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

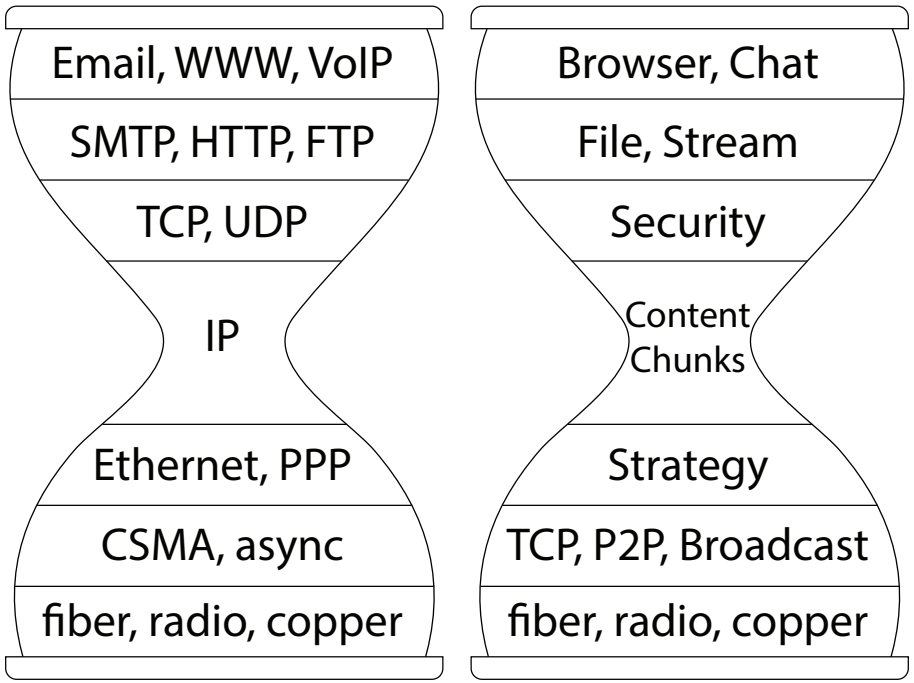

CCN: Naming & Packets

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

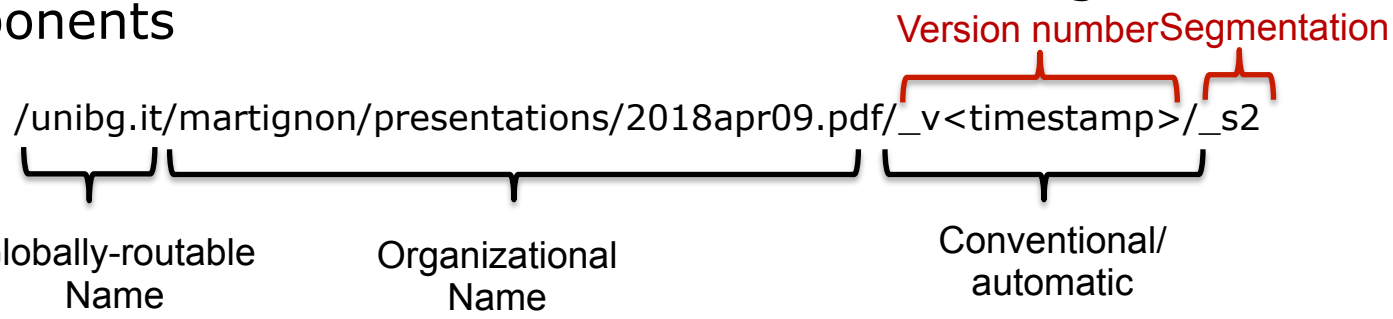

- Two type of packets are envisioned:
  - Interest Packet
  - Data Packet

Interest packet

Content Name
Selector (order preference, publisher filter, scope, ...)
Nonce

Data packet

Content Name
Signature (digest algorithm, witness, ...)
Signed Info (publisher ID, key locator, stale time, ...)
Data

Reliability and Flow Control

- In CCN, **one Interest** retrieves ***at most* one Data packet**.
- This basic rule ensures that *flow balance* is maintained in the network and allows efficient communication between varied machines over networks of widely different speeds.

- Just as in TCP, however, it is possible to **overlap** data and requests.

- **Multiple Interests** may be issued at once, before Data arrives to consume the first.
  - Hence, the Interests serve the role of window advertisements in TCP. A recipient can dynamically vary the window size by varying the Interests that it issues.
  - Since CCN packets are independently named, the pipeline does not stall on a loss – the equivalent of TCP SACK is intrinsic.

Van Jacobson et Al.,"Networking Named Content", CoNEXT'09

Reliability and Flow Control

- In a large network, the end-to-end nature of TCP conversations means there are **many points** (bottlenecks, nodes…) between the sender and the receiver where **congestion** can occur from conversation (flow) aggregation even though each conversation is operating in flow balance.
- The effect of this congestion is delay and packet loss.

- The TCP solution is for end-points(TCP sources) to dynamically adjust their window sizes to keep the aggregate traffic volume below the level where congestion occurs.
- The need for this congestion control is a result of TCP's flow balance being end-to-end.

Van Jacobson et Al.,"Networking Named Content", CoNEXT'09

Reliability and Flow Control

- In CCN, on the other hand, all communication is local so there are *no points* between sender and receiver that are *not* involved in their balance.
- Since CCN flow balance is maintained at each hop, there is no need for additional techniques to control congestion in the middle of a path.

- This is not the same as hop-by-hop flow control, where backpressure between adjacent nodes is used to adjust resource sharing among continuous flows. CCN does not have FIFO queues between links but rather an LRU memory (the cache) which decouples the hop-by-hop feedback control loops and damps oscillations.

Van Jacobson et Al.,"Networking Named Content", CoNEXT'09