# Caching and Mobility Support in a Publish-Subscribe Internet Architecture

George Xylomenos, Xenofon Vasilakos, Christos Tsilopoulos, Vasilios A. Siris, and George C. Polyzos

Presented by Max Nicosia

# Motivation

- Need for Mobility
- Change in usage model
  - From sharing Resources -> Info (regardless of end-point hosting it)

PSI claims:

- Scaling not affected by addressing or routing assumptions
- Supports native multicasting
- Asynchronous nature and caching help with MN handoff

# Architecture Overview

Naming:

- Scope Identifier (SID)
- Rendezvous Identifier (RID)

Network Primitives:

- Publish
- Subscribe

# Architecture Overview (cont)

Functions (3):

- Rendezvous
- Topology Management
- Forwarding

# Example Retrieval

Steps:

1. Producer publishes item
2. Consumer subscribes
3. Subscription is propagated to the Rendezvous Function
4. Topology Manager constructs forwarding path -> Forwarding Identifier (FID) and returns it to Publisher (notification)
5. Publisher starts transmission

# Network Design

4 Types of Elements:

1. Hosts
2. Forwarding Nodes (FNs)
3. Rendezvous Nodes (RNs)
4. Topology Managers (TNs)

Key Benefits (Hosts behind gateway RN):

- Increased security and privacy
- Reduced control overhead and rendezvous traffic

# Forwarding

- Using LIPSIN
  - On-path links encoded in fixed size bit string (Bloom Filter) -> FID
  - Included in packet
- Packets can be delivered over multicast trees by adding tree links to the FID
  - No extra state in the network is required
- Topology Managers can enforce routing policies

# RENE and Topology Management

- Made out of multiple RNs
  - Built from FNs and RNs talking to TMs
  - Link State Announcements (LSA)
- Topologies built used for intra-domain routing
- TMs send FIDs with paths

# Transport

- RNs maintain list of subscribers (state)
- New subscribers are added to FID
- Documents (reliable delivery)
  - TMs compute bidirectional FIDs
- Push-based

- Pull-based

# Caching and Replication

- On-path Caching
- Off-path Caching
- Content Replication

# On-path Caching

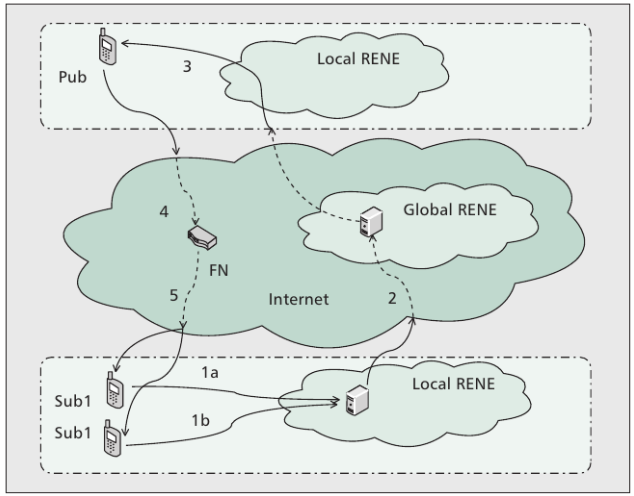


- Requests can be served by FNs on the path
- Prevent feedback implosion of multicast errors
- Possible to use on the Internet at the Application layer
  - Requires proxy server for Browser

# Off-path Caching

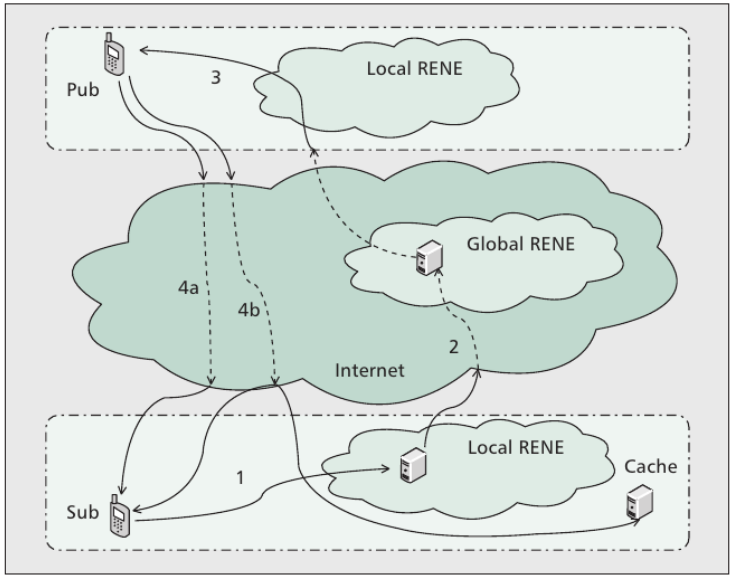


- Take advantage that any node can serve request
- RENE can decides to cache or not to cache info
- Additional overhead -> announce to local RENE
- Hard to achieve on the Internet as it needs coordination between various application-level caches

# Content Replication

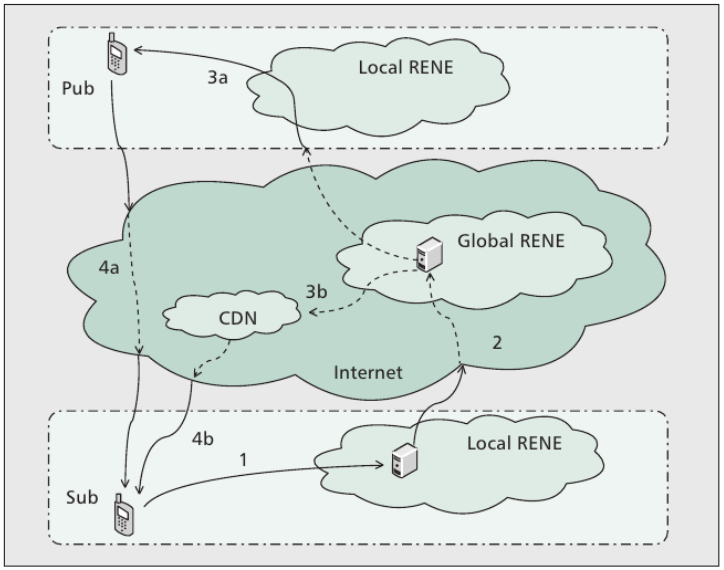


- Used to minimise inter-domain traffic and improve user experience
- Does not rely on DNS tricks like on the Internet
- Can also exploit off-path caching

# Mobility and Caching

- Decoupling from Time and Space
- Re-issue subscriptions for missed data
  - RENE takes care of redirection to caches

- Micro-mobility:
  - Use off-path caches to support MNs -> requests forwarded after handoff to cache
  - Assigned based on topology or forecasting
- Macro-mobility:
  - Set caches in neighbouring networks (expensive)

# Conclusions

- Information-centric architecture
- Individually identified items can be cached at Network Layer
- Asynchronous nature of PSI make it mobile-ready
- Constantly evolving architecture

# Critique

- No performance analysis about the impact on the network on various conditions
  - Possibly not very applicable in many cases
- Privacy concerns about data being replicated without any control

# Questions?