIGCOMM CCR* (also *NDN Technical Report 0019*)
2. D. Kulinski, J. Burke, L. Zhang. Video Streaming over Named Data Networking. *IEEE COMSOC MMTC E-Letter*, 2013.
3. Z. Zhu, C. Bian, A. Afanasyev, V. Jacobson, and L. Zhang. Chronos: Serverless multi-user chat over NDN. *Technical Report NDN-0008*, NDN Project, October 2012.
4. A. Afanasyev, Z. Zhu, L. Zhang, The story of ChronoShare, or how NDN brought distributed file sharing back, under review
5. J. Burke, P. Gasti, N. Nathan, and G. Tsudik. Securing instrumented environments over Content-Centric Networking: the case of lighting control. In *IEEE INFOCOM NOMEN Workshop*, Apr. 2013.
6. W. Shang, Q. Ding, A. Marianantoni, J. Burke, and L. Zhang. Securing building management systems using named data networking. *IEEE Network Special Issue on Information-Centric Networking*, April 2014.
7. G. Grassi, D. Pesavento, G. Pau, R. Vuyyuru, R. Wakikawa, and L. Zhang. VANET via Named Data Networking. In *IEEE INFOCOM NOMEN Workshop*, Apr. 2014.
8. A. Hoque, S. O. Amin, A. Alyyan, B. Zhang, L. Zhang, and L. Wang. Named-data link state routing protocol. In *ACM SIGCOMM ICN Workshop*, 2013.

# References (cont'd)

9. Z. Zhu, A. Afanasyev, and L. Zhang. Let's ChronoSync: Decentralized dataset state synchronization in NDN. In *ICNP*, 2013.
10. H. Yuan, T. Song, and P. Crowley. Scalable NDN forwarding: Concepts, issues and principles. In *ICCCN*, 2012.
11. H. Yuan and P. Crowley. Scalable pending Interest table design: From principles to practice. *IEEE INFOCOM*, 2014.
12. W. So, A. Narayanan, and D. Oran. Named data networking on a router: Fast and DoS-resistant forwarding with hash tables. In *ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, Oct 2013.
13. M. Varvello, D. Perino, and J. Esteban. Caesar: A content router for high speed forwarding. In *ACM SICCMM Workshop on Information-centric Networking*, 2012.
14. C. Olschanowsky, S. Shannigrahi, C. Papadopoulos. Supporting climate research using named data networking. In *IEEE International Workshop on Local & Metropolitan Area Networks (LANMAN)*, 2014
15. S. Shannigrahi *et. al.*. Named Data Networking in Climate Research and HEP Applications. In *21 International Conference on Computing in High Energy and Nuclear Physics*, 2015