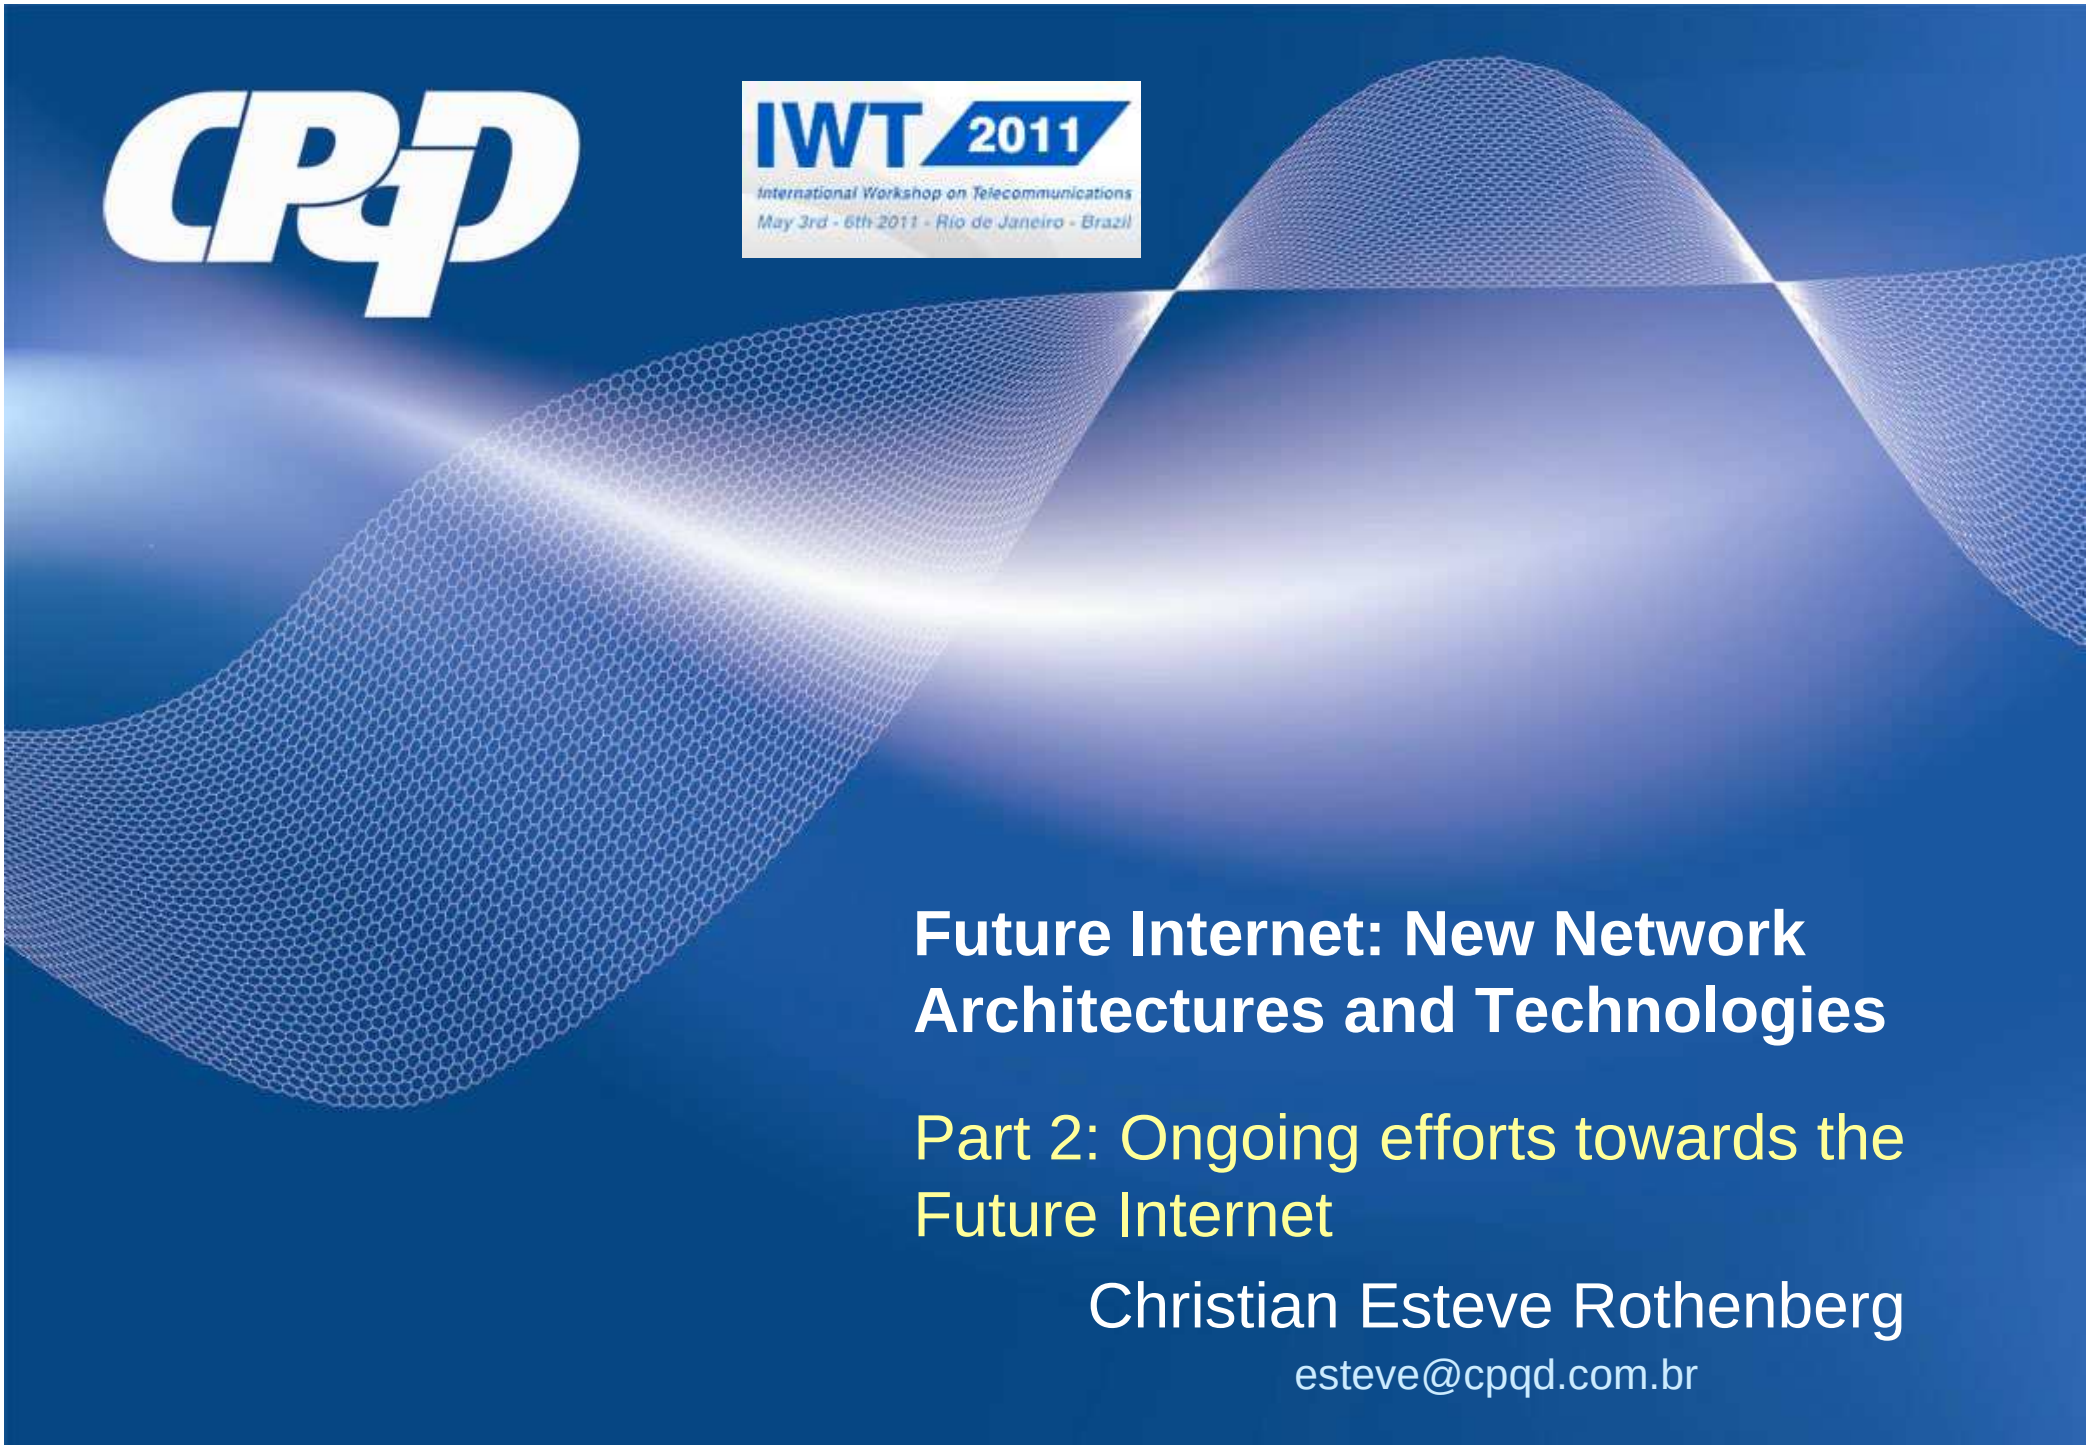

Future Internet: New Network Architectures and Technologies

Part 2: Ongoing efforts towards the Future Internet

Christian Esteve Rothenberg
esteve@cpqd.com.br

Innovation and Ossification in the Internet

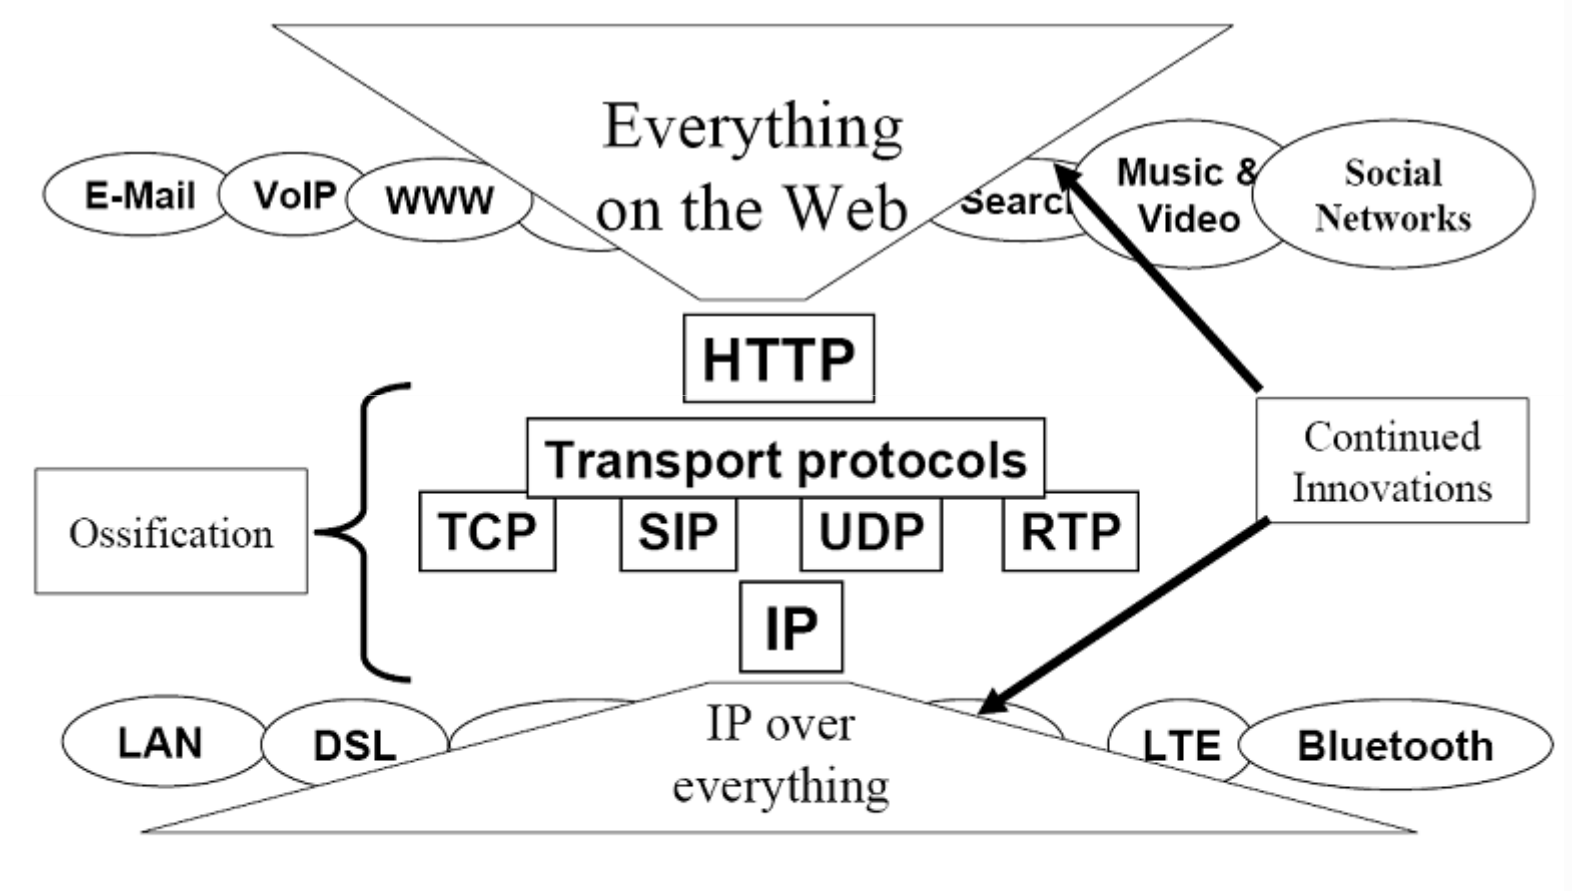

Source: Peter Stuckmann and Rainer Zimmermann, "European Research on Future Internet Design", IEEE Wireless Communications Magazine, October 2009

What happened with the All-IP dream?

- Wait a minute, a few years ago all was about IP convergence ...
  - see IWT´XY, SBRT XY, SBRC XY, etc.
- Now that the Telecom world has adopted IP,
  we don’t want IP anymore?
  - IP is not good, it does not scale, security, etc.
- We are researchers,
  - our job is to question paradigms
  - our job is to start the debate on a post-IP scenario
- Besides the researcher’s duty,
  there is a lot of rationale behind re-thinking the Internet architecture

Issues

Experienced by User

- Security
- Reliability and QoE

Attackers

- Denial of Service
- Intrusion, Session Capturing, Phishing
- Worms, Viruses, Spammers

Pain for the operators

- Limited Address Space
- Mobility
- Multi-homing
- Routing table explosion
- Scalable management
- Too much P2P traffic?
- Business model with over-the-top services?

Why the research boom on Future Internet?

- There is a big momentum on Future Internet research
- The Internet has invaded most aspects of life and society
  - Changing life, work, communication, social interaction, …
- It brings many benefits but also threats
  - Governments are concerned about it (critical infrastructures... e-war, cybercrime)
  - Funding Internet research is considered one important contribution to dealing with the situation

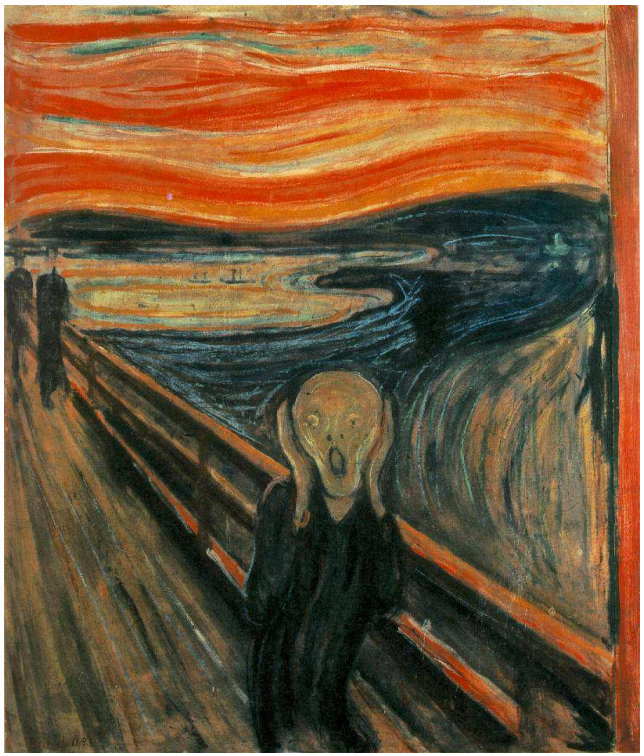

“The Internet Is Broken”

“The Internet will Collapse”

“The Internet Is Ossified”

So, what?

- There is a common consensus that the Internet needs improvement
- There is no shared vision on how this may happen
  - Not even a rough direction can be outlined
  - Popular (misleading) discussion item: *incremental* or *clean slate*?
- Consequence:

Let’s do a broad search instead of intensive research

  - US: FIND initiative plus GENI building a big playground
  - EU: funding many projects with different approaches

Future Internet research projects popping up everywhere

- GENI/FIND, USA
- Future Internet Cluster, EU
- New Generation Network / AKARI, Japan
- Future Internet Forum, Korea
- CNGI, China Next Generation Internet Project
- RNRT, France
- G-Lab Initiative, Germany
- SHOK, Finland
- Ambient Sweden Initiative, Sweden
- Internet del Futuro, Spain
- CANARIE, Canada
- ...

Future Internet Research across the globe

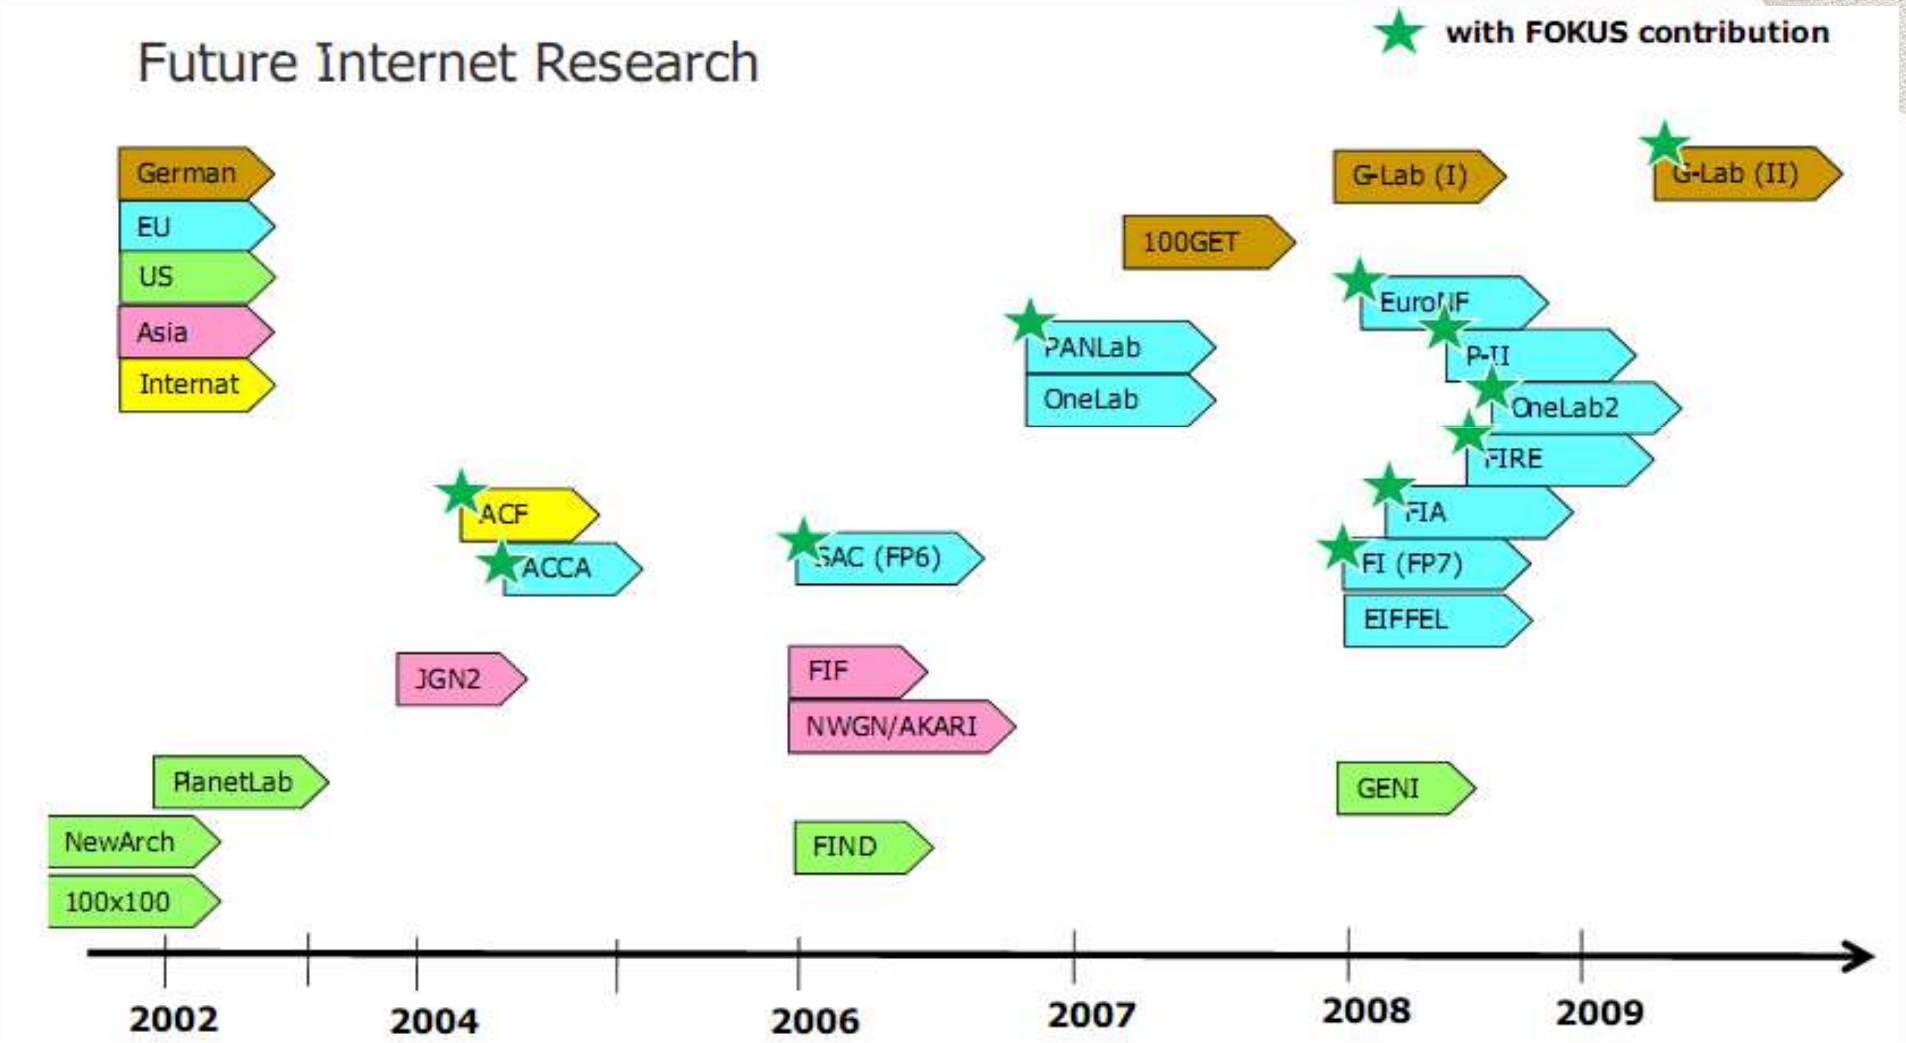

Source: T. Tseby, “Future Internet Technologies”

EU Future Networks Project Portfolio and Clustering

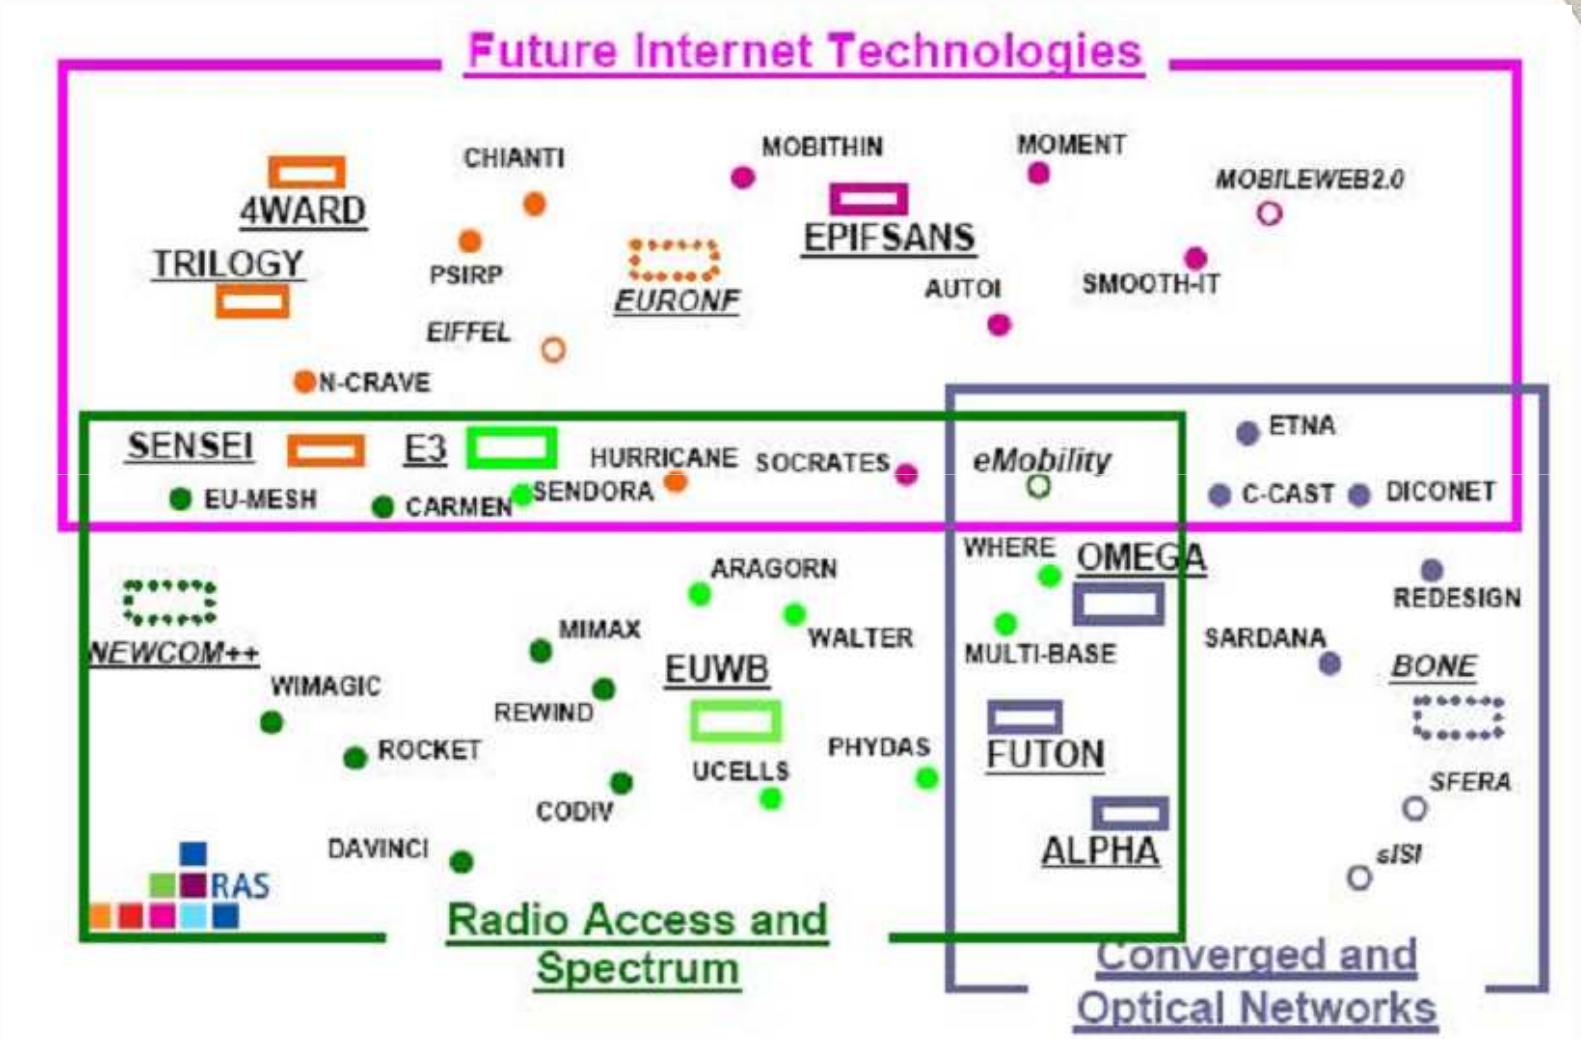

Big momentum on Future Internet research

We may have only one bullet,
so we better use it right!

What do we need?

1. We need visions for the Future Internet
   - Re-thinking fundamentals (transport, routing, addressing, identity, new Internet waist)
   - Defining goals and requirements for the FI
2. We need experimentally-driven research for validation at scale and under realistic scenarios
   - E.g., GENI, FIRE, Federica, OneLab
3. We need business incentives for adoption
   - Think IPv6, MobileIP, IP Multicast, etc.
   - EIFELL, MIT CFP, BIRD, socio-economics market evaluations, Industrial engagement, etc.

Visions through Clean Slate Designs

Stanford University

Clean Slate Design for the Internet

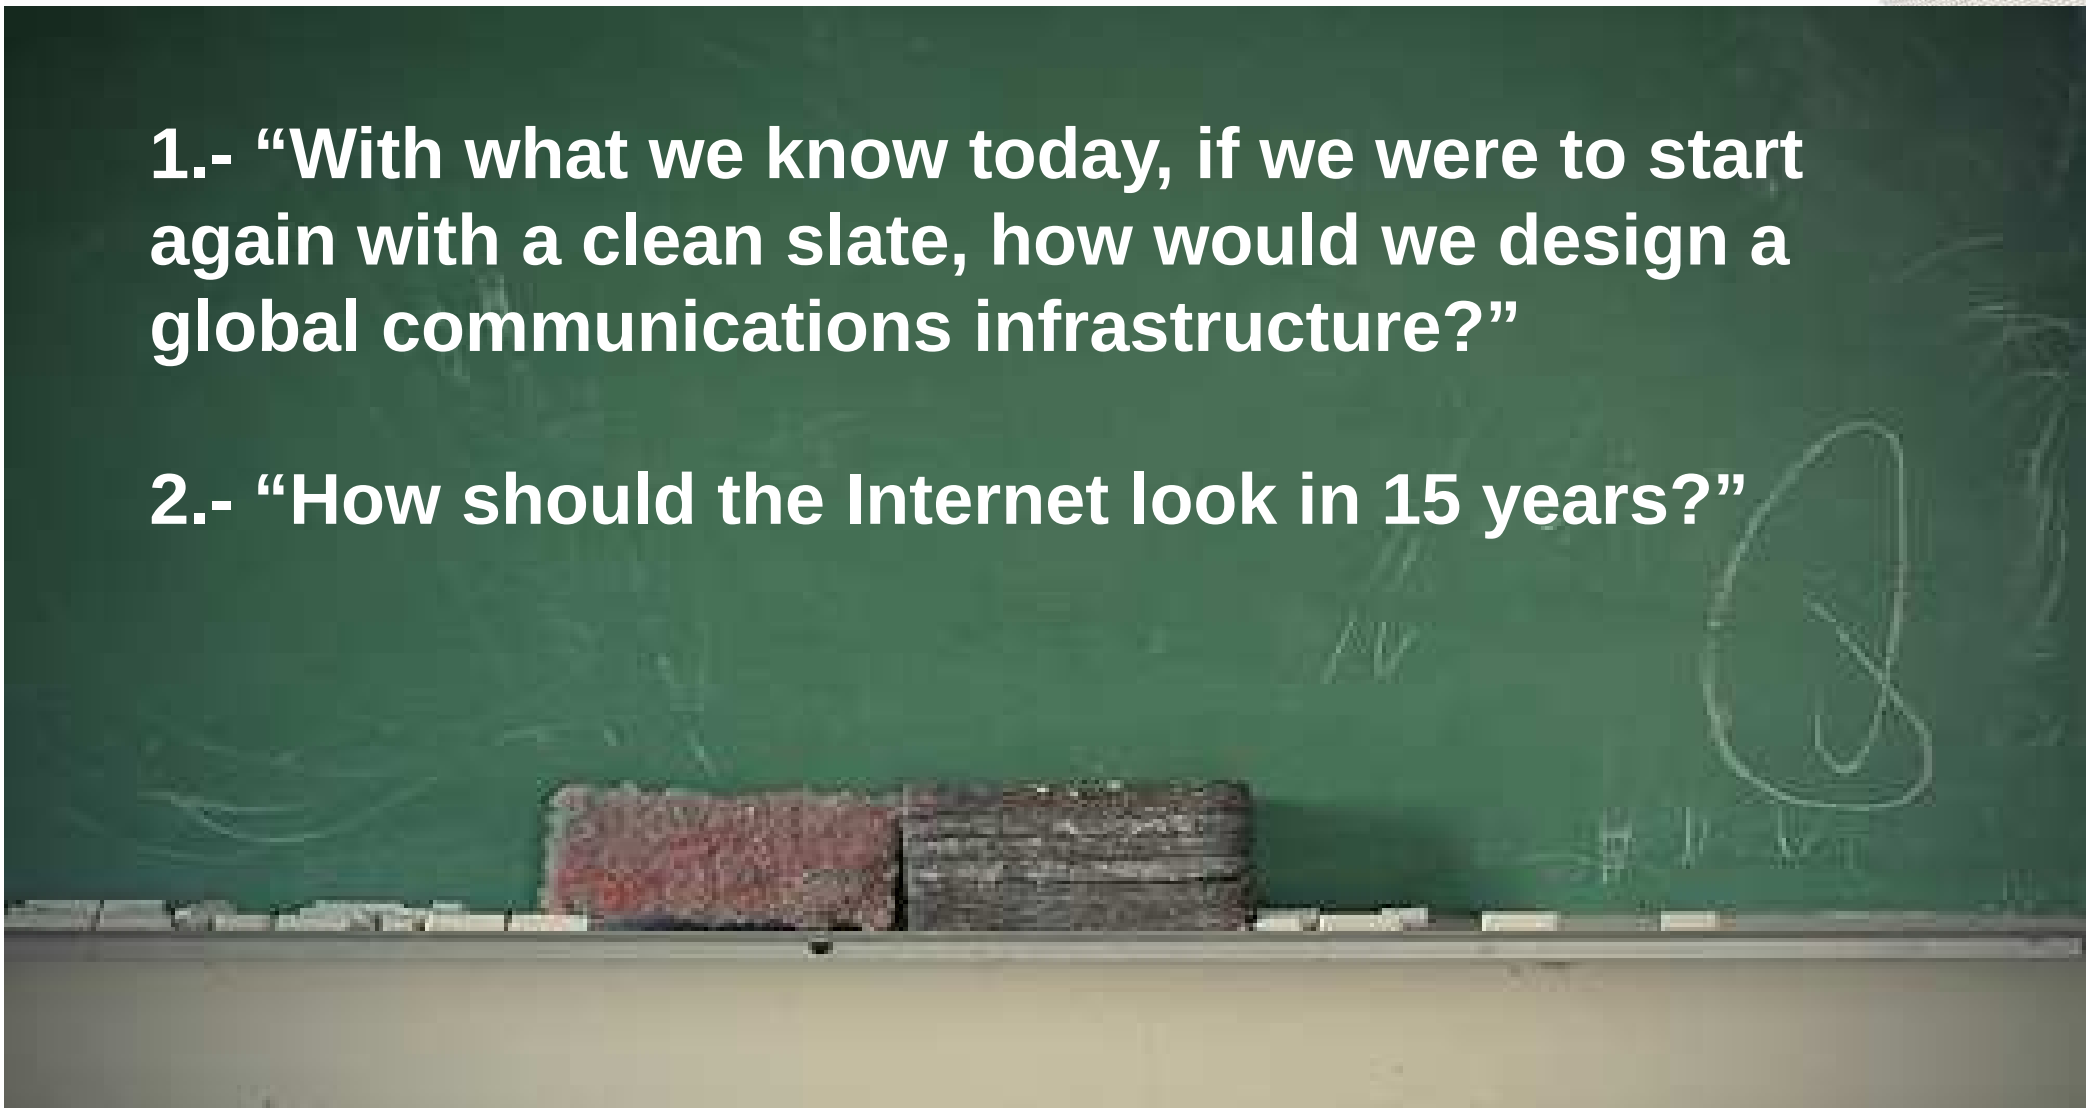

Disclaimer Notice:
Clean slate design does not presume clean slate deployment.

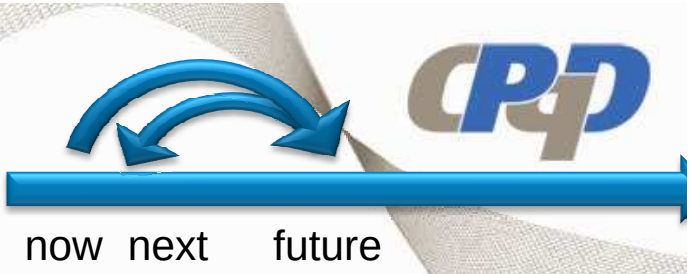

- Late binding to reality -

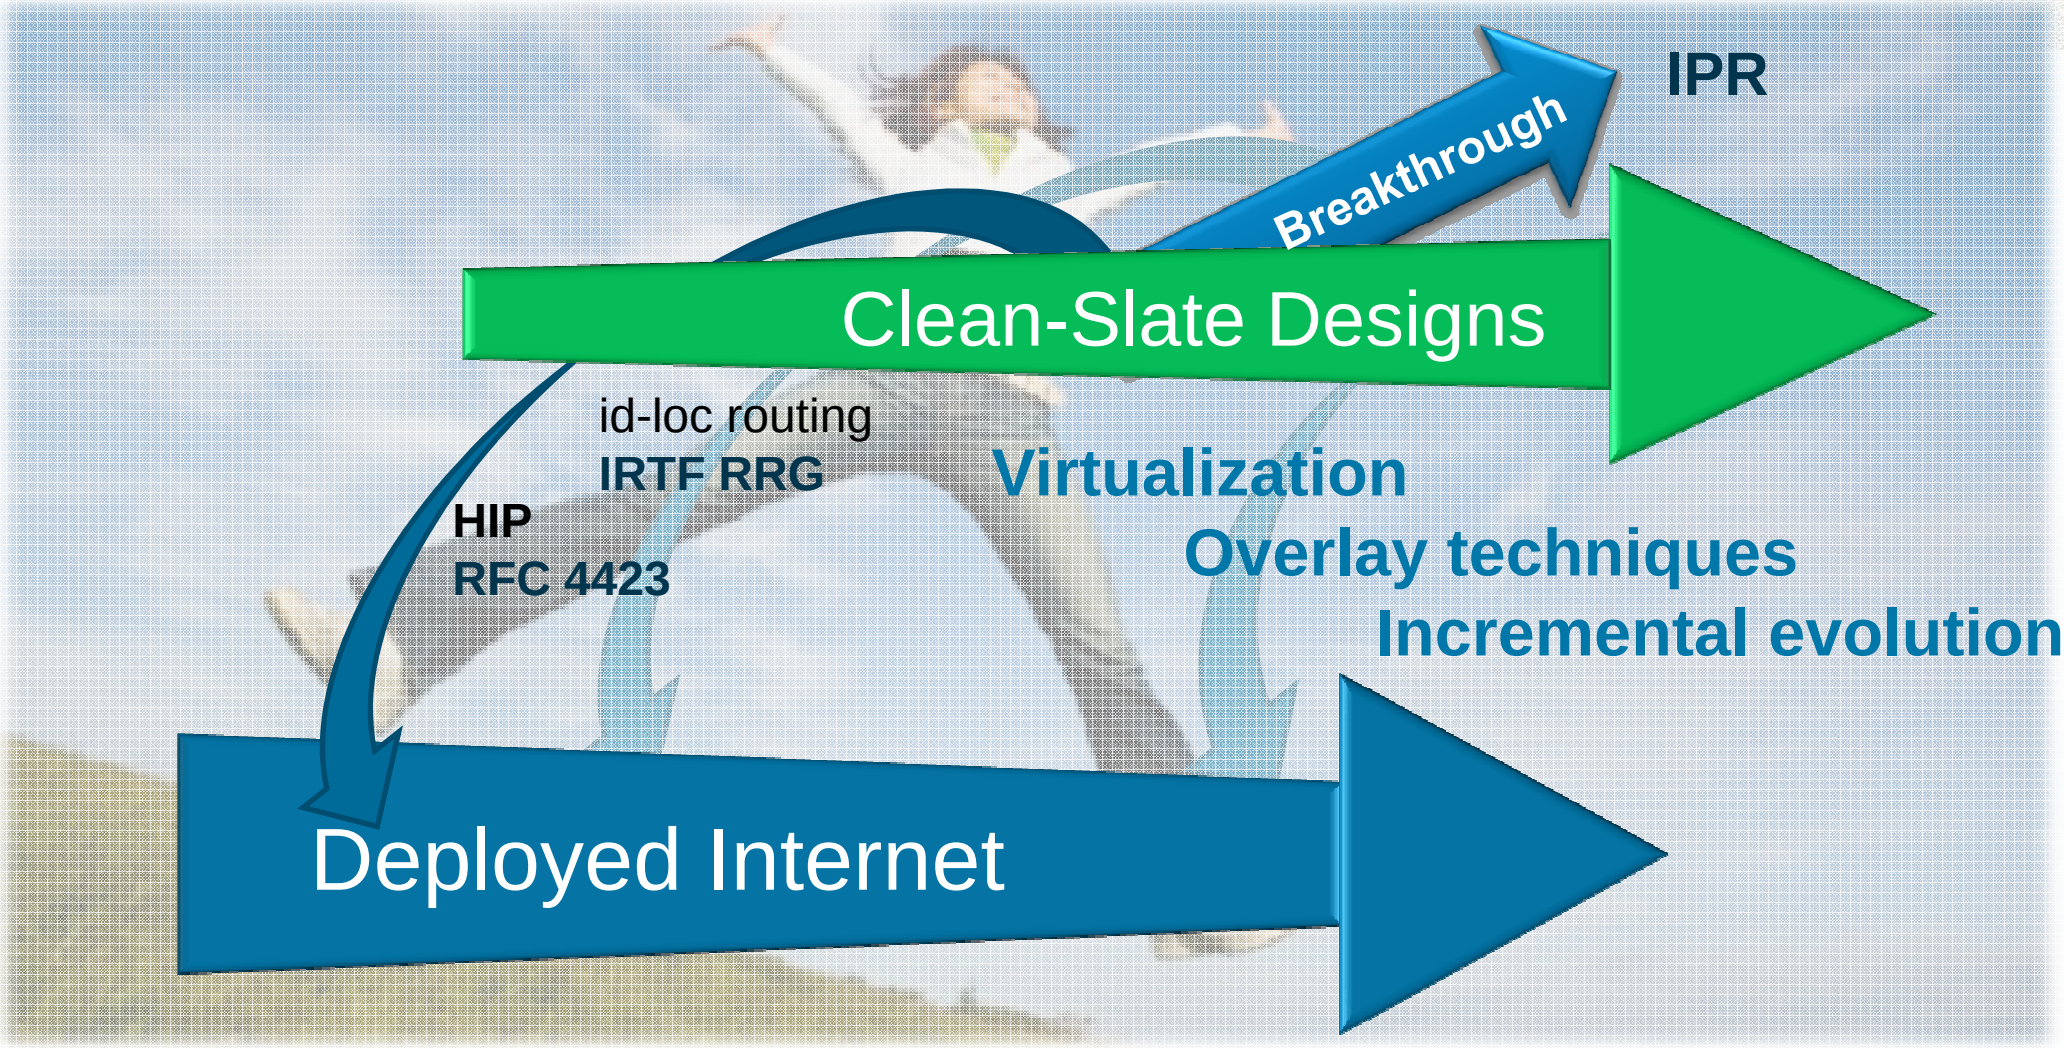

Approaches and (visionary) ideas

- There are many ideas out there,
  - Some are already several years old
  - Few are fundamentally new
- On the following slides some will be presented
  - Subjective selection based on potential of networking revolution
- Credits and references:
  - D. Clark, “Moving FIND to the next stage”, Jul. 2009
  - http://groups.csail.mit.edu/ana/People/DDC/Working%20Papers.html
  - V. Jacobson, “Networking Named Content” to appear at CoNEXT 2009 http://www.ccnx.org
  - EU FP7 PSIRP – Publish Subscribe Internet Routing Paradigm, http://psirp.org
  - J. Quittek, “The Future Internet, Is it time to look for a new one?”
  - T. Tseby, “Future Internet Technologies -- A Technical Overview of Evolutionary and Revolutionary Ideas”, 2010

Approaches and (visionary) ideas

- **New control architectures**
- **Overlay networks**
- **Network virtualization**
- **Software-defined networking**
- **Locator-identifier split**
- **Information-oriented networking**
- **Self-management**
- **Revisiting networking fundamentals**
  - What defines an architecture? There is no networking science.
  - Addressing, Routing, Security, Management, Availability

    [D.Clark]
- Many more….

Multiplexing - a basic issue

- Old (1960's) idea: packets.
  - Seems to have worked out well.
- New ideas:
  - Integrated management of packets and circuits (aggregates).
    - Integrated management.
    - Fault recovery, routing/traffic engineering.
    - Integrate future concepts in optics (routing vs. TE)
  - Virtualization of routers and links
    - Avoid need to have one design.
    - Needs assessment and practical validation

Source: D. Clark

Virtualization Today

- **Virtual Memory**
  - Virtual Machines Virtual Machines
  - Sharing system resources (e.g. vmware)
- **Virtual Paths**
  - Virtual Private Networks (VPNs)
  - Multiprotocol Label Switching (MPLS)
  - Generalized MPLS (GMPLS)
- **Virtual Application Layer Networks**
  - Overlay Networks

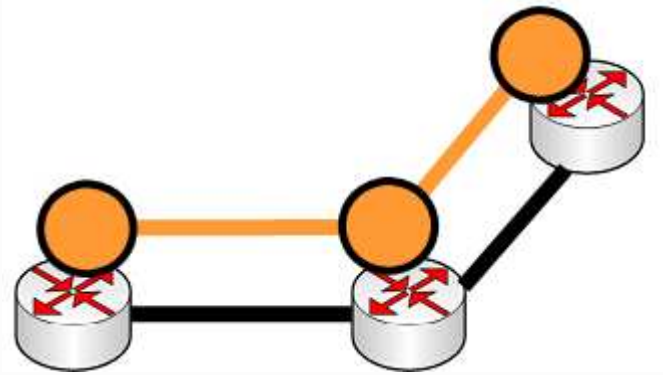

Source: T. Tseby, “Future Internet Technologies”

Network Virtualization

- Virtualization of all network resources
  - Network nodes
  - Links

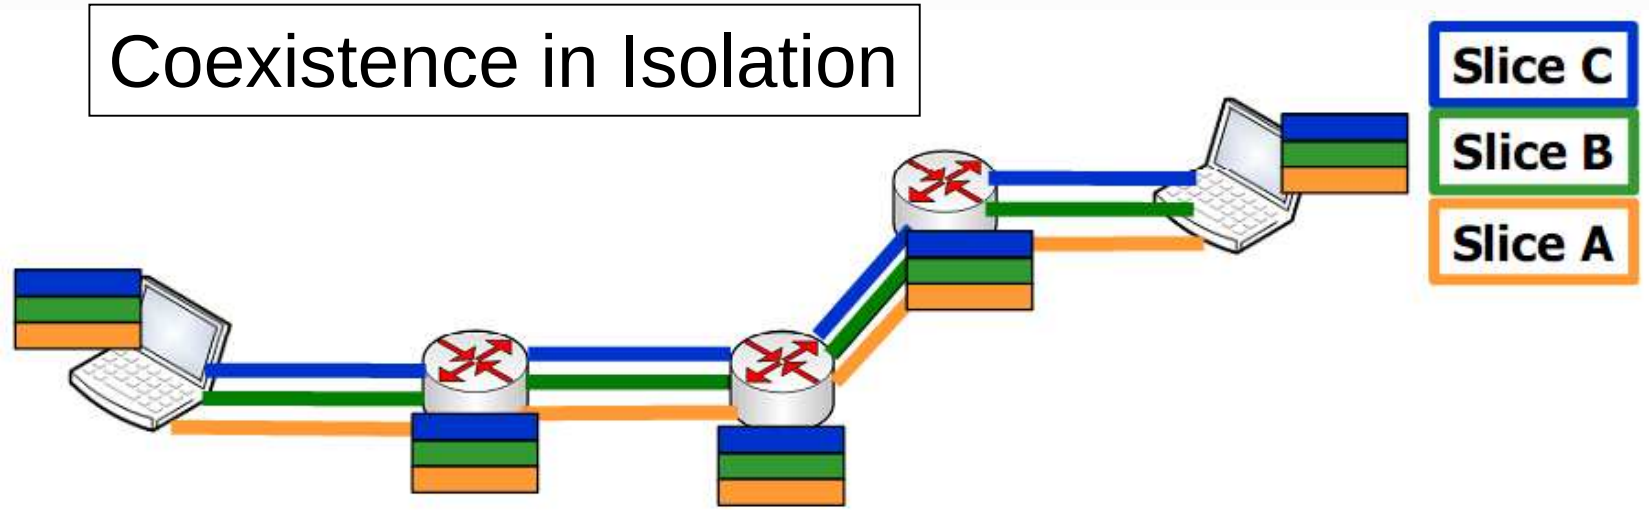

Source: T. Tseby, “Future Internet Technologies”

Benefits of Network Virtualization

- **Customization**
  - Applications can adapt network to their needs (e.g. buffer management)
- **Economic Refactoring**
  - Separate infrastructure providers from service providers
- **Optimized Resource Utilization -> Cost Reduction**
  - Resource sharing, Load balancing, power saving, etc.
- **Isolation**
  - Controlled resource access and usage
  - Provide Isolated networks for experiments
- **Security**
  - Set up secure environment among trusted entities

Virtualized Testbeds for Future Internet Research
– Provide separated resources over shared infrastructure for experiments
Virtualized Future Internet
– Virtualization as general approach for the Future Internet

Connection establishment

- Old idea:
  - Minimize the round trips.
- New ideas:
  - Need a phase for exchange of identity.
    - May need a “cross-layer” initial exchange.
    - Re-modularize TCP to be less layered.
  - Need to diffuse attacks.
    - Adding a round trip or two (esp. if not always) worth the cost in order to allow an E2E (indentity) check.
    - Part of availability framework.
  - Fit this thinking into the DTN paradigm.

Source: D. Clark

Addressing

- Old view:
  - Designed for efficient forwarding.
- New view: take into account
  - Security issues
    - Accountability, privacy, deterrence, hiding.
  - Management issues
    - Re-numbering
  - Multi-homing
  - Do you really want to address physical nodes?
    - How about services? Information? Anycast?
    - But consider lower-layer management issues.

Source: D. Clark

Routing

- Old view:
  - Find the lowest cost route
  - Load-based dynamics lead to instability.
- New ideas:
  - Random route selection (oblivious routing avoids link DoS and TE)
  - User route selection (P2P, Multi-homing)
  - Multi-path routing. (TCP multipath, IETF, Trilogy)
  - Energy/cost-aware routing (SIGCOMM 09)
  - Machine learning to achieve high-level policies (self-optimization)
  - Move route computation out of forwarders (4D, OpenFlow)
  - Multiple simultaneous routing schemes (virtual network slices)
  - ID-loc separation (HIP, LISP)
  - Routing on flat identifiers (Pasquini et al.)

Source: D. Clark

Identifier-Locator Split

- **Main issue: IP semantic overload**
  - Transport layer: IP is an identifier (naming)
  - Network layer: IP is a locator (addressing)
- **Consequences**
  - Lack of a stable identifier for end-to-end communication
  - Mobility/Multihoming
  - Heterogeneity
  - Security
- **Solution**
  - Identifier/locator separation
    - HIP, IETF RRG LISP, NodeID

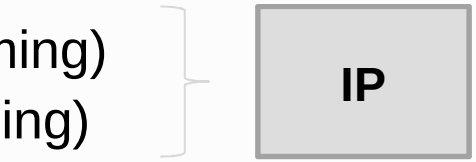

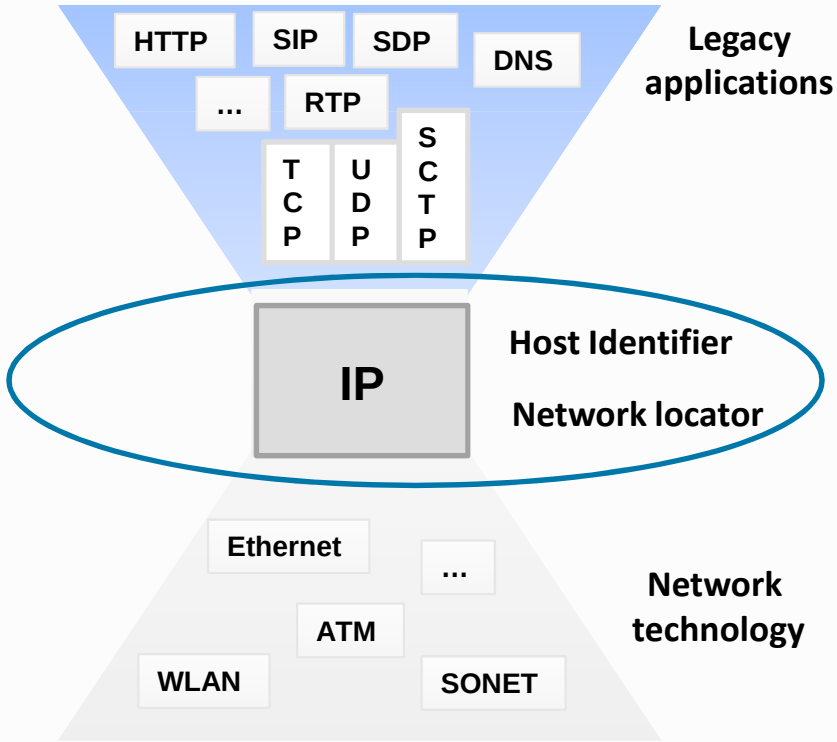

Identifier-Locator Split

Identifier/locator separation

- Introduction of an **identification layer** between the network and transport layers
  - E.g. Identifiers are 32-bit (128-bit in IPv6) flat (topology-free), persistent and unique node IDs
- Issue: Mapping/binding of identifier to locators

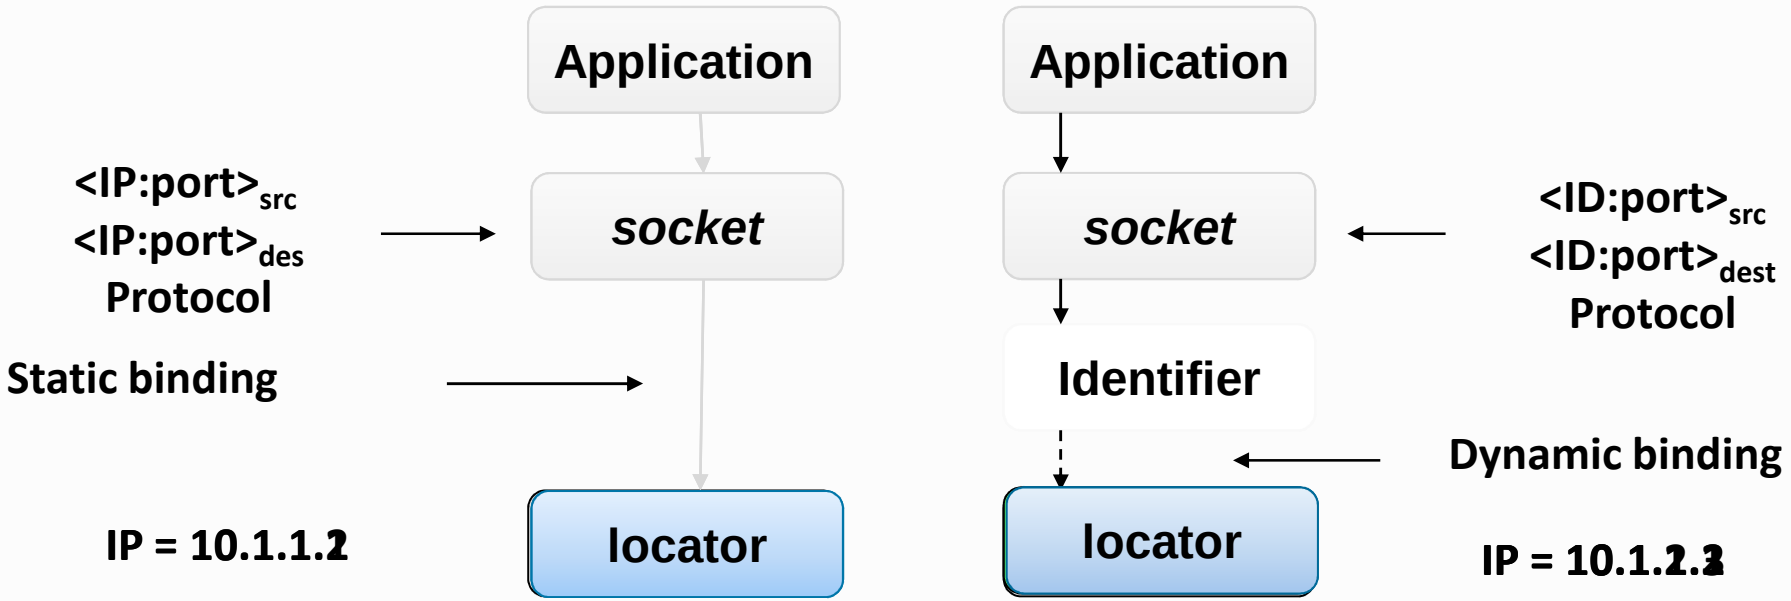

Application design

- Old view (simplistic): our machines talk.
  - Host-to-host conversation
- New view:
  - Lots of servers and services (resource pooling in cloud DCs)
  - Need for cross-application core services
  - Identity management, social networks
  - Modulate behavior based on trust.
  - Name-oriented socket API [cf. C. Vogt]
  - Linked Data (cf. Semantic web)
- Application design patterns and building blocks should be part of the future network.

Source: D. Clark

Information-layer

- Old idea: an application issue (ignore it.)
- New idea: need a framework
  - Naming and identity of information.
    - Independent of how you get it.
  - Dissemination
    - Swarms, P2P: (heterogeneous).
    - Improves availability of information if information is pushed into the network.
  - Economics: one service or many competing?
    - Competitive info dissemination “on top of” lower-layer transport.
  - Information-Centric Networking
    - Can we create a network architecture based on naming data instead of naming hosts?

Source: D. Clark

Overview of the 4 “flagship” NSF Projects on Future Internet Architectures

Future Internet Projects

Named Data Networking

- **Principal Investigator**: Lixia Zhang, UCLA

- **Collaborating Institutions**: Colorado State University, PARC, University of Arizona, University of Illinois/Urbana-Champaign, UC Irvine, University of Memphis, UC San Diego, Washington University, and Yale University

- http://www.named-data.net/

Re-Architecturing the Internet
- Information-centric approaches -

information-centrism

PSIRP

DONA

content-centric networking

New ID spaces

TRIAD

ROFL

id-loc

CDN

clean-slate

3 ***Interconnecting information***

DPI

P2P

overlays

IPv6

NAT

middleboxes

networking (r)evolution

patching

TCP/IP

2 ***Interconnecting hosts***

Telephony

1 ***Interconnecting wires***

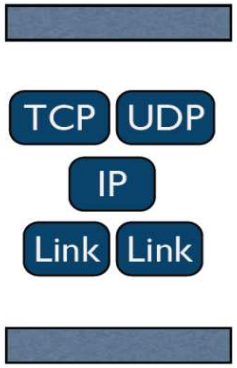

Information-centric Networking
- Rethinking fundamentals -

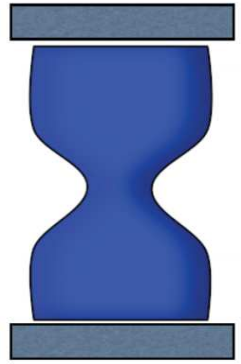

- **Send / Receive → Publish / Subscribe**
- **Sender-driven → Receiver-driven**
- **Host names → Data names**
- **Host reachability → Information scoping**
- **Channel security → Self-certified metadata**
- **Unicast → Multicast**

CCN: A New Layering

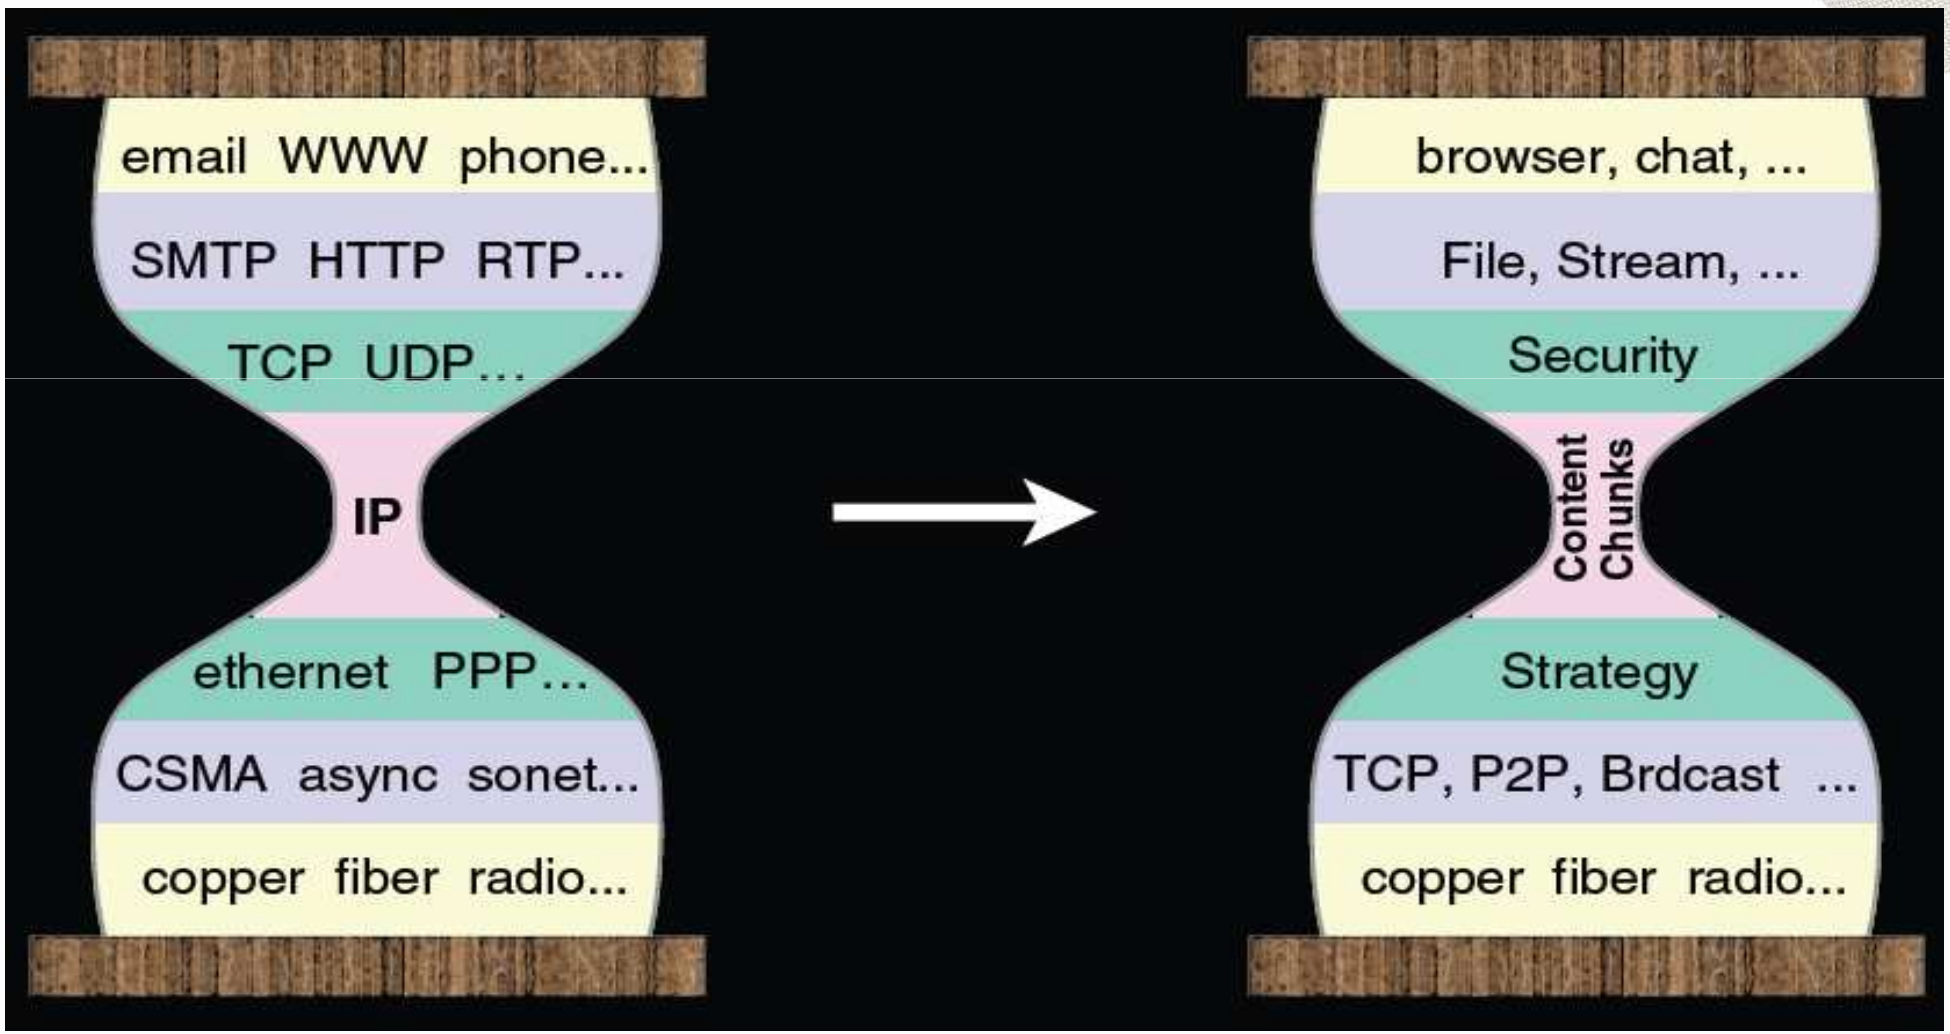

What's in a Name (user/app view)

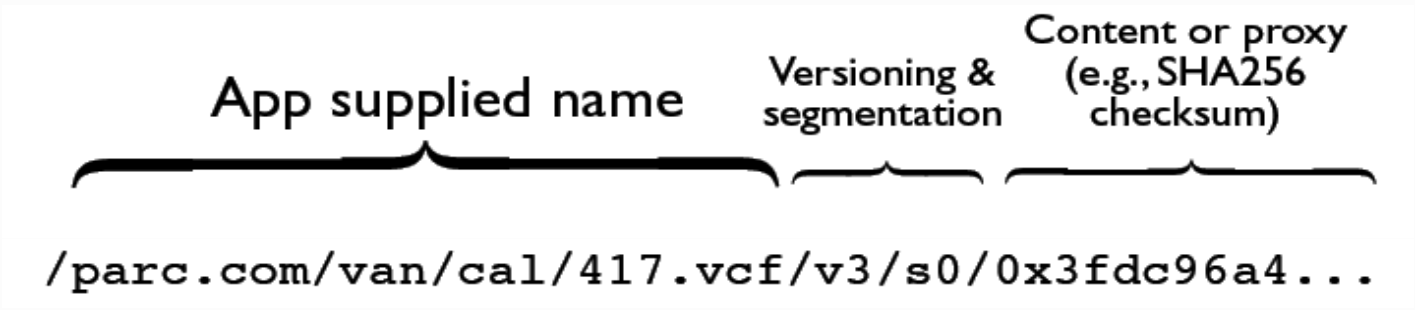

- Note that this binding is *immutable*
  - the data associated with the name can't change

Built-in security through self-certified data

- Metadata contains encrypted cryptographic checksum and locator for the public key of the producer.
- Producer's key is typically hierarchically structured.

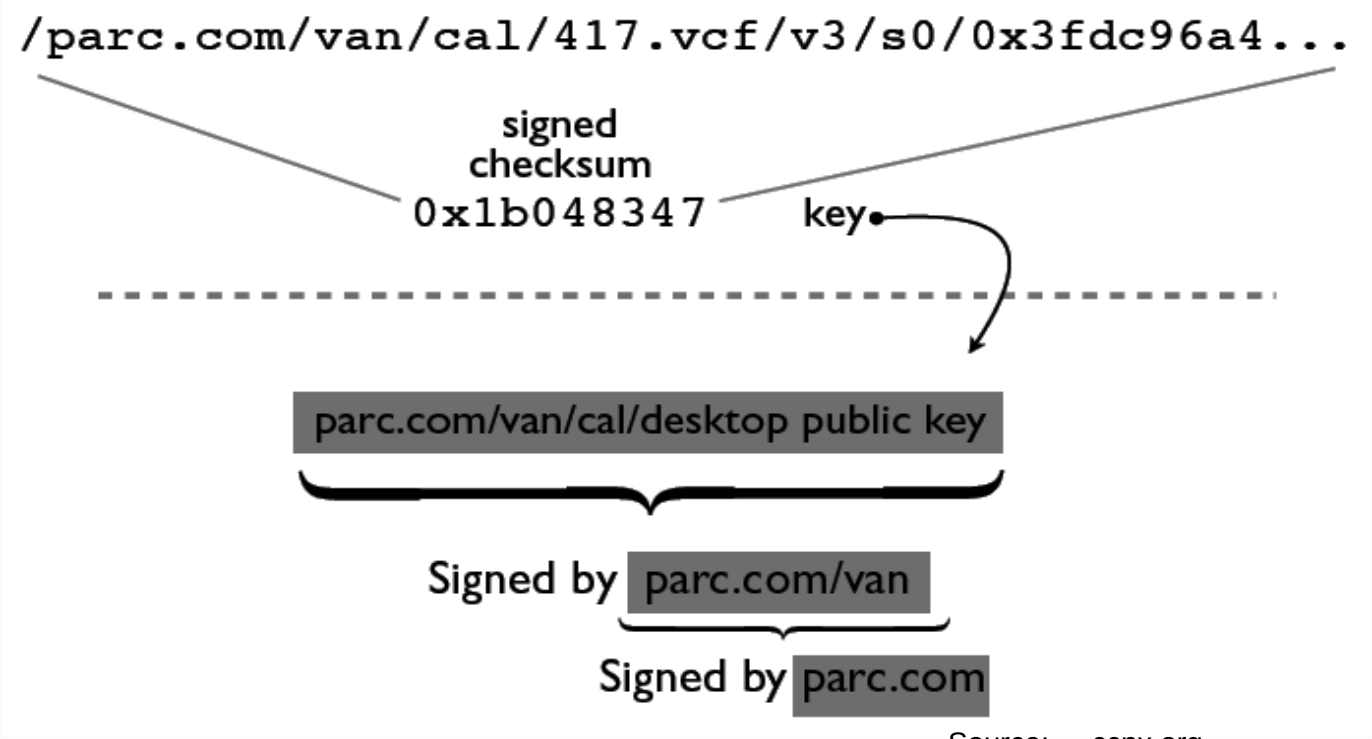

Source: ccnx.org

Two basic primitives

- There are just two CCN packet types -
  interest (similar to “http get” or “subscribe”)
  data (similar to “http response” or “publish”).

Interest packet

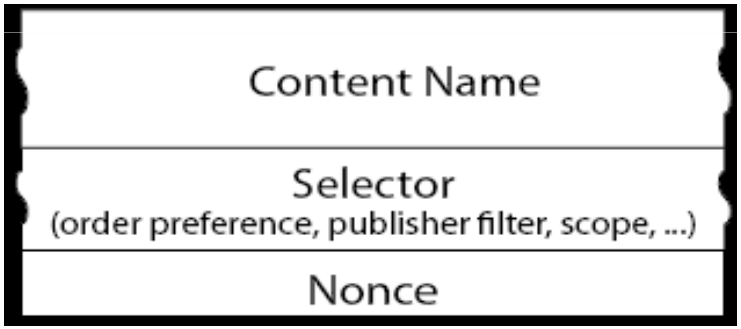

Data packet

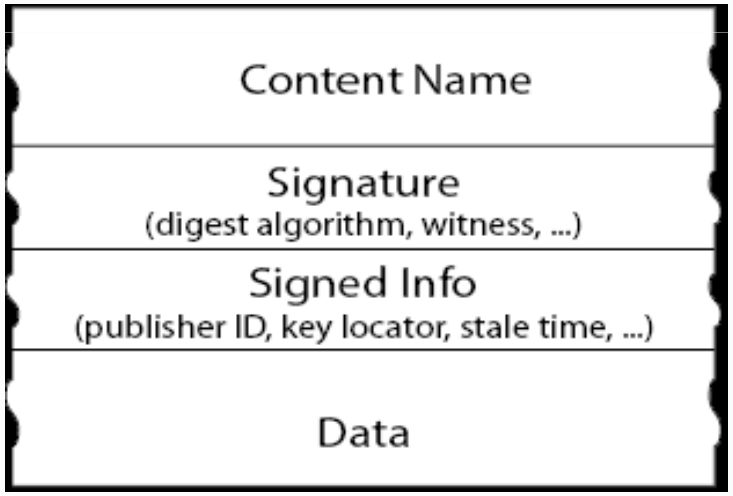

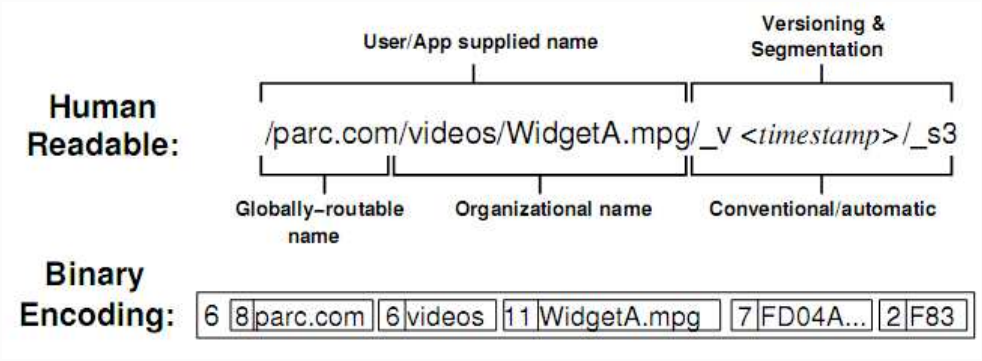

Content name

Name-oriented routing and forwarding

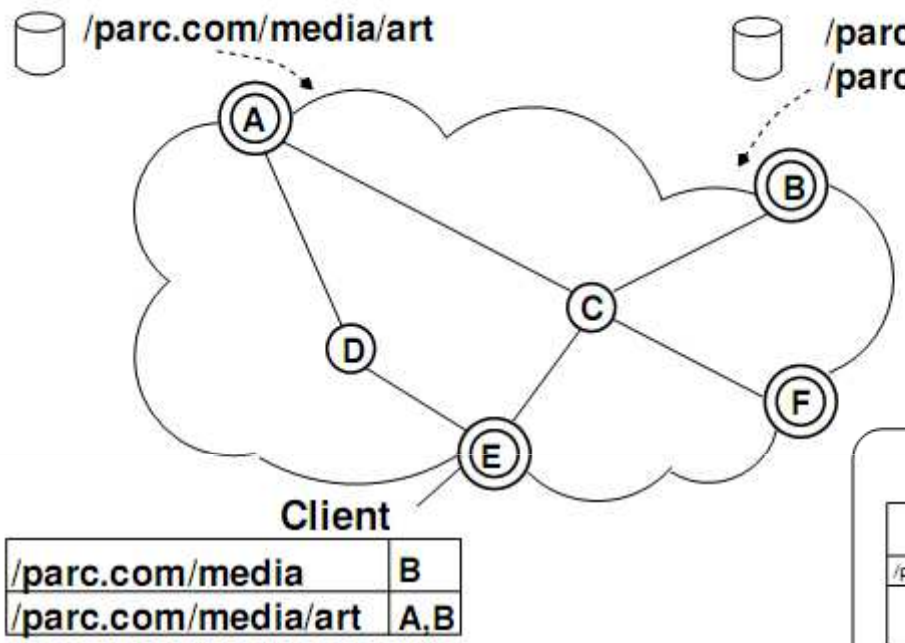

CCN forwarding engine model

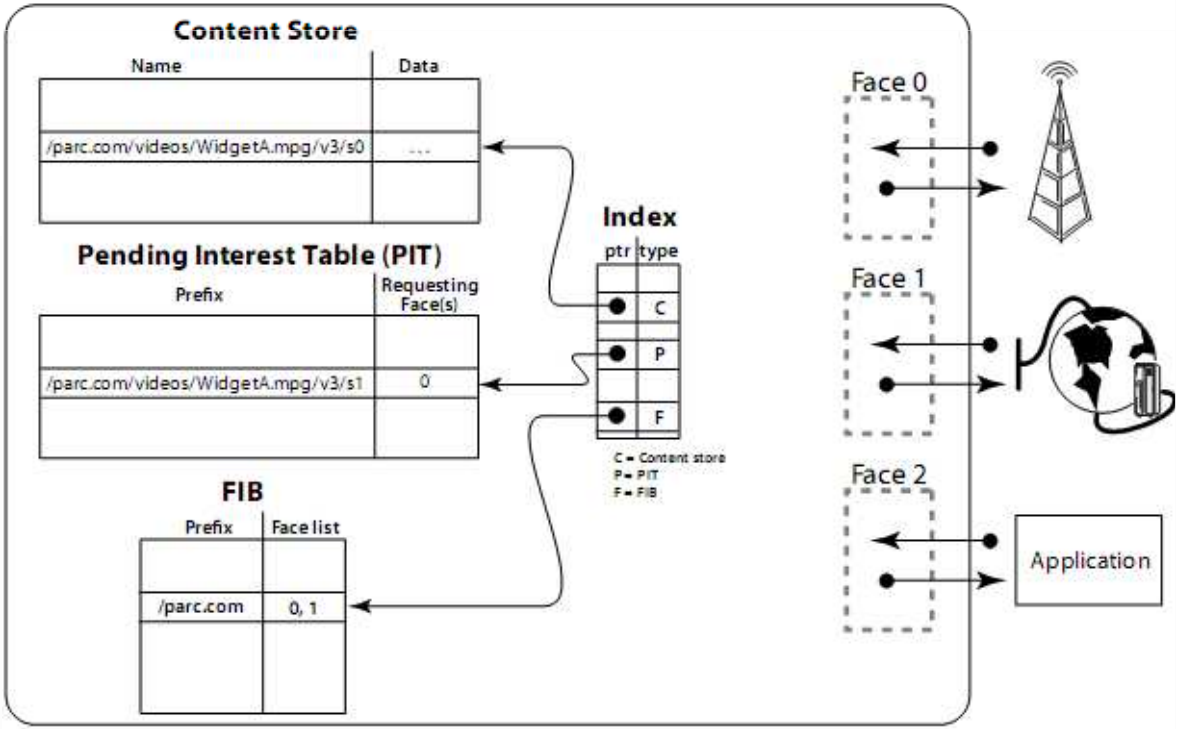

- FIB populates based on name aggr.
- Interests go to pending interest table and is forwarded based on the FIB
- Data packets remove PIT entries
- Content Store are opportunistic caches
- Flow-balance and loop-free

At a minimum, same hardware req. as IP

IP networking

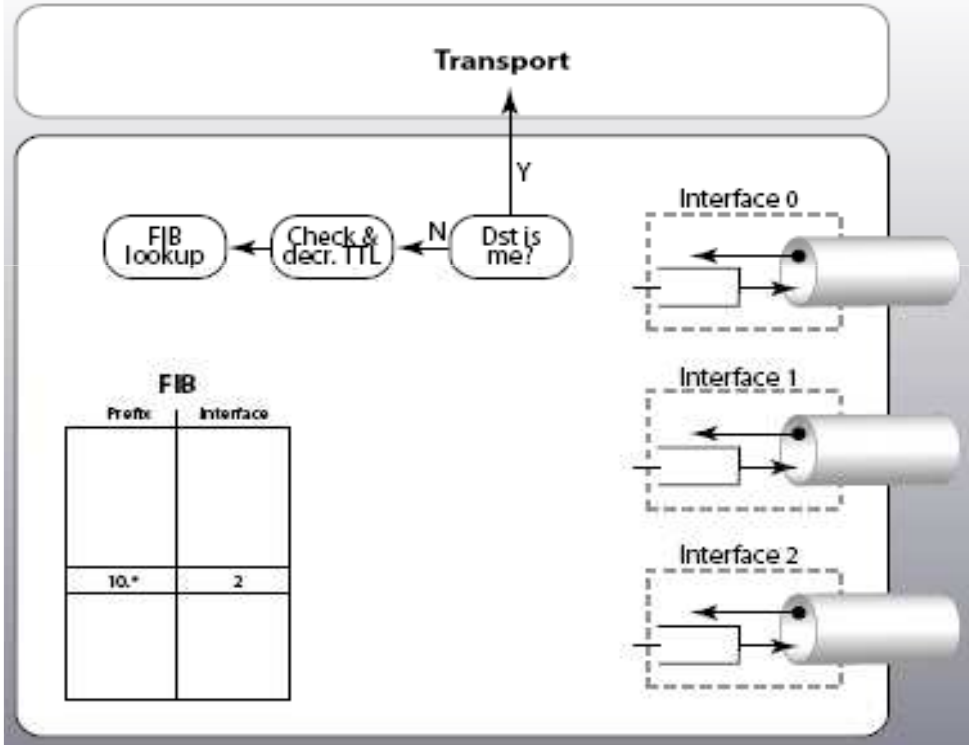

Content-Centric Networking

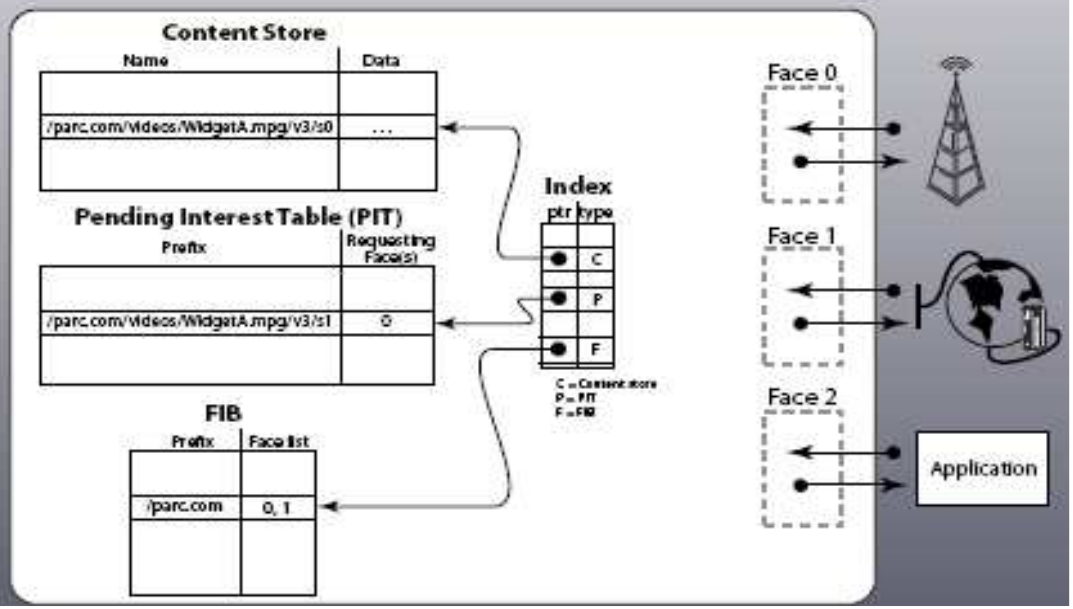

EU FP7 PSIRP Project

PSIRP
PUBLISH-SUBSCRIBE INTERNET ROUTING PARADIGM

- Redesign the Internet architecture from the pub/sub point of view, taking nothing (not even IP) for granted.
  - *Make “information” the centre of attention*
  - *Remove the “location-identity split” that plagues current networks*
  - *Innovative multicasting and caching features to optimize performance and efficiency*
  - *Security functionality as a native core component of the architecture*

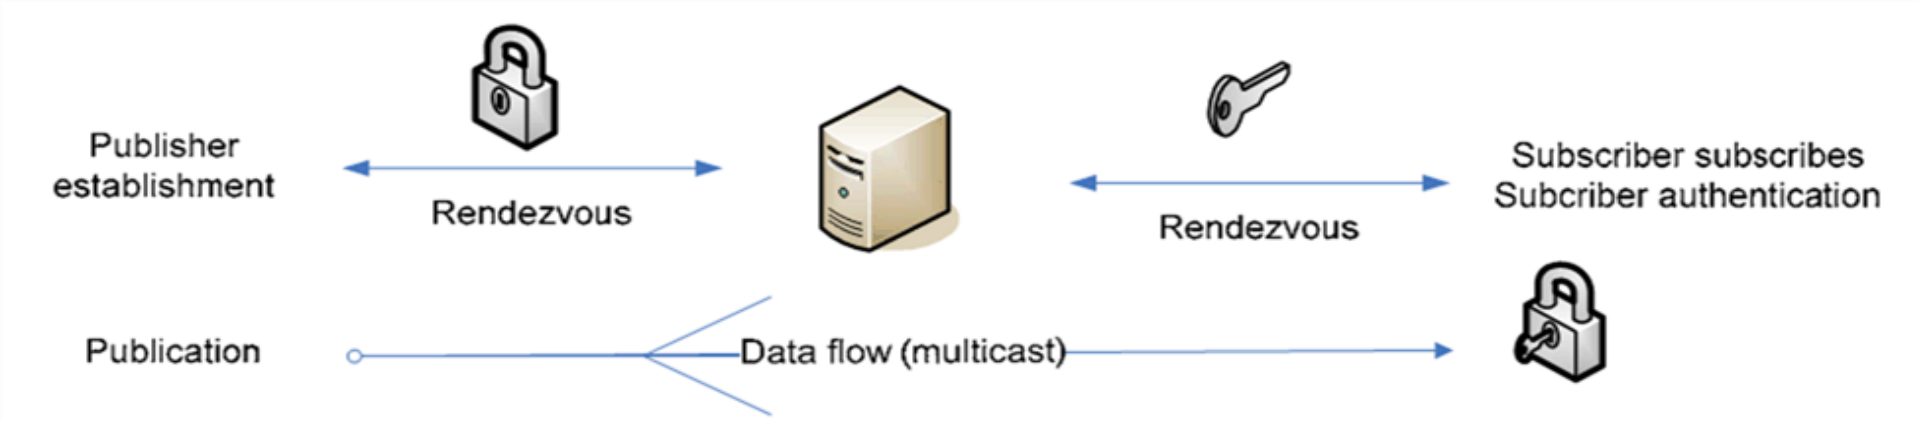

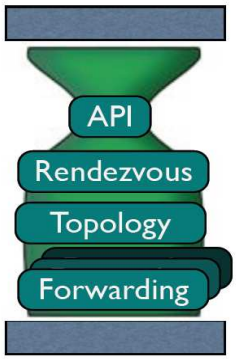

RTFM

PSIRP
PUBLISH-SUBSCRIBE
INTERNET ROUTING
PARADIGM

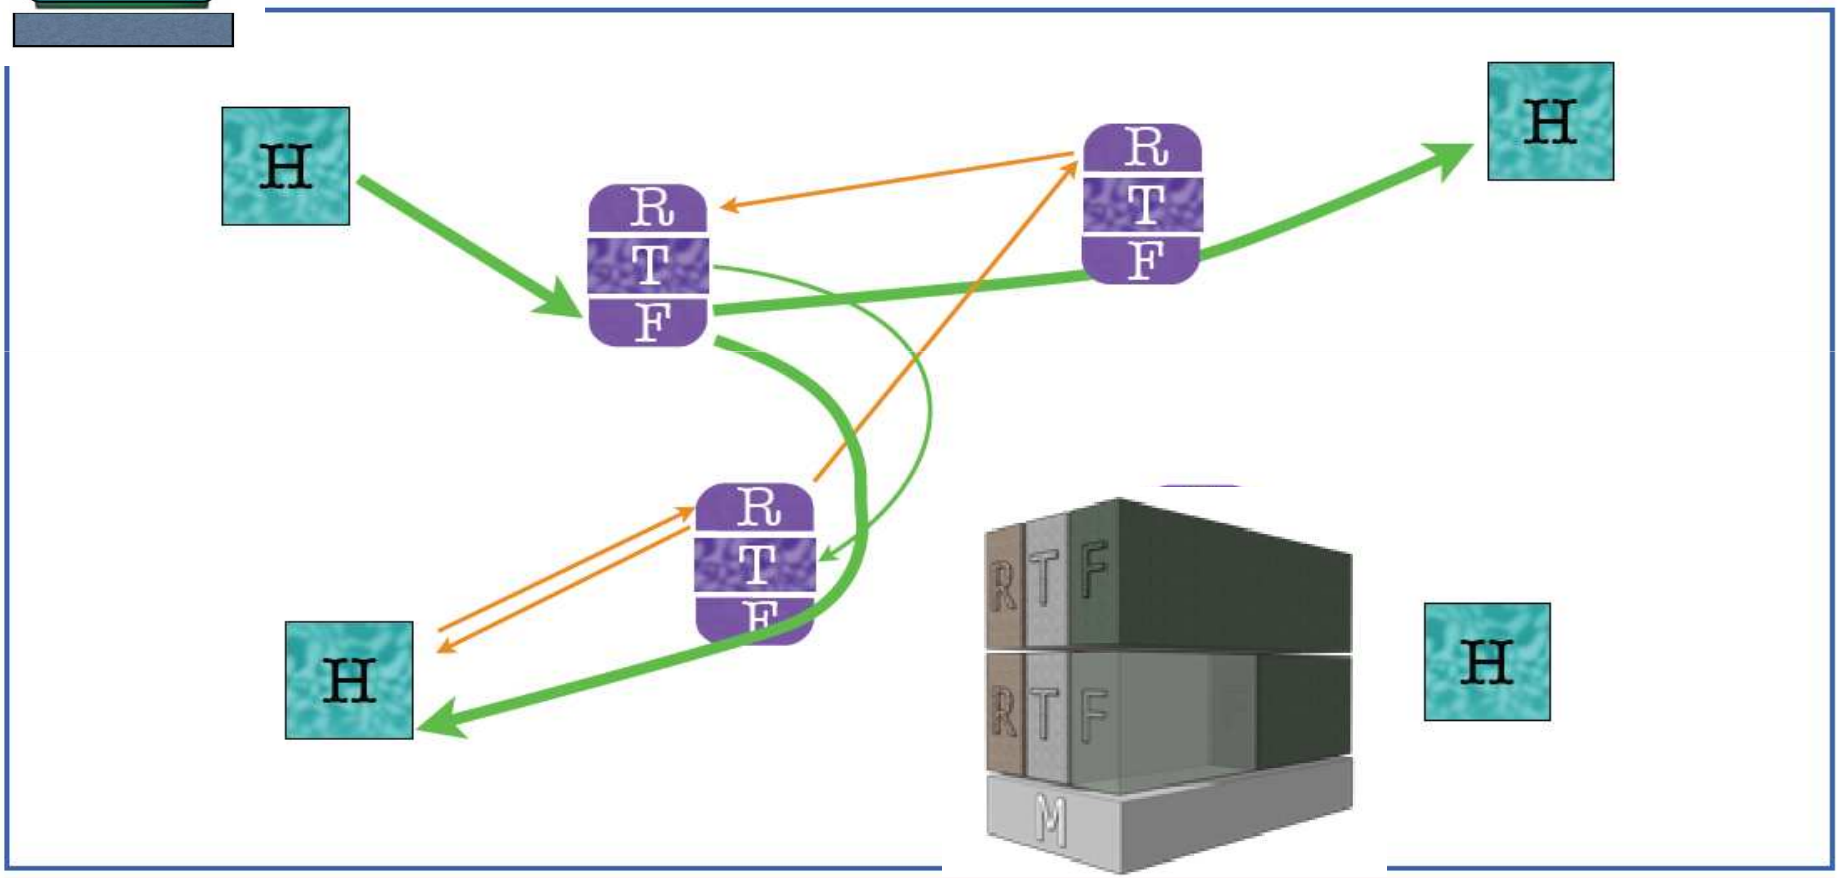

Source: EU FP7 PSIRP Project, http://psirp.org

eXpressive Internet Architecture

- **Principal Investigator:** Peter Steenkiste, Carnegie Mellon University

- **Collaborating Institutions:** Boston University, University of Wisconsin/Madison

- http://www.cs.cmu.edu/~xia/

Vision of the eXpressive Internet Architecture

XIA envisions a future Internet that:

- Is trustworthy
  - Security broadly defined is the biggest challenge
- Supports long-term evolution of usage models
  - Including host-host, content retrieval, services, …
- Supports long term technology evolution
  - Not just for link technologies, but also for storage and computing capabilities in the network and end points
- Allows all actors to operate effectively
  - Despite differences in roles, goals and incentives

XIA - Evolvability

- Narrow waist of the Internet has allowed the network to evolve significantly network to evolve significantly
- But need to evolve the waist as well!
  - Can make the waist smarter

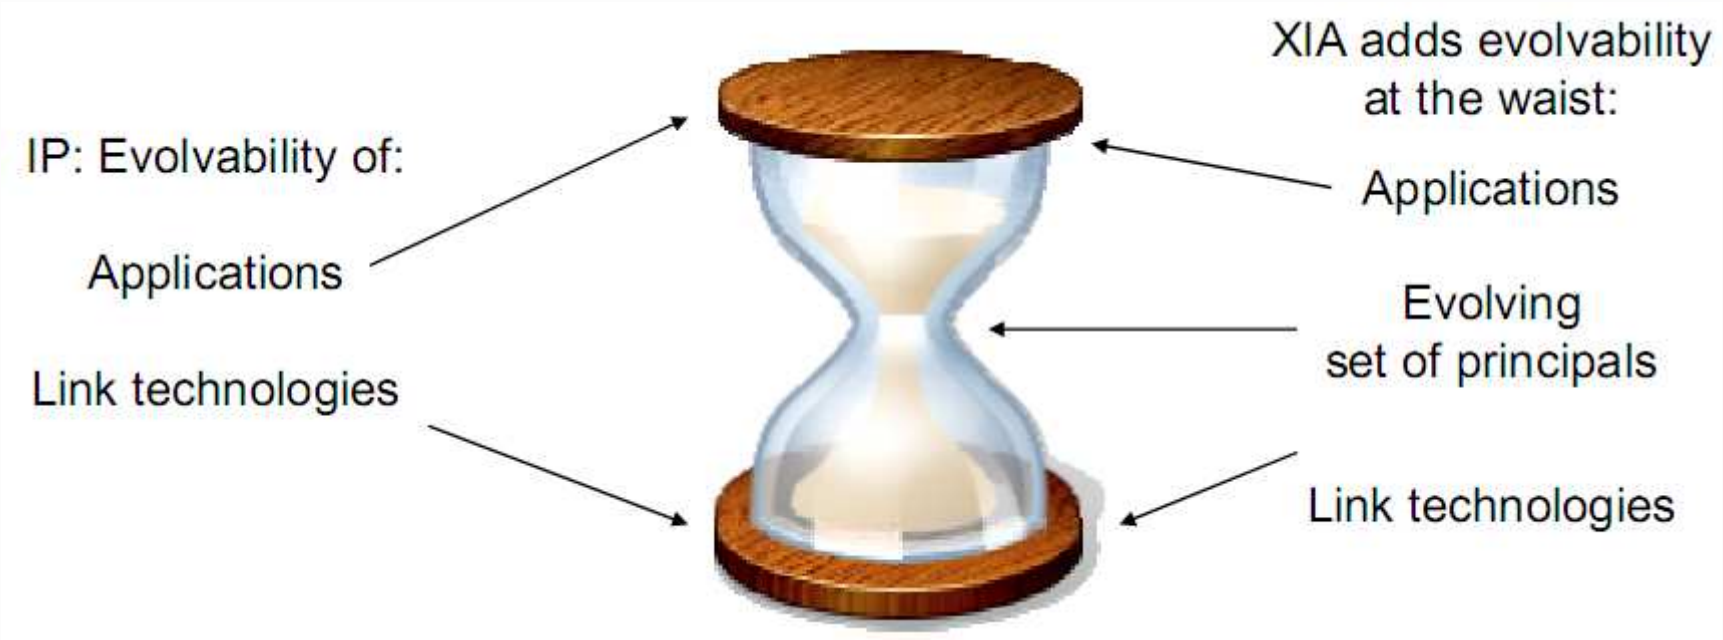

P1: Evolvable Set of Principals

- Identifying the intended communicating entities reduces complexity and overhead
  - No need to force all communication at a lower level (hosts), as in today's Internet
- Allows the network to evolve

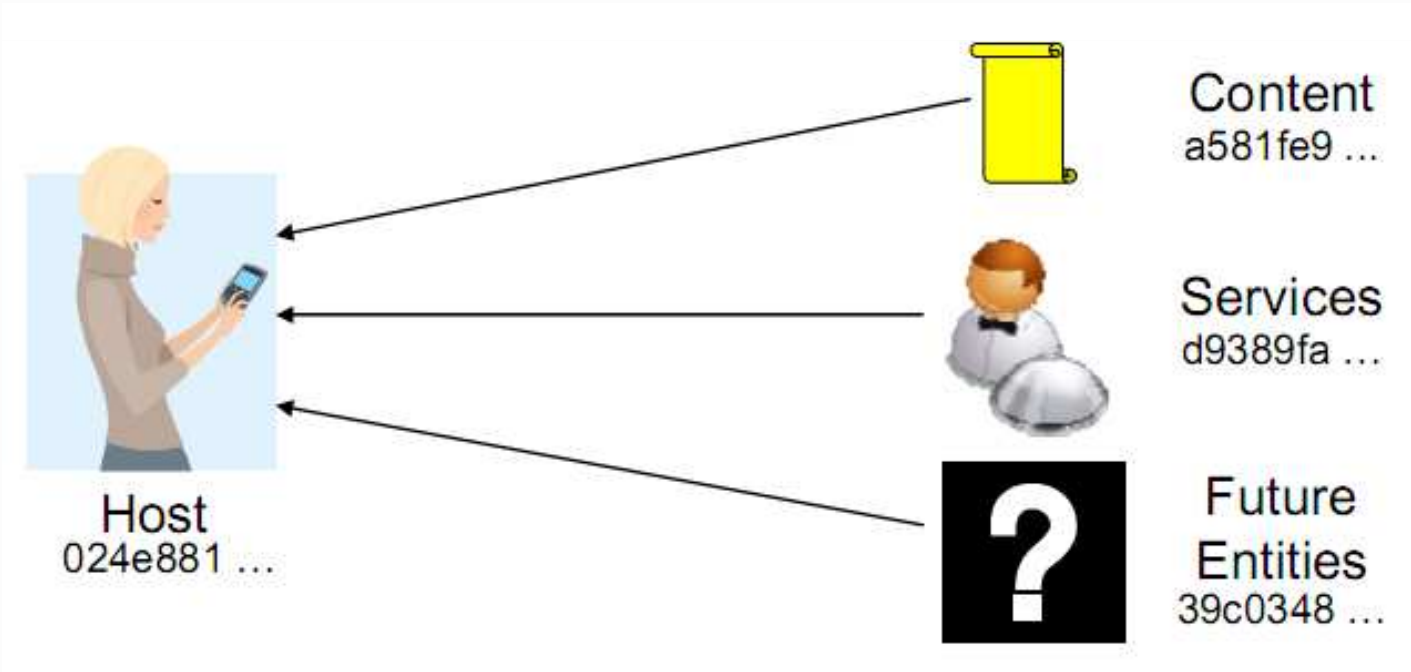

P2: Security as Intrinsic as Possible

- Security properties are a direct result of the system design
  - Do not rely on correctness of external configurations, actions, databases
  - Malicious actions can be easily identified
- XIA uses self-certifying identifiers that guarantee security properties for c ommunication operation

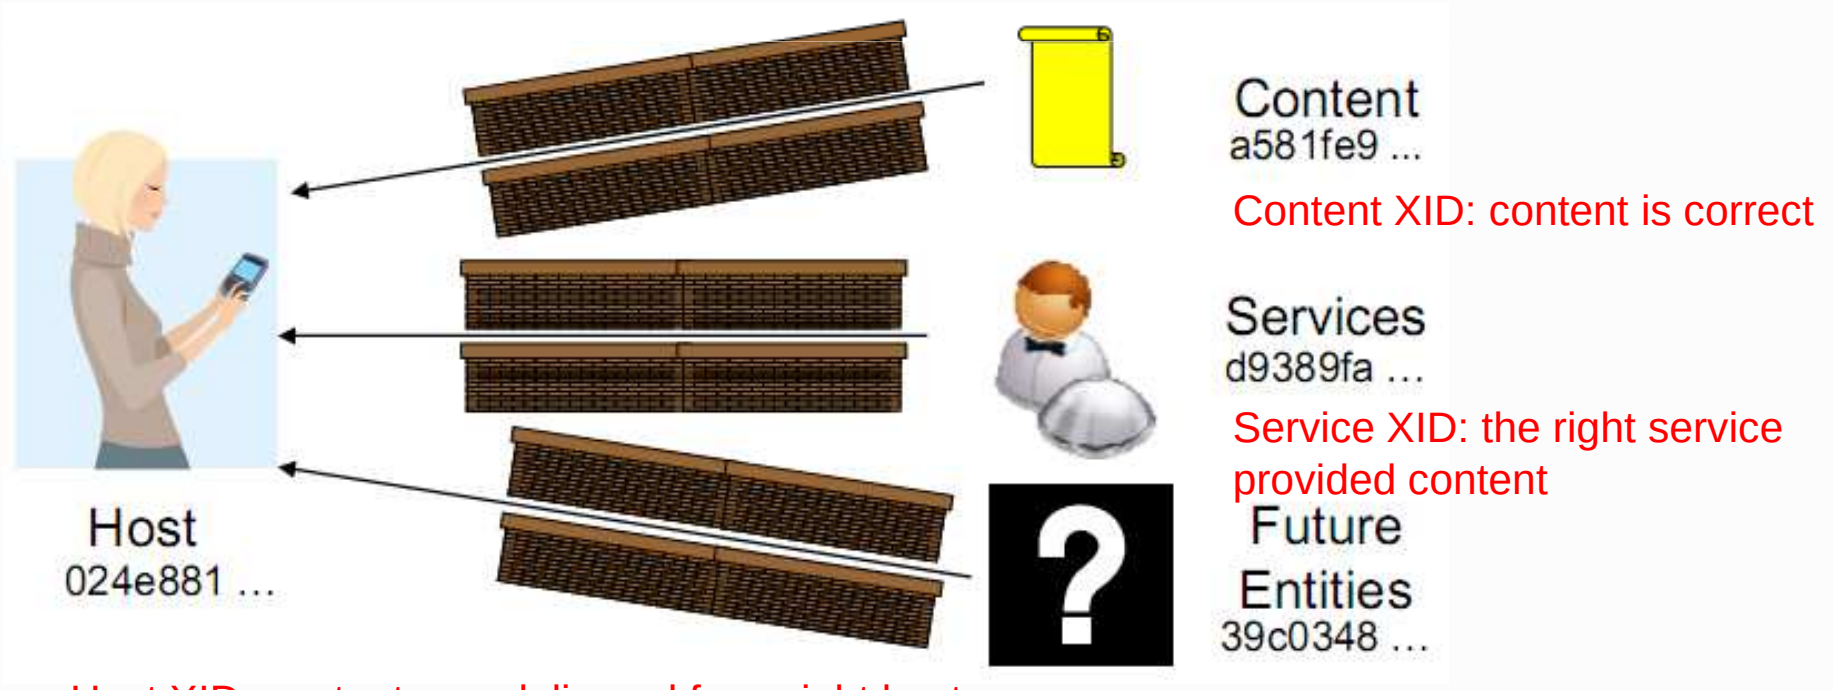

Intrinsic Security in XIA

- XIA uses self-certifying identifiers that guarantee security proper ties for communication operation
  - Host ID is a hash of its public key – accountability (AIP)
  - Content ID is a hash of the content – correctness
  - Does not rely on external configurations
- Intrinsic security is specific to the principal type
  - Important – guarantees depend on principal type
- Example: retrieve content using …
  - Content XID: content is correct
  - Service XID: the right service provided content
  - Host XID: content was delivered from right host

Other XIA Principles

- Narrow waist for trust management
- Ensure that the inputs to the intrinsically secure system match th e trust assumptions and intensions of the user
  - Narrow waist allows leveraging diverse mechanisms for trust management: CAs, reputation, personal, …
- Narrow waist for all principals
  - Defines the API between the principals and the network protocol mechanis ms
- All other network functions are explicit services
  - XIA provides a principal type for services (visible)
  - Keeps the architecture simple and easy to reason about

XIA: eXpressive Internet Architecture

- Each communication operation expresses the intent of the operation intent of the operation
  - Also: explicit trust management, APIs among actors
- XIA is a single inter-network in which all principals are connecte d
  - Not a collection of architectures implemented through, e.g., virtualization o r overlays
  - Not based on a “preferred” principal (host or content), that has to support a ll communication

Multiple Principal Types

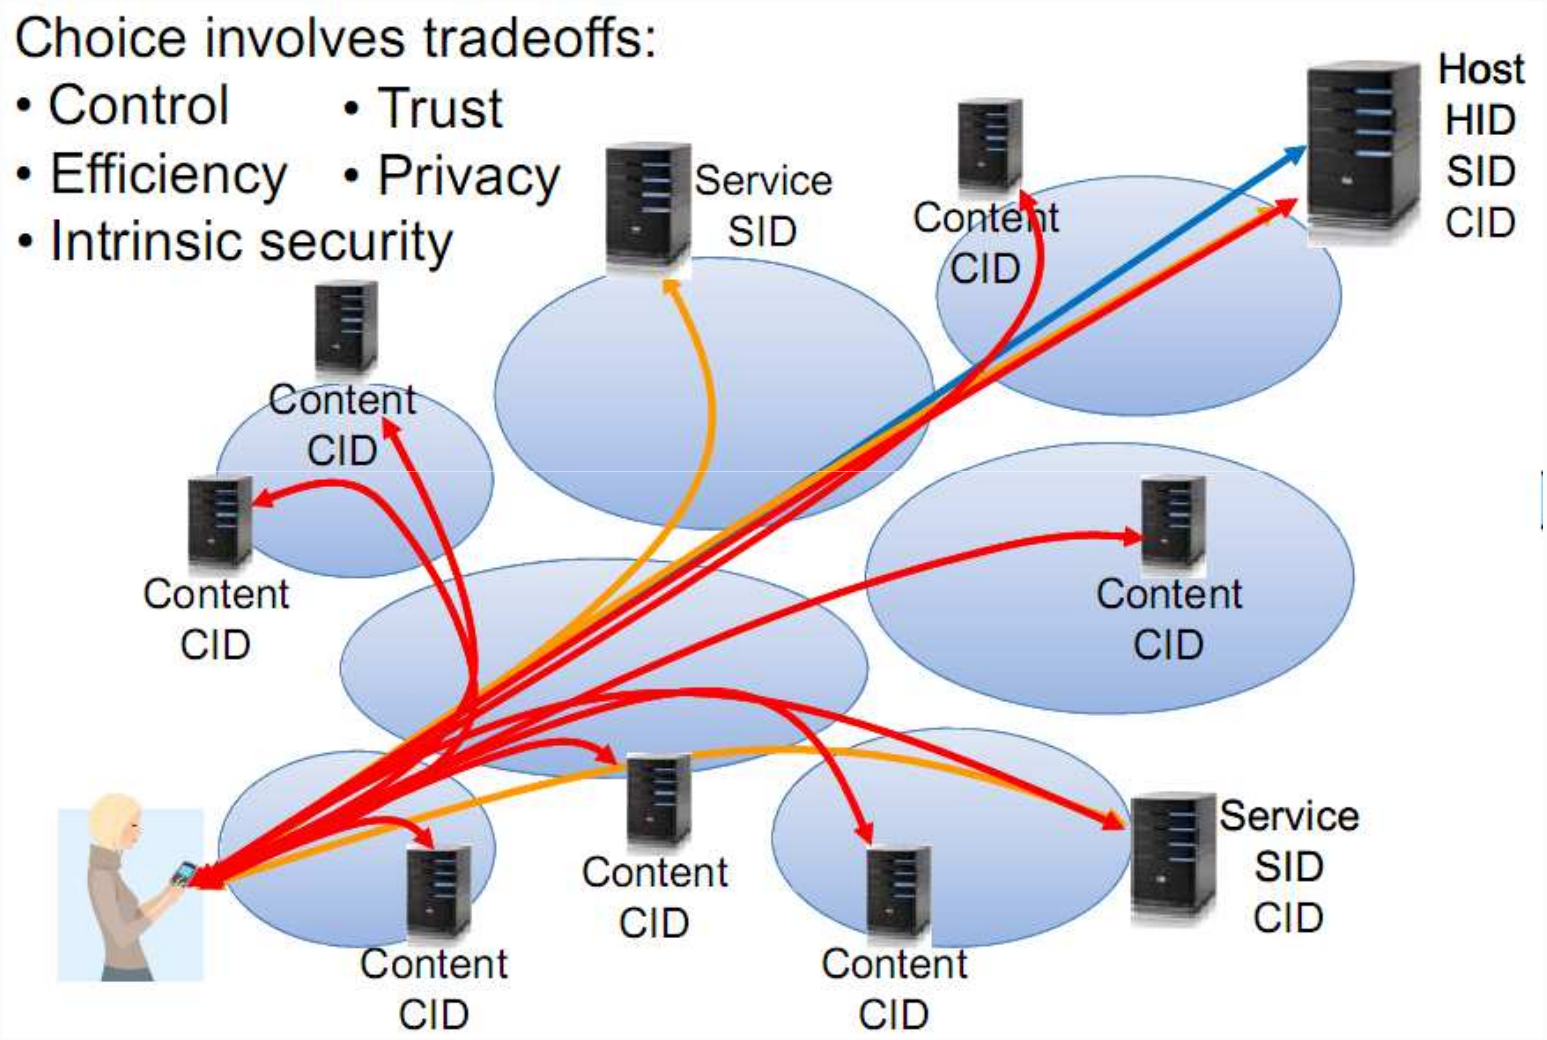

XIA: eXpressive Internet Architecture

- Each communication operation expresses the intent of the operation intent of the operation
- XIA is a single inter-network in which all principals are connected
- Multiple Principal Types:

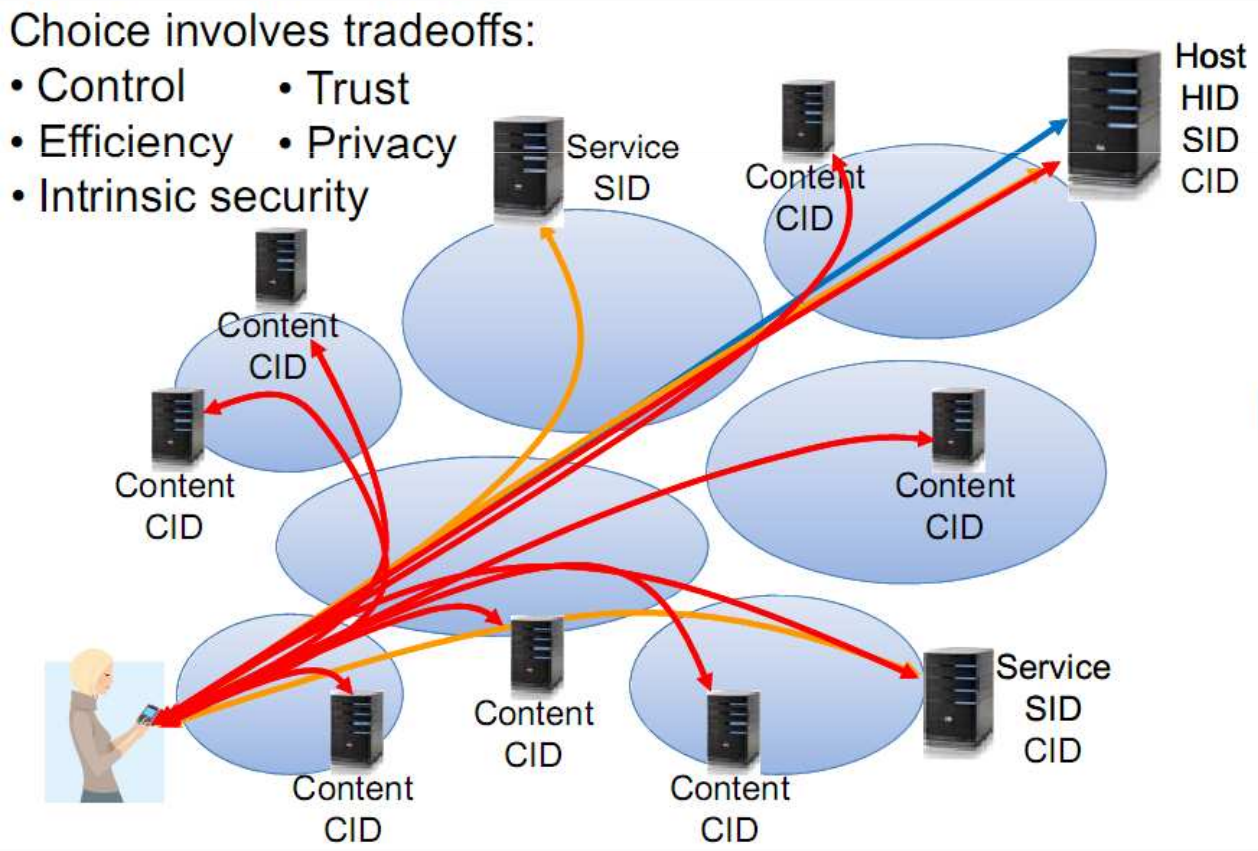

MobilityFirst

- **Principal Investigator**: Dipankar Raychaudhuri, Rutgers University/New Brunswick

- **Collaborating Institutions**: Duke University, Massachusetts Institute of Technology, University of Massachusetts/Amherst, University of Massachusetts/Lowell, University of Michigan, University of Nebraska/Lincoln, University of North Carolina/Chapel Hill

- http://mobilityfirst.winlab.rutgers.edu/TechApproach.html

MobilityFirst

The Internet does not satisfy **support host and network mobility**, as the assumption that end-hosts and networks are mostly stationary is deeply embedded in its architecture:

- **Naming and addressing**:
  - DNS binds a host name to an IP address.
  - DNS is designed to be queried frequently but updated slowly.
  - Caching enables scalability at the cost of update-propagation delay.
- **Routing:**
  - Aggregation of IP addresses into prefixes is central to routing scalability.
  - Host mobility is assumed to be infrequent enough to be relegated to an indirect routing approach that is inefficient and requires additional infrastructure not universally available.
  - Network mobility (e.g., mobile networks of vehicles, planes, or VMs) scales poorly with the Internet's interdomain routing protocol.
- **Disruption-tolerance:**
  - TCP/IP stack assumes that hosts and routers are stationary long enough so that the routing protocol can compute a contemporaneous end-to-end path and the transport protocol can respond to end-to-end loss and congestion signals.
  - TCP/IP stack falls short in emerging wireless scenarios (e.g., vehicular, sensor).

MobilityFirst - Requirements

The *MobilityFirst* architecture is based on the following high-level requirements:

- **Mobility as the norm:** Seamless host and network mobility at scale; multi-provider mobile network access; heterogeneous wireless technologies.
- **Robustness:** with respect to intrinsic properties of wireless medium (disconnection, varying bandwidth, high error rates, scarce spectrum).
- **Trustworthiness:** Enhanced security for mobile networks and wired infrastructure (strong authentication, enhanced trust models, privacy, DDoS resistance, secure routing).
- **Usability:** Architectural support for context-aware pervasive mobile services; evolvable core network services; network manageability; economic viability, regulability and universal access.

MobilityFirst - Design goals

- Unlike the current Internet however, mobility and trustworthiness form the starting point for these design goals.
- **Host and network mobility (G1)**: End-to-end communication must continue (i) despite frequent mobility of end-hosts or networks; (ii) despite the absence of a contemporaneous end-to-end path.
- **No global root of trust (G2):** Correct network behavior must not depend on a single root of trust.
- **Intentional data receipt (G3)**: An end-host must receive data only if the transmission is consistent with its receipt policy.
- **Byzantine robustness (G4)**: End-to-end communication must continue despite the compromise of (a small fraction of) end-hosts or infrastructural nodes.
- **Content addressability (G5)**: The network should assist with content retrieval in addition to enabling host-to-host communication.
- **Evolvable network services (G6)**: The architecture should allow for the co-existence or rapid creation of new and different network services.

MobilityFirst - Design principles

- **Visibility and choice (P1)**: Networks and end-users should have visibility and choice in determining what resources are available and how they are used.
- **Usability (P2):** The architecture should be easy to use for end-users, operators and app developers, suggesting plug-and-play interfaces and socket APIs that simplify app development.
- **Manageability (P3):** Networks should be easy to manage. This principle suggests a well-instrumented management plane that provides necessary aggregate views to operators.
- **Simplicity (P4):** As a general principle, the design should favor simple methods that do not involve significant control complexity or maintain per-flow or per-packet state in the network.
- **Regulability (P5):** The architecture should confirm to specifiable public policies and laws.
- **Commercializability (P6):** The architecture must be economically viable.
- **Technology-awareness (P7)**: The architecture should be cognizant of foreseeable technology trends.

MobilityFirst key protocol features

- Separation of naming & addressing
- Fast global naming service
- Storage-aware (GDTN) routing
- Hop-by-hop (segmented) transport
- Self-certifying public key names
- Support for content/context/location
- Programmable computing layer
- Separate network mgmt plane

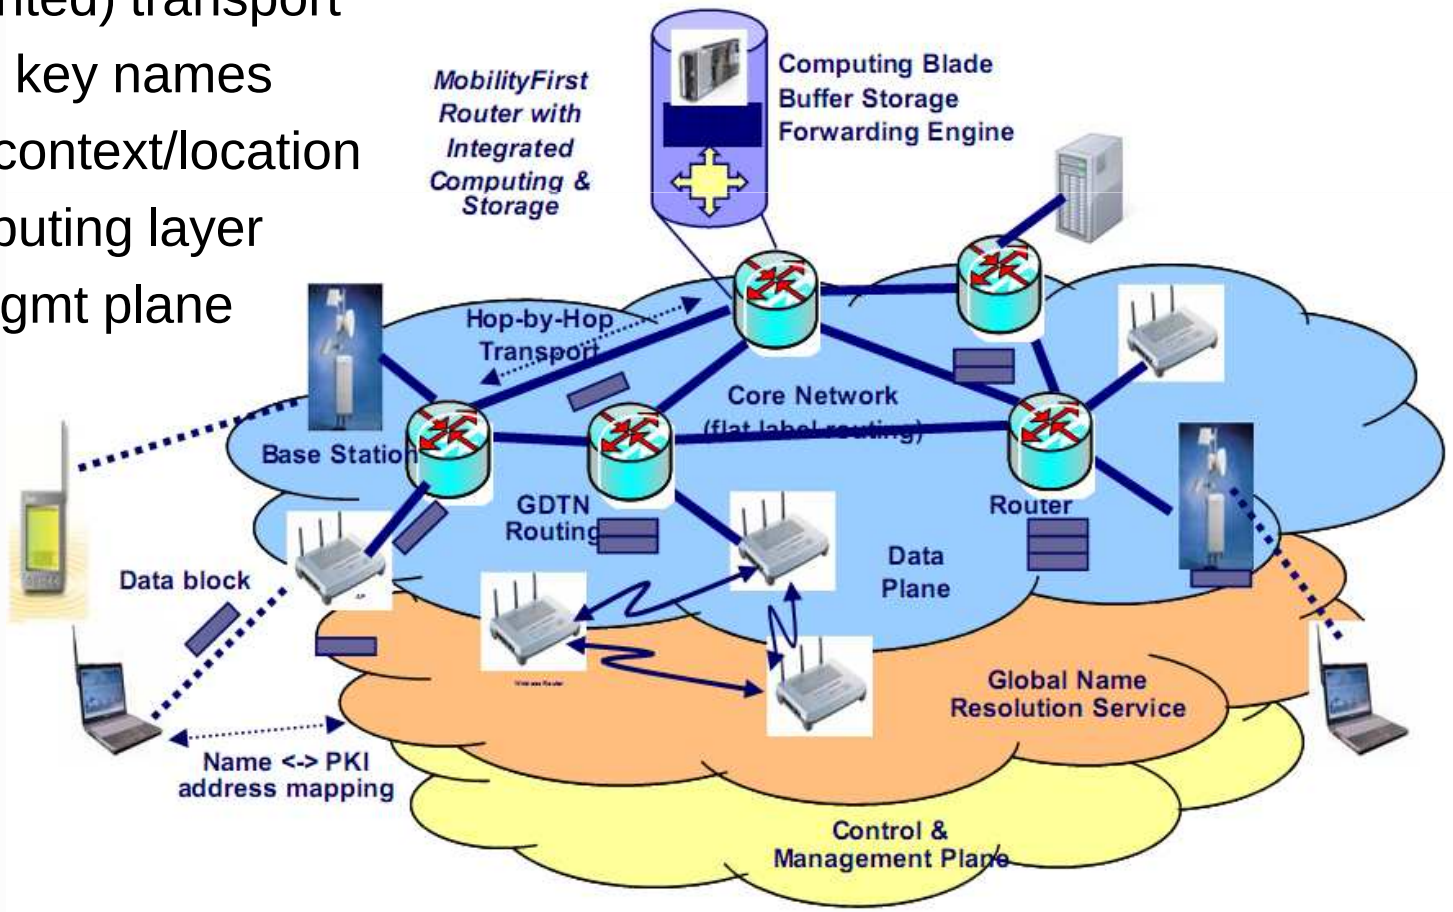

MobilityFirst - Overview of Component Architecture

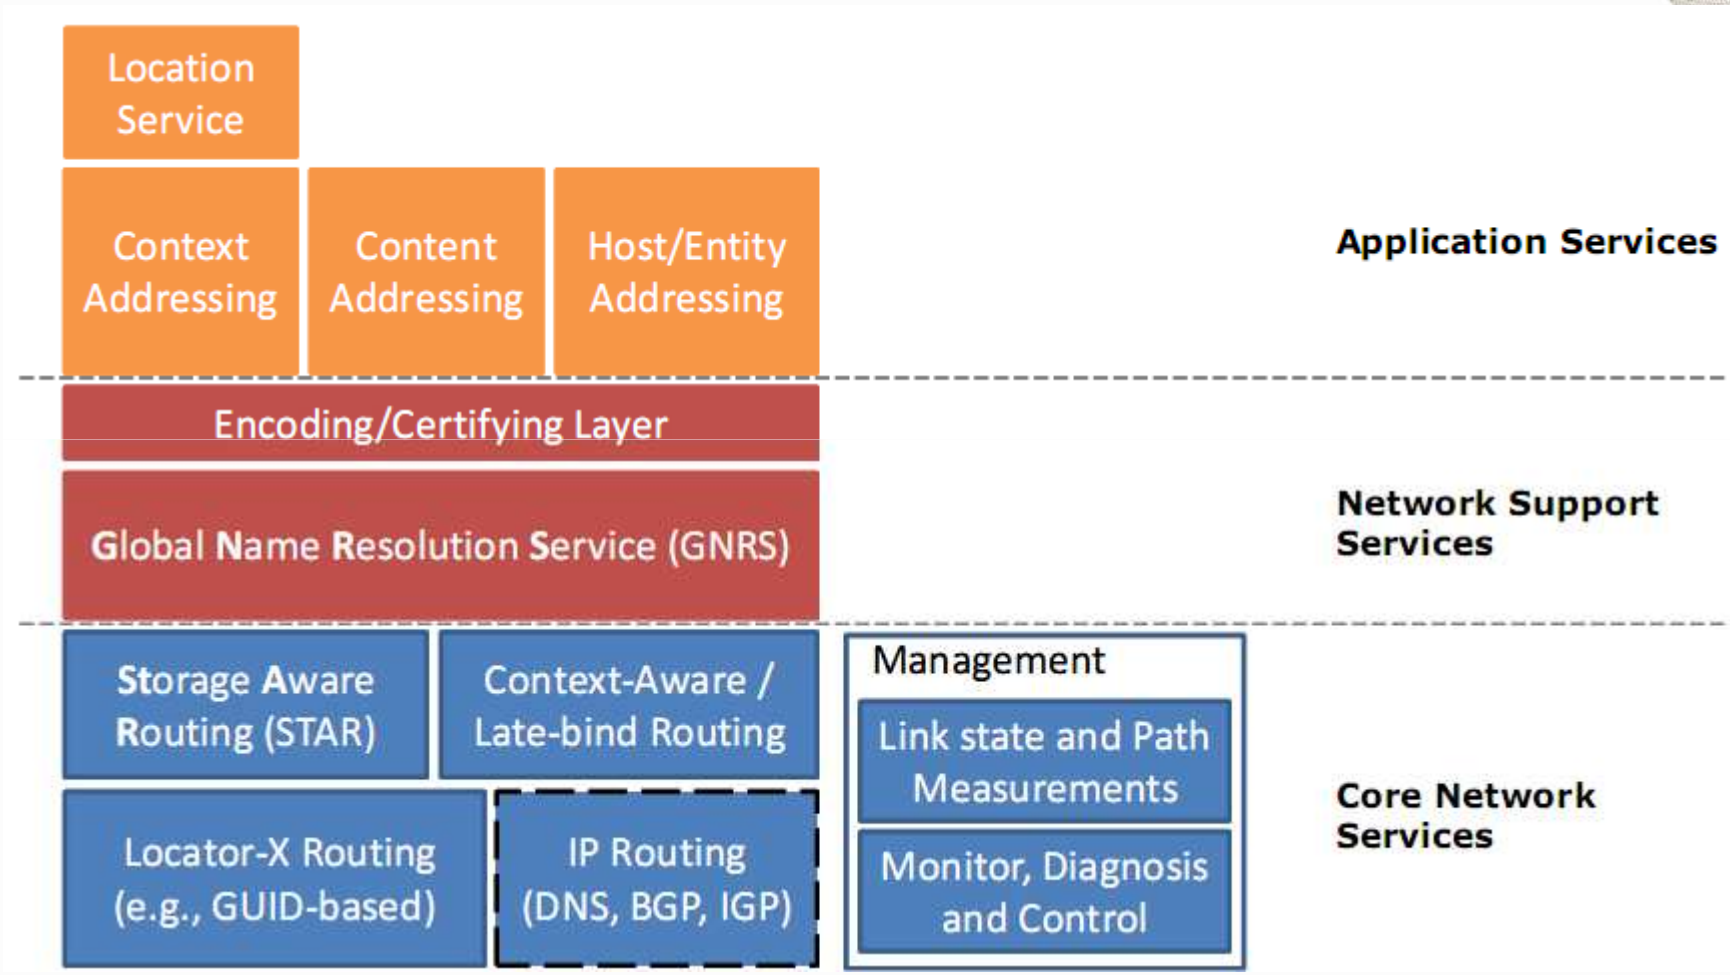

NEBULA

- **Principal Investigator**: Jonathan Smith, University of Pennsylvania

- **Collaborating Institutions**: Cornell University, Massachusetts Institute of Technology, Princeton University, Purdue University, Stanford University, Stevens Institute of Technology, University of California/Berkley, University of Delaware, University of Illinois/Urbana-Champaign, University of Texas, University of Washington

- http://nebula.cis.upenn.edu/NEBULA_brief

NEBULA

- NEBULA is an architecture for the Cloud-based Future Internet
  - More secure and reliable
  - Deployable and evolvable
  - Truly clean-slate
- Technology, Economics and Policy continue to evolve
  - NEBULA co-design with Economist and Lawyer on team
- Motivation for a new architecture:
  - Availability: At risk of network outages
  - Security:
    - Poor endpoint authentication
    - HIPAA policy restrictions not expressible with existing routing protocols
  - Consistency:
    - Communications end point focused, not data focused
    - Cloud systems have embraced weak consistency CAP Theorem)!

NEBULA – Architecture and Principles

Architecture:

- Services provided by cloud data centers
- Multiple cloud providers, that each use replication
- Agreements for "roaming" allow user to connect to nearest center
- Variety of access mechanisms (wired & wireless)
- Transit networks to connect access to data centers

Principles:

- Ultra reliable, high-speed core interconnecting data centers
- Parallel paths between data center and core router
- Secure access and transit
- Policy-based path selection
- Authentication during connection establishment

NEBULA: A Network Architecture to Enable Security

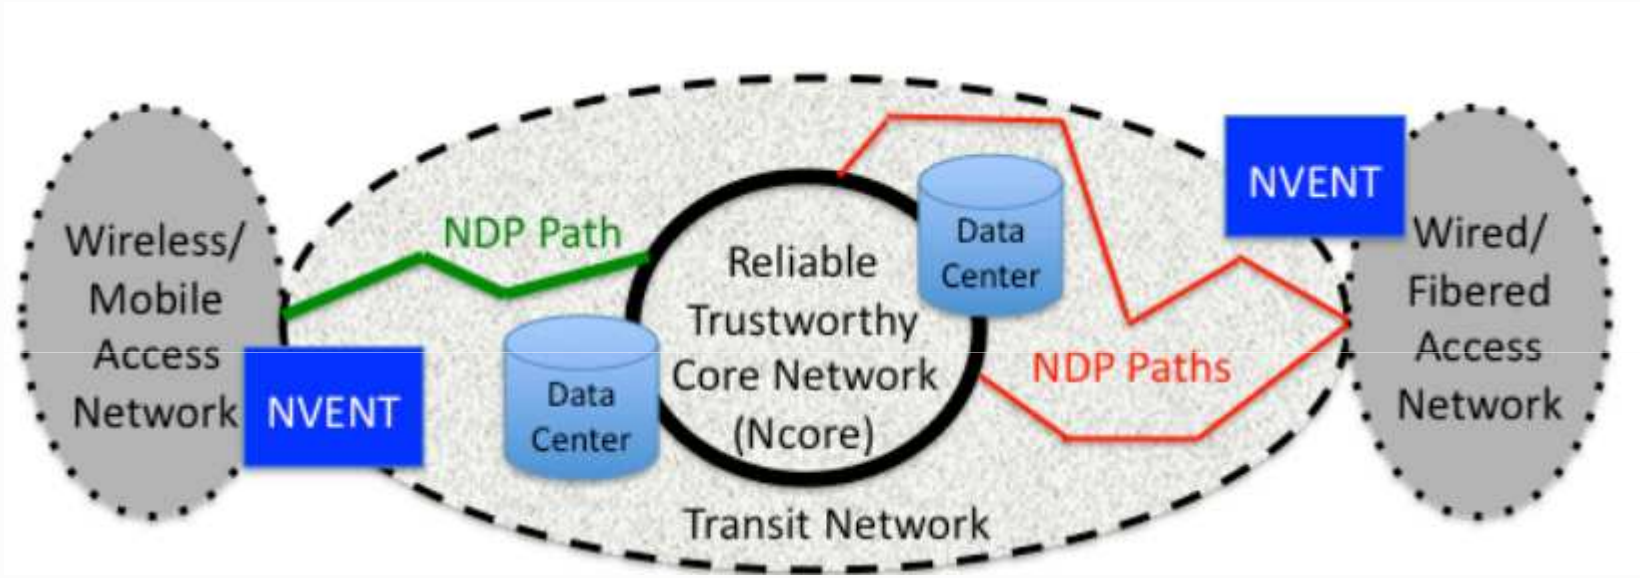

NDP – NEBULA Data Plane – distributed path establishment with guarantees
NVENT – NEBULA Virtual and Extensible Networking Techniques
– extensible control plane
NCore – NEBULA Core – redundantly connected highFavailability routers

NEBULA Data Plane (NDP) in a nutshell

- Use cryptography for:
  - Proof of consent (PoC) – route authorized?
  - Proof of path (PoP) – route followed?

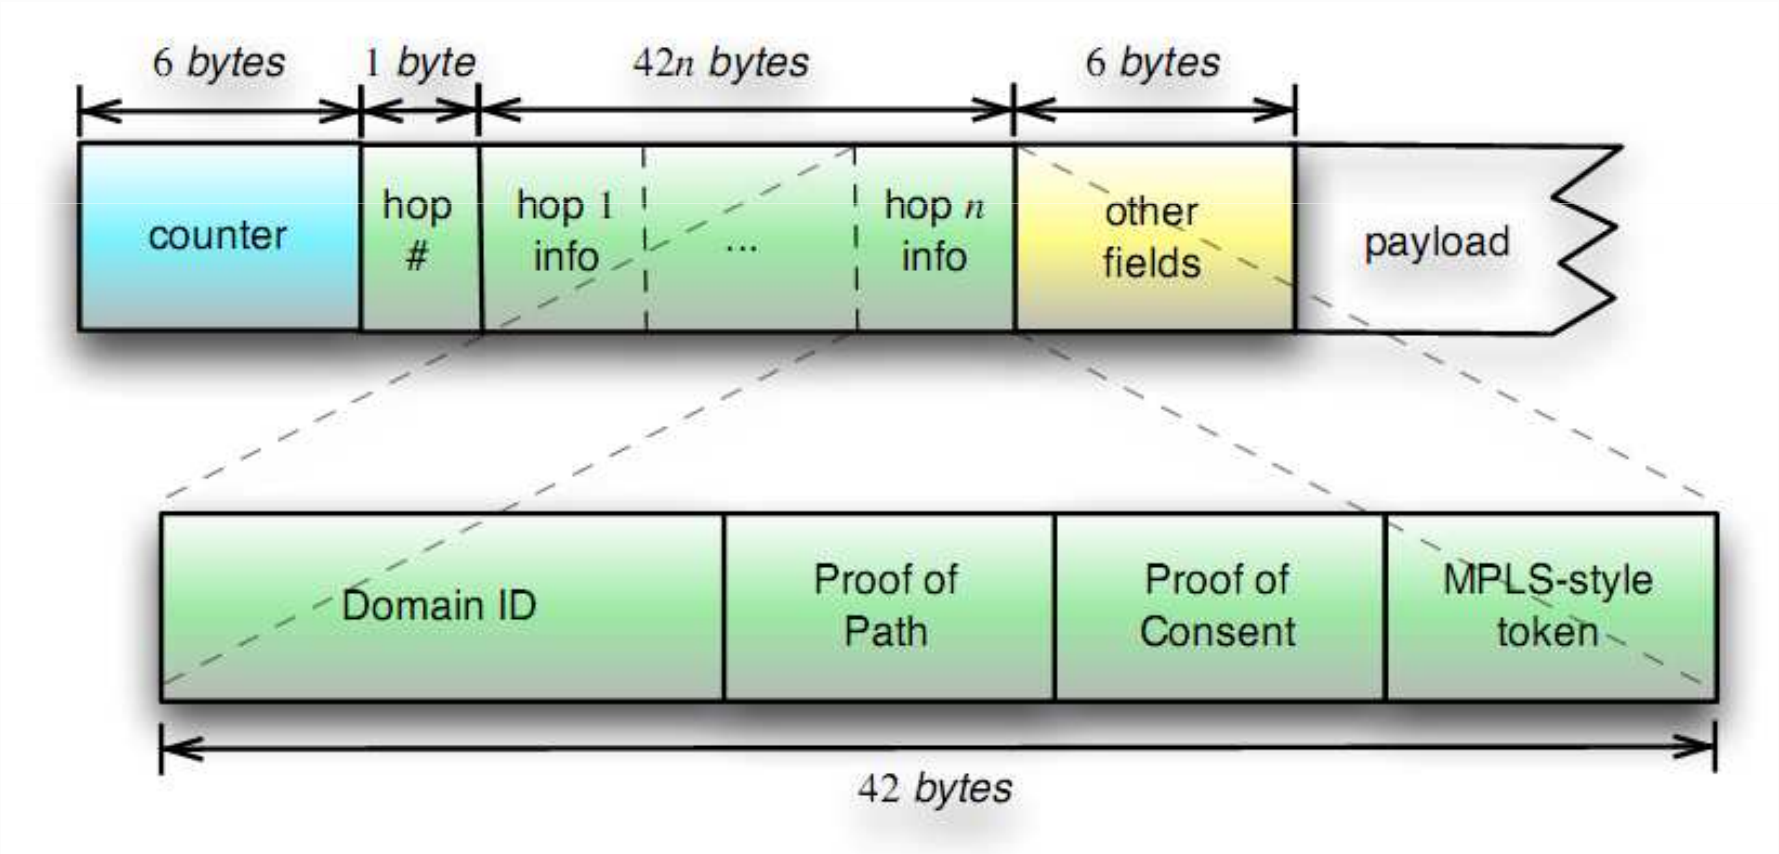

NEBULA Virtual and Extensible Network Techniques (NVENT)

- Service discovery
- New service injection
- Secure control plane for naming, border gateways, etc.
- Generalized path discovery for specifying policies, multiple paths and dynamic path construction via NDP

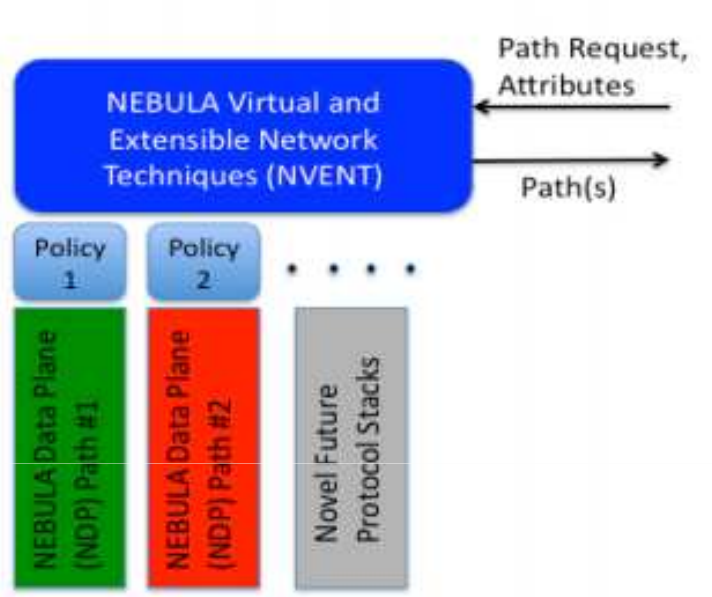

Extensible: Paths over new substrates
Deployable: Linux implementation of reliable BGP

NEBULA Core (NCore)

- High availability via redundant high-throughput links and novel high-availability router control software
- Reliable distributed software builds a routing-complex-from multiple chassis

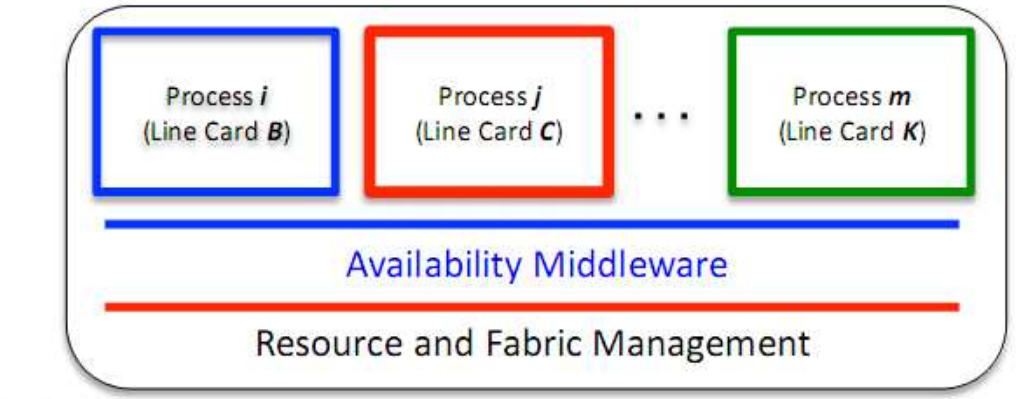

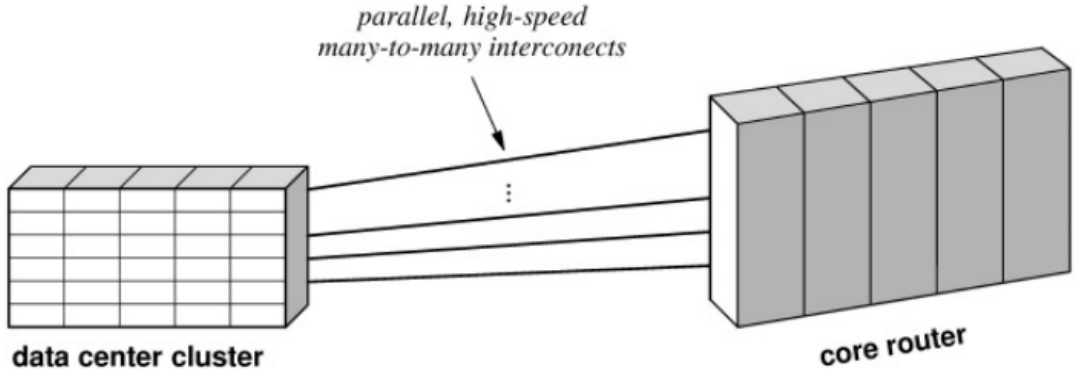

NEBULA Architectural Criteria

Design Goal	NEBULA
Communication must continue despite loss of networks, links, or gateways.	NEBULA uses multiple dynamically allocated paths and reliable transport.
Allow host attachment and operation with a low level of effort	NVENT/NDP is as easy to automate and use as DHCP/IP.
Support secure communication (authentication, authorization, integrity, confidentiality) among trusted nodes.	Mutually suspicious NDP nodes self-select paths exhibiting cryptographic proofs of properties required for security.
Provide a cost-effective communications infrastructure	NCORE places resources where architecturally needed; regulatory/policy analysis.
Implement network and user policies	Policies implemented with NDP and NVENT.
The architecture must accommodate a variety of networks.	NDP sends packets by encapsulation, NVENT networks by virtualization
The architecture must permit distributed management of its resources.	NDP path establishment decentralized, NVENT

Trends and Experimental Research

Future Internet Technology Trends

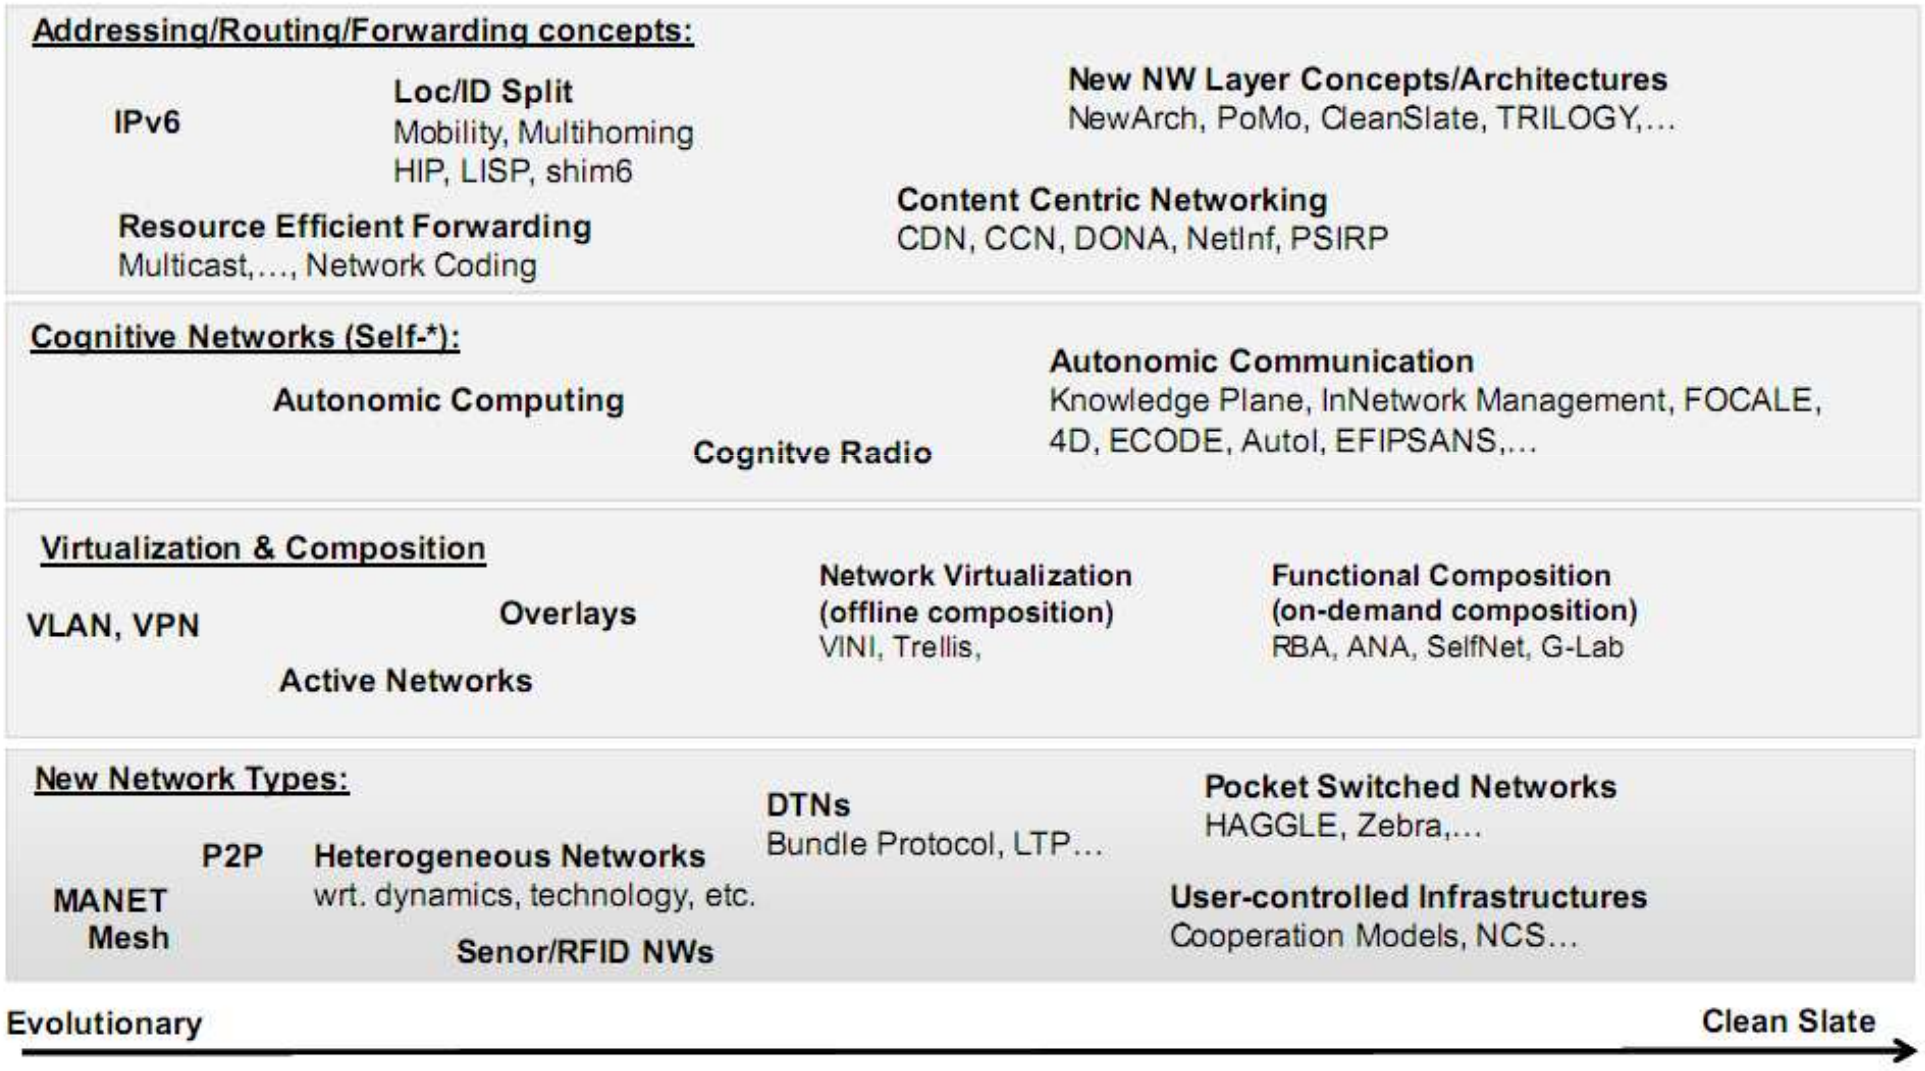

Source: T. Tseby, “Future Internet Technologies”

Future Internet Research and Evolution

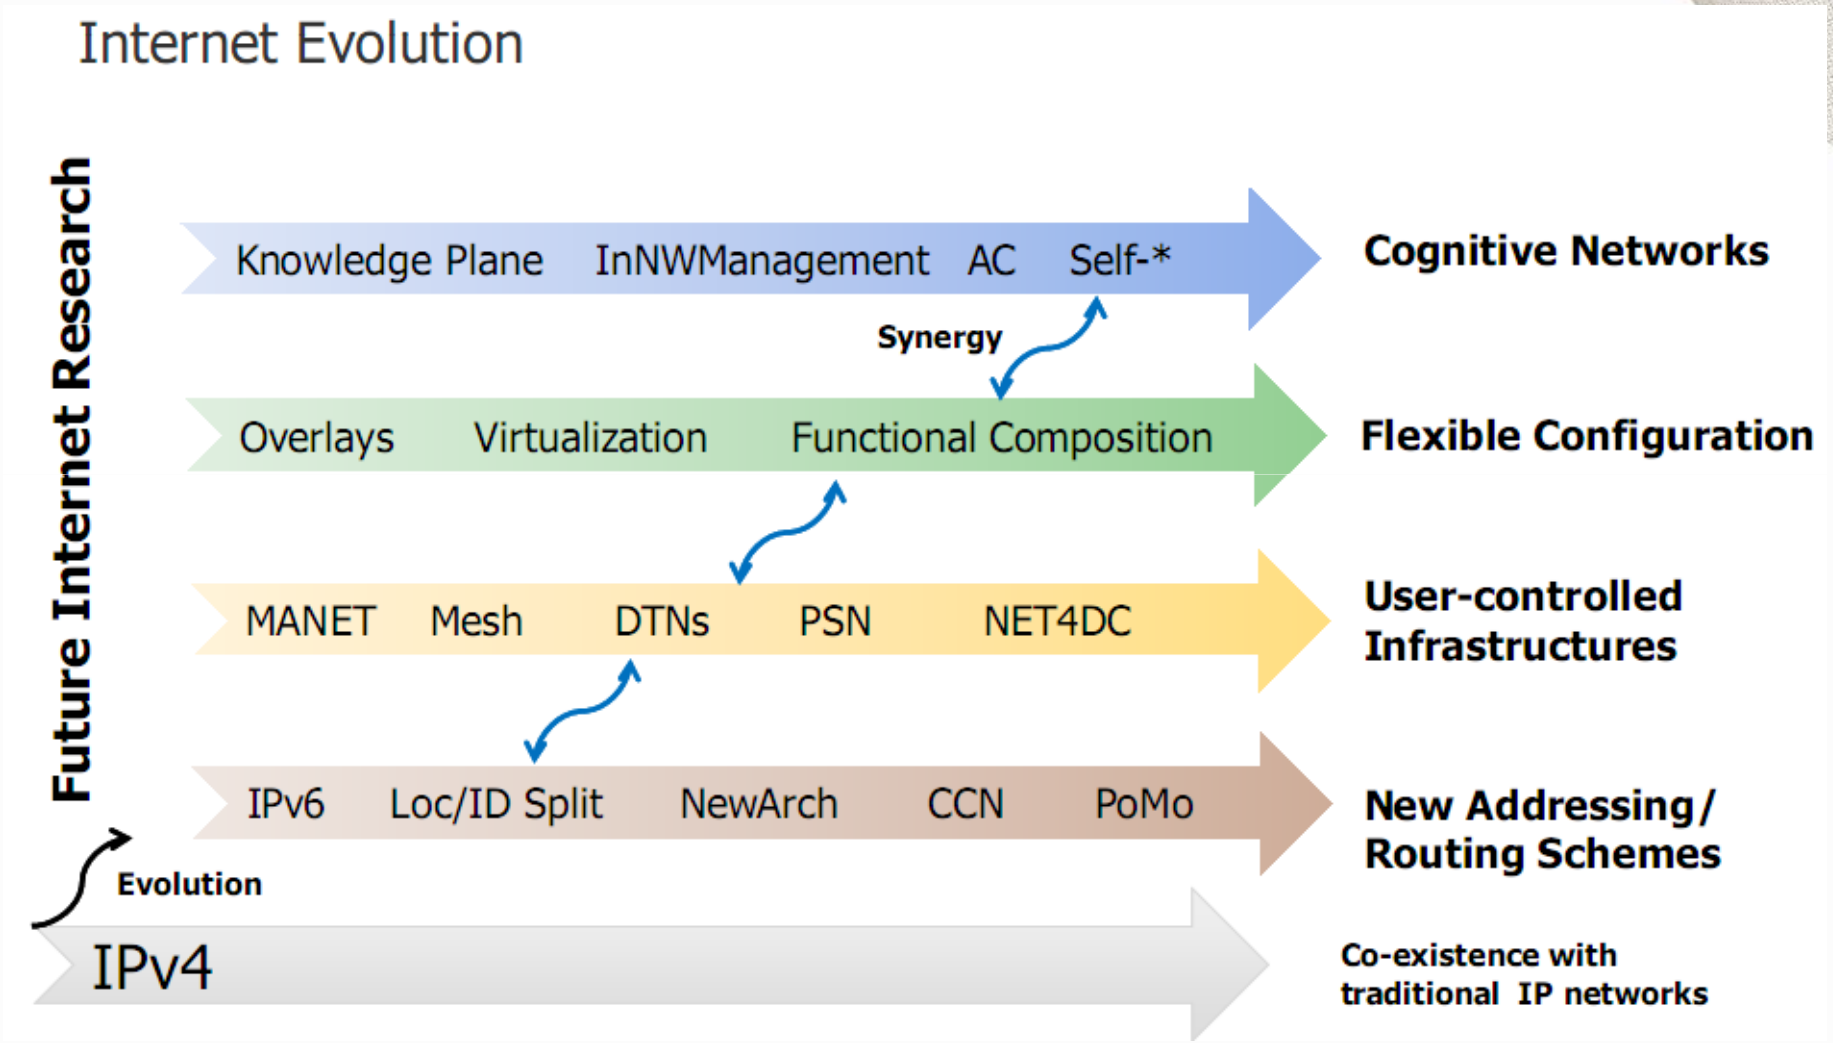

Source: T. Tseby, “Future Internet Technologies”

EU Future Internet Research Environments

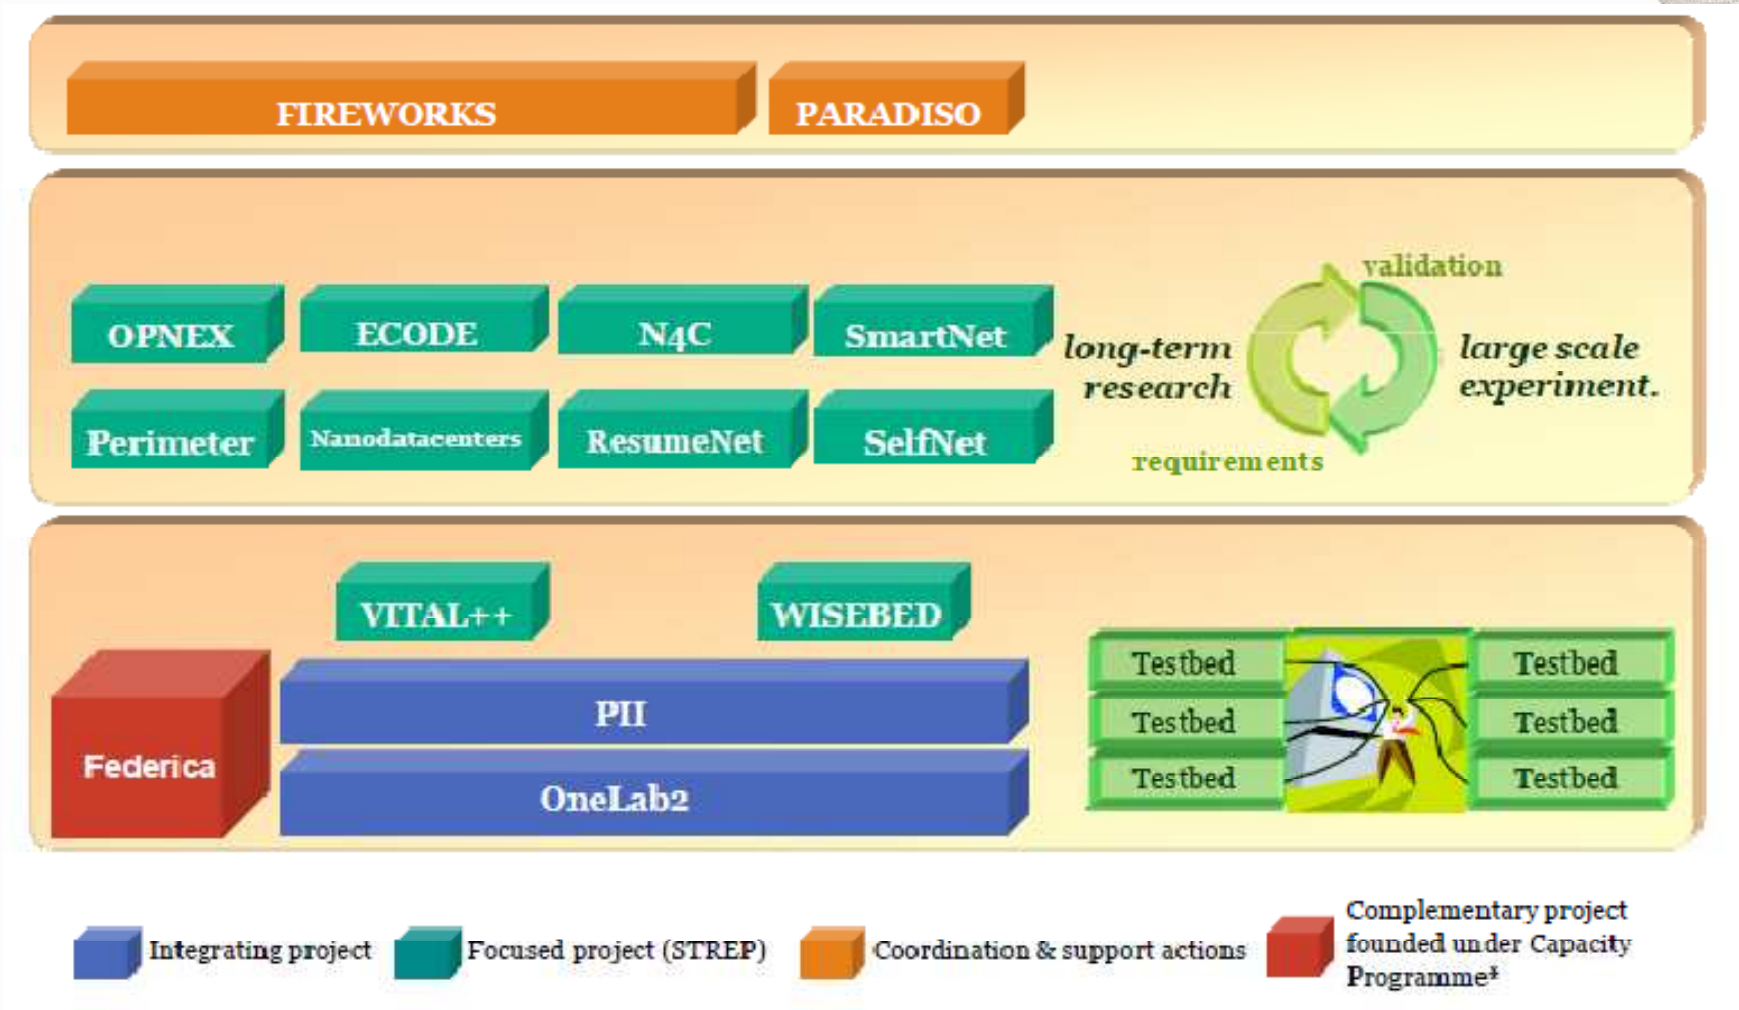

http://cordis.europa.eu/fp7/ict/fire/fire-fp7_en.html

GENI

Global Environment for Network Innovations

- Initiative of the US National Science Foundation (NSF)
- **Goal:** build an open, large-scale, realistic experimental facility
  - for evaluating new network architectures
  - create customized virtual network
  - opportunity to experiment independent of today's Internet (e.g. conditions, assumptions)
  - slices of resources in space and time

GENI Design Principles

- Physical network ‘substrate’
  - building block components
  - elements / nodes / links / subnets
- Software control & management framework
  - knits building blocks together
  - allows many parallel experiments (slices)
  - creates arbitrary logical topologies (virtualization)
- Programmable for ‘Clean Slate’ research
- Instrumented for accurate analysis
- Flexible and Phased Design
  - Support Technology Introduction during GENI Lifetime

geni
Exploring Networks of the Future

Slicing and Virtualization

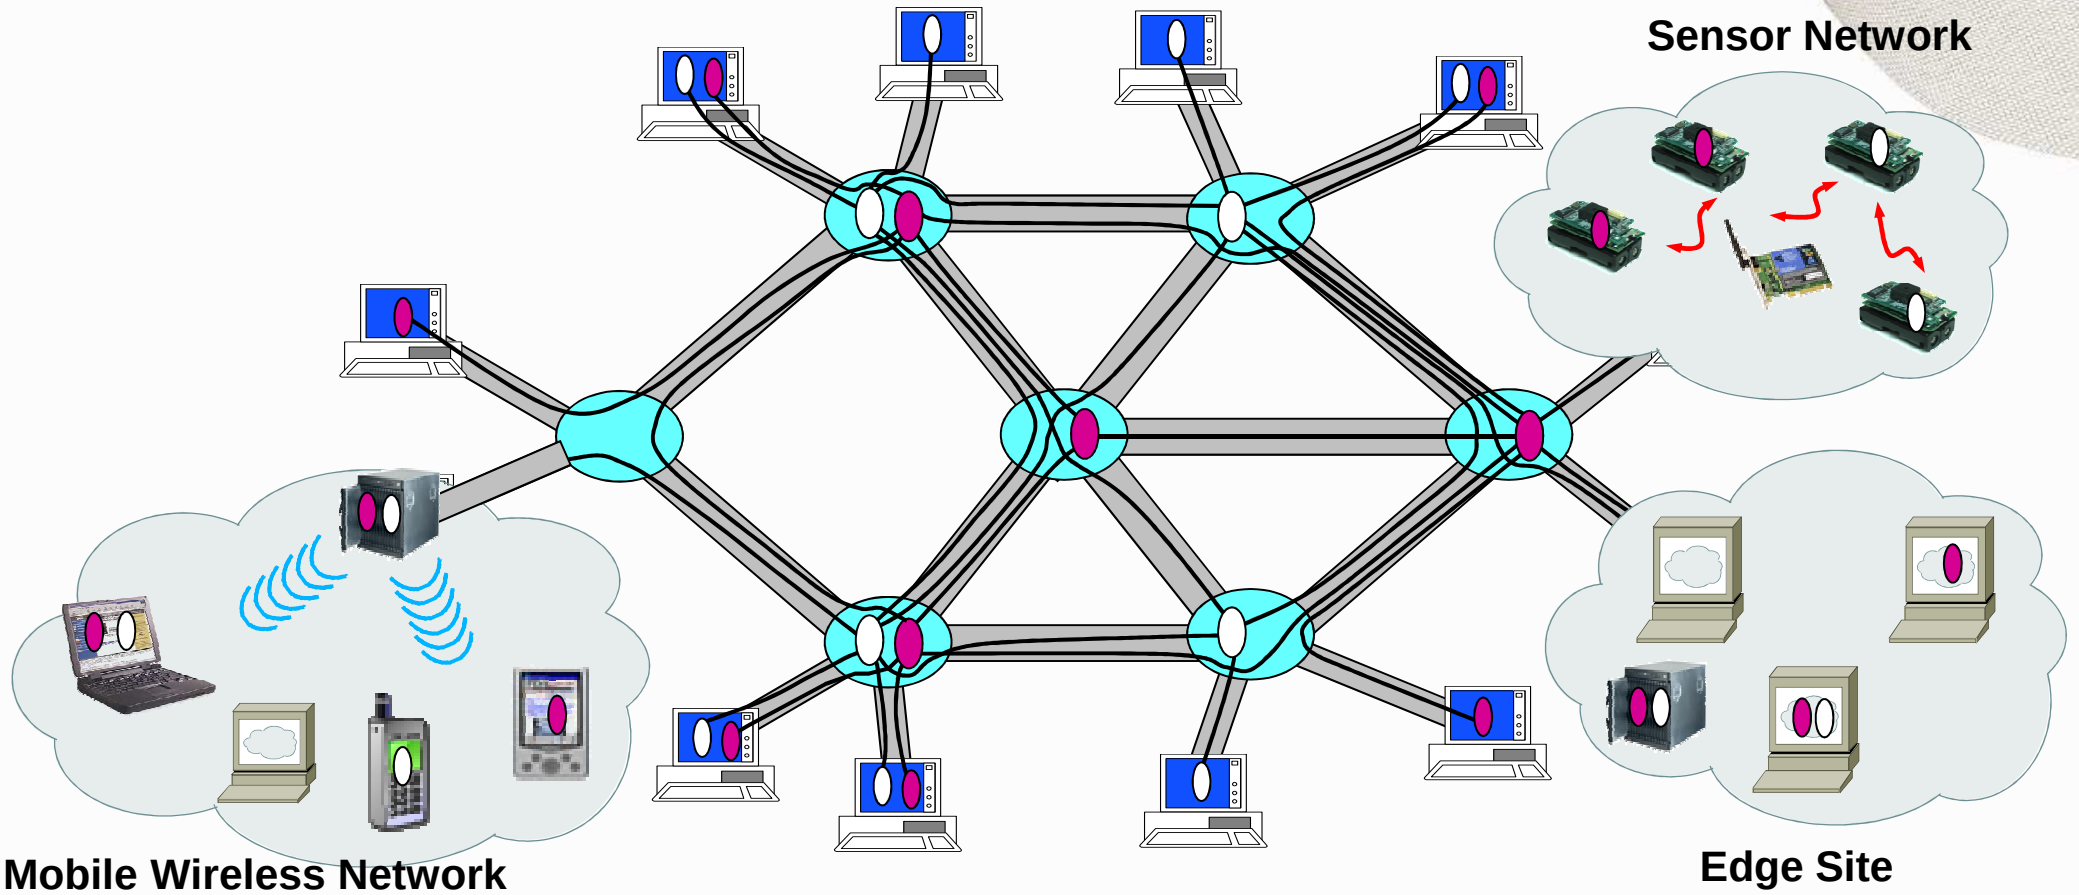

- share resources to support many simultaneous experiments

A testbed with Network Programmability

Network elements programmable via open interfaces and/or downloadable user code

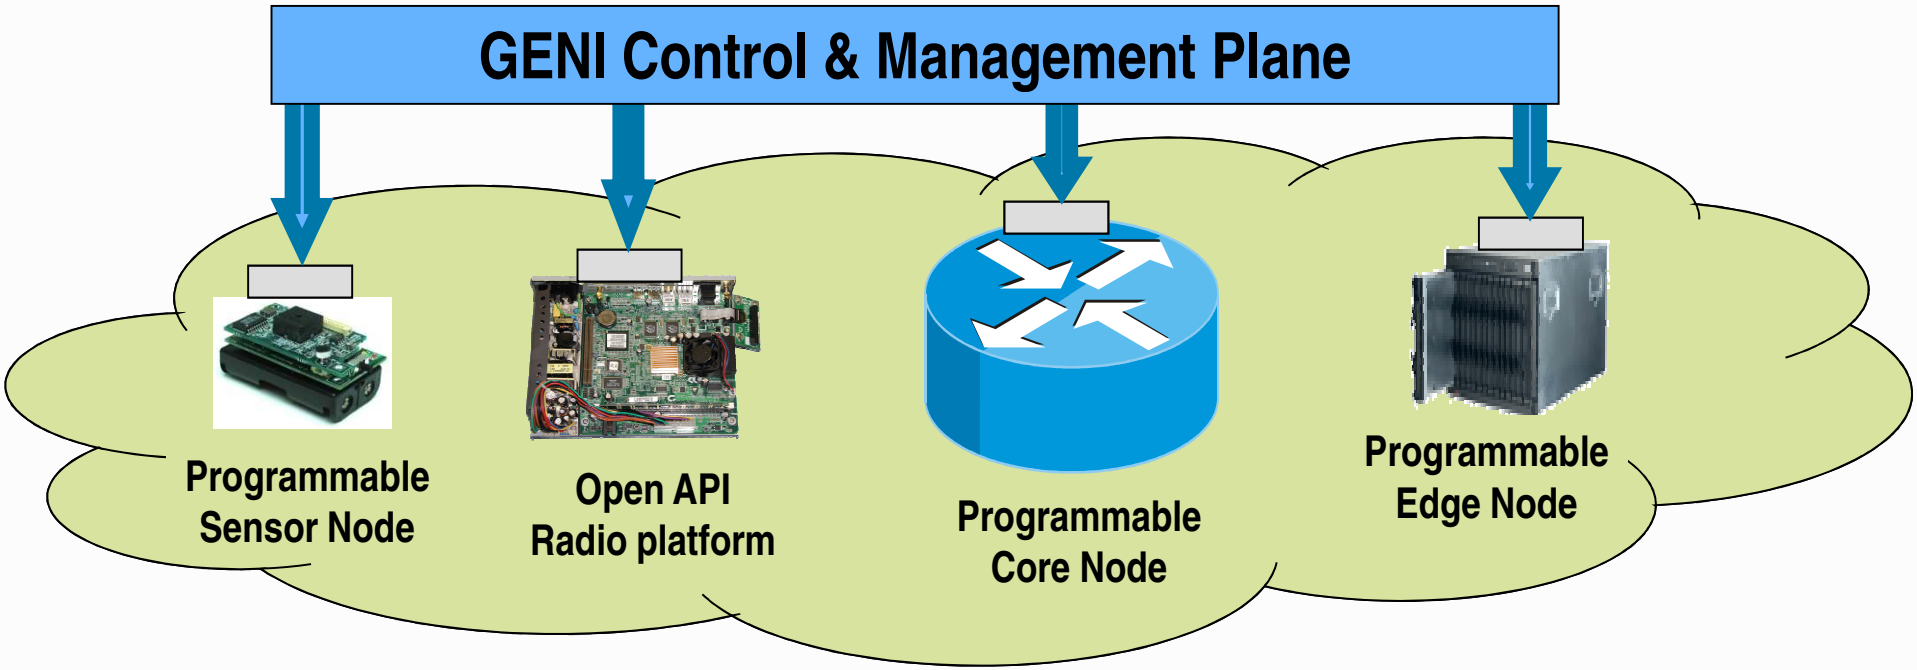

Software-defined networking

- A pragmatic approach to provide a substrate for switch programmability
- Open-source control software
  - Leads to innovation
- Out-source intelligence to commodity PCs
  - Leads to lightweight, inexpensive, commoditized but customizable switches
- Flow model
  - Simple
  - Plumbing
  - Control
  - Rewrite
- External open API to flow-table

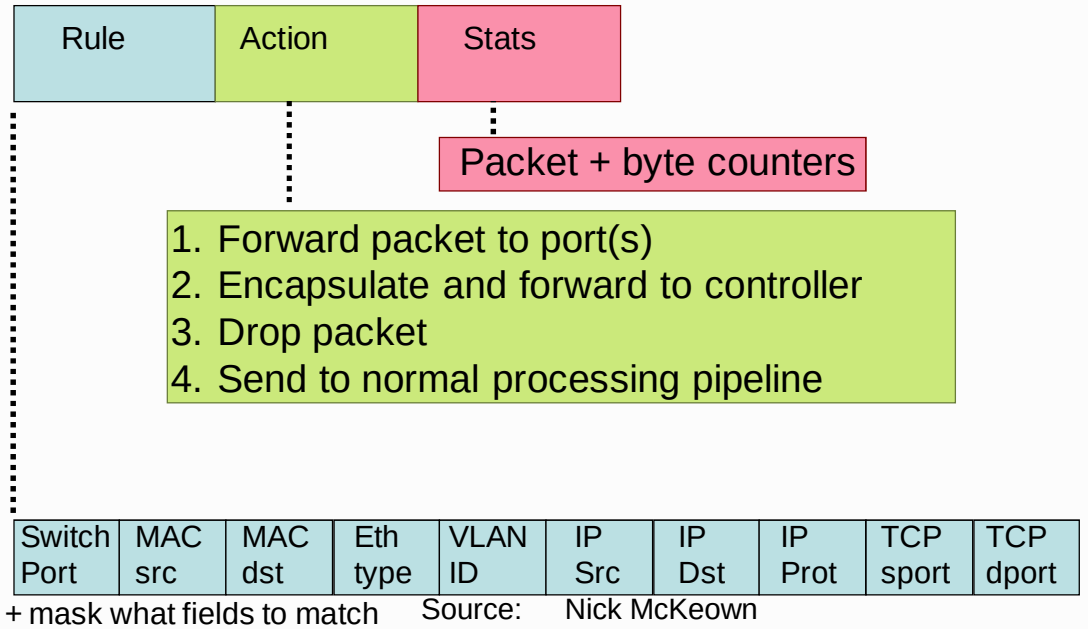

Source: Nick McKeown

Conclusions

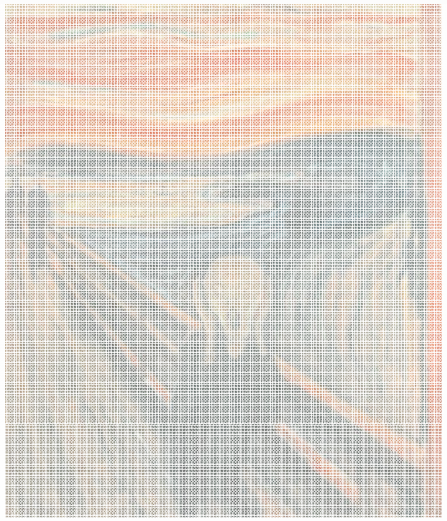

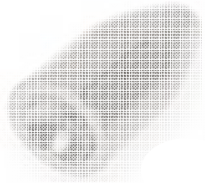

Future Internet

Obrigado!

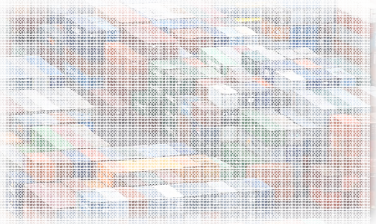

Perguntas?

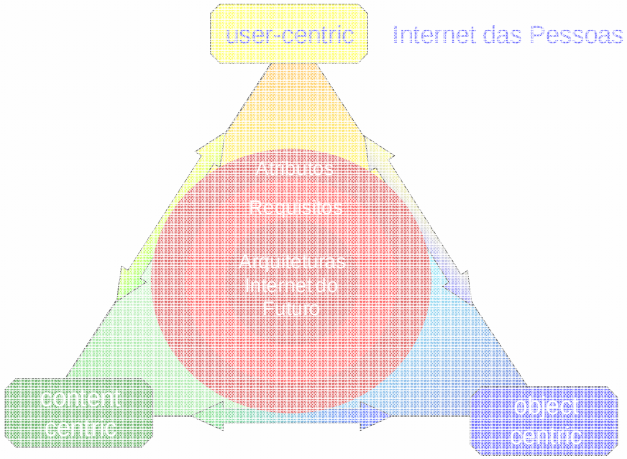

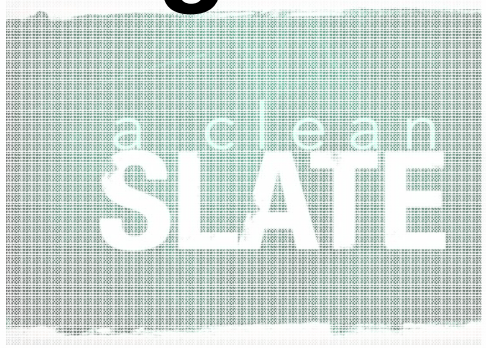

REFERENCES AND FURTHER READING

- T. Tronco (ed.) New Network Architectures, Studies in Computational Intelligence, 2010, Volume 297/2010
- B. Leiner, V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, and S. Wolff, “A brief history of the internet.” SIGCOMM Computer Communication. Rev. 39, 5 (Oct. 2009)
- D. Clark, “Moving FIND to the next stage”, Jul. 2009 http://groups.csail.mit.edu/ana/People/DDC/Working%20Papers.html
- T. Tseby, “Future Internet Technologies -- A Technical Overview of Evolutionary and Revolutionary Ideas”, 2010
- A. Feldmann, “Future Internet Revisiting the Internet architecture?”
- J. Rexford and C. Dovrolis, "Future Internet architecture: Clean-slate vs. evolutionary research," Communications of the ACM, September 2010.
- E. Mikóczy, I. Kotuliak, O. Deventer. Evolution of the Converged NGN Service Platforms Towards Future Networks. Future Internet 2011, 3, 67-86.

Source:

BACK-UP