UCL DEPARTMENT OF ELECTRONIC AND ELECTRICAL ENGINEERING
COMMUNICATIONS AND INFORMATION SYSTEMS GROUP

IEEE ISCC 2011 Keynote

Information-Centric Networking: Overview, Current State and Key Challenges

Prof. George Pavlou
http://www.ee.ucl.ac.uk/~gpavlou/
Communications and Information Systems Group
Dept of Electronic & Electrical Engineering
University College London, UK

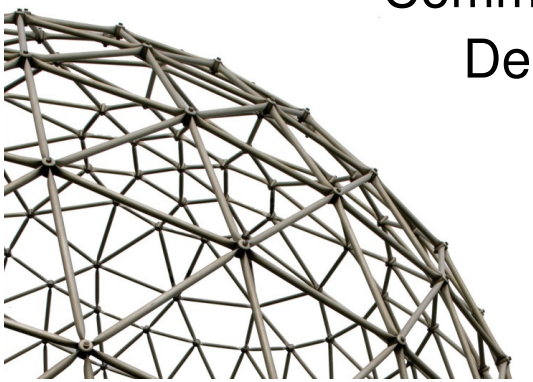

Internet-based Content

- The Internet plays a central role in our society
  - Work and business, education, entertainment, social life, …
- The vast majority of interactions relate to content access
  - P2P overlays (e.g. BitTorrent, eMule, live streaming)
  - Media aggregators (e.g. YouTube, GoogleVideo)
  - Over-the-top video (e.g. Hulu, iPlayer)
  - Content Delivery Networks (e.g. Akamai, Limelight)
  - Social Networks (e.g. Facebook, MySpace)
  - Photo sharing sites (e.g. Picasa, Flickr)
- New approaches are required to cater for the explosion of video-based content and for creating novel use experiences
- **Continue throwing more capacity cannot work anymore!**

The Emerging Content-Oriented Internet

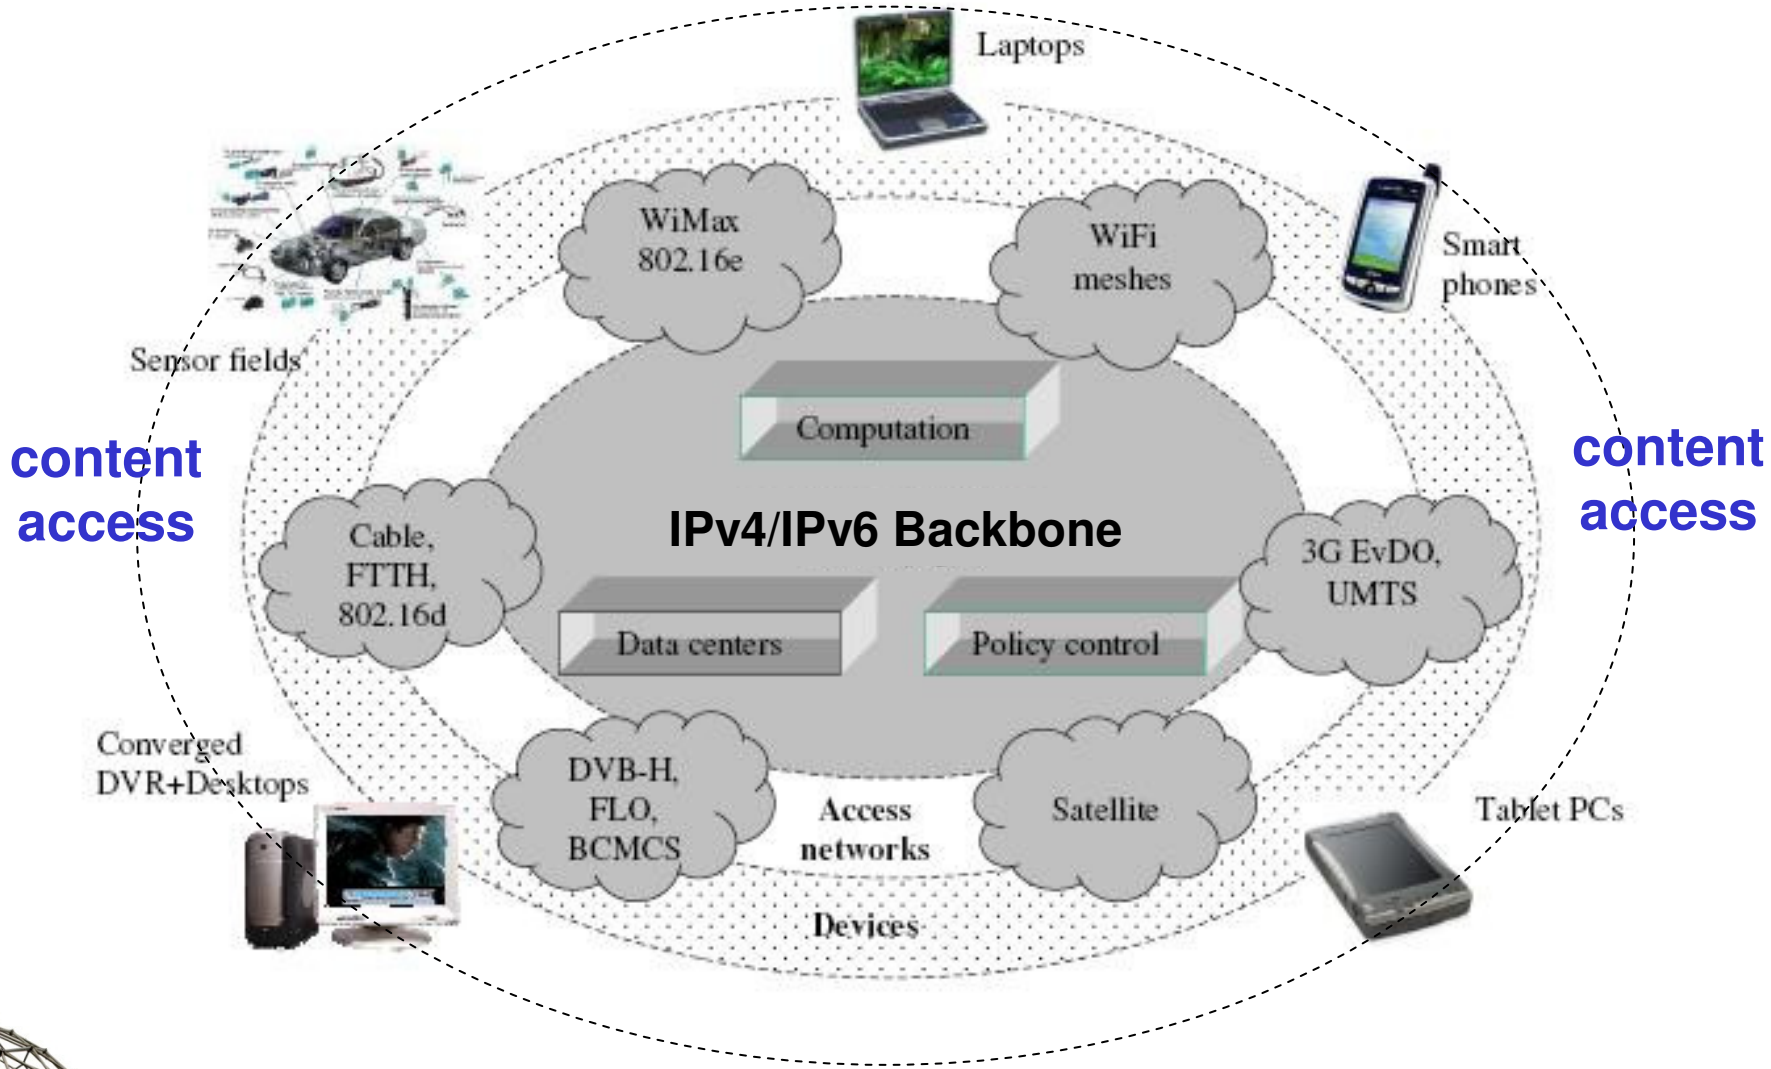

Expected IP Traffic Growth 2009-2014

- According to the Cisco Visual Networking Index 2010:
  - Global IP traffic will quadruple every year until 2014
  - 64 exabytes per month is expected by 2014
  - Global Internet video traffic will surpass P2P traffic in 2010
  - Approx. 55% of the overall Internet traffic will be video by 2014
  - Global mobile data traffic will double every year until 2014
  - Approx. 65% of the overall mobile traffic will be video by 2014
- It will take over 2 years to watch the amount of video that will cross global IP networks every second in 2014!
- **Infrastructure evolution needs to be partnered with novel approaches and associated business models**

Expected IP Traffic Growth 2009-2014 (cont'd)

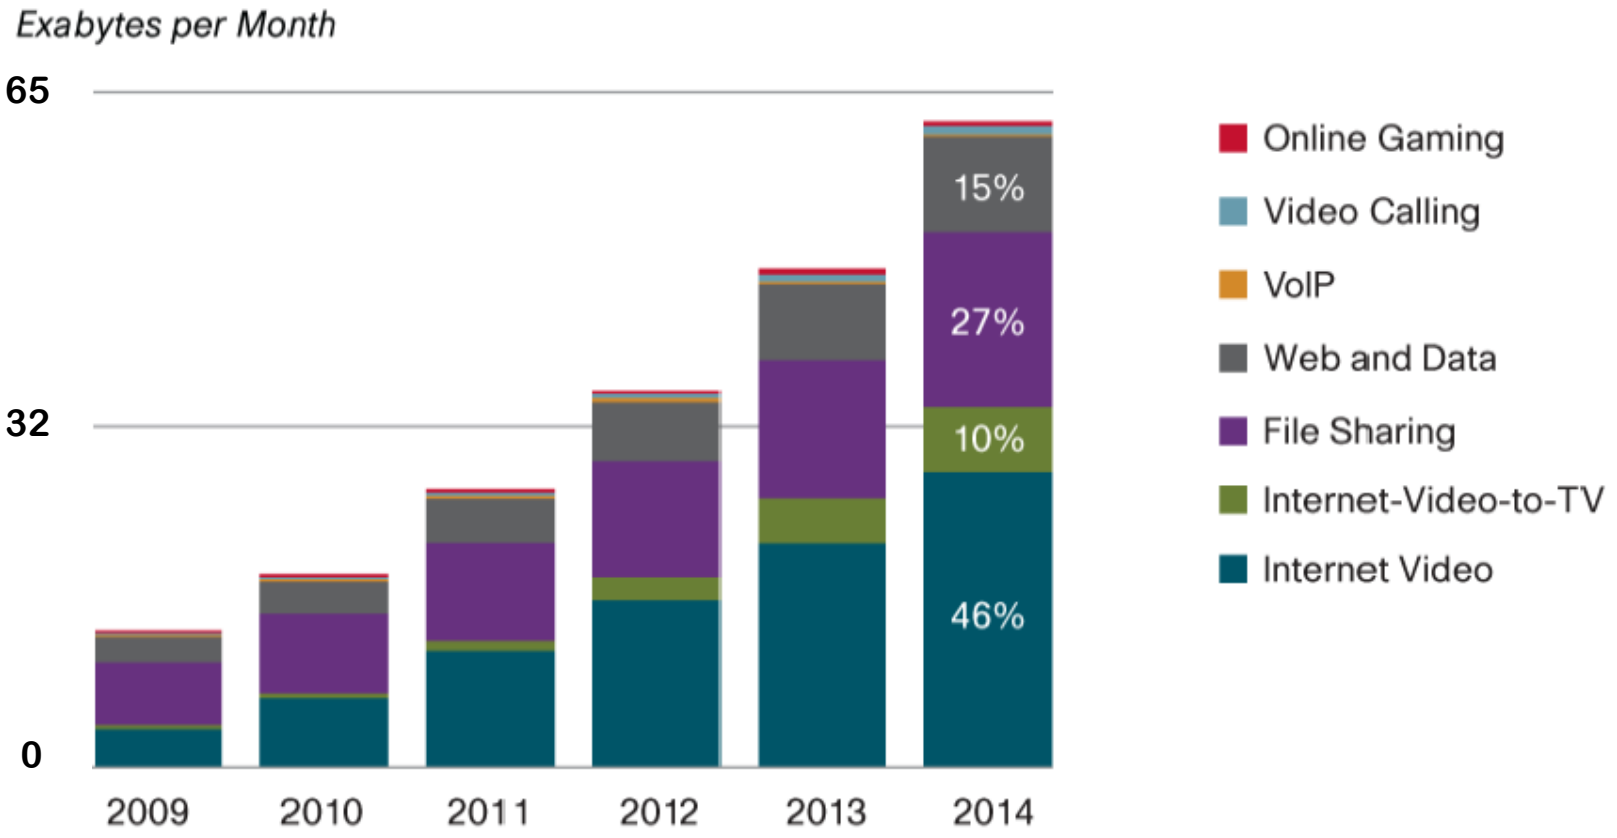

P2P Overlays and CDNs

- Peer-to-Peer (P2P) Overlays: started from file sharing and evolved to multicast-streaming real-time video through overlay nodes
  - Self-organized, adaptive, fault-tolerant content distribution
  - Content object names are resolved to candidate peers
- Content Distribution Networks (CDNs): pioneered by Akamai, they support anycast by choosing the most appropriate (i.e. topologically close) content replica to maximise user QoE
  - Use DNS-based redirection
  - Mostly offline content replica placement based approach

- Both P2P overlays and CDNs make the content server transparent for accessing "named content", allowing access to cached copies
  - A first step towards an information-oriented communication model

Current Content Naming and Security Problems

- Content URIs are effectively object locators, resolving to the IP address of the hosting server i.e. location-dependent
  - Binding breaks when object moves or when site changes domain
  - Replicas all have different URIs, appearing as different objects
  - Unique, persistent, location-transparent naming is required

- The current Internet security model provides connection endpoint as opposed to content object authentication
  - Once an object copy has left the origin server, its authenticity cannot be verified anymore, which is a problem for caching
  - In an information-centric approach it is important to be able to authenticate content objects as opposed to connection endpoints

Current Paradigm Shift

Node-centric design: sharing network resources

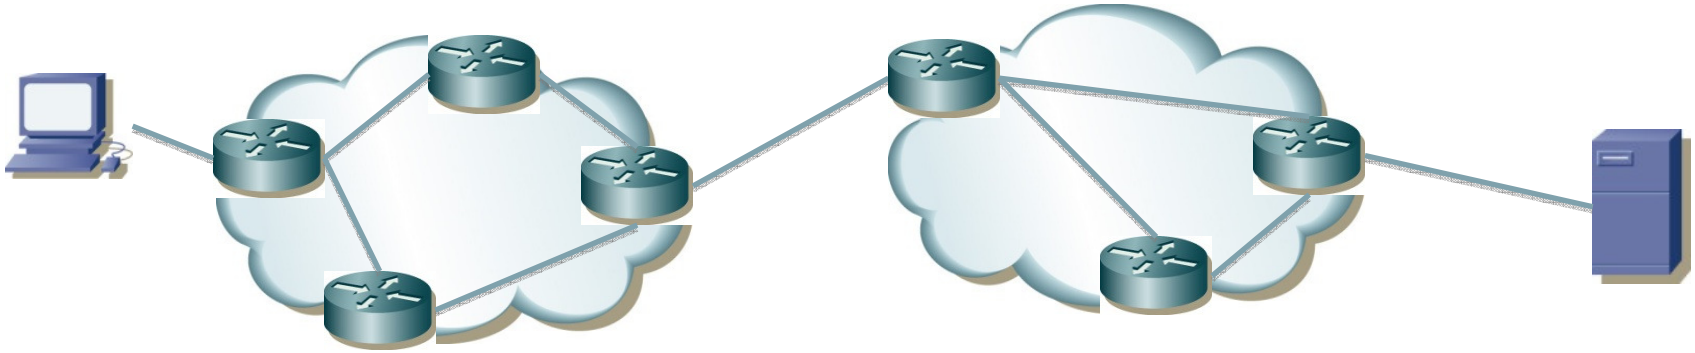

Information-centric design: content access and distribution

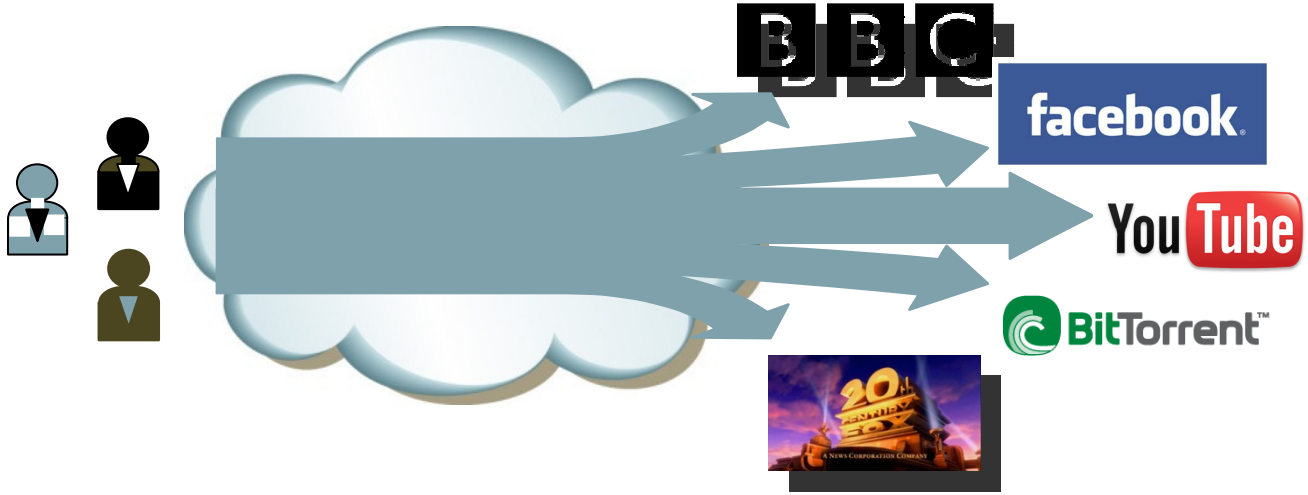

Information-Centric Networking

- Given that users are interested in named content and not in node endpoints, is there a clean architectural approach to address the relevant requirements?
  - All encompassing instead of add-ons to specific domains
  - Provide an enhanced P2P/CDN-like paradigm within the network

- **Information-Centric Networking (ICN)** targets general infrastructure that provides in-network caching so that content is distributed in a scalable, cost-efficient & secure manner
  - Receiver-driven model – subscribe/get objects of interest
  - Support for location transparency, mobility & intermittent connectivity
  - Needs also to be able to support interactivity (e.g. voice) and node-oriented services (e.g. telnet)

Flash-Crowd Effect Due to Content Popularity

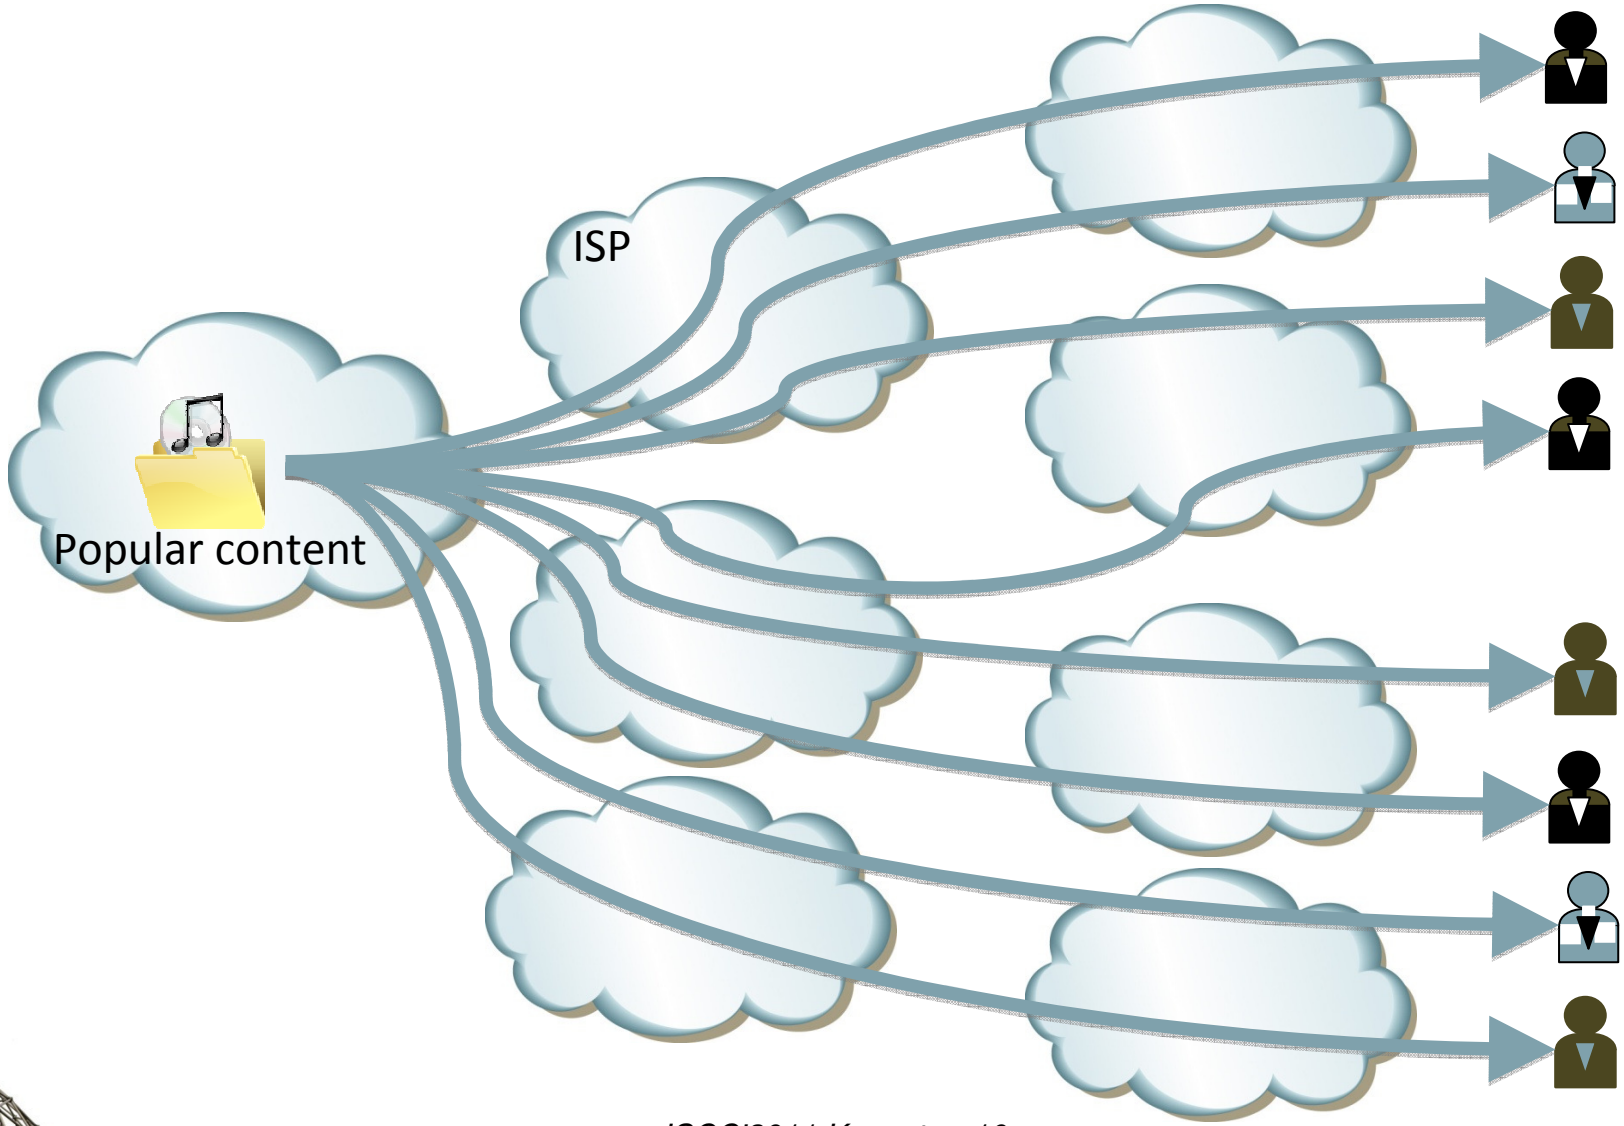

Scalable Cache-based Content Distribution

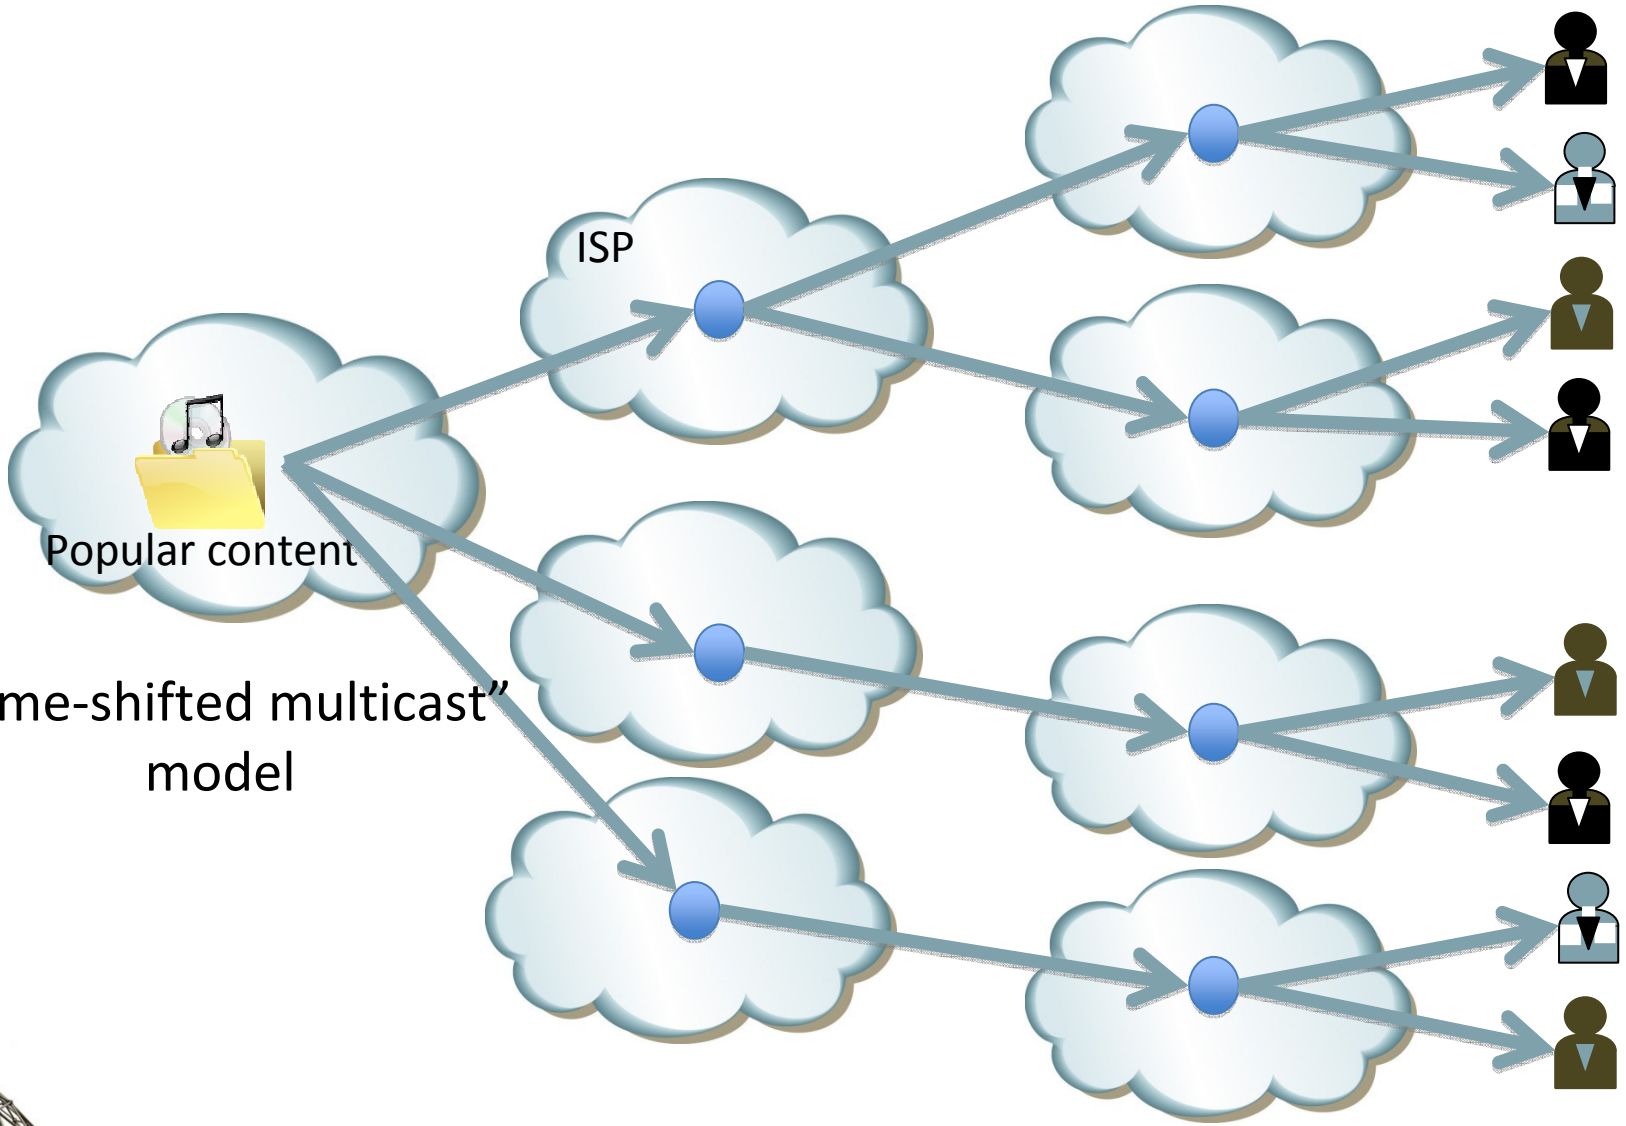

"Time-shifted multicast" model

Caching Approaches

- Two general approaches: offline proactive (as in CDNs) and dynamic reactive (as in P2P overlays)
- Different options for the granularity of caching:
  - Object-level: caching whole information objects
  - Chunk-level: caching information chunks
  - Packet-level: caching individual packets (yes, this is a possibility!)

- Coordinated intelligent decision making is required w.r.t. what/where to cache/drop for maximizing gain

Information Objects

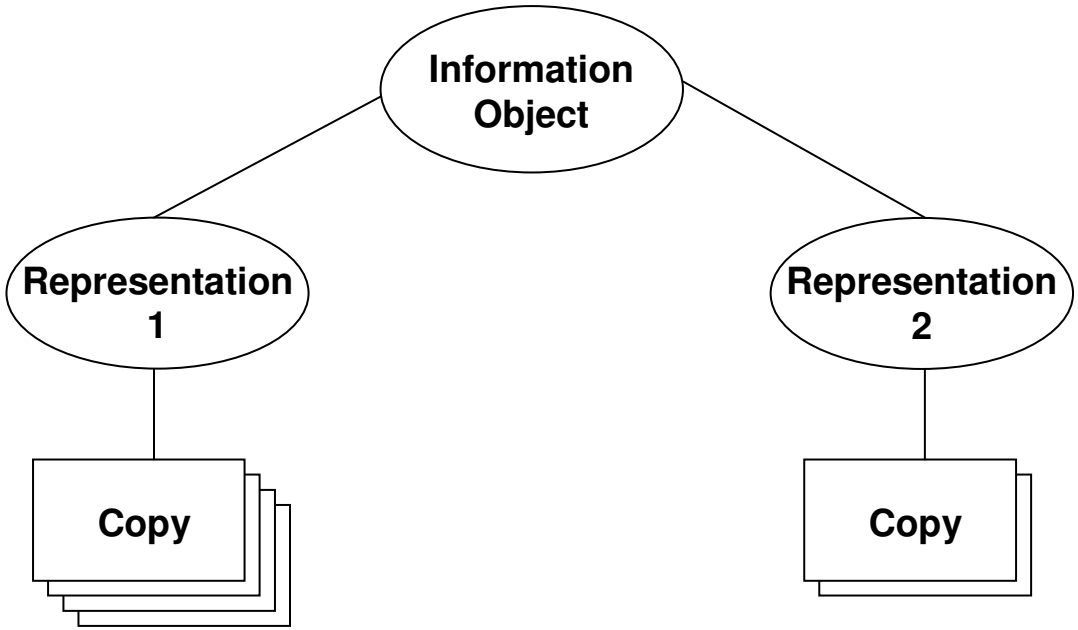

Relationship between information object, its representations and copies of the latter – all these share the same ID

Content Naming Issues

- Information objects are identified by location-independent IDs, with all the object copies sharing a unique ID
- Given that in ICN security applies to information, object IDs in many ICN architectures incorporate security
  - Non human-friendly IDs
  - Human-friendly names can also be associated with IDs
- Flat, hierarchical or combined ID schemes

- Scalability a concern in particular for flat naming schemes

Naming Scalability

- A vast amount of information objects
  - Currently more than 1 trillion unique URLs (Google 2008)
  - 26 billion web pages (www.worldwidewebsize.com)
  - 119 million 2nd level domain names in the DNS (end of 2010)
- Possible to operate DHTs with >2 million nodes
  - For 1000 trillion objects (2^{15}) with 100 bytes per record and no replication, 50Gb of DRAM is necessary
  - With 10 times replication and 1Kb per record 5Tb of RAM is necessary and can be supported with SSD, albeit expensively
  - 4WARD/SAIL experiments indicate 100ms per resolution is possible
- Aggregation at the publisher level may possibly allow a DNS-like solution

Name Resolution and Routing Issues

- Two general approaches: two-phase and one-phase
  - Approach heavily dependent on namespace/ID properties
- In the two-phase approach, name resolution takes place first by mapping the ID to locators, with the most suitable one selected (anycast)
  - Content name resolution servers are required e.g. DNS++
  - Routing to the content source and subsequent content delivery simply use locators i.e. IP addresses
  - The locator is typically not visible to the application which uses a *Get(ID)* API abstraction

Name Resolution and Routing Issues (cont'd)

- In the one-phase approach, in-network content ID-based routing to the source is used
  - Content-ID based routing uses a "structured" ID, content state in the network ("breadcrumps") and includes anycast
- The content delivery path can be the reverse path of the request or (user) ID-based routing can be used

- Different characteristics of the two approaches:
  - The two-phase one can be incrementally deployed over the current Internet given that locator-based routing is used
  - The one-phase ID-based routing is radical

Application Programming Interface

- All ICN approaches use information-centric APIs
  - Location-independence a key feature
  - A pull-based or receiver-driven approach
- *Get(ID)* and *Put(ID)* are the key primitives
  - Get/Subscribe can even request content of certain type which is not yet in place
  - Put/Publish places content in the global content space

- The publish/subscribe API semantics has led researchers to also consider a pub/sub routing paradigm

Key Projects

- UCB **DONA** - Data-Oriented Network Architecture
- 4WARD/SAIL **NetInf** - Network of Information
- PSIRP/PURSUIT **PubSub** - Publish Subscribe Routing
- Xerox PARC **CCN** - Content-Centric Networking
- COMET **CMP** - Content Mediation Plane

- Also other projects and research efforts worldwide

Data-Oriented Network Architecture (DONA)

- Originated at University of California Berkeley
  - Follow on to the Routing on Flat Labels (ROFL) first effort
- One-phase approach through Resolution Handlers (RHs) that exhibit a hierarchical structure
  - IDs are also hierarchical and incorporate security
  - *Query*/*Response* packets, with the closest object copy returned
  - In pure data-oriented fashion, content delivery uses the reverse path
- DONA was the first ICN approach and has had significant influence on other approaches

Network of Information (NetInf)

- Started in the EU project 4WARD and is currently continued in the follow-on project SAIL
- Both one-phase and two-phase approaches
  - One-phase approach uses a hierarchy of DHTs
  - Two-phase approach uses “late locator construction” that targets dynamic environments with high mobility
  - Cache-aware transport protocol

- Significant European industry support

Publish Subscribe Routing

- Started in the EU project PSIRP and is currently continued in the follow-on project PURSUIT
- Two-phase resolve/retrieve model but a radical revolutionary approach
  - Resolvers are called Rendezvous points
  - After content matching resolves to a rendezvous ID, *Subscription*/*Data* packets fetch the content
  - Data packets use source routing with Bloom filters
- A high-level data-oriented architecture with potentially different instantiations (two current implementations)

Content-Centric Networking (CCN)

- Originated by Van Jacobson
- One-phase approach through *Interest*/*Data* packets flowing in a "reverse ack/data TCP-style"
  - Data packets are cached everywhere along the delivery path as they may be useful to other consumers
  - Least Recently Used (LRU) packet discard policy implements the "time-shifted multicast"
  - Hierarchical naming scheme

- CCNx implementation is publicly available while the recently started NSF NDN project looks at more general CCN-related research issues

CCN-like In-Network Content Caching

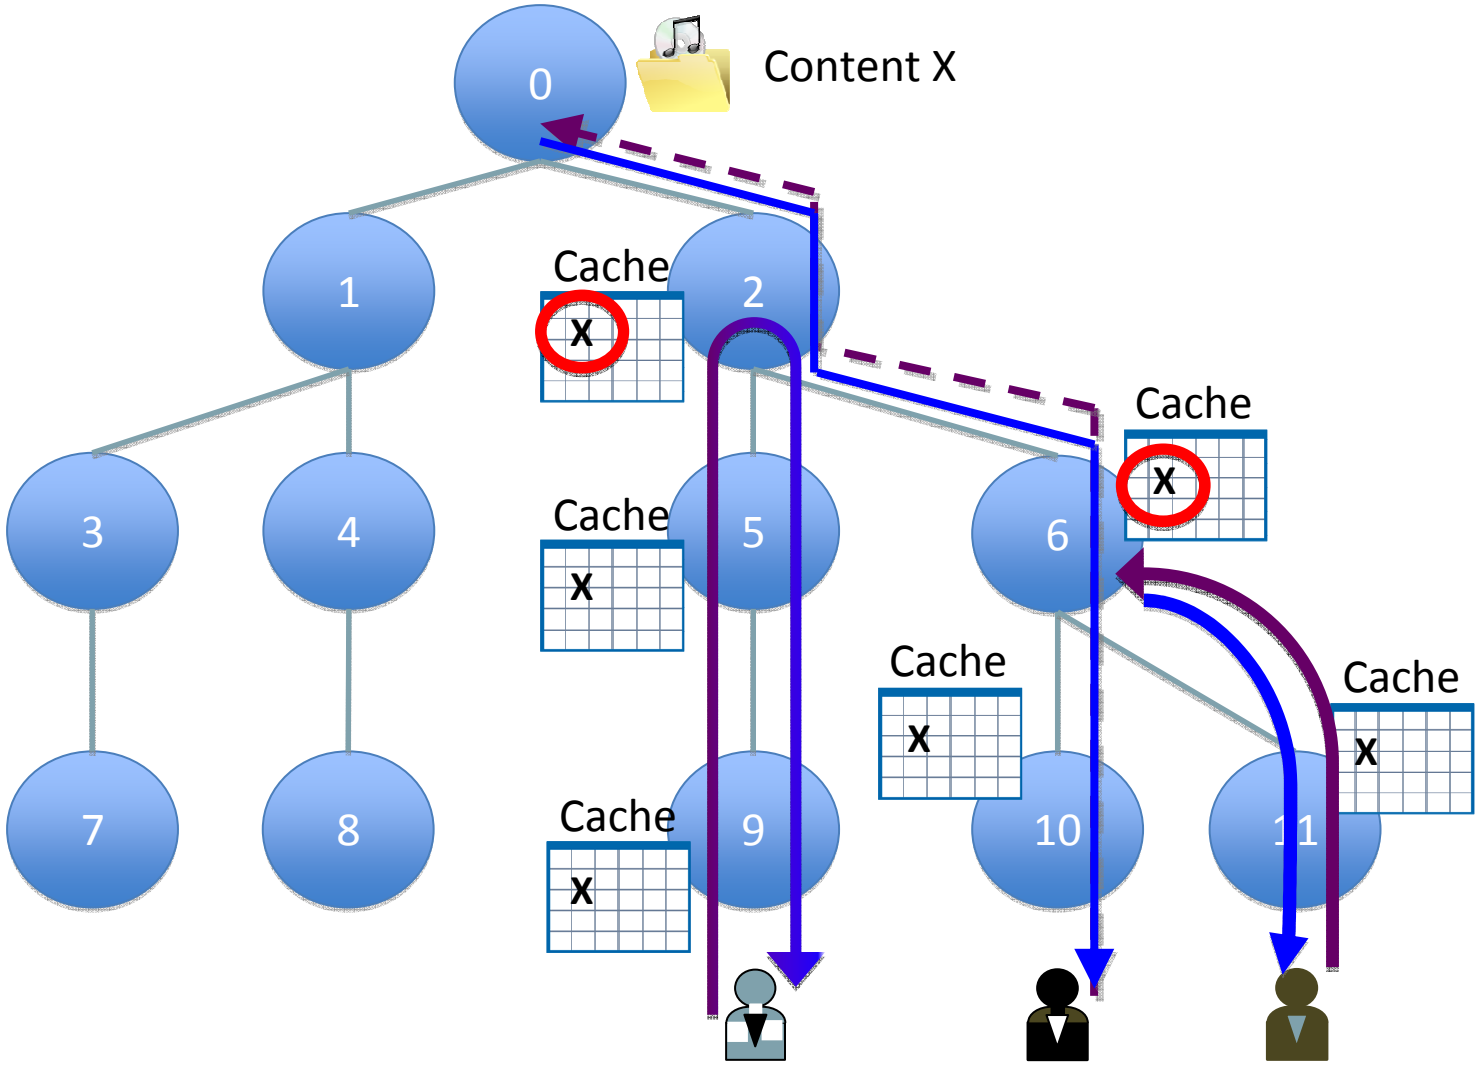

Content Mediation Plane (COMET)

- EU project COMET
- Two-phase approach, with resolution through Content Mediation Servers (CMSs) and delivery influenced by them
  - DONA-style resolution but adds information scoping/filtering and also anycast based on server load and network conditions
  - Delivery can use paths configured by the CMSs for better user QoE
  - Proactive caching at the network-edge ("content-aware routers")

- Evolutionary approach with minimal network modifications for better-than-best-effort content delivery

Content Mediation Plane (cont'd)

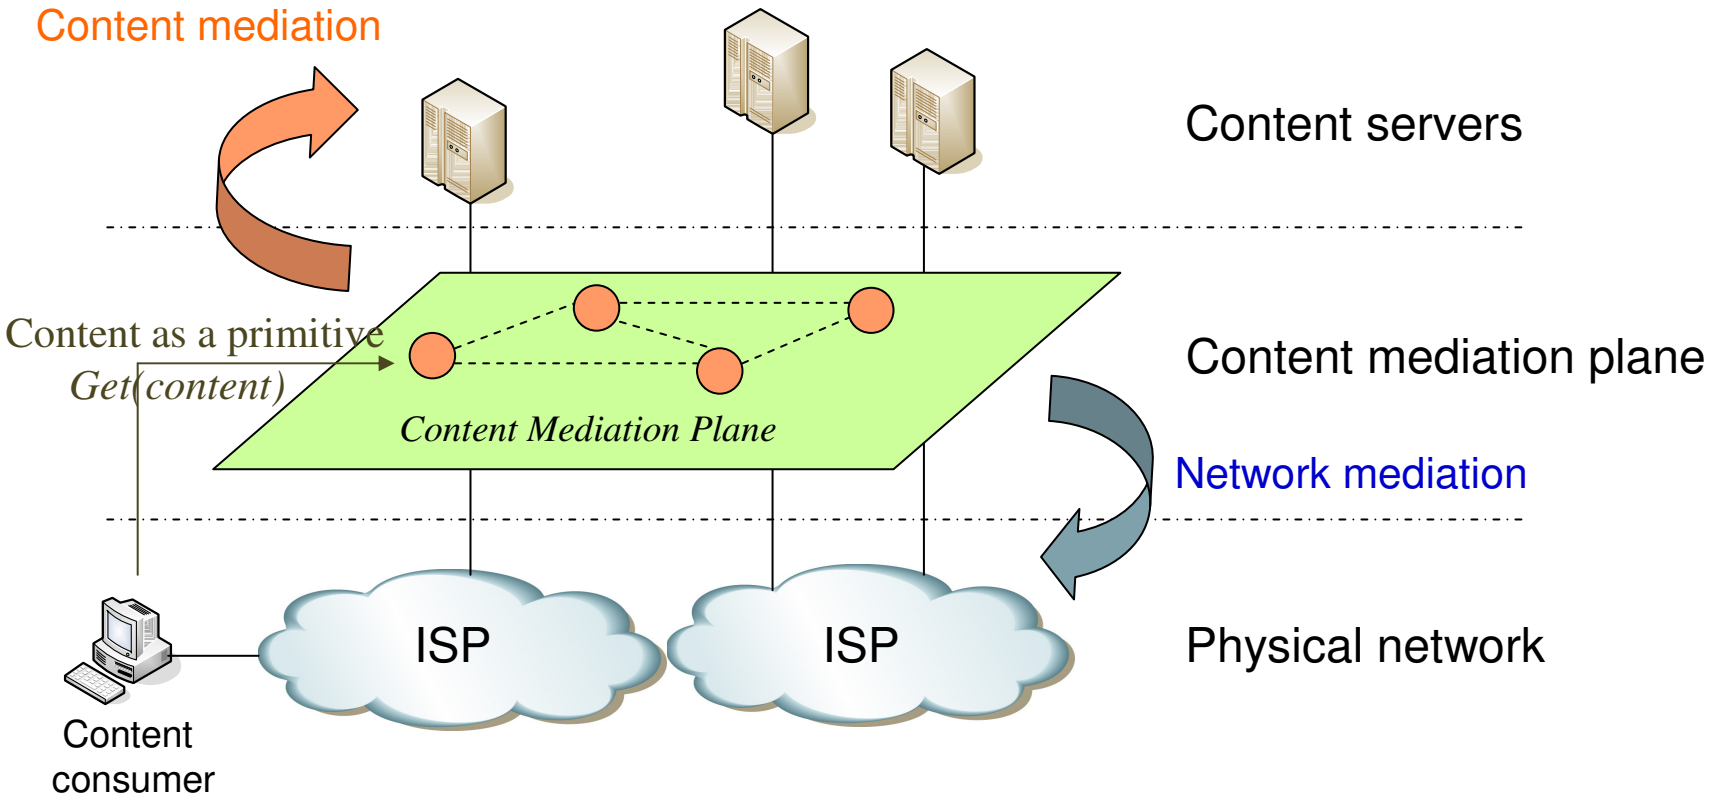

The content mediation plane can be also implemented in a radical manner within the network

Coupled Content Resolution and Routing in COMET

- Follows domain-level ***hop-by-hop gossip-like communication***
  - Content resolution is driven by ISP business relationships, BGP routing and content consumer preferences
  - Requires a content ID based on aggregatable labels which can be sequentially ordered
- *Register/Publish* and *Consume* messages
- Content can be only published to specific areas using INCLUDE e.g. BBC iPlayer content to be only available from within the UK
  - The same applies to content consumption i.e. from specific areas

- Pre-established state in the network is required ("breadcrumps")

Content Publication in COMET

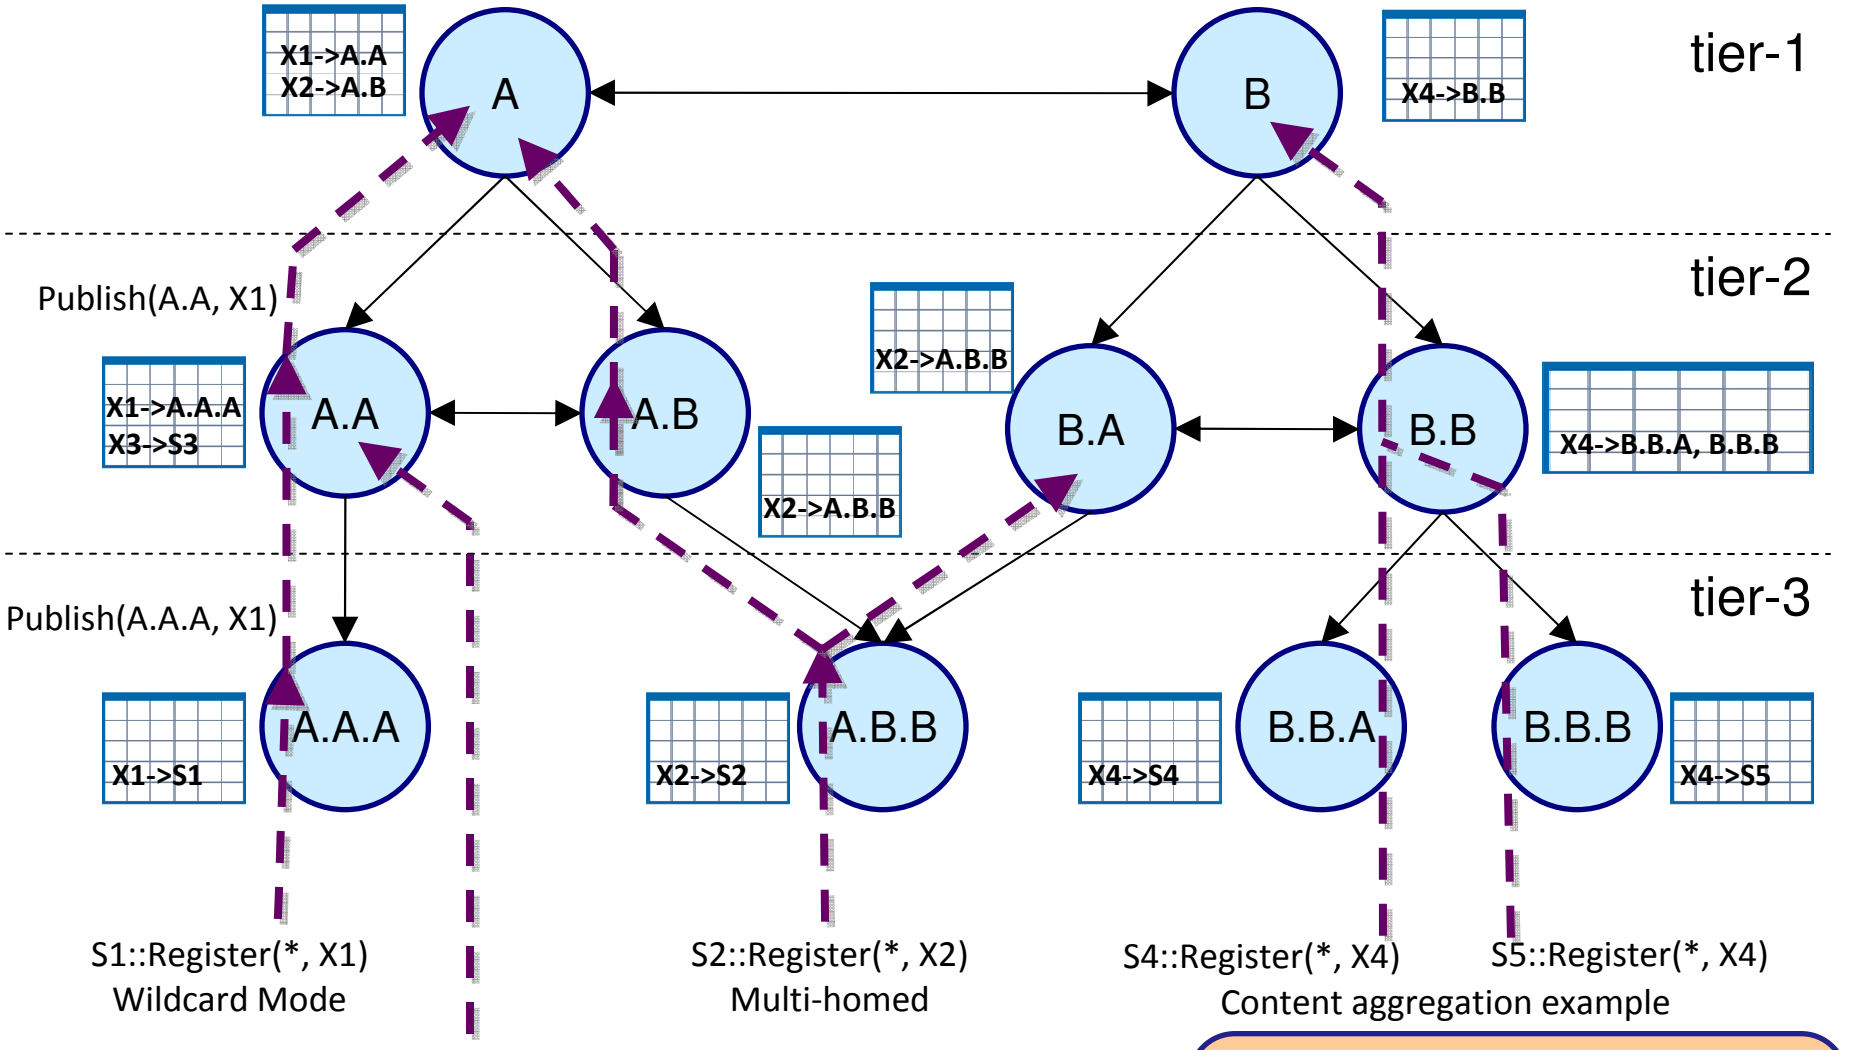

ICN Research Group in the IRTF

- Proposal in the forthcoming IETF to bring ICN researchers together, exchange research results, create a common ICN framework and feed input to existing IETF WGs
  - Longer term plan an IETF ICN WG
- Possible research topics to be addressed:
  - ICN naming schemes
  - Scalable name resolution for flat names
  - Scalable routing
  - Protocol framework
  - Security
  - API / application design
  - Business, legal and regulatory framework

Future Internet Requirements…

- Better mobility support
  - Impact on addressing
- More flexible and reliable routing
  - Multi-path as opposed to current single path
- Better service-aware resource control
  - Service-aware mapping of traffic to resources => better QoE
- Better security and spam protection
  - Possibly other paradigms of identity/presence, e.g. default-off

…to which ICN could be the Answer

- ICN can deal with:
  - Mobility - content/user ID not bound to location
  - Multi-path routing – anycast through in-network caching
  - Content-aware resource mapping – using metadata
  - Security – integrated with the content
  - Spam protection - receiver-driven model

Key ICN Challenges

- Naming – intricately linked with resolution and ID-based routing, so essential to get it right
- Scalability - cope with at least 10^{15} information objects
- Security per object, privacy concerns given that the network "sees" the information objects, spam control
- Manageability, real-time usage data to drive e.g. opportunistic caching through closed loop control
- Incremental deployment, the ability to gradually migrate without obliterating existing IPv4/v6 infrastructure
- Incentives and novel business models to engage involved stakeholders

ICN Could Make This Much Better!

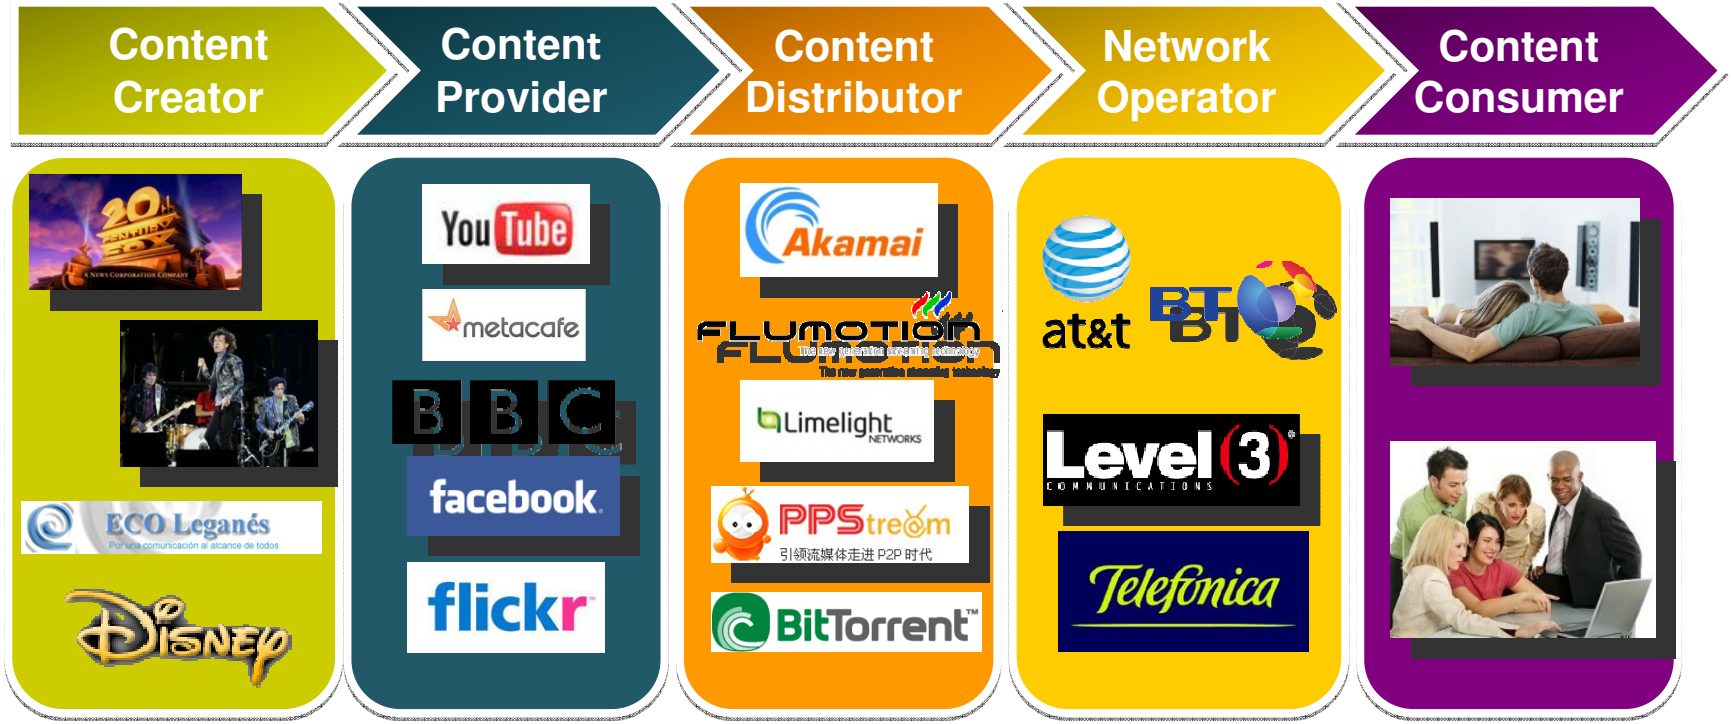

- ICN can provide tangible benefits to most stakeholders in an Internet that will be engineered according to its prevailing use
- Pave the way towards new media applications and user experiences