Prototyping of the MobilityFirst Future Internet Architecture using the ORBIT and GENI Testbeds

GREE Workshop at GEC-16

Salt Lake City, March 21, 2013

D. Raychaudhuri
WINLAB, Rutgers University
ray@winlab.rutgers.edu

Introduction

Introduction: NSF Future Internet Architecture (FIA) Program

- **FIA program started in Oct 2010, with 5 teams funded:**
  - XIA (led by CMU) – project aims to develop very flexible architecture which can evolve to meet new requirements
  - NEBULA (led by UPenn) – project aims to design fast/managed flows to cloud services at the core of the Internet
  - NDN (led by UCLA/PARC) – project aims to re-design Internet to handle named content efficiently
  - ChoiceNet (led by RENCI/UNC) – project aims to enable choice and competition at each layer of protocol stack
  - **MobilityFirst** (led by Rutgers) – project aims to develop efficient and scalable architecture for emerging mobility services
- Scope of all these FIA projects includes architecture/design, protocol validation and comprehensive evaluation of usability and performance (using real-world applications in later stages)

MobilityFirst Project: Collaborating Institutions

RUTGERS (LEAD)

D. Raychaudhuri, M. Gruteser, W. Trappe, R, Martin, Y. Zhang, I. Seskar, K. Nagaraja

University of Massachusetts Amherst

A. Venkataramani, J. Kurose, D. Towsley

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

M. Reiter

THE UNIVERSITY of WISCONSIN MADISON

S. Bannerjee

Massachusetts Institute of Technology

W. Lehr

MICHIGAN

Z. Morley Mao

Duke UNIVERSITY

X. Yang, R. RoyChowdhury

UMASS LOWELL

G. Chen

UNIVERSITY OF Nebraska Lincoln

B. Ramamurthy

Project Funded by the US National Science Foundation (NSF)
Under the Future Internet Architecture (FIA) Program, CISE

+ Also industrial R&D collaborations with AT&T Labs, Bell Labs, NTT DoCoMo,, Toyota ITC, NEC, Ericsson and others

Introduction: Mobility as <u>*the*</u> key driver for the future Internet

- Historic shift from PC's to mobile computing and embedded devices…
  - ~4 B cell phones vs. ~1B PC's in 2010
  - Mobile data growing exponentially – Cisco white paper predicts 3.6 Exabytes by 2014, significantly exceeding wired Internet traffic
  - Sensor/IoT/V2V just starting, ~5-10B units by 2020

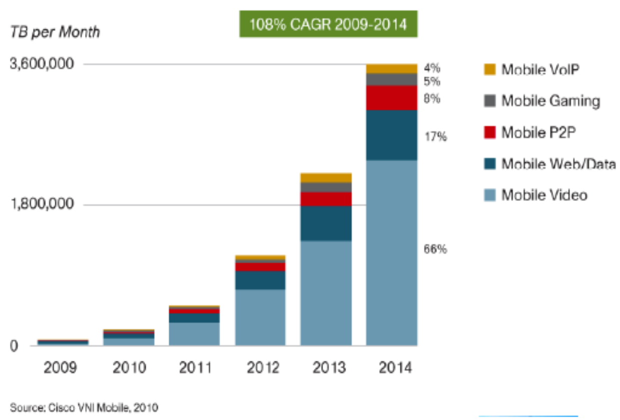

WINLAB

Introduction: Cellular-Internet convergence

- Technology disparity today
  - Two sets of addresses (cell number & IP), protocols (3GPP, IP), and protocol gateways (GGSN)
  - Poor scalability, fragile, difficult to manage
  - More complications with heterogeneous radio access

Single unified architecture can simplify and speed up mobile Internet application development across diverse networks and platforms

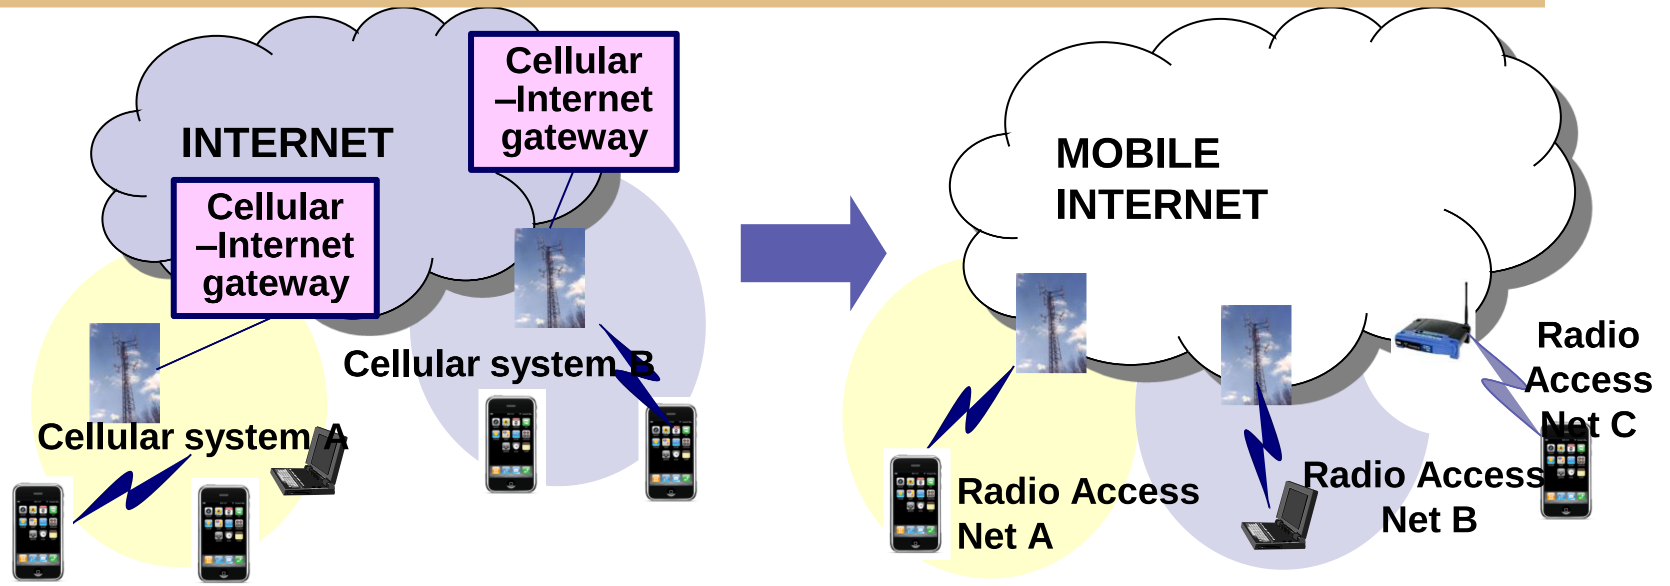

Introduction: Opportunistic Content Delivery

- Opportunistic, delay tolerant (DTN-type) delivery modes and content caching increasingly important for mobile/sensor devices
  - Heterogeneous access; network may be disconnected at times
  - Significant performance gain via content caching and opportunistic delivery
  - P2P mode of content delivery can also play a useful role …

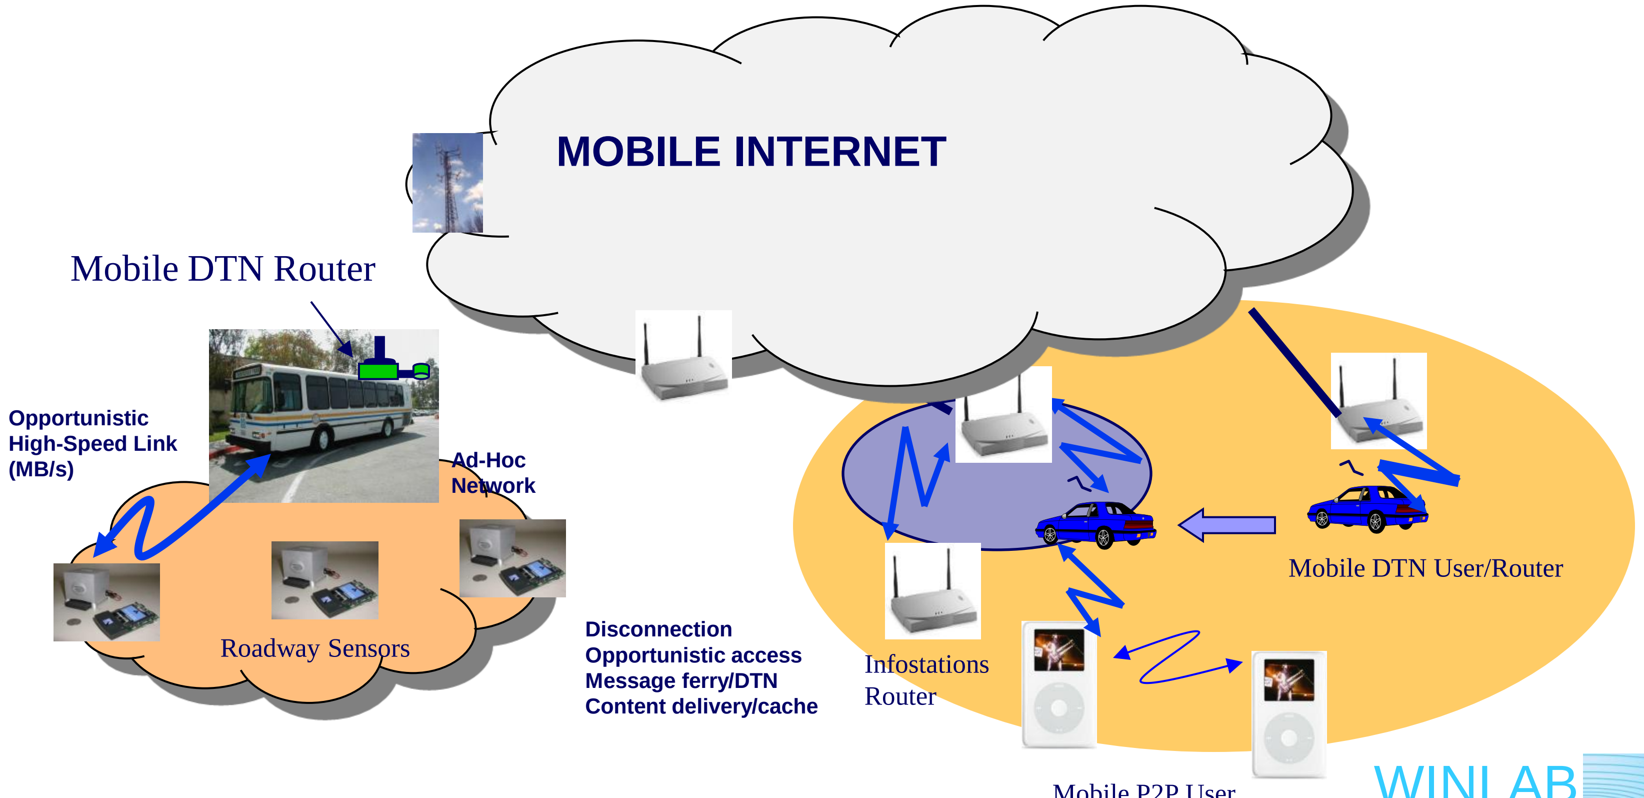

Introduction: Vehicular Networks

- 100's of millions of cars with radios by ~2015
  - Vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) modes
  - Support for dynamic network formation, network mobility, geo-routing, etc.
  - Critical new security and privacy requirements

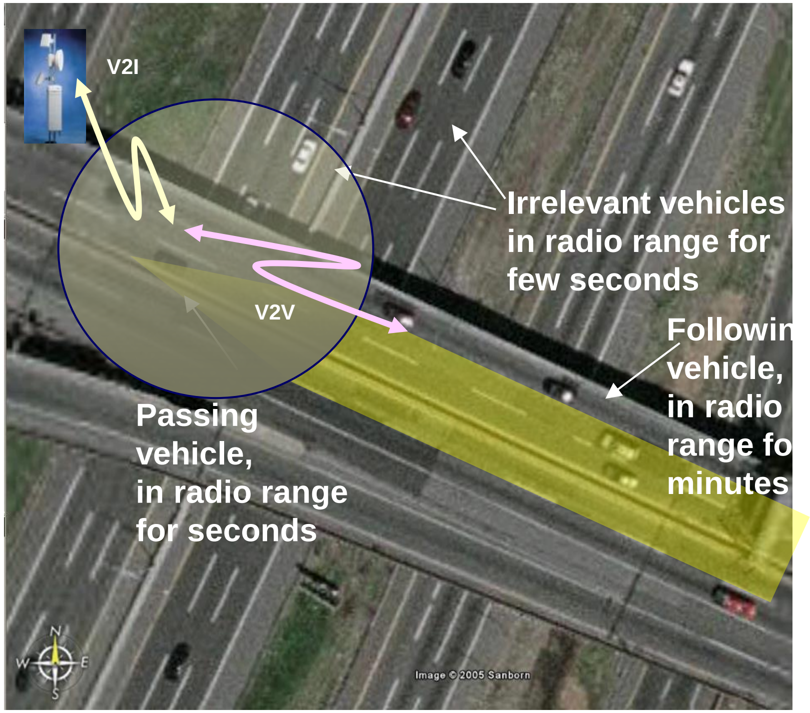

Introduction: Pervasive Cyber-Physical Systems/Internet-of-Things (IoT)

- Next-generation Internet applications will interface human beings with the physical world, e.g.,
  - Machine-to-machine (M2M), cyber-physical systems, smart grids, ..
  - Location and context-aware embedded computing
  - Secure and flexible network computing ('edge cloud") model

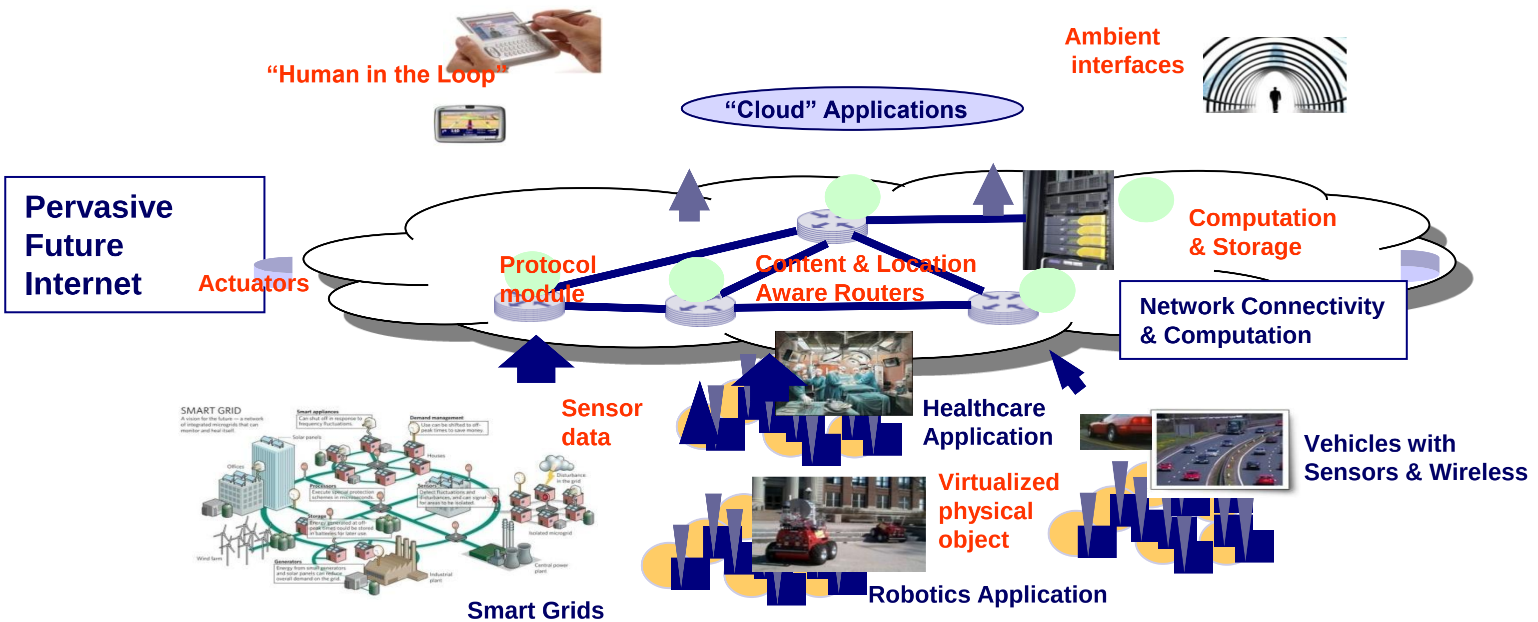

MobilityFirst Protocol

MobilityFirst Protocol: *Feature Summary*

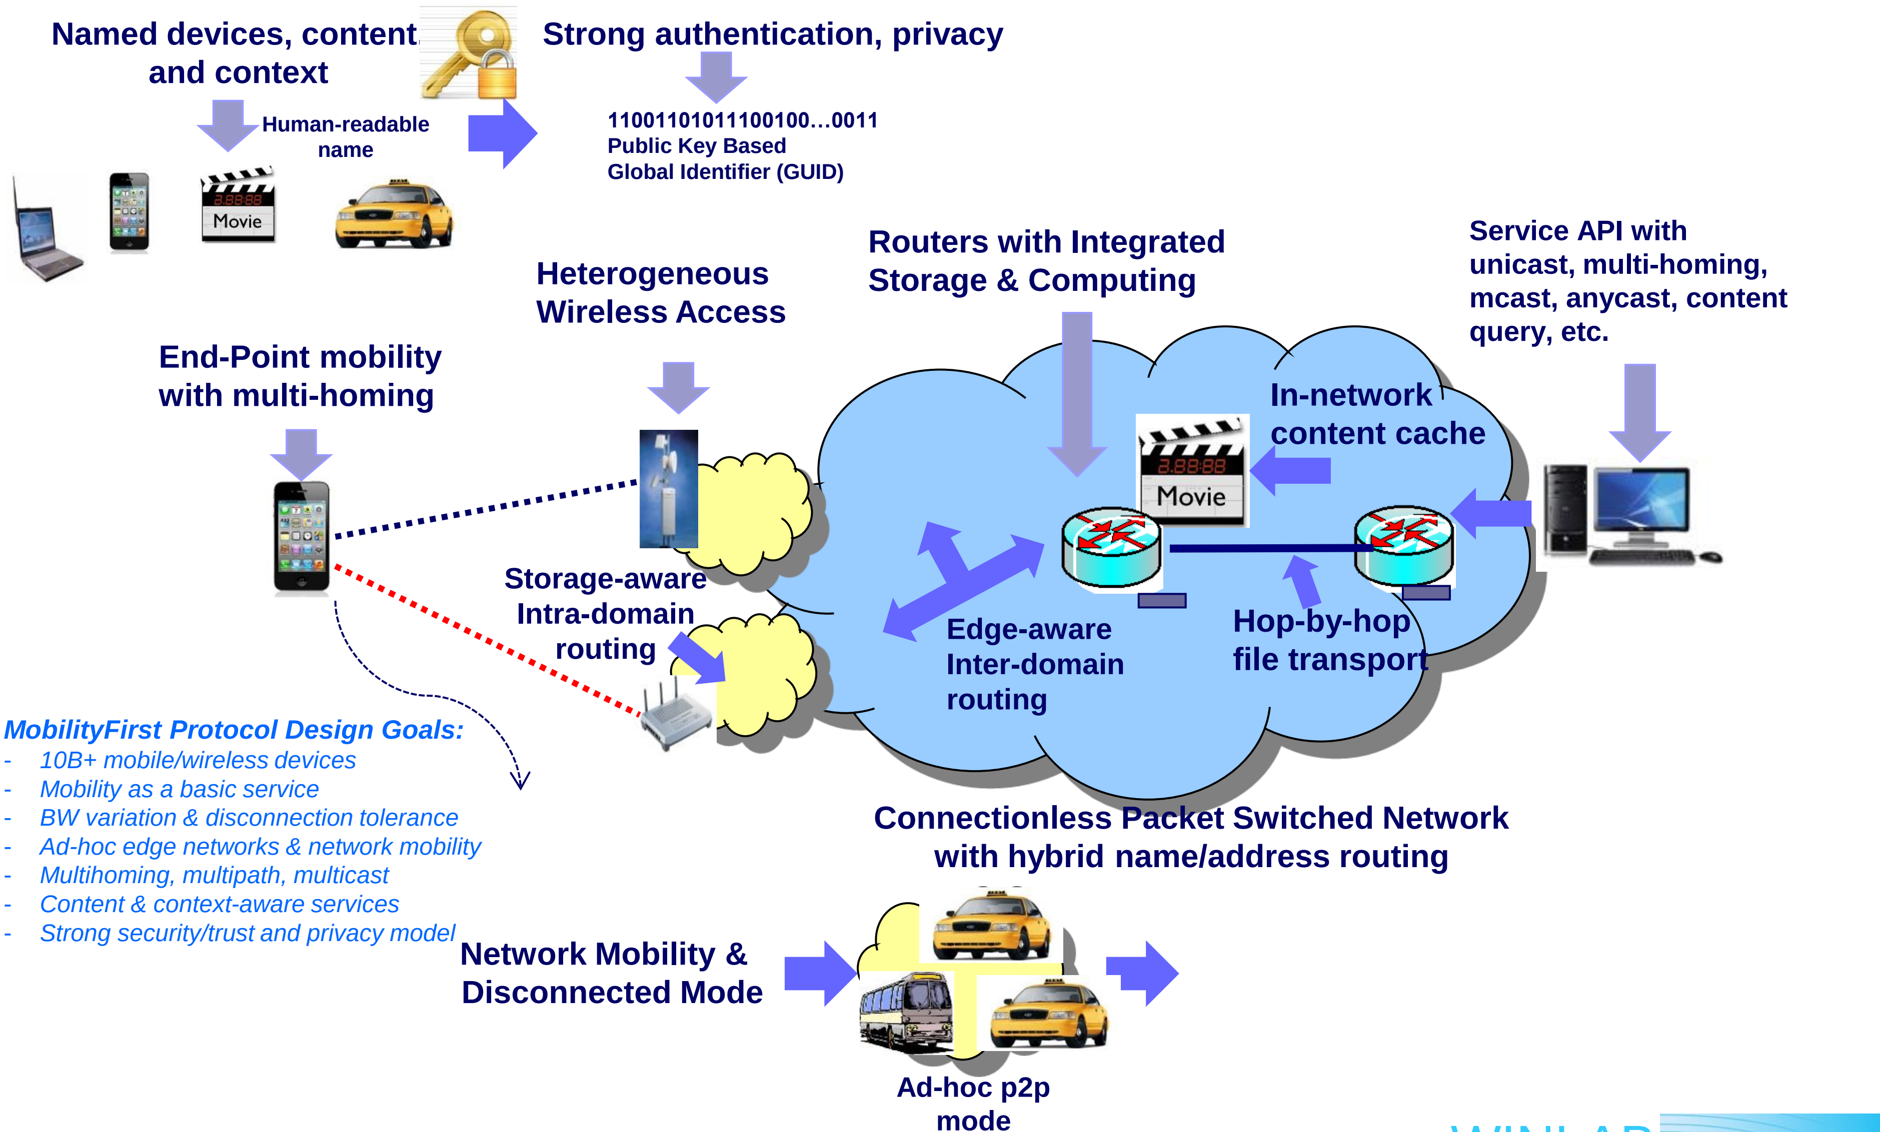

MobilityFirst Protocol Design Goals:

- *10B+ mobile/wireless devices*
- *Mobility as a basic service*
- *BW variation & disconnection tolerance*
- *Ad-hoc edge networks & network mobility*
- *Multihoming, multipath, multicast*
- *Content & context-aware services*
- *Strong security/trust and privacy model*

MobilityFirst Protocol: *Technology Solution*

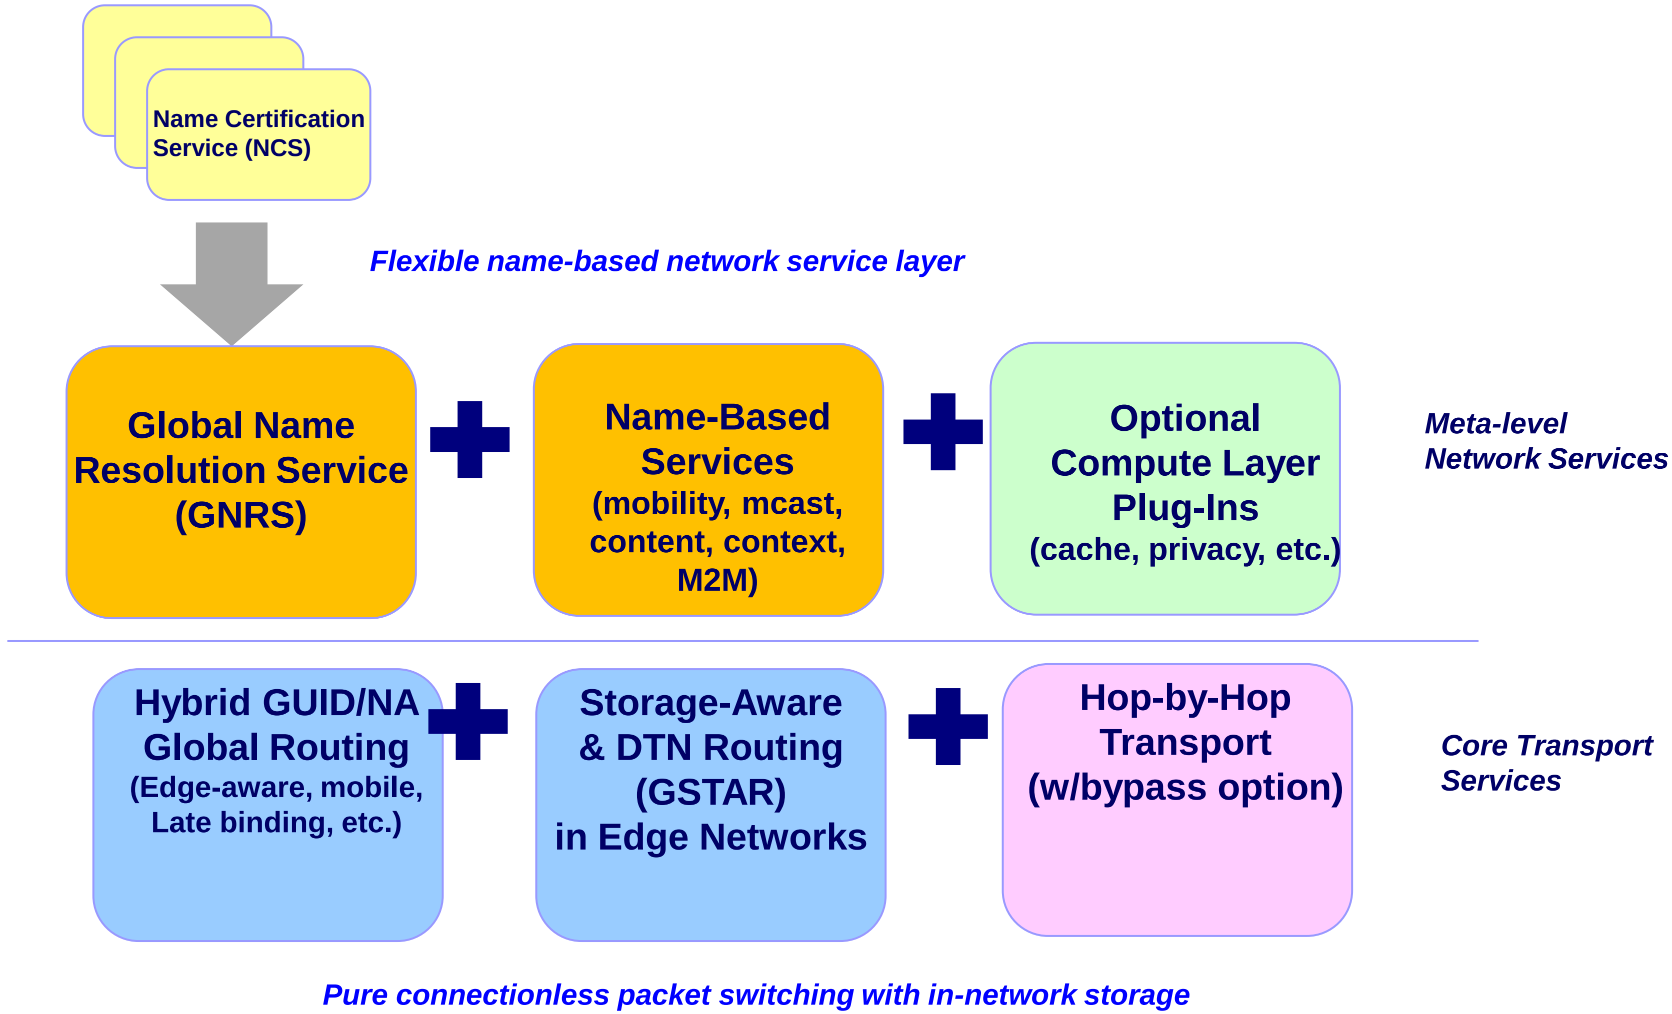

MobilityFirst Protocol: *The Protocol Stack*

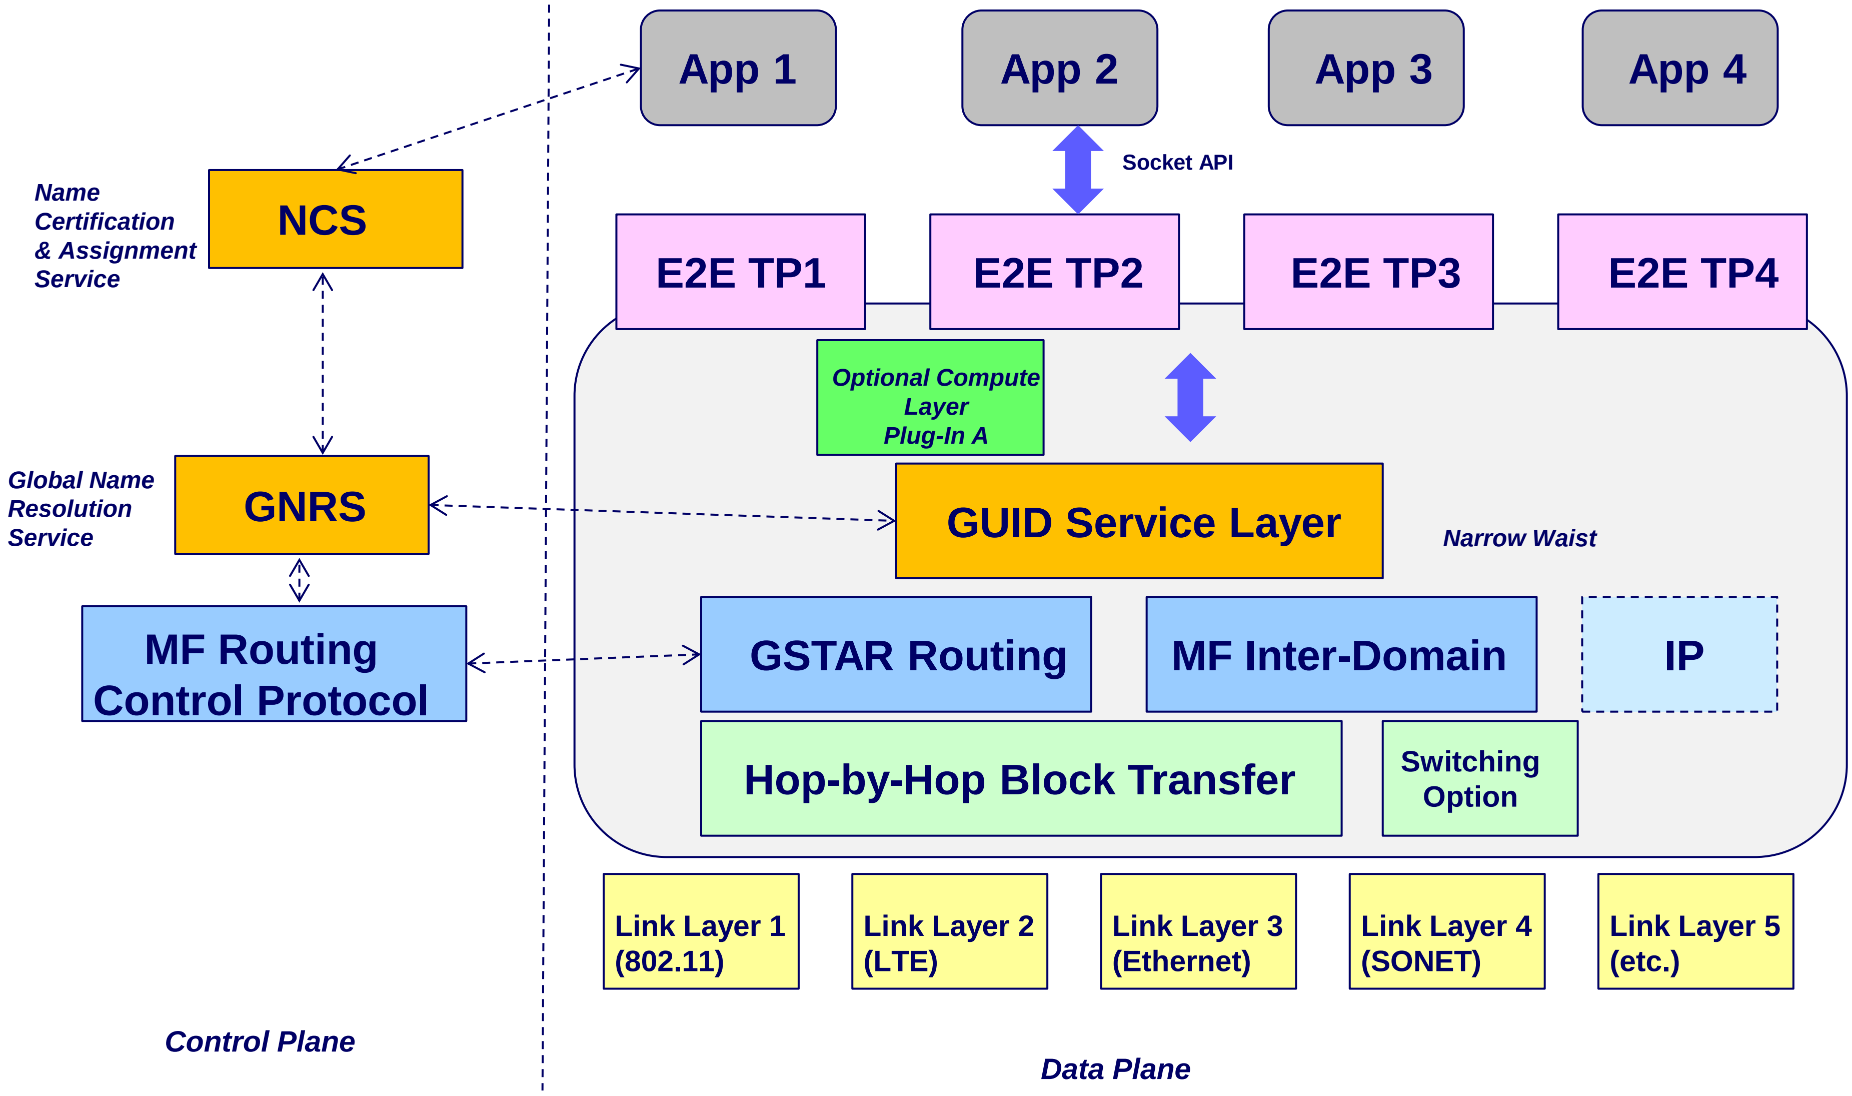

WINLAB

MF Protocol: Name-Address Separation → GUIDs

- Separation of names (ID) from network addresses (NA)
- Globally unique name (GUID) for network attached objects
  - User name, device ID, content, context, AS name, and so on
  - Multiple domain-specific naming services
- Global Name Resolution Service for GUID → NA mappings
- Hybrid GUID/NA approach
  - Both name/address headers in PDU
  - "Fast path" when NA is available
  - GUID resolution, late binding option

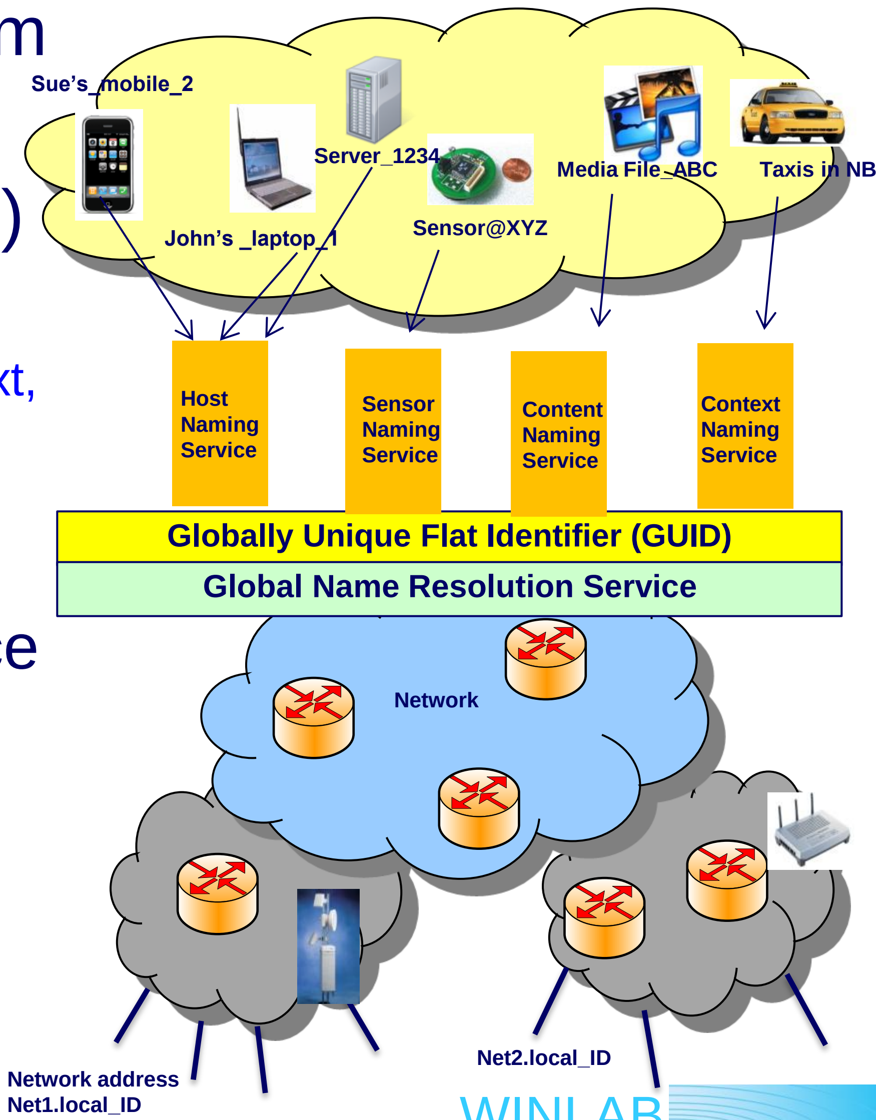

WINLAB

Protocol Example: Mobility Service via Name Resolution at Device End-Points

Service API capabilities:
- *send (GUID, options, data)*
Options = anycast, mcast, time, ..
- *get (content_GUID, options)*
Options = nearest, all, ..

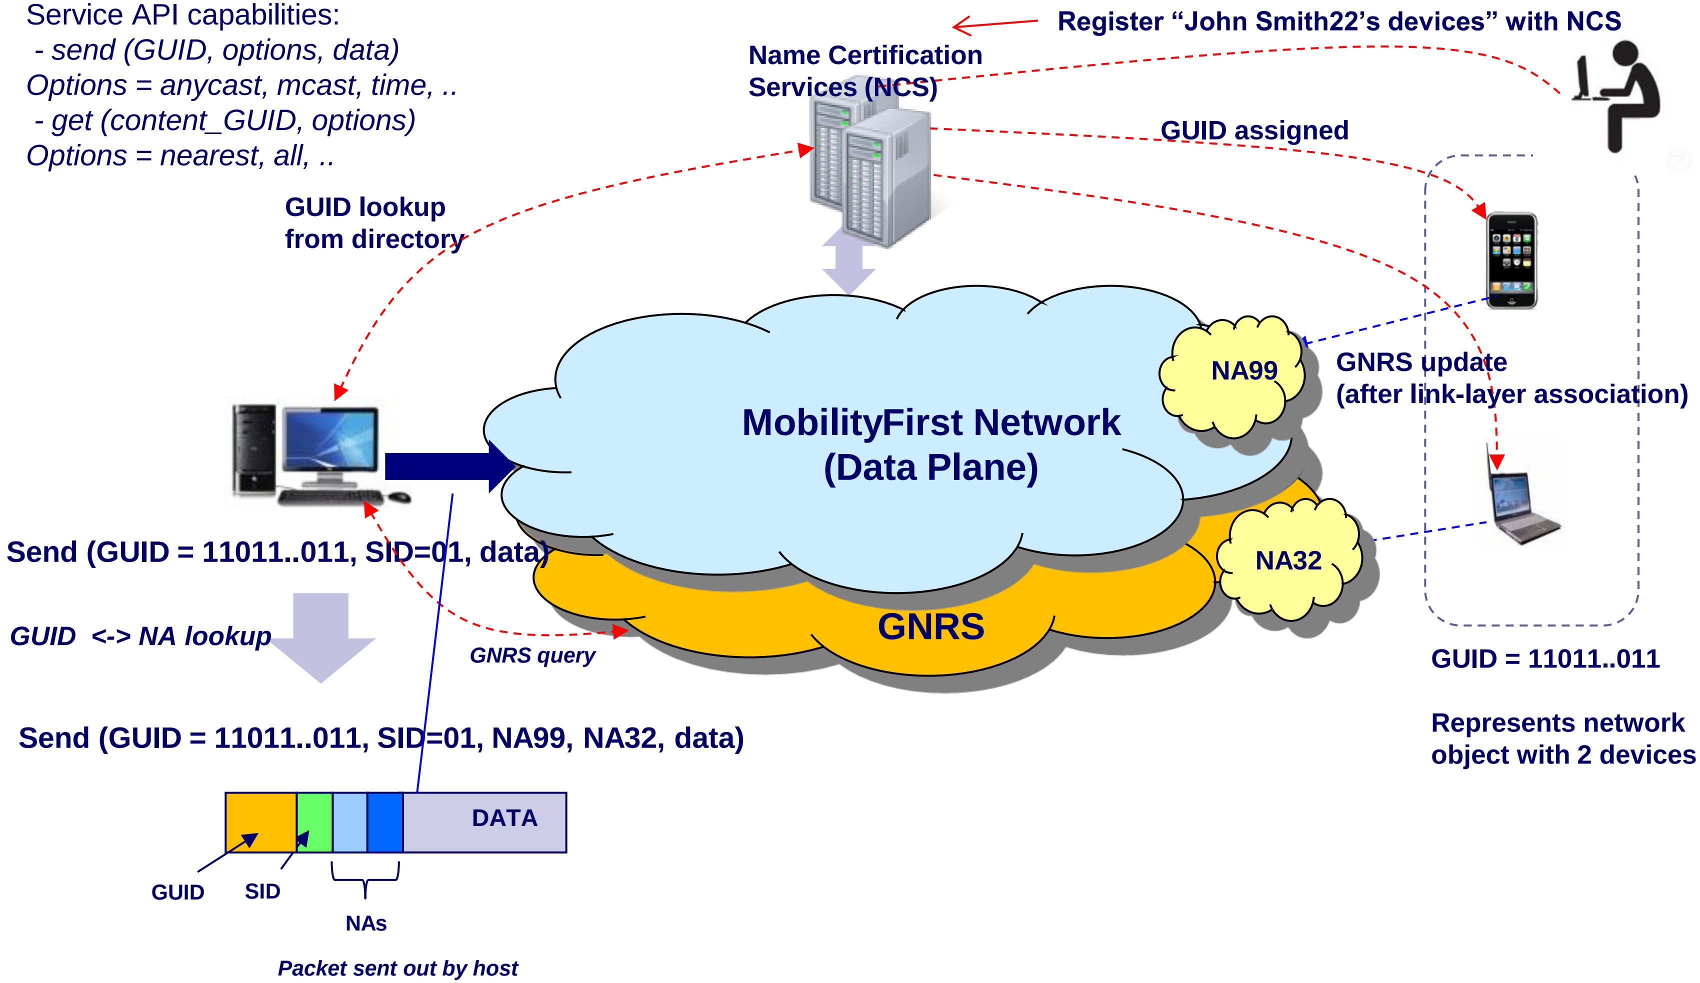

MF Protocol: Hybrid GUID/NA Storage Router in *MobilityFirst*

- **Hybrid name-address based routing in MobilityFirst requires a new router design with in-network storage and two lookup tables:**
  - "Virtual DHT" table for GUID-to-NA lookup as needed
  - Conventional NA-to-port # forwarding table for "fast path"
  - Also, enhanced routing algorithm for store/forward decisions

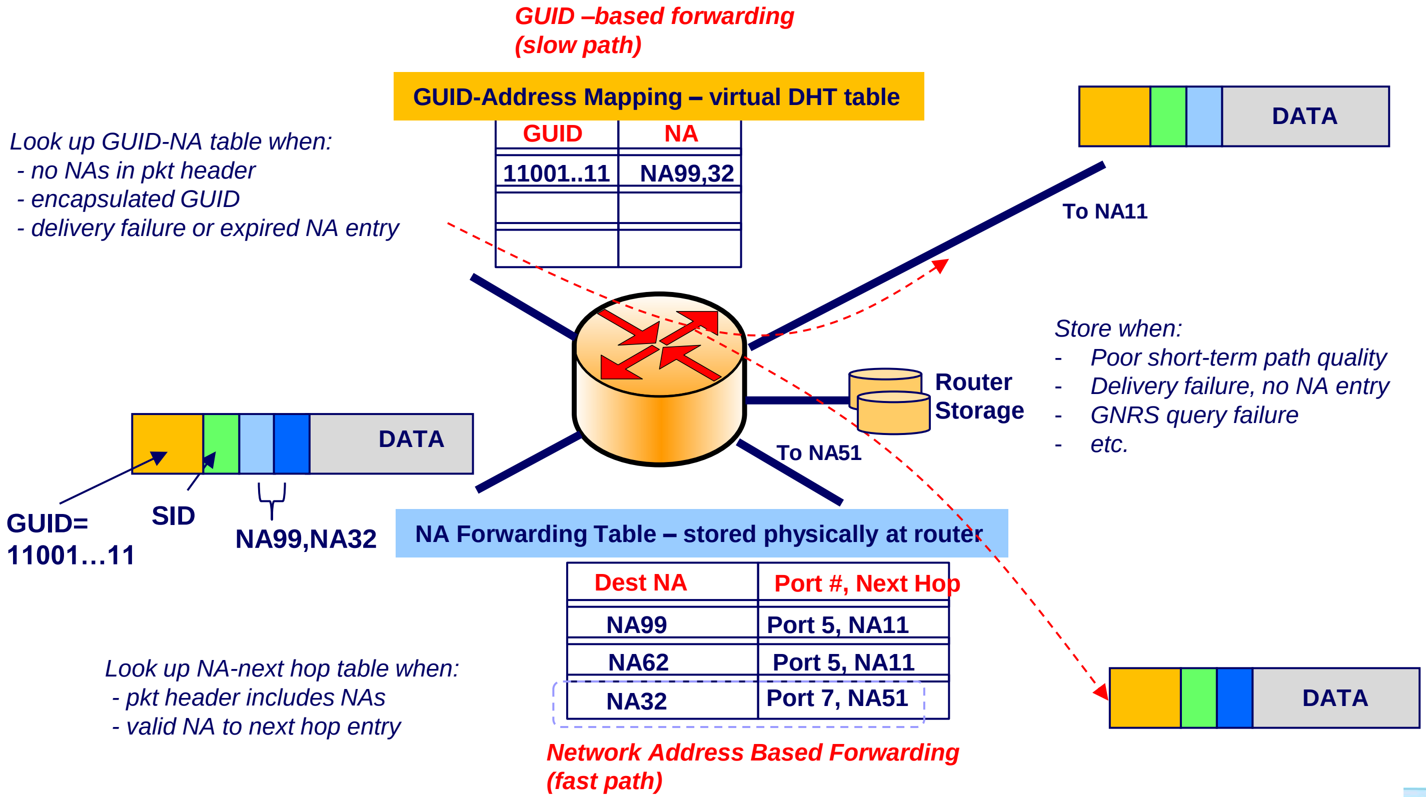

Protocol Design: Realizing the GNRS

- Fast GNRS implementation based on DHT between routers
  - GNRS entries (GUID <-> NA) stored at Router Addr = hash(GUID)
  - Results in distributed in-network directory with fast access (~100 ms)

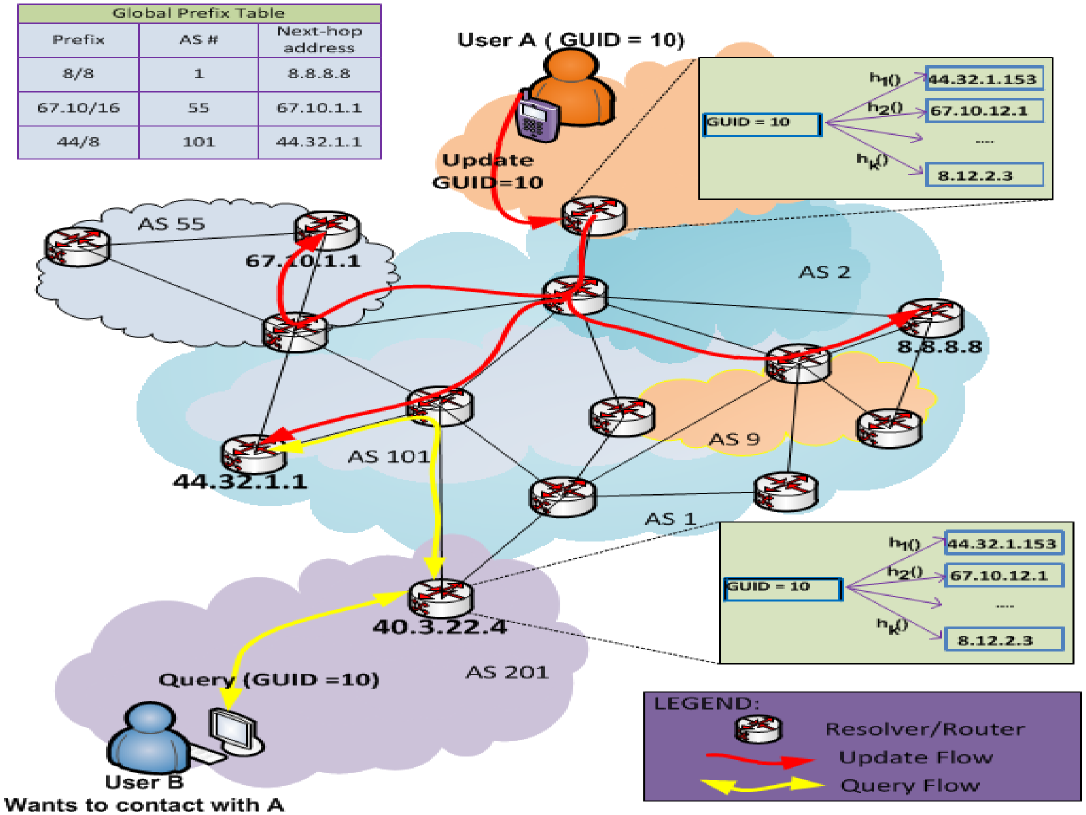

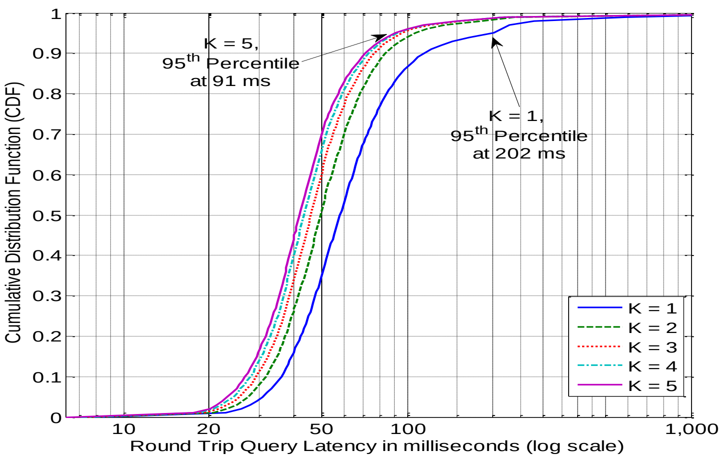

Internet Scale Simulation Results Using DIMES database

WINLAB

Protocol Design: Storage-Aware Routing (GSTAR)

- Storage aware (CNF, generalized DTN) routing exploits in-network storage to deal with varying link quality and disconnection
- Routing algorithm adapts seamlessly adapts from switching (good path) to store-and-forward (poor link BW/short disconnection) to DTN (longer disconnections)
- Storage has benefits for wired networks as well..

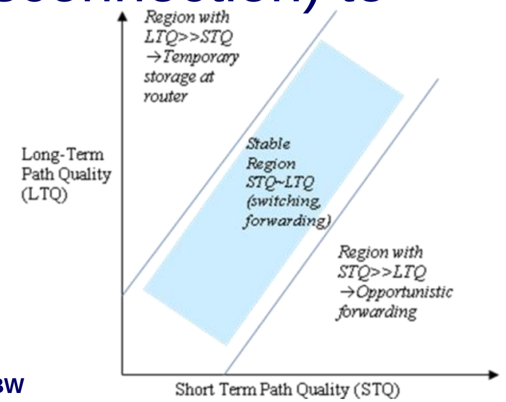

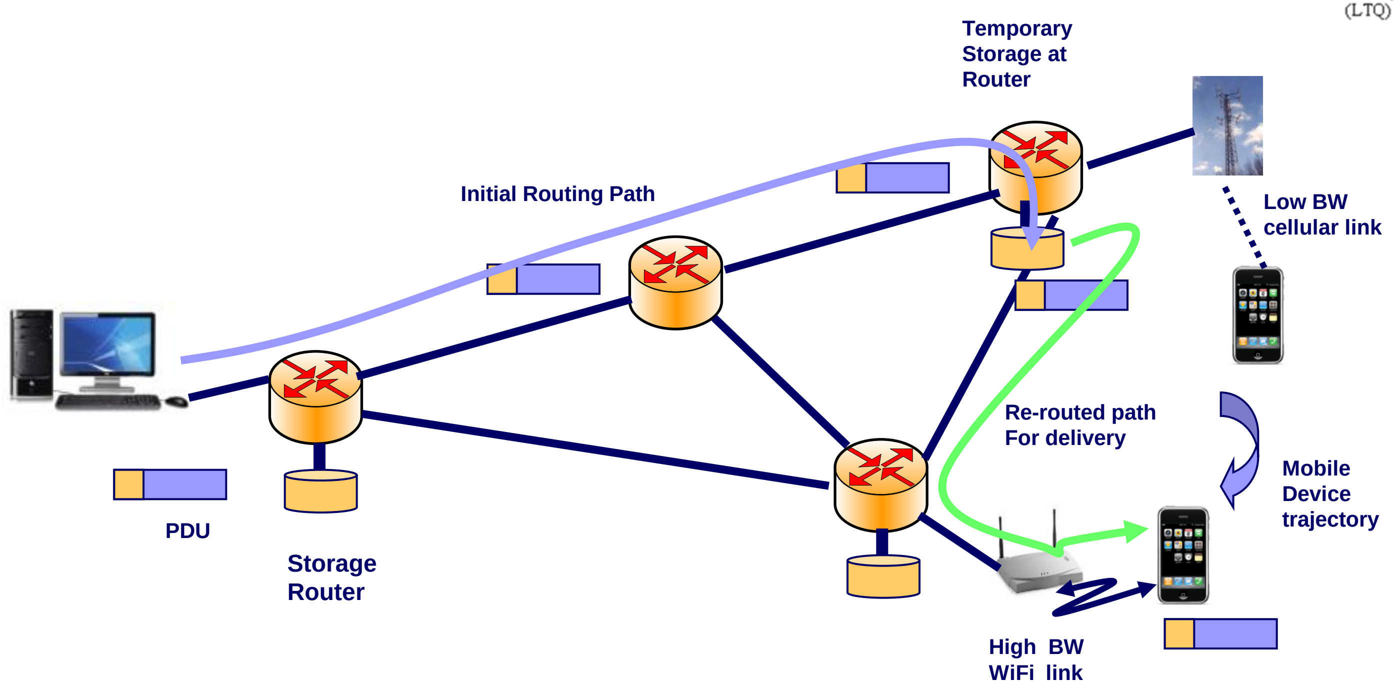

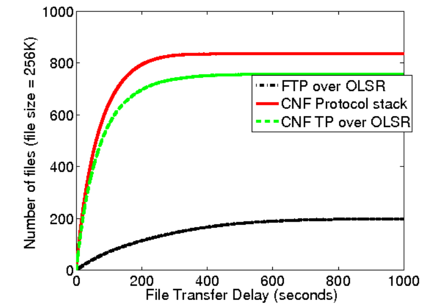

Sample CNF routing result

MF Protocol: Service Abstractions (1)

- MobilityFirst offers a named-object service API that supports mobility, disconnection, multi-homing, multicast in a natural way
- Replaces the point-to-point virtual link abstraction of IP …
- Example: Sending to a mobile device with multiple interfaces

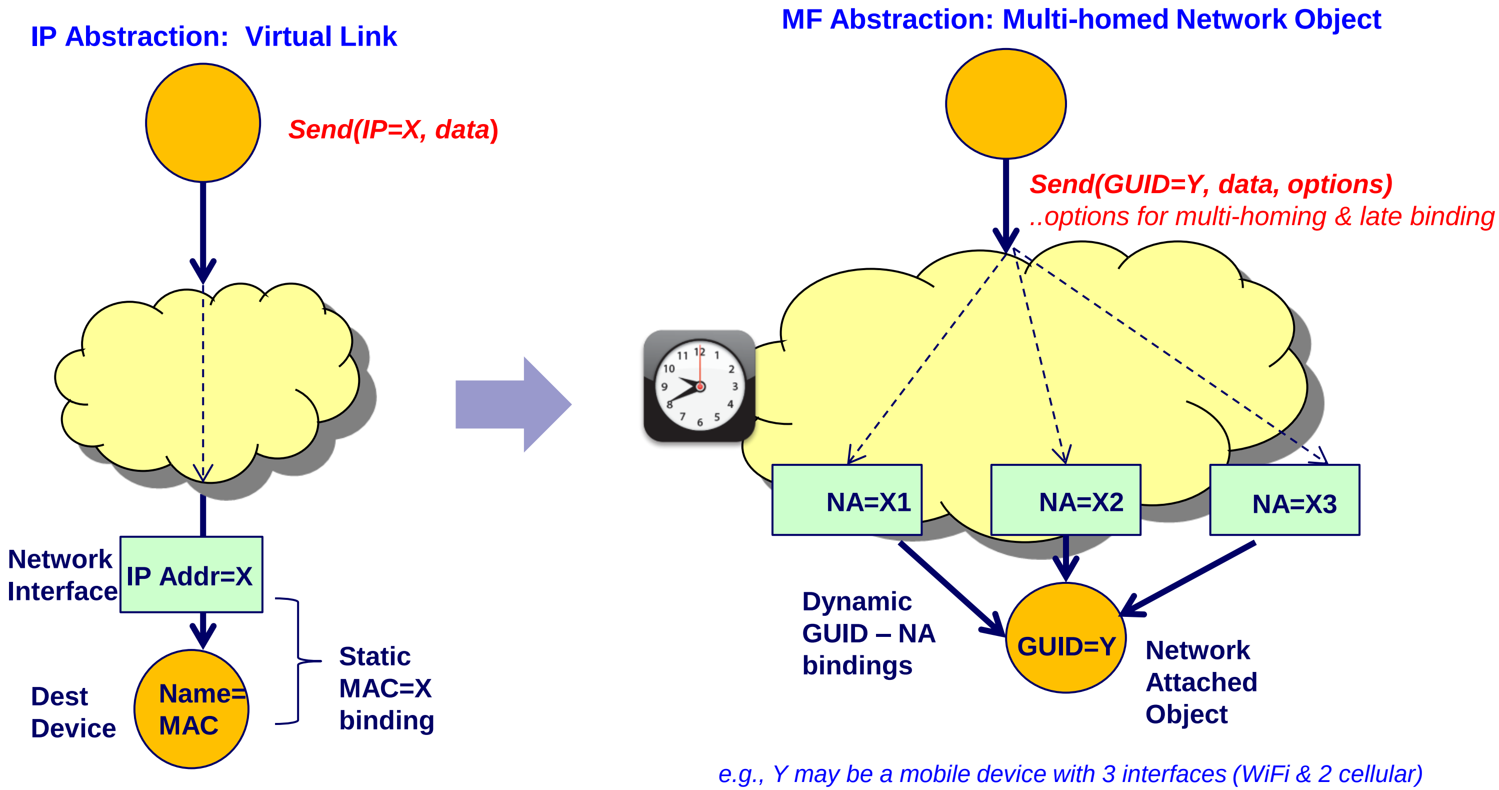

e.g., Y may be a mobile device with 3 interfaces (WiFi & 2 cellular)

MF Protocol: Service Abstractions (2)

- Use of MF Service API for content retrieval and dynamic group multicast (..membership may be specified by context)

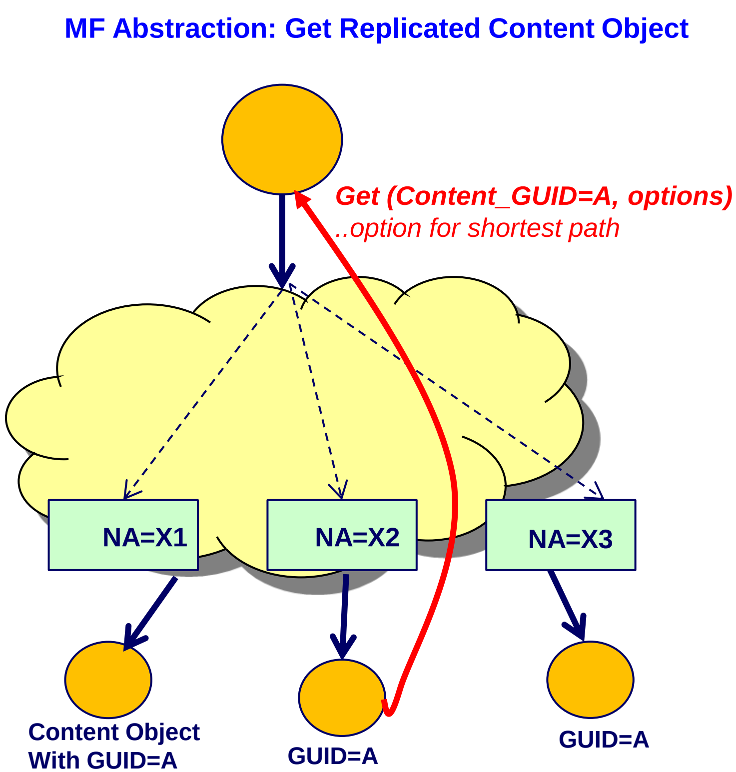

e.g., A is a replicated content object at multiple network locations

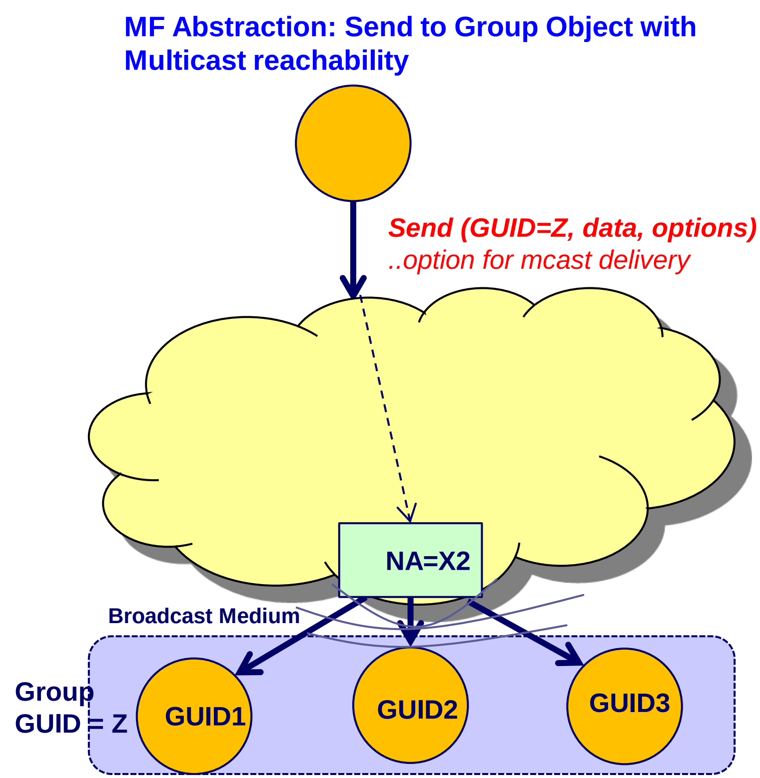

e.g., Z may be a context group of M2M devices or a cloud service

MobilityFirst Protocol Stack: Service API

- MobilityFirst API (or “socket” interface) provides a new set of services corresponding to the MF protocol stack’s capabilities
- Key features are:
  - GUID as the unique identifier right up to application level (no need for port #)
  - Service identifier choice including: basic unicast/mcast, anycast, dual-homing, late binding mobile delivery, delayed delivery, content query, context delivery, plug-in computing service, etc. (others TBD)
  - Unicast vs multicast determined
  - Lightweight transport protocol choices for end-to-end reliability and flow control
  - Socket interface examples:
    - Open(URL, myGUID, TP=A, TP_options) - identity specification and stack customization
    - Send(GUID=X, SvcFlags/SID=anycast, data, len)
    - Send(GUID=X, SvcFlags/SID=unicast, TP=B, data, len)
    - Send(GUID=X, SvcFlags/SID=late binding, options, data)
    - Send(GUID=X, SvcFlags/SID=anycast+, data)
    - Recv(data_buffer, len, GUID_allow={X, Y, X})
    - Get(GUID=X, SvcFlags/SID=anycast+content query, options, data_buffer, len)
    - ….

Wireless/Mobile Use Case

Wireless/Mobile Use Case

- MobilityFirst enables a stitched-together architecture for mobile networks

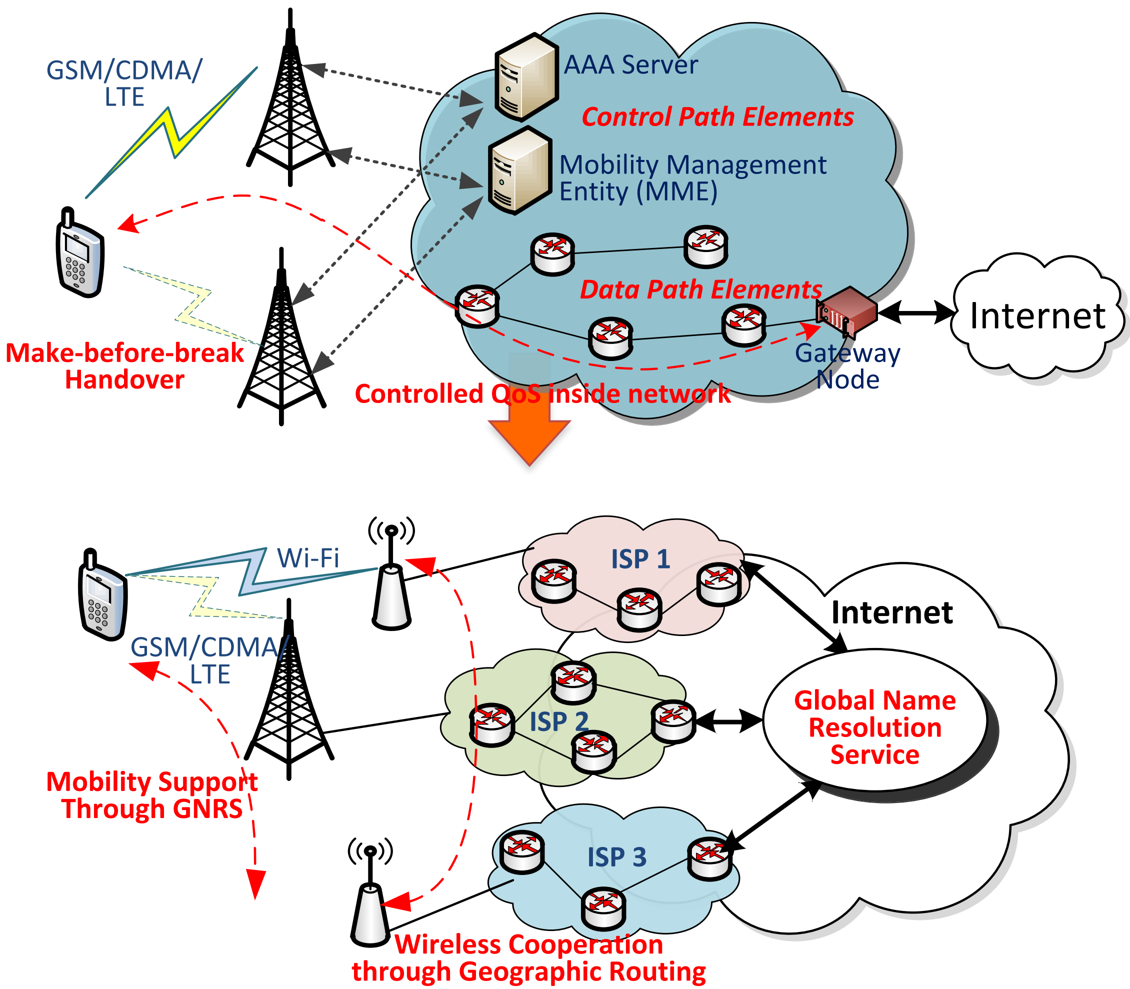

Current Mobile Networks

- Planned Deployment
- Licensed Spectrum
- Fine-grained Managed QoS
- Centralized Mobility Support
- Homogeneous Topology
- Network-wide Authentication

MF Network-of-Networks

- Ad-hoc Deployment
- Unlicensed Spectrum
- Coarse-grained Managed
- In-network Mobility Support
- Heterogeneous topology
- Authentication at APs

Wireless/Mobile Use Case:
Service Capability Examples

- MF provides several useful capabilities for mobile networks, e.g. mobility, multi-homing, multi-network access, late binding, multicast, ..

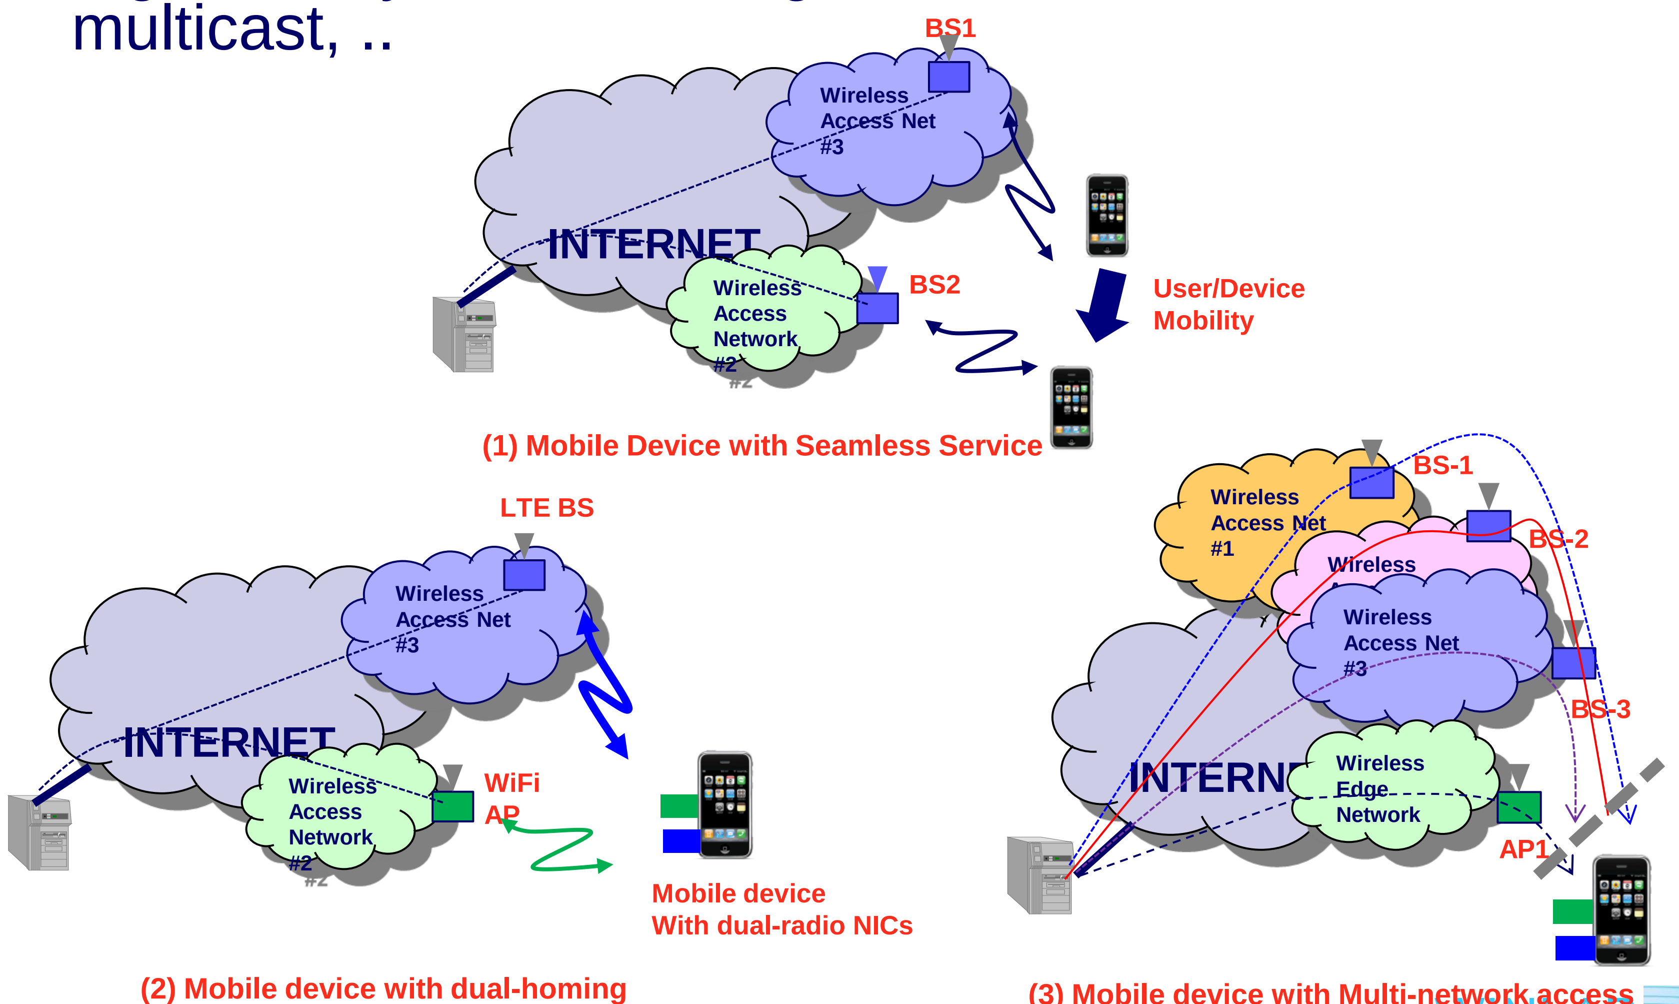

WINLAB

Protocol Example: Mobility Service via Name Resolution at Device End-Points

Service API capabilities:
- *send (GUID, options, data)*

Options = anycast, mcast, time, ..
- *get (content_GUID, options)*

Options = nearest, all, ..

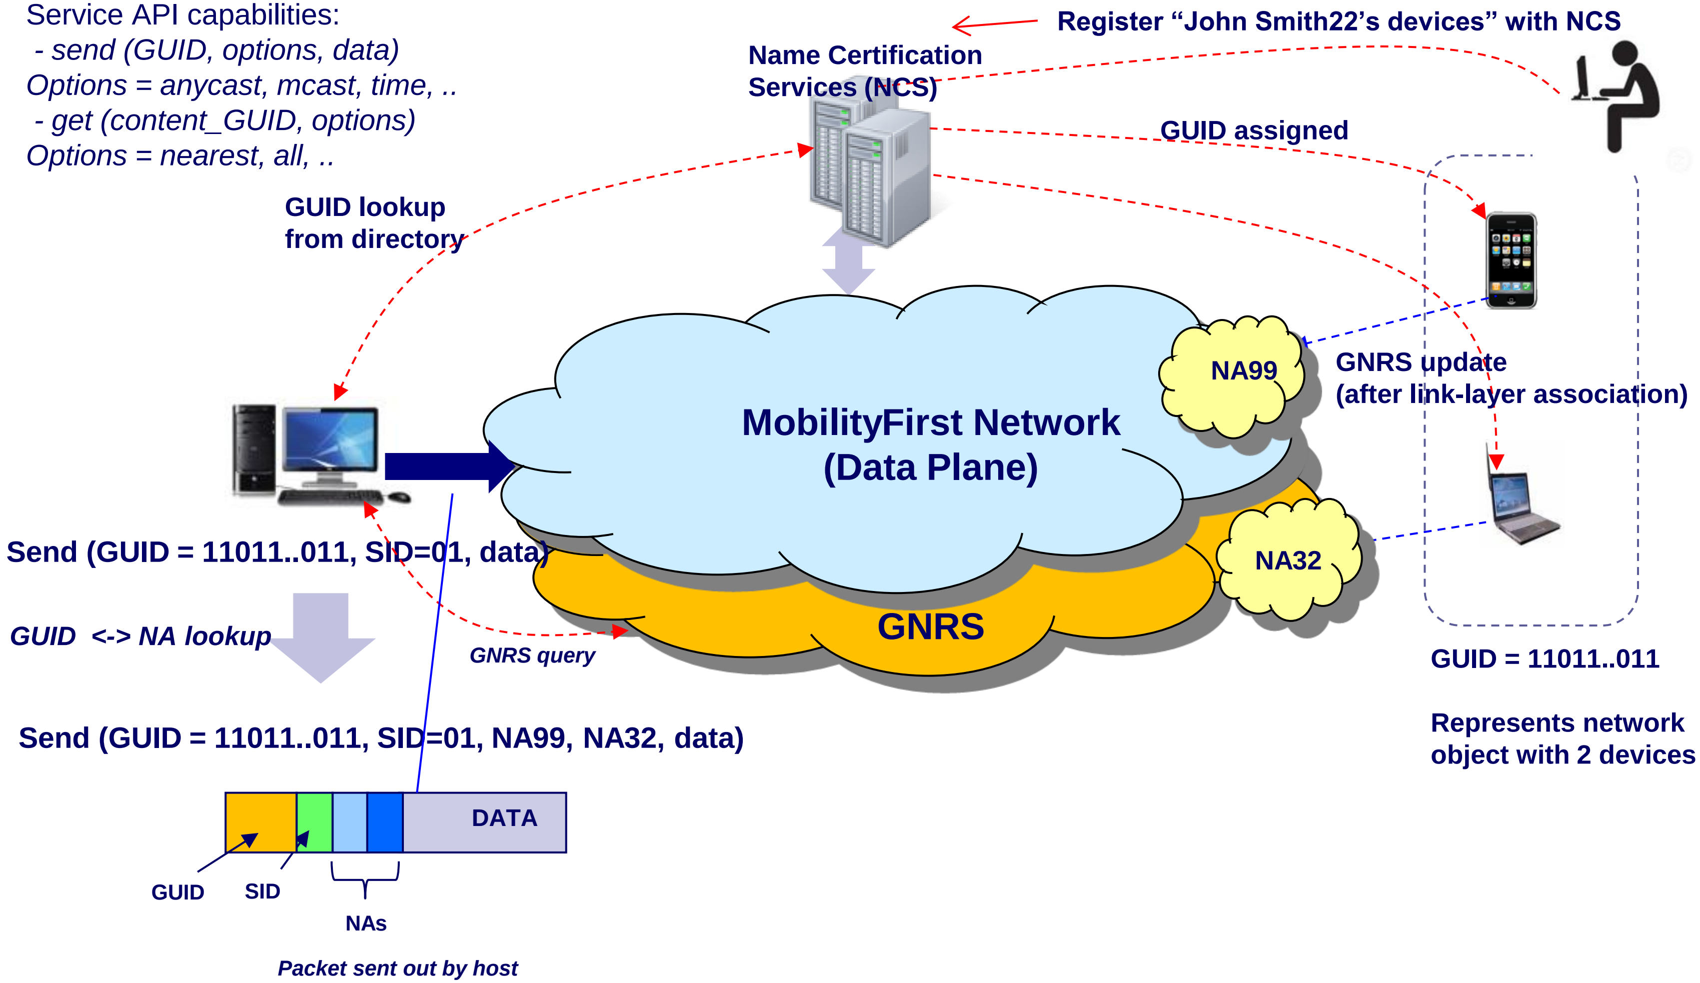

Protocol Example: Dual Homing Service via Named-Object / GNRS

Multihoming service example

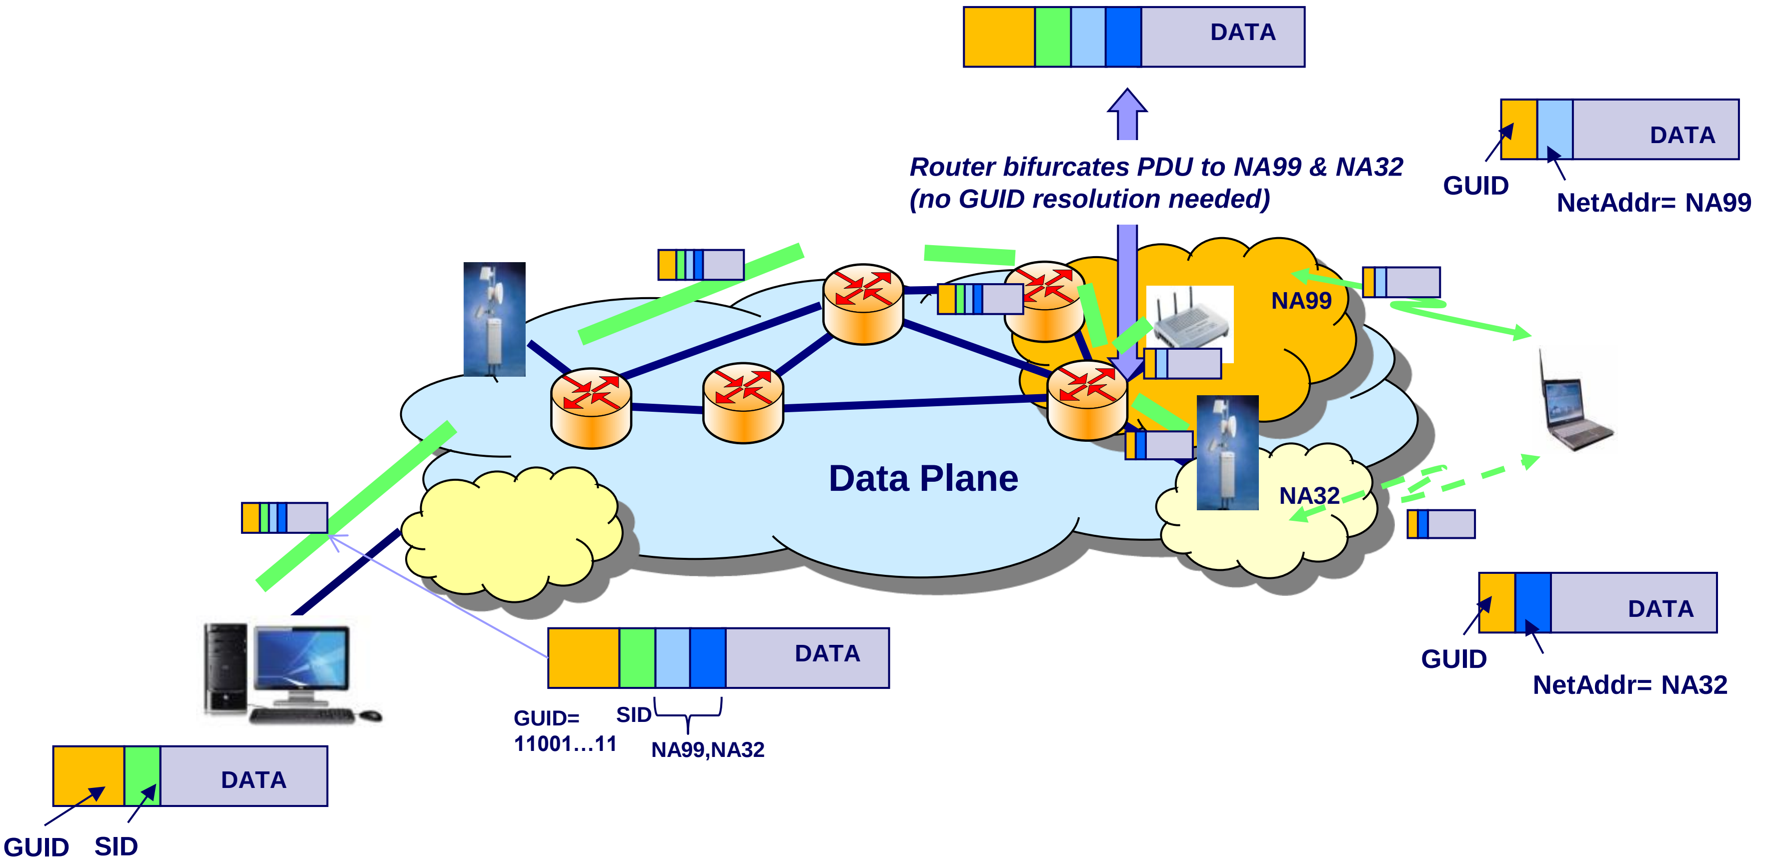

Send data file to "John Smith22's laptop", SID= 129 (multihoming – all interfaces)

WINLAB

Protocol Example: Handling Disconnection via Late Binding

Store-and-forward mobility service example

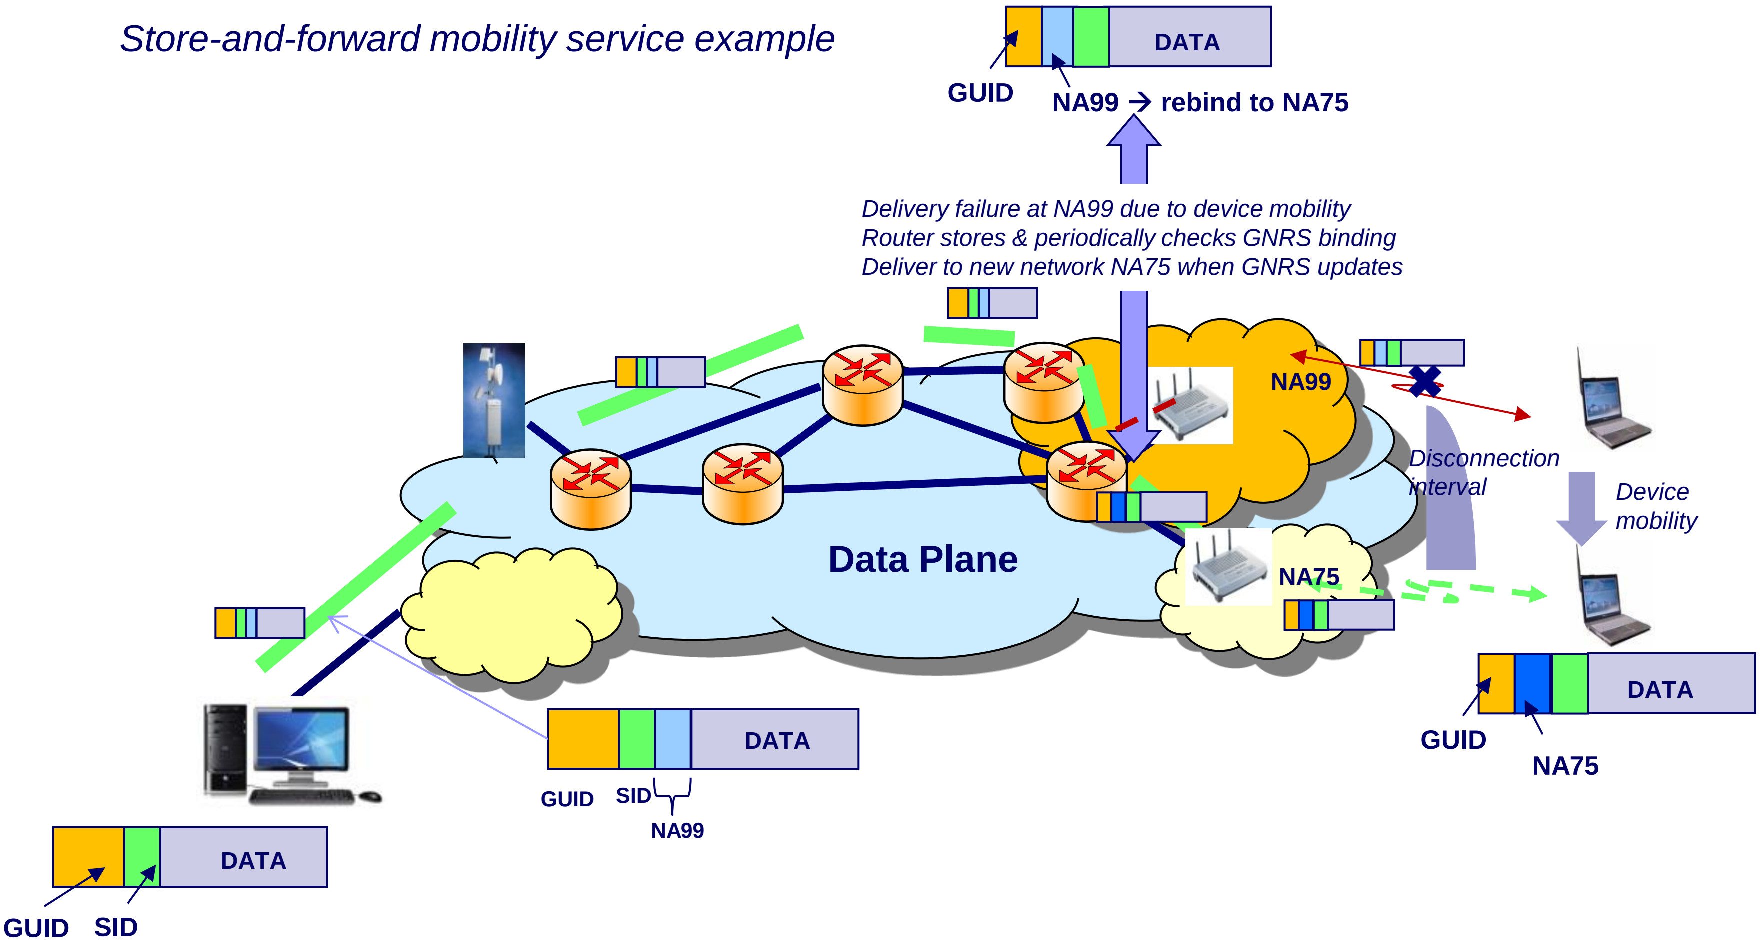

Send data file to "John Smith22's laptop", SID= 11 (unicast, mobile delivery)

WINLAB

Sample Results: MF vs TCP/IP for WiFi Roaming

- Quantifying the benefits of in-network mobility management & storage aware routing

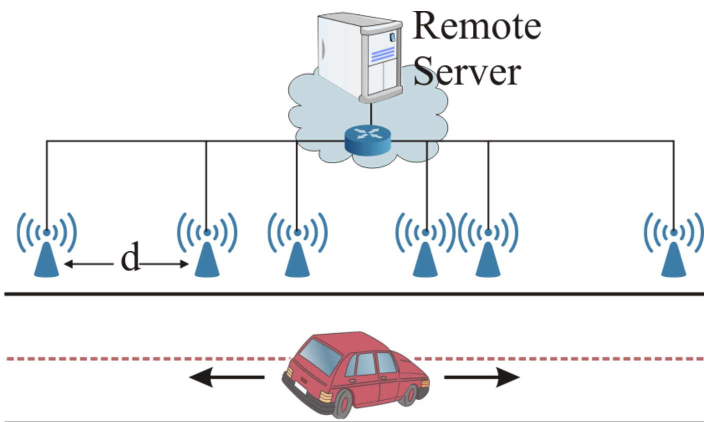

- NS3 Simulation with full 802.11 stack, different vehicular speeds & offered load
- Comparing TCP/IP with MF

Transmission of stored packets → Throughput Jumps

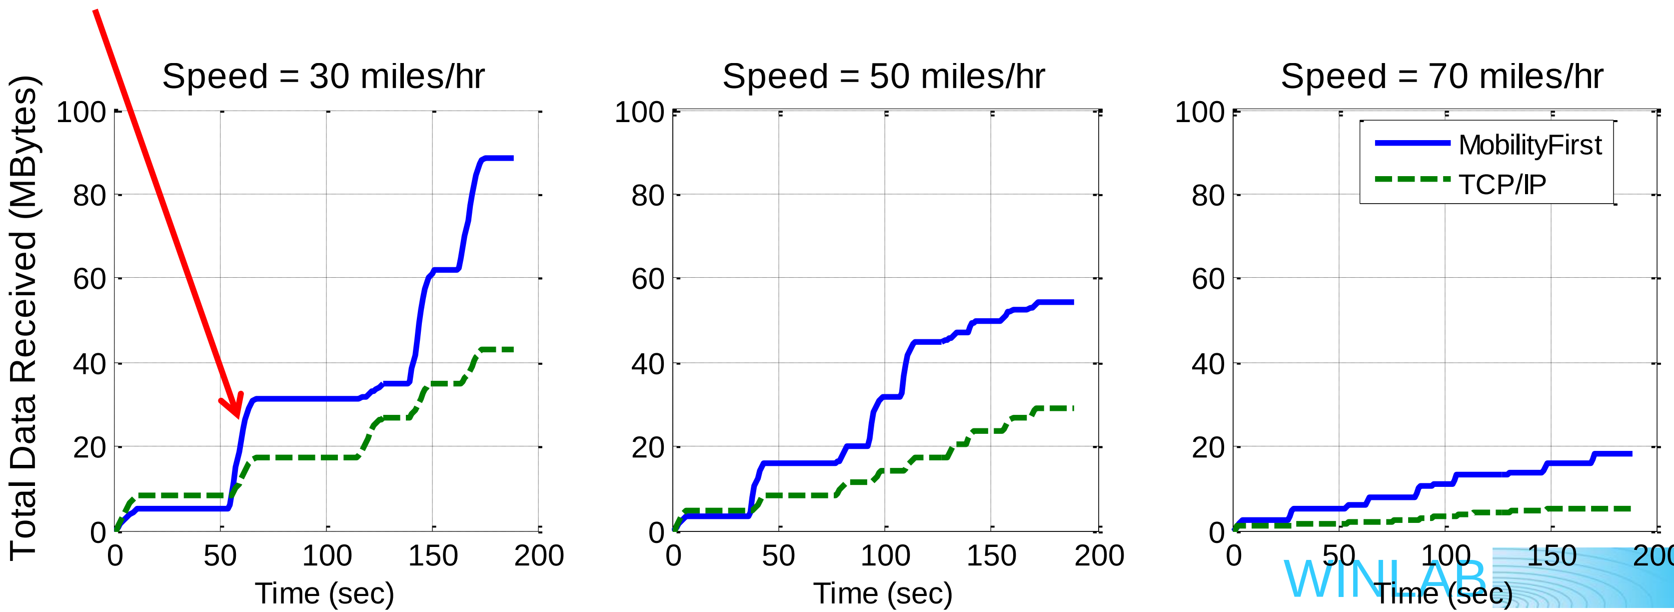

WINLAB

Sample Results: MF vs. TCP/IP for LTE/WiFi Switching

- MF provides several benefits in a heterogeneous wireless environment:
  - Mobility Mgmt. is not specific to each network
  - Automatic storage of packets in transition
  - Simultaneous use of multiple networks

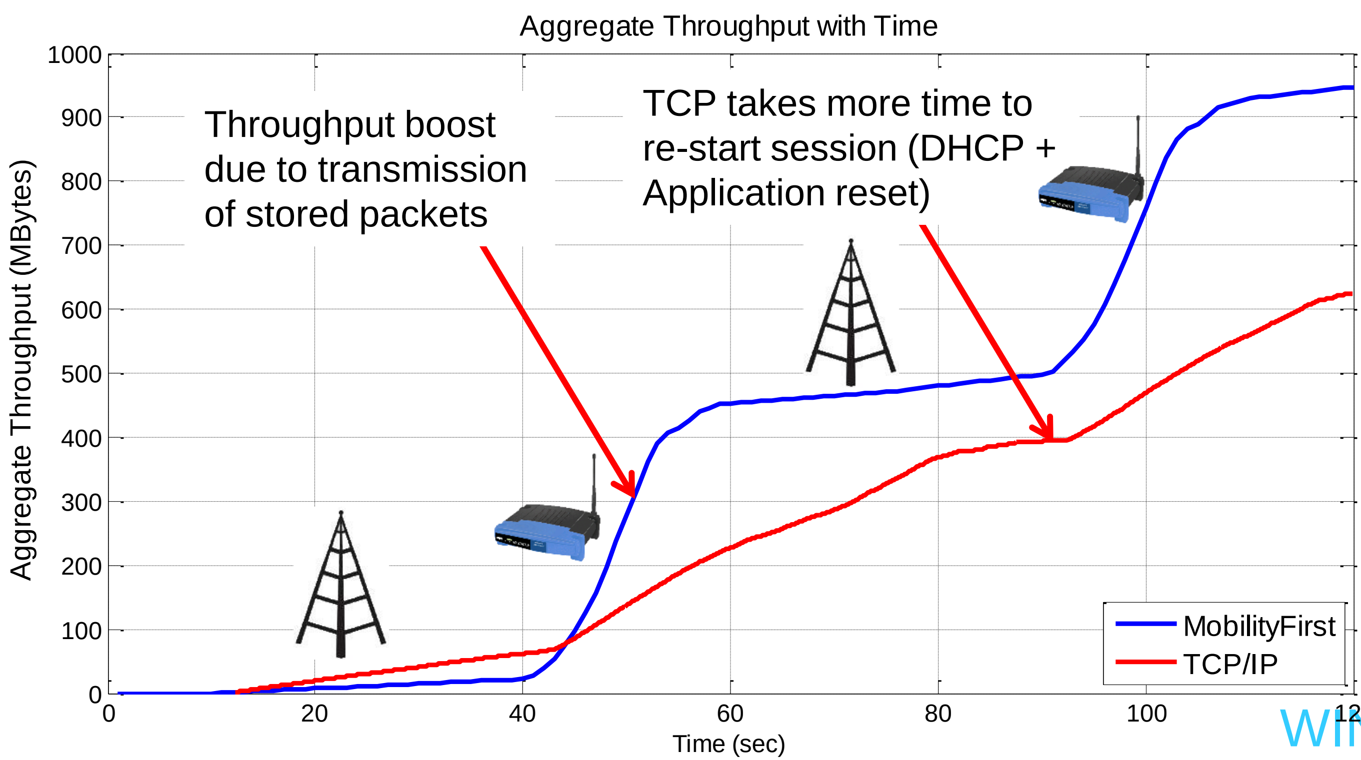

Content & Context/M2M Use Cases

Content Delivery Use Case:
Content Retrieval via GUID Name

- Network support for content addressability and caching → service primitives such as *get(content-ID, ..)*
- Optional computing layer plug-in for enhanced services

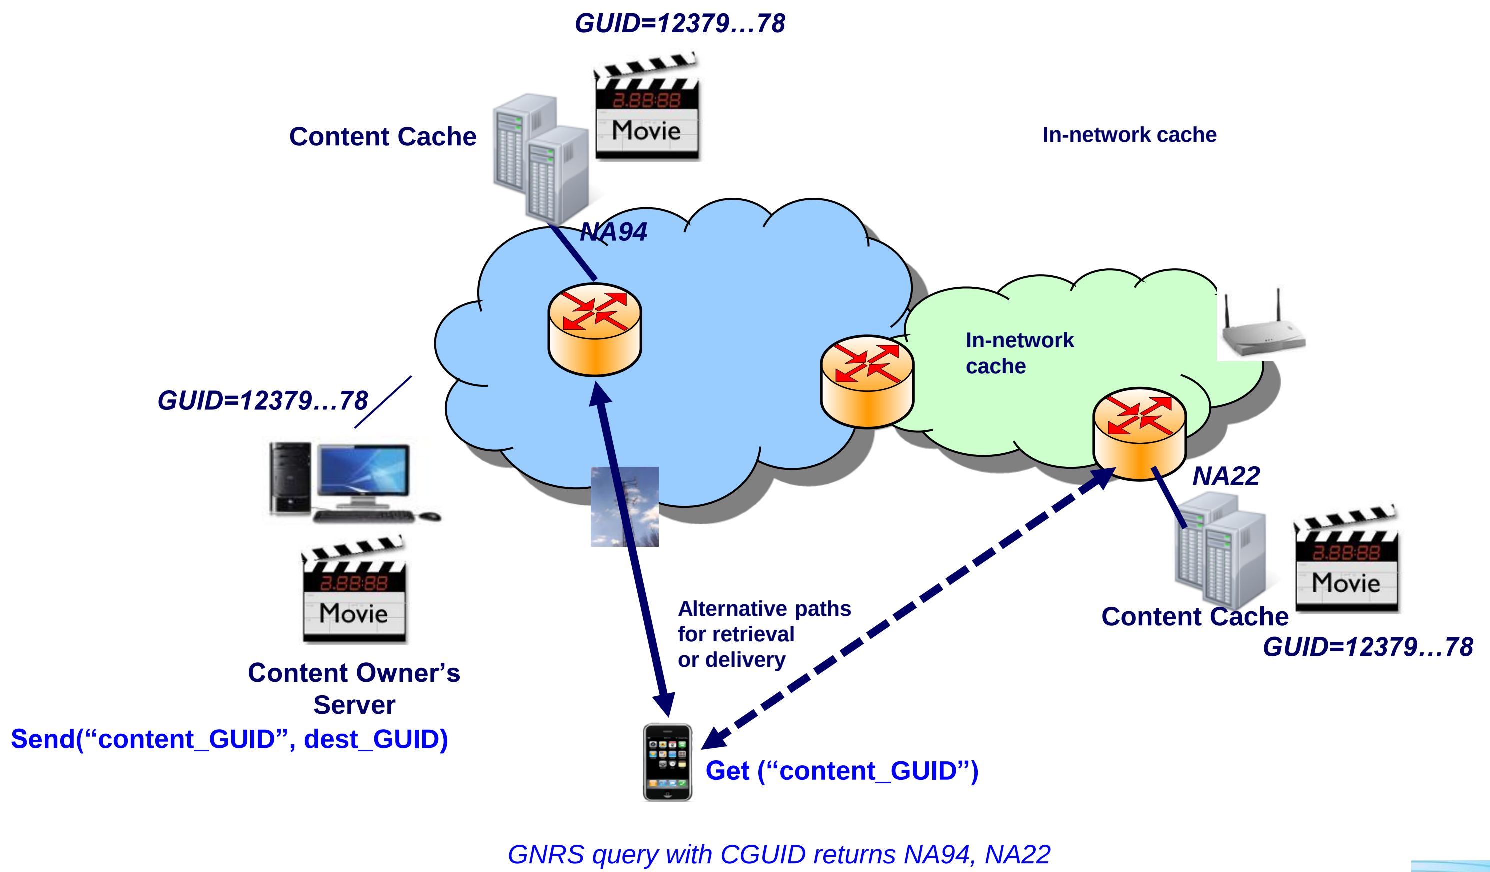

Content Delivery Use Case: Enhanced CDN Service

Enhanced service example – content delivery with in-network caching & transcoding

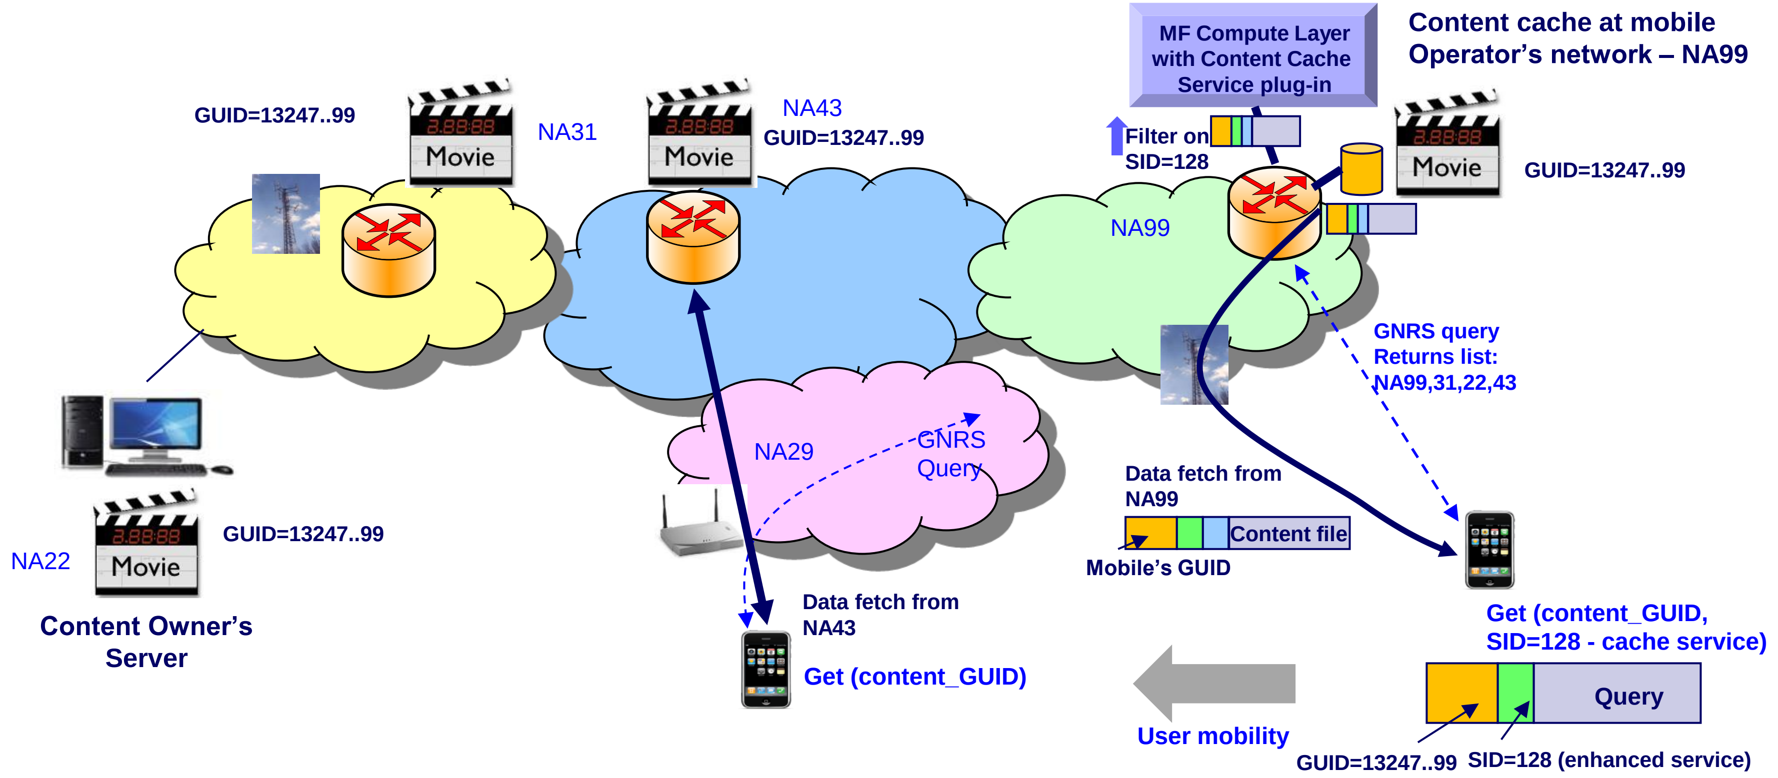

M2M Use Case:
Supporting Context-Aware Services

- Context-aware delivery often associated with M2M services
  - Examples of context are group membership, location, network state, …
  - Requires framework for defining and addressing context (e.g. “taxis in New Brunswick”)
  - Anycast and multicast services for message delivery to dynamic group

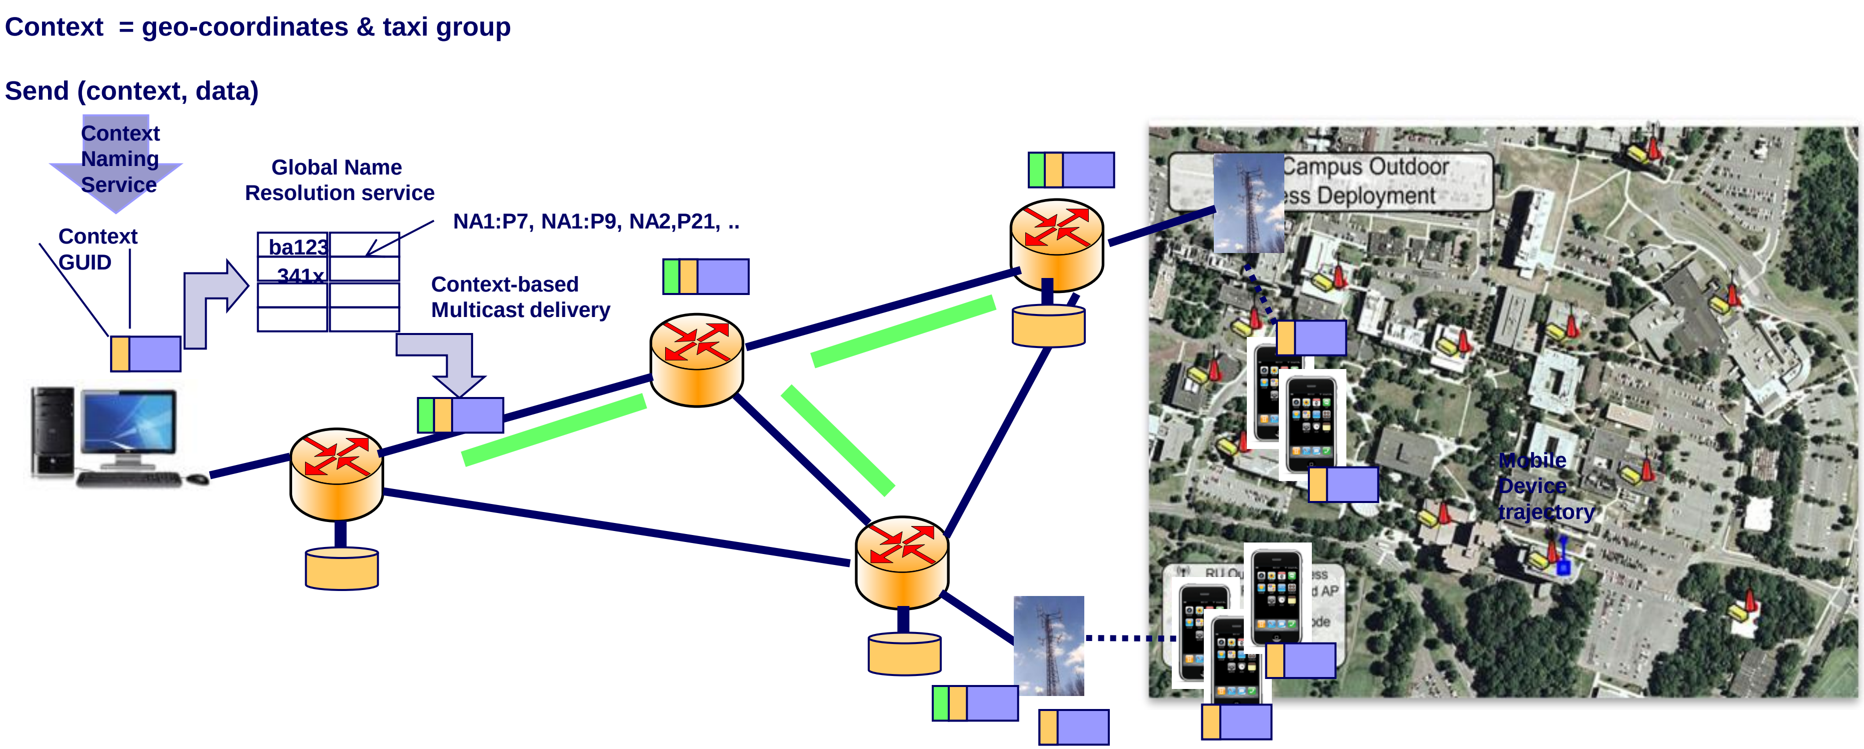

M2M Use Case: Protocol Example

- M2M: S1 data is picked up by Mobile GW and sent to MF for A1 that subscribes it
- Context: T1, conf@WINLAB, is subscribed by P1 …P4;A2 sends a file to T1 reaches P1..P4 via MF
- M2M/Context server updates MF-GNRS for mappings: S1→ A1 and T1 → P1…P4

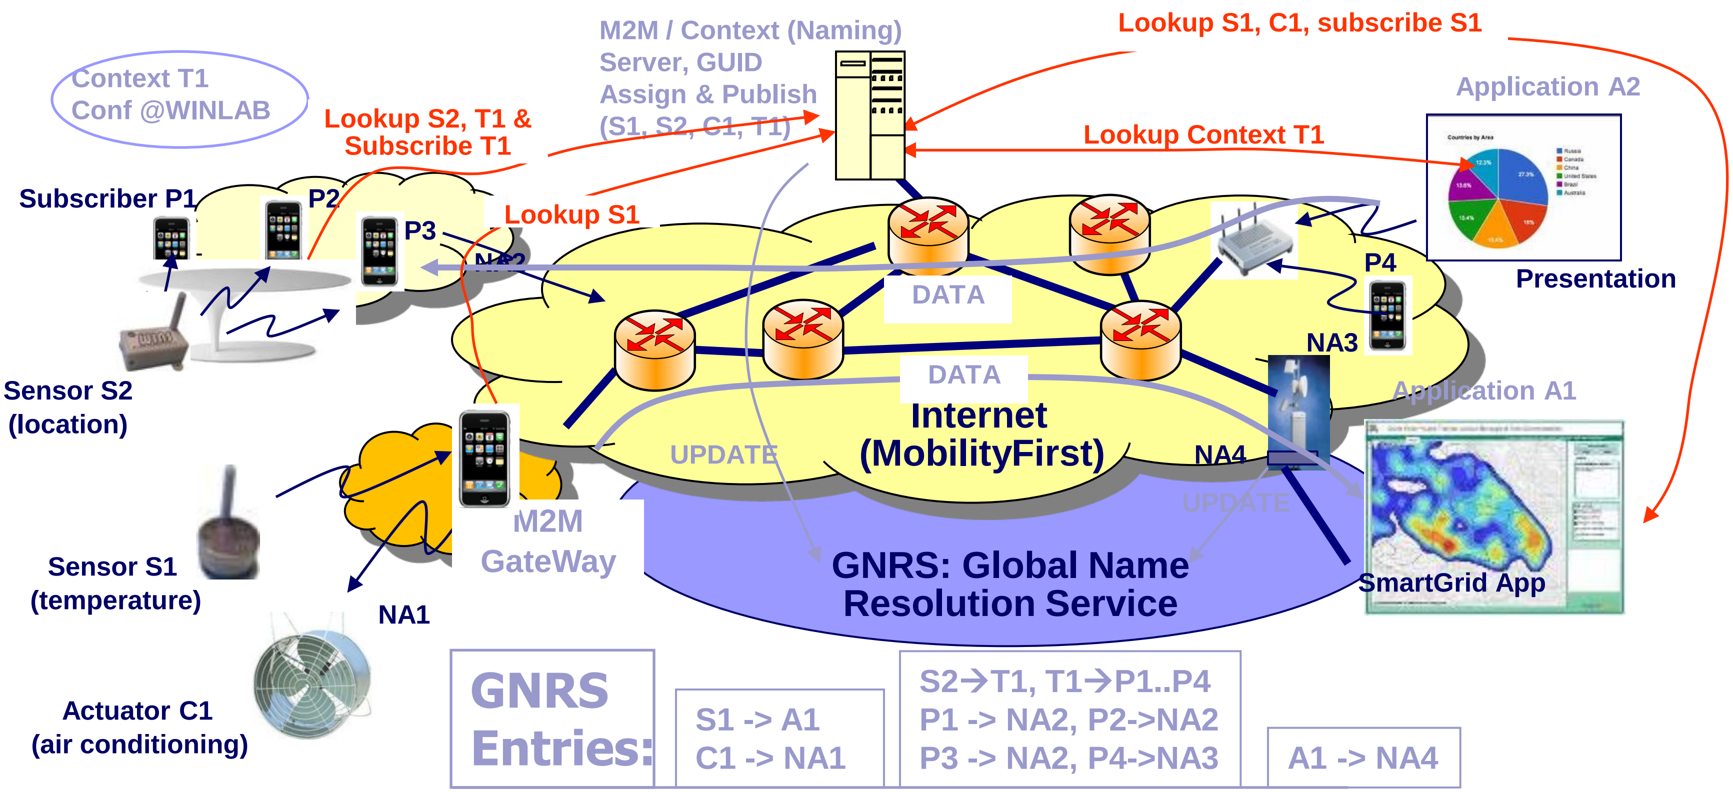

Mobility First Prototyping & Deployment Status

MobilityFirst Prototyping: Phased Strategy

Phase 1 | **Phase 2** | **Phase 3**

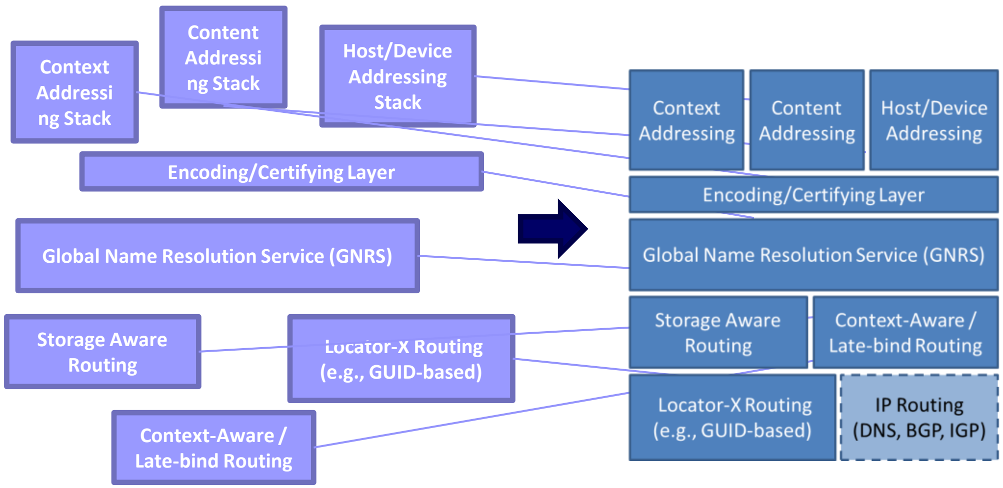

Prototype

Standalone Modules

Integrated MF Protocol Stack and Services

Deployable s/w pkg., box

Evaluation

Simulation and Emulation

Smaller Scale Testbed

Distributed Testbed
E.g. 'Live' on GENI

Experimental Platforms: ORBIT Testbed

- ORBIT is an NSF-supported open-access network research testbed in operation at WINLAB, Rutgers University since 2005
- Up to 400 radio/router nodes, with emulated network topology
- Options for software defined network (SDN) and software defined radio (SDR) technologies

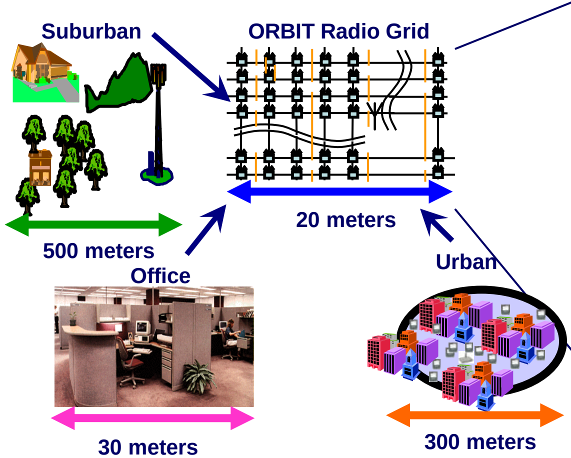

Radio Mapping Concept for ORBIT Emulator

400-node Radio Grid Facility at WINLAB Tech Center

Current ORBIT sandbox with GNU radio

Planned upgrade (2007-08)

Programmable ORBIT radio node

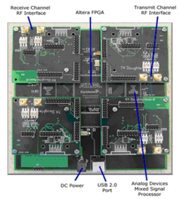

URSP2 CR board

Experimental Platform: ORBIT OpenFlow Sandbox

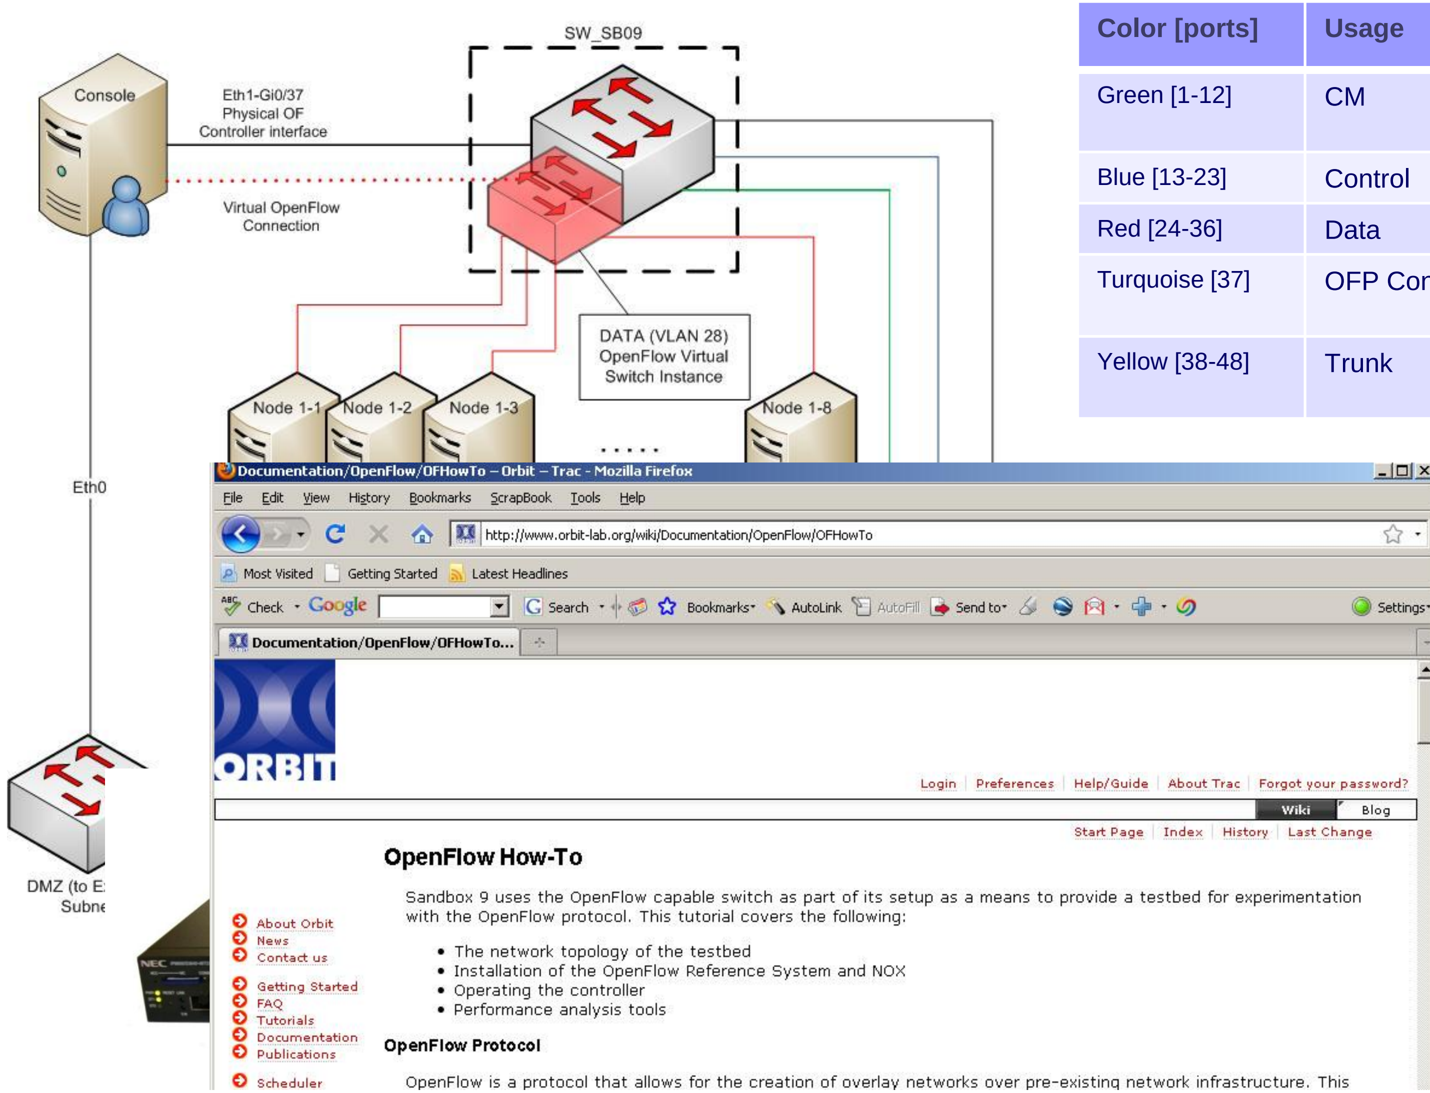

Color [ports]	Usage	VLAN
Green [1-12]	CM	3
Blue [13-23]	Control	27
Red [24-36]	Data	28
Turquoise [37]	OFP Controller	-
Yellow [38-48]	Trunk	-

ched to IP8800 switch
orts of the OF virtual switch
als:
low reference
OX-core
vireless experimentation

Experimental Platforms: GENI WiMAX Deployment at Rutgers

- Rt.1 campus deployment Q1/09
- Performance evaluation in progress

RF Module
(sector)

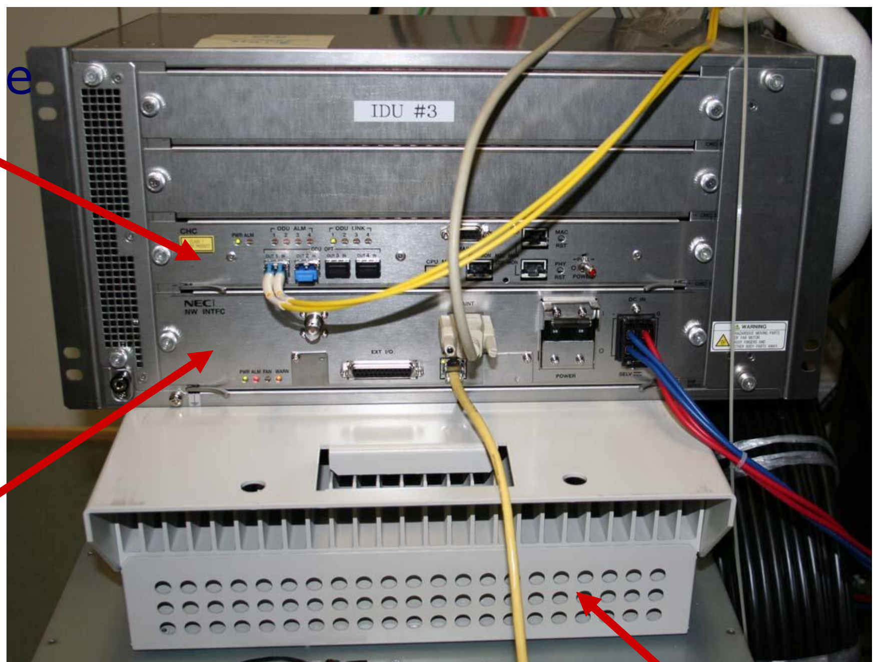

Base
Module

Outdoor Unit (ODU)

Omni-directional antenna
(elev. < 6ft above roof!)

MF Click Software Router

Lightweight, scalable multicast

- GNRS for maintenance of multicast memberships
- Heuristic approaches to reduce network load, limit duplicated buffering, and improve aggregate delivery delays
- Click prototype, with SID for multicast flows
- Evaluating hail a cab application as a example multipoint delivery scenario

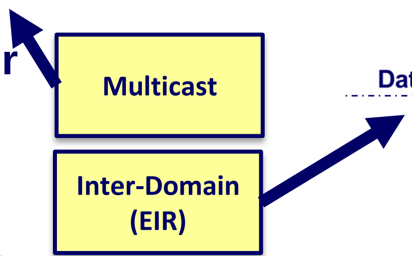

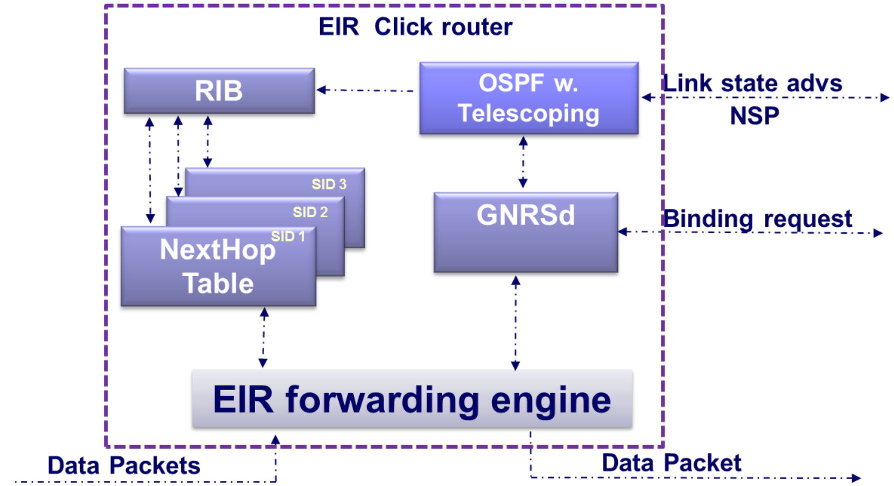

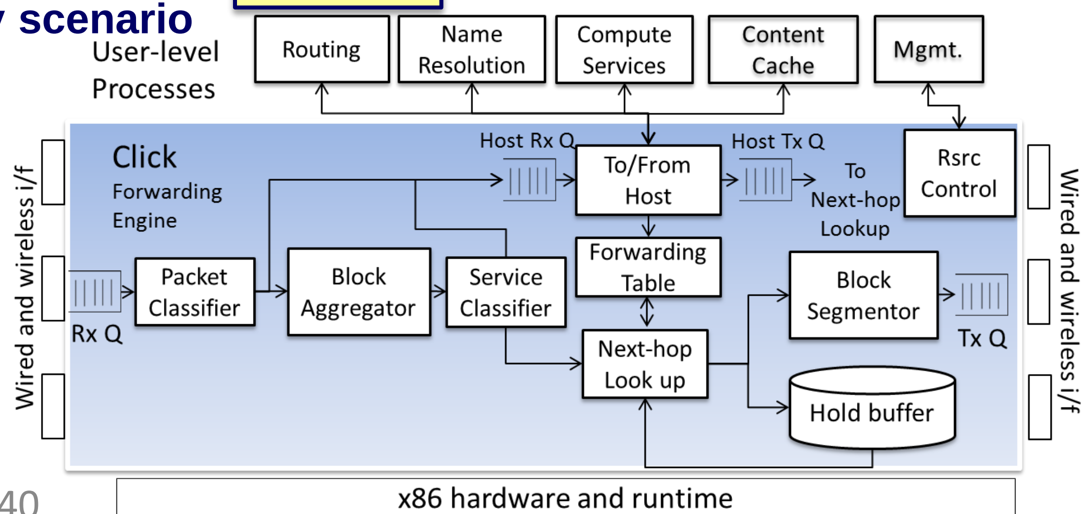

MobilityFirst Prototyping: Host Protocol Stack

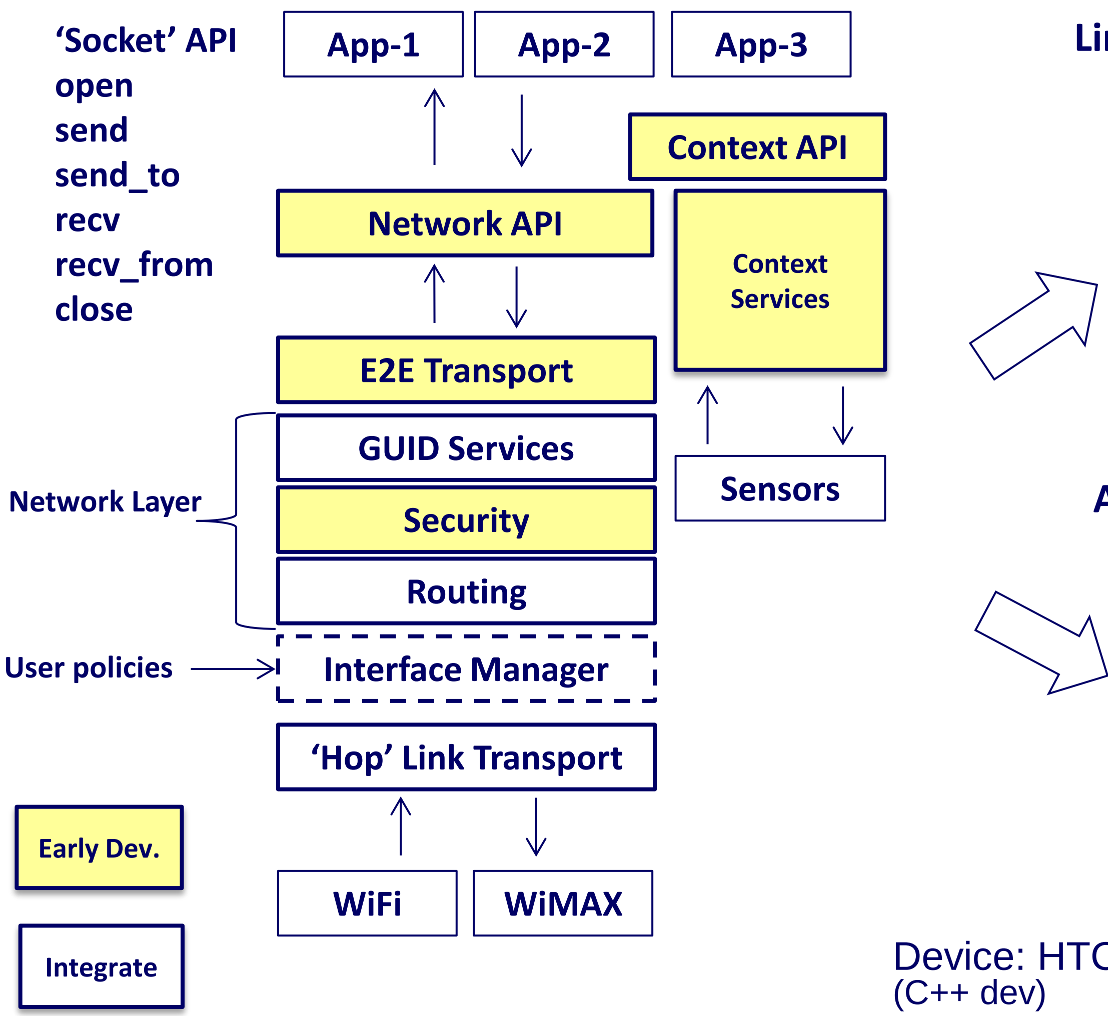

Linux PC/laptop with WiMAX & WiFi

Android device with WiMAX & WiFi

Device: HTC Evo 4G, Android v2.3 (rooted), NDK (C++ dev)

MF Routing Prototype on ORBIT

- Click-based prototyping of edge-aware inter-domain routing (EIR) on Orbit nodes
  - Implementation on 200+ nodes on the grid
  - Routing protocol uses "aNode" concept to disseminate full topology and aggregated edge network properties
  - Telescopic NSP (network state packet) advertisement for scalability

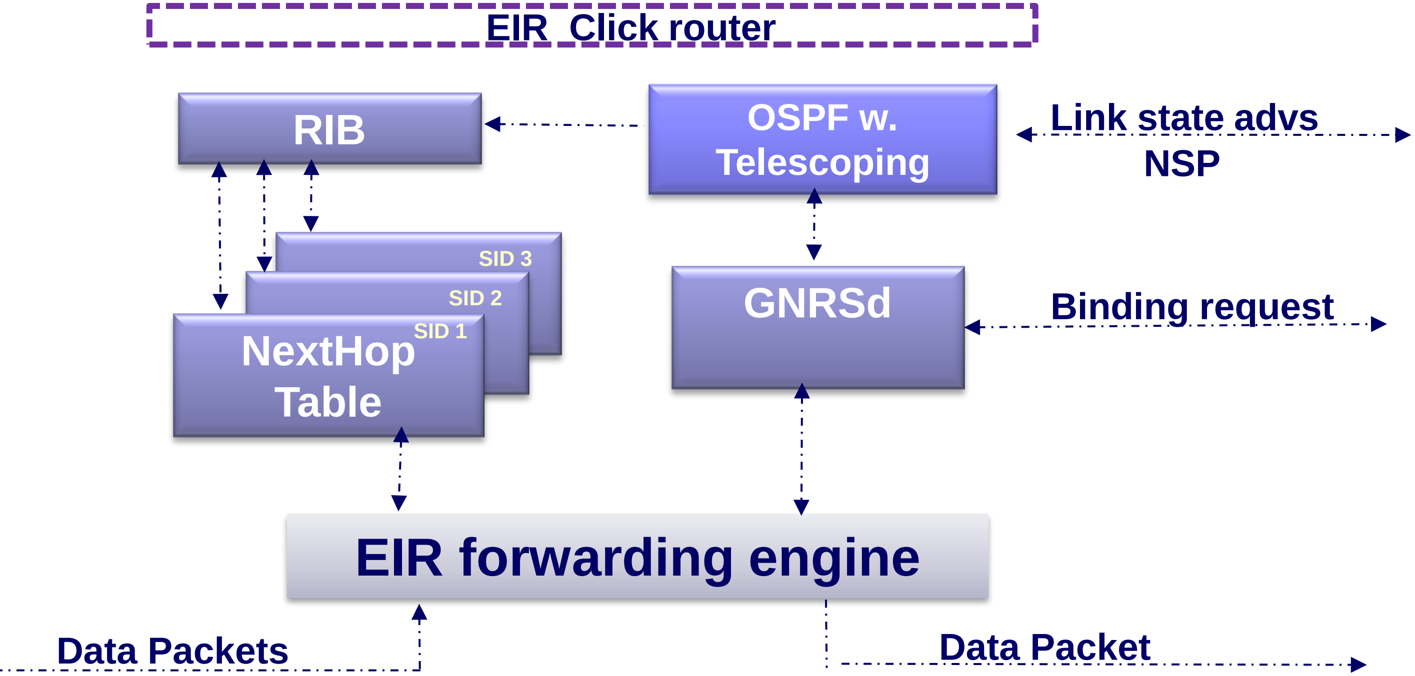

MF Routing Prototype on ORBIT (2)

- Click-based EIR prototype on Orbit nodes
  - Message delivery with late binding
  - Storage-aware routing
  - Efficient multicast & multipath data delivery

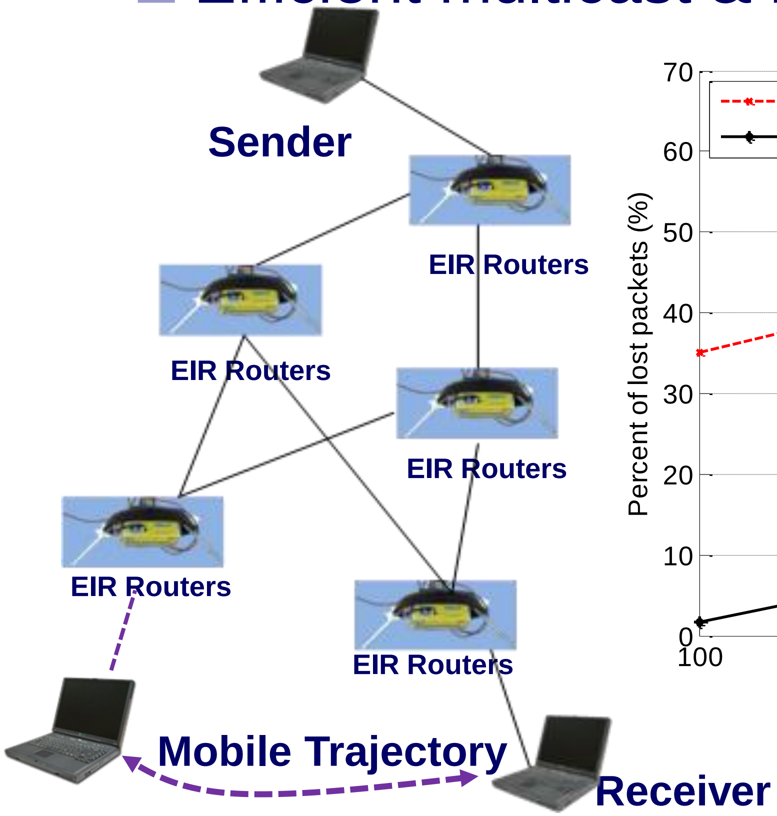

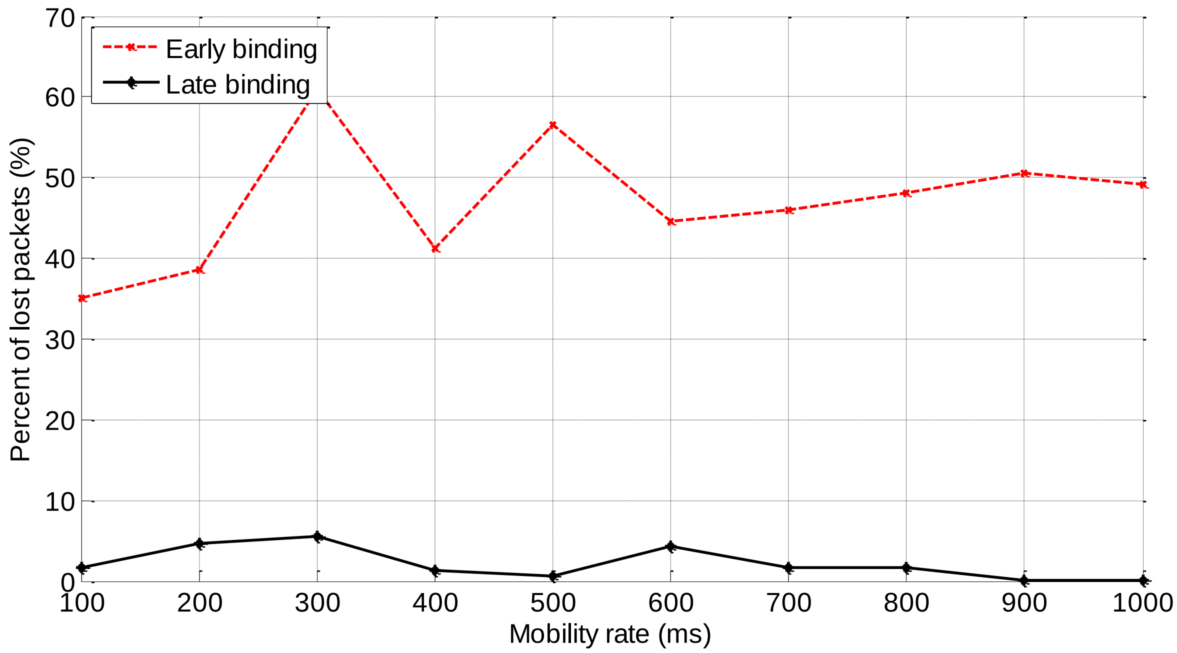

WINLAB

MF GNRS Implementation

- Two alternate designs:
  - network-level one-hop map service; co-located with routing (Dmap)
  - Locality-aware, cloud-hosted service (Auspice)
- Three alternate evaluation platforms:

1. GENI Wide Area Deployment

2. ORBIT lab with net. emulation

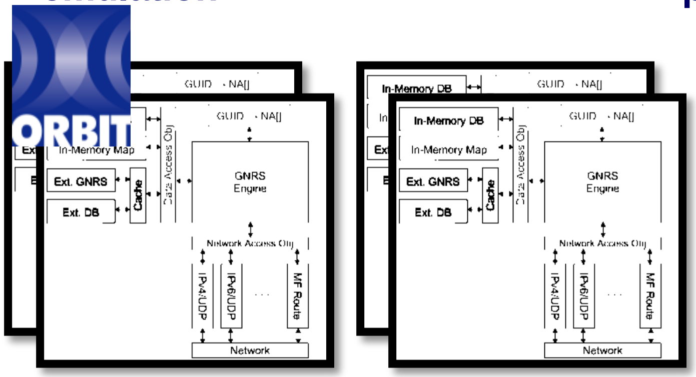

Network Emulation

3. Commercial cloud platform

amazon webservices™

OpenFlow/SDN Implementation of MF

- **Protocol stack embedded within controller**
- Label switching, NA or GUID-based routing (incl. GNRS lookup)
- **Controllers interact with other controllers and network support services such as GNRS**
- **Flow rule is set up for the remaining packets in the chunk based on Hop ID (which is inserted as a VLAN tag in all packets)**
  E.g., SRC MAC = 04:5e:3f:76:84:4a, VLAN = 101
  => OUT PORT = 16

MF Protocol Stack

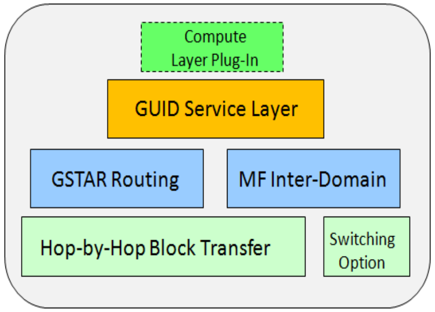

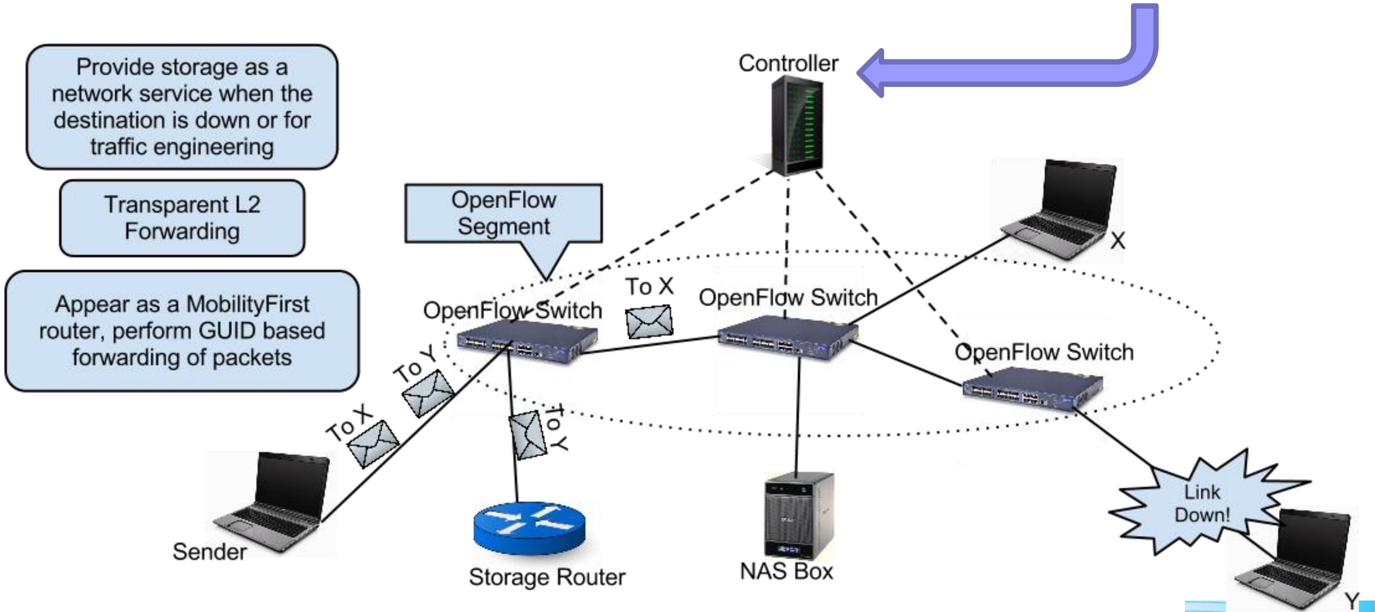

MF Router Prototype on FLARE SDN Platform from U Tokyo (Nakao)

- Objectives
  - Multi-site deployment of MobilityFirst routing and name resolution services
    - Impact of large RTTs on MobilityFirst network protocols
  - High performance evaluation of MobilityFirst delivery services on FLARE - 1Gbps, 10Gbps
    - Augmented Click router elements compiled down to FLARE native
  - Evaluation of FLARE platform for design and evaluation of next-generation network protocols
  - Demo at GEC-16, March 2013

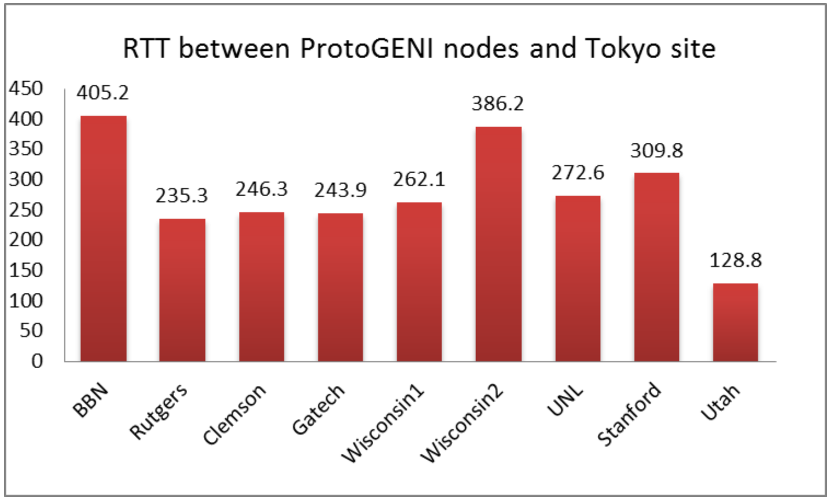

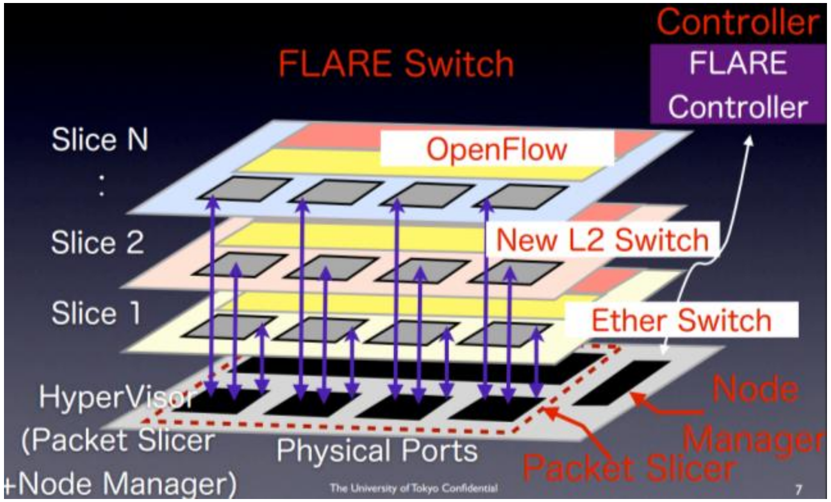

MF FPGA Router Prototype

- Designing Architecture and hardware (RTL) for NetFPGA platform to prototype Mobility First router
  - Supporting 4 Gigabit Ethernet Channels, standard 1500 Ethernet packets.
  - Supporting 2 DMA Channels for Host packet transfer (over PCI bus).
  - Designed to be configurable through Host computer
  - Designed for Fast NA (2-level cache) and GUID lookups.
- Evaluating different lookup strategies for GUID and NA
  - 2-level caches (BCAM for L 1 and SRAM for L2)
  - Compare and contrast of hardware Binary heap search Vs hardware Bloom Filters for L2 cache.

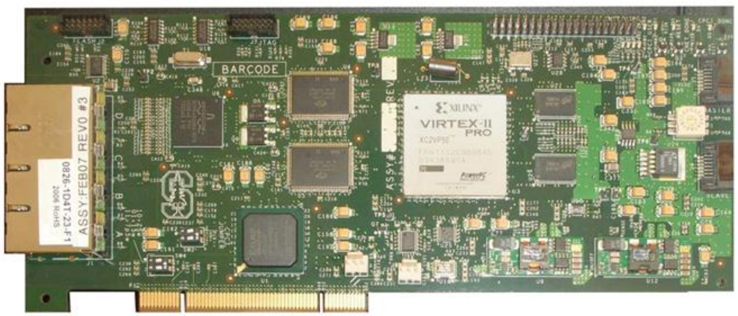

NetFPGA – 1st Generation:
Xilinx Virtex-II Pro 50
4 x 1Gbps bi-directional ports
SRAM: 4.5 MB DDR2 DRAM: 64MB

MobilityFirst Prototyping: GENI Deployment

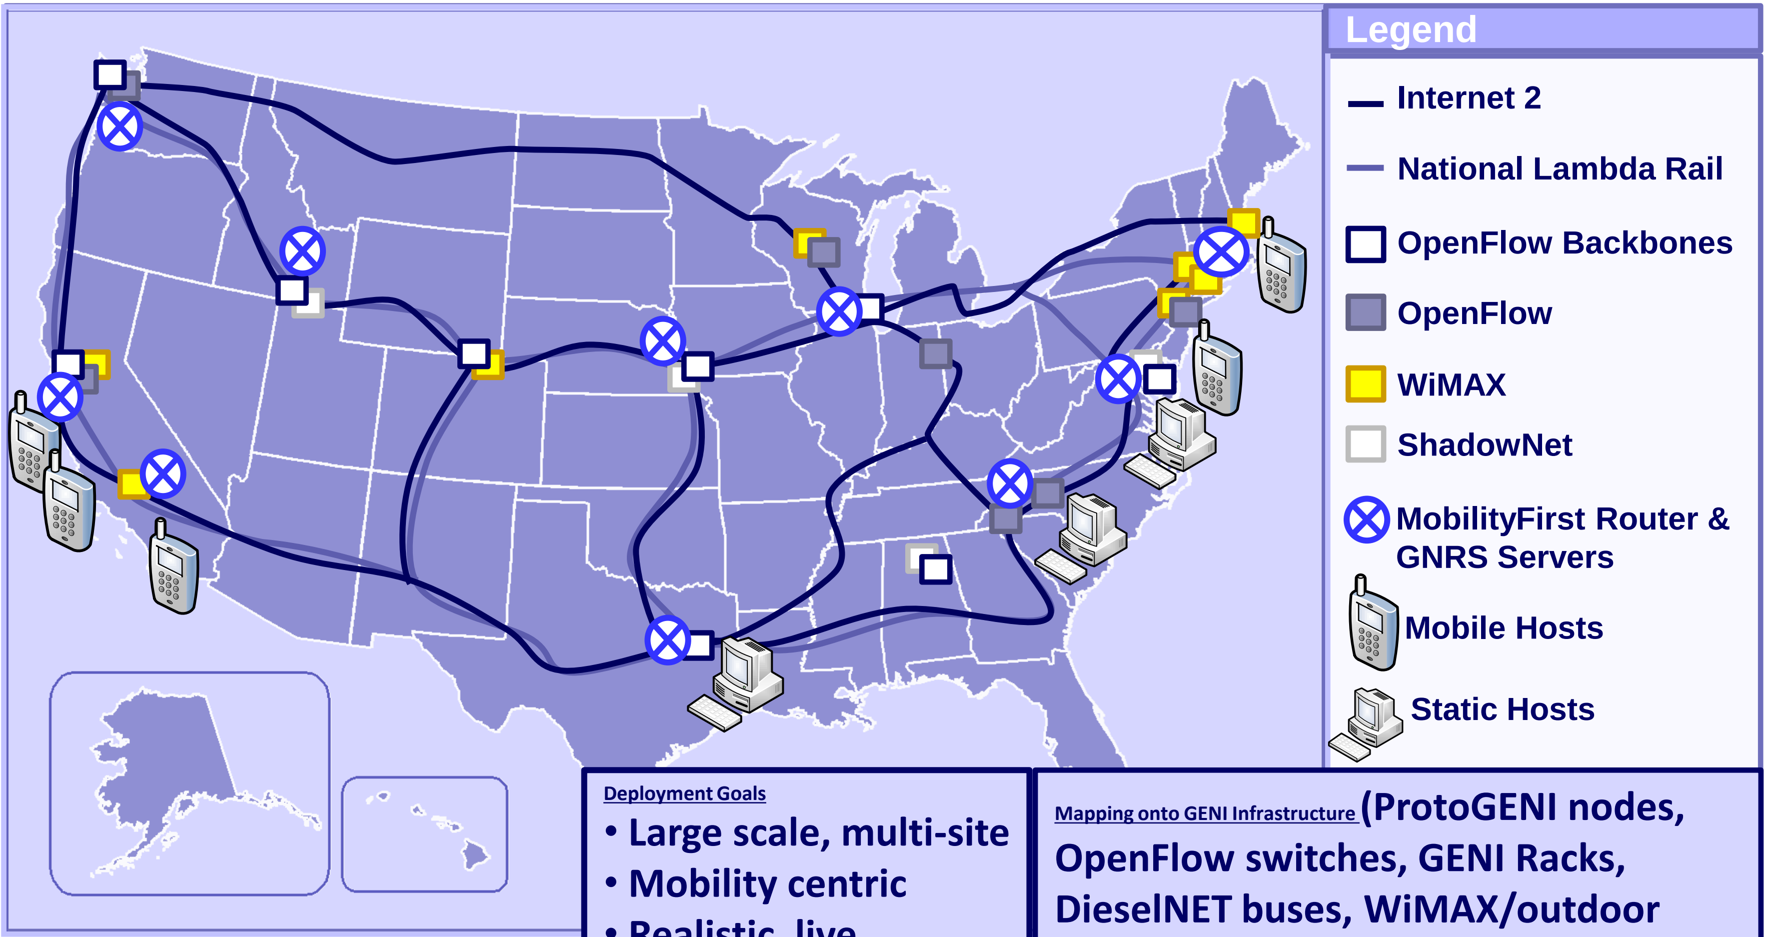

<u>Deployment Goals</u>

- Large scale, multi-site
- Mobility centric
- Realistic, live

<u>Mapping onto GENI Infrastructure</u> (ProtoGENI nodes, OpenFlow switches, GENI Racks, DieselNET buses, WiMAX/outdoor ORBIT nodes)

MF Prototype Deployment on GENI (GEC-12 Recap)

Physical Topology

- MF Routers at 7 campuses across US on ProtoGENI hosts
- Layer2 links: Internet2 & NLR backbones, OpenFlow switches
- Edge networks: WiFi and WiMAX access at BBN & Rutgers

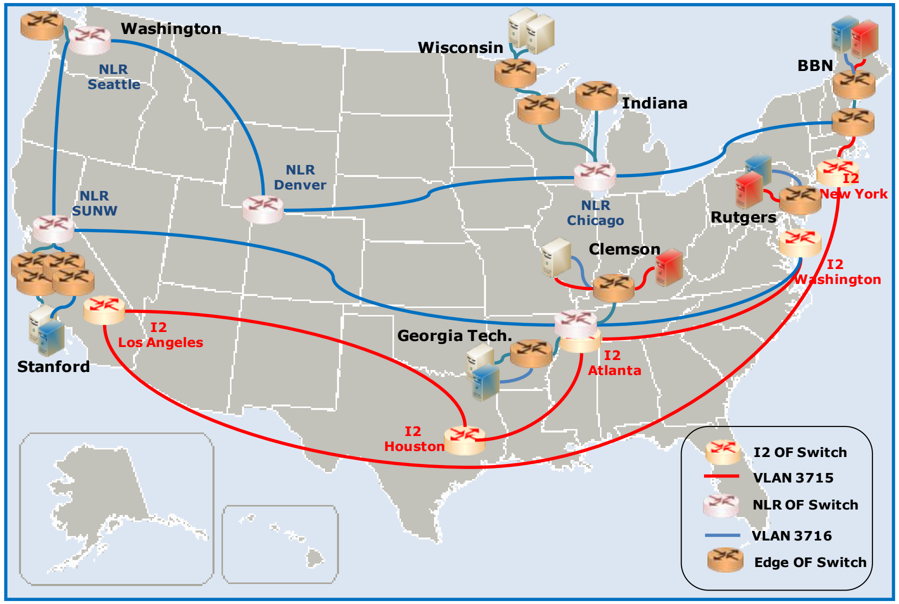

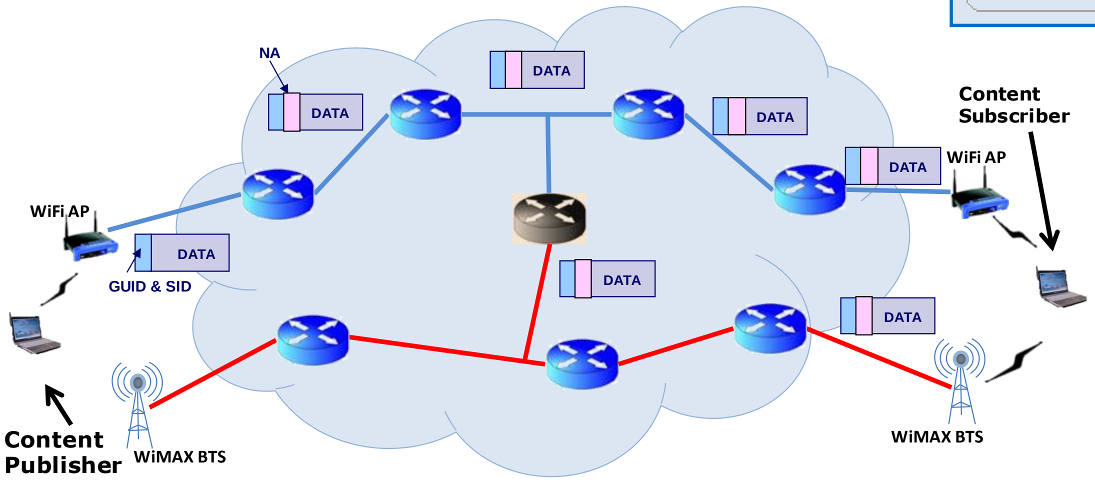

Robust Content Delivery

- Dual-homed mobile subscriber with WiFi + WiMAX
- Adapt to disconnection and variable link quality (GSTAR)
- Dual-homed delivery

WINLAB

MobilityFirst Prototyping: Hot Mobile 2012

Delivery Services for Multi-Homed Devices with User preference of delivery interface

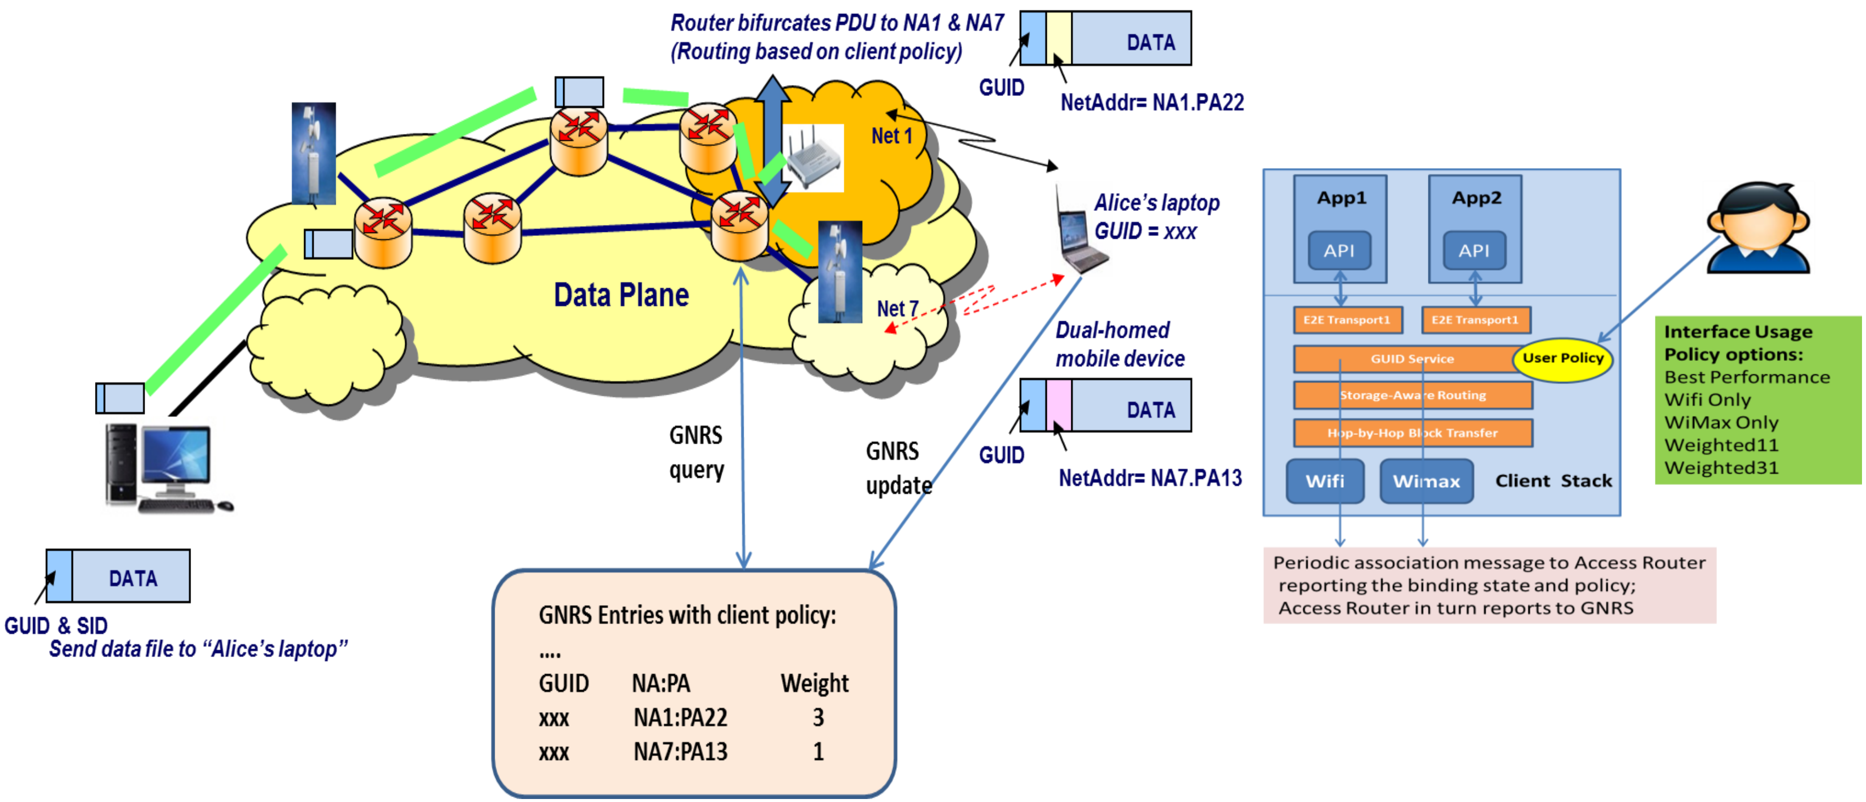

MobilityFirst Prototyping: GEC-13 Demo (Mobility, Multi-homing), ~3/12

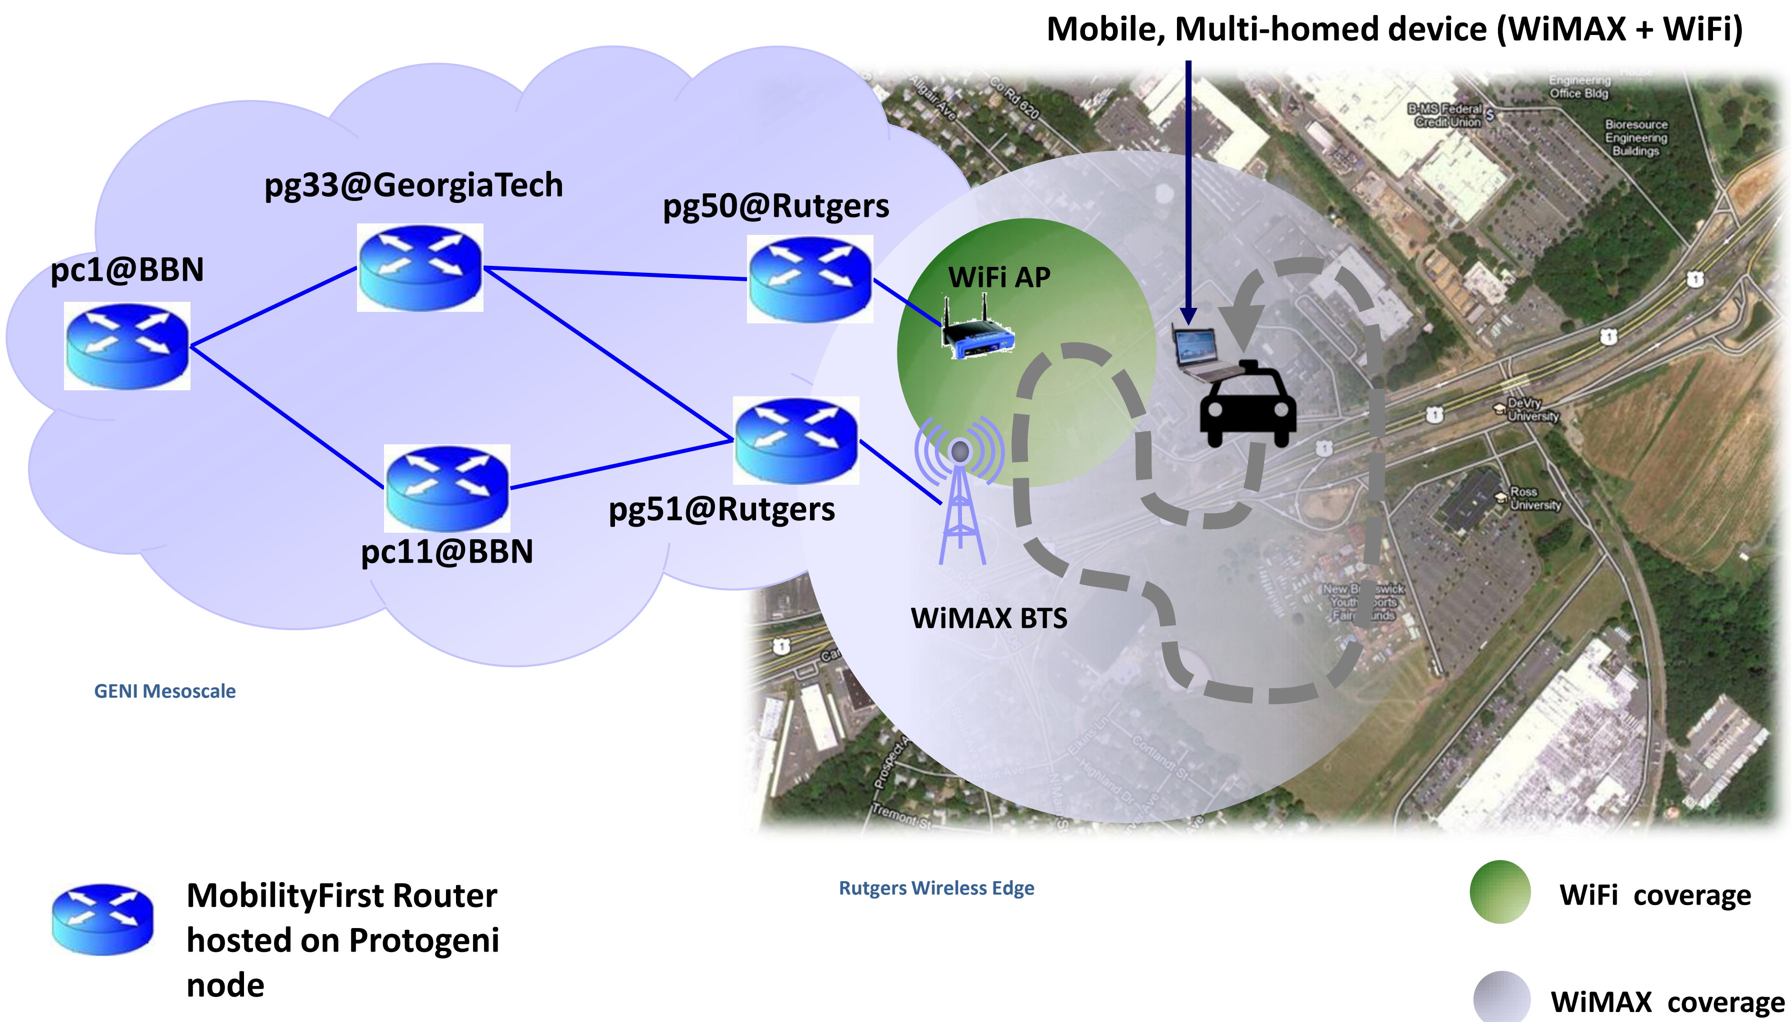

MF Multi-Site Deployment (GEC-16)

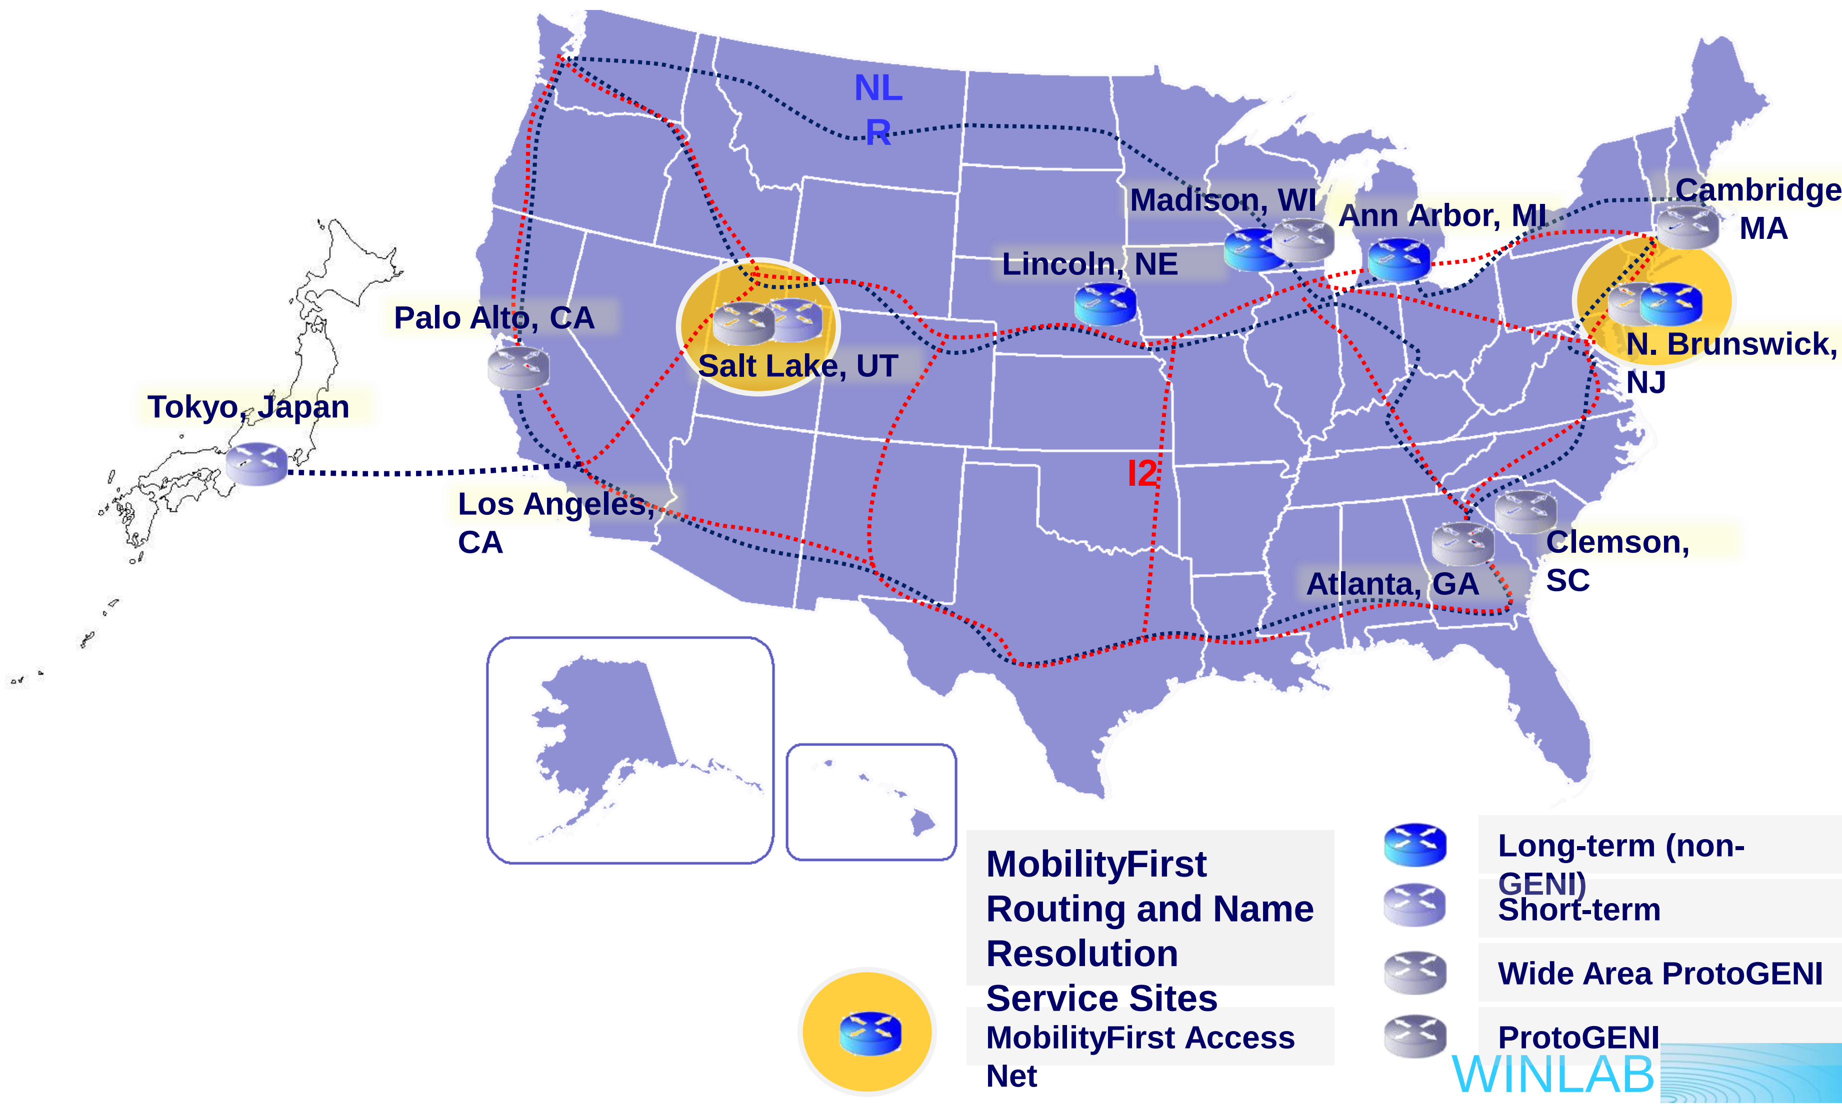

WINLAB

MF/GENI Connectivity: VLAN Stitching

Wide Area ProtoGENI nodes connect on VLAN 3716 at: Internet2 PoP
OR
NLR PoP

MobilityFirst Routing and Name Resolution Service Sites

MobilityFirst Access Net

Long-term (non-GENI)

Short-term

Wide Area ProtoGENI

ProtoGENI

WINLAB

MF/GENI Setup: Emulated Topology

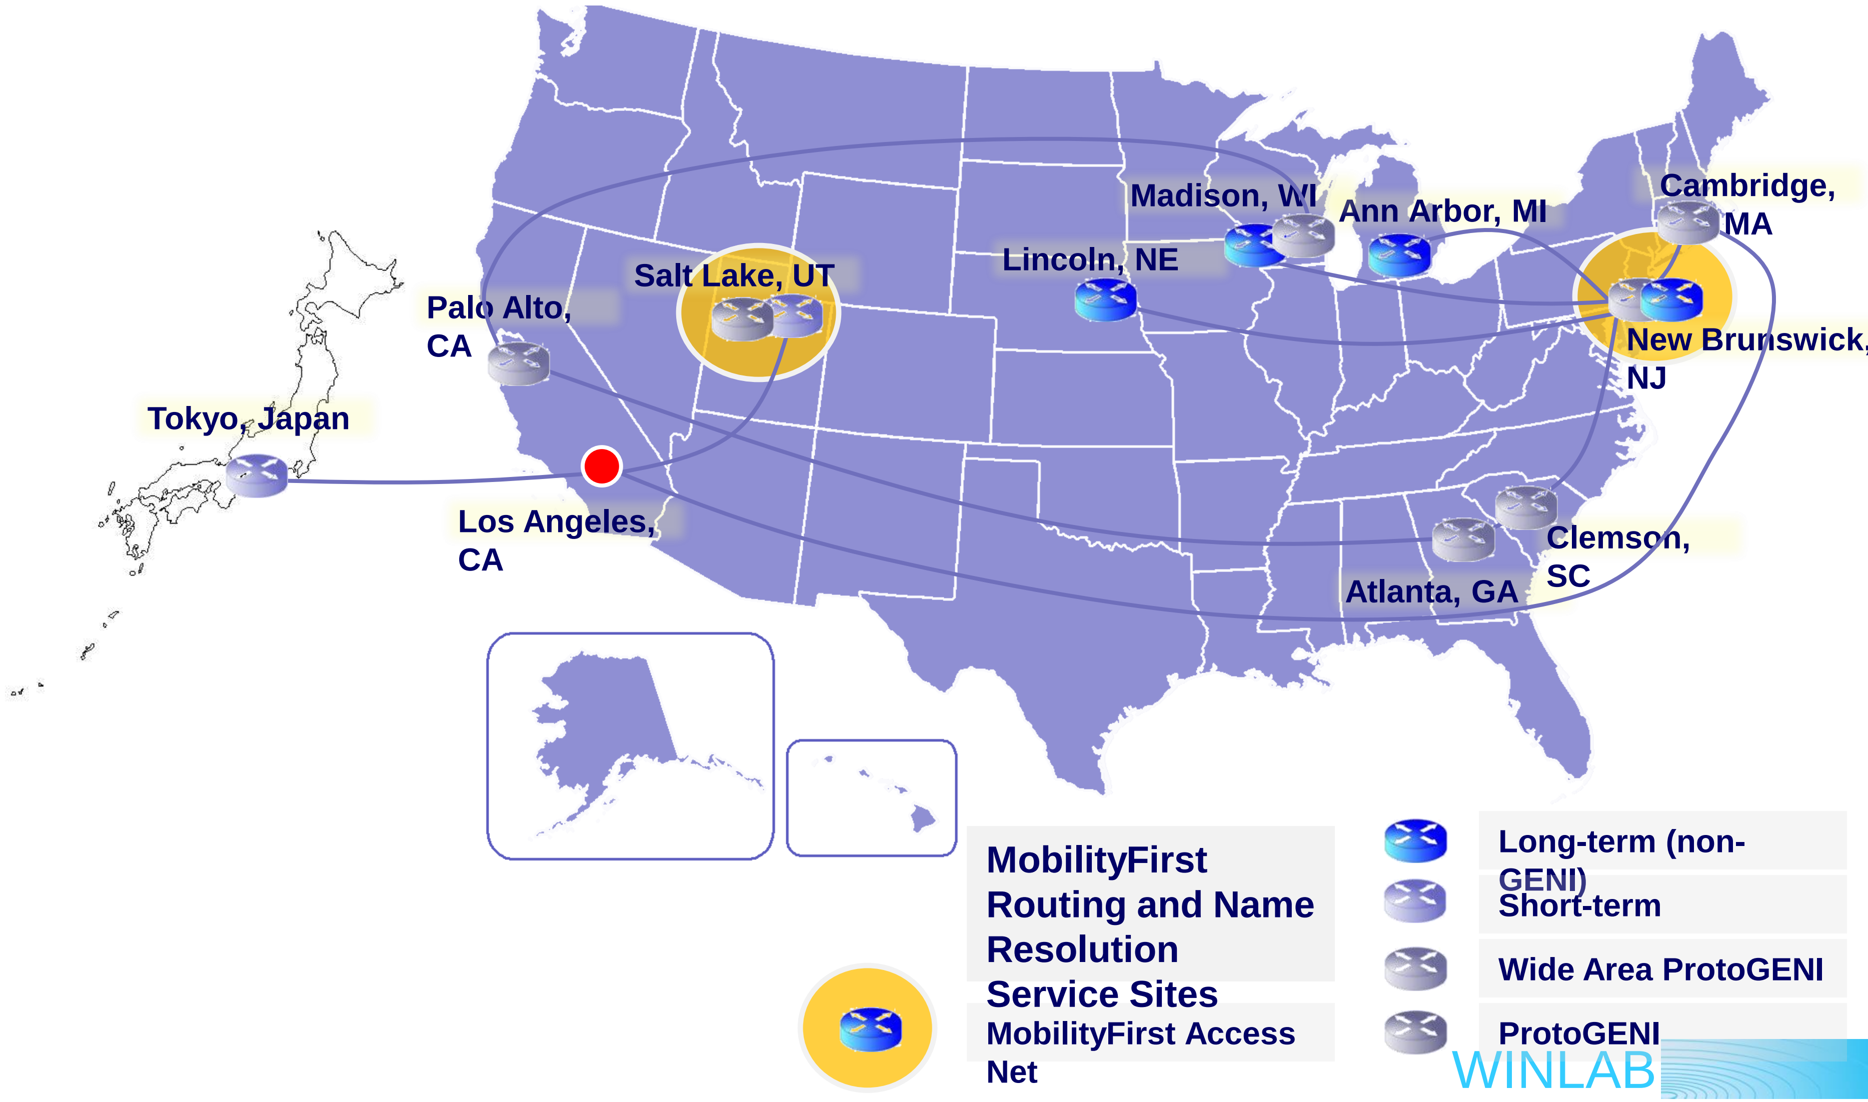

WINLAB

MobilityFirst GENI Demo Topology

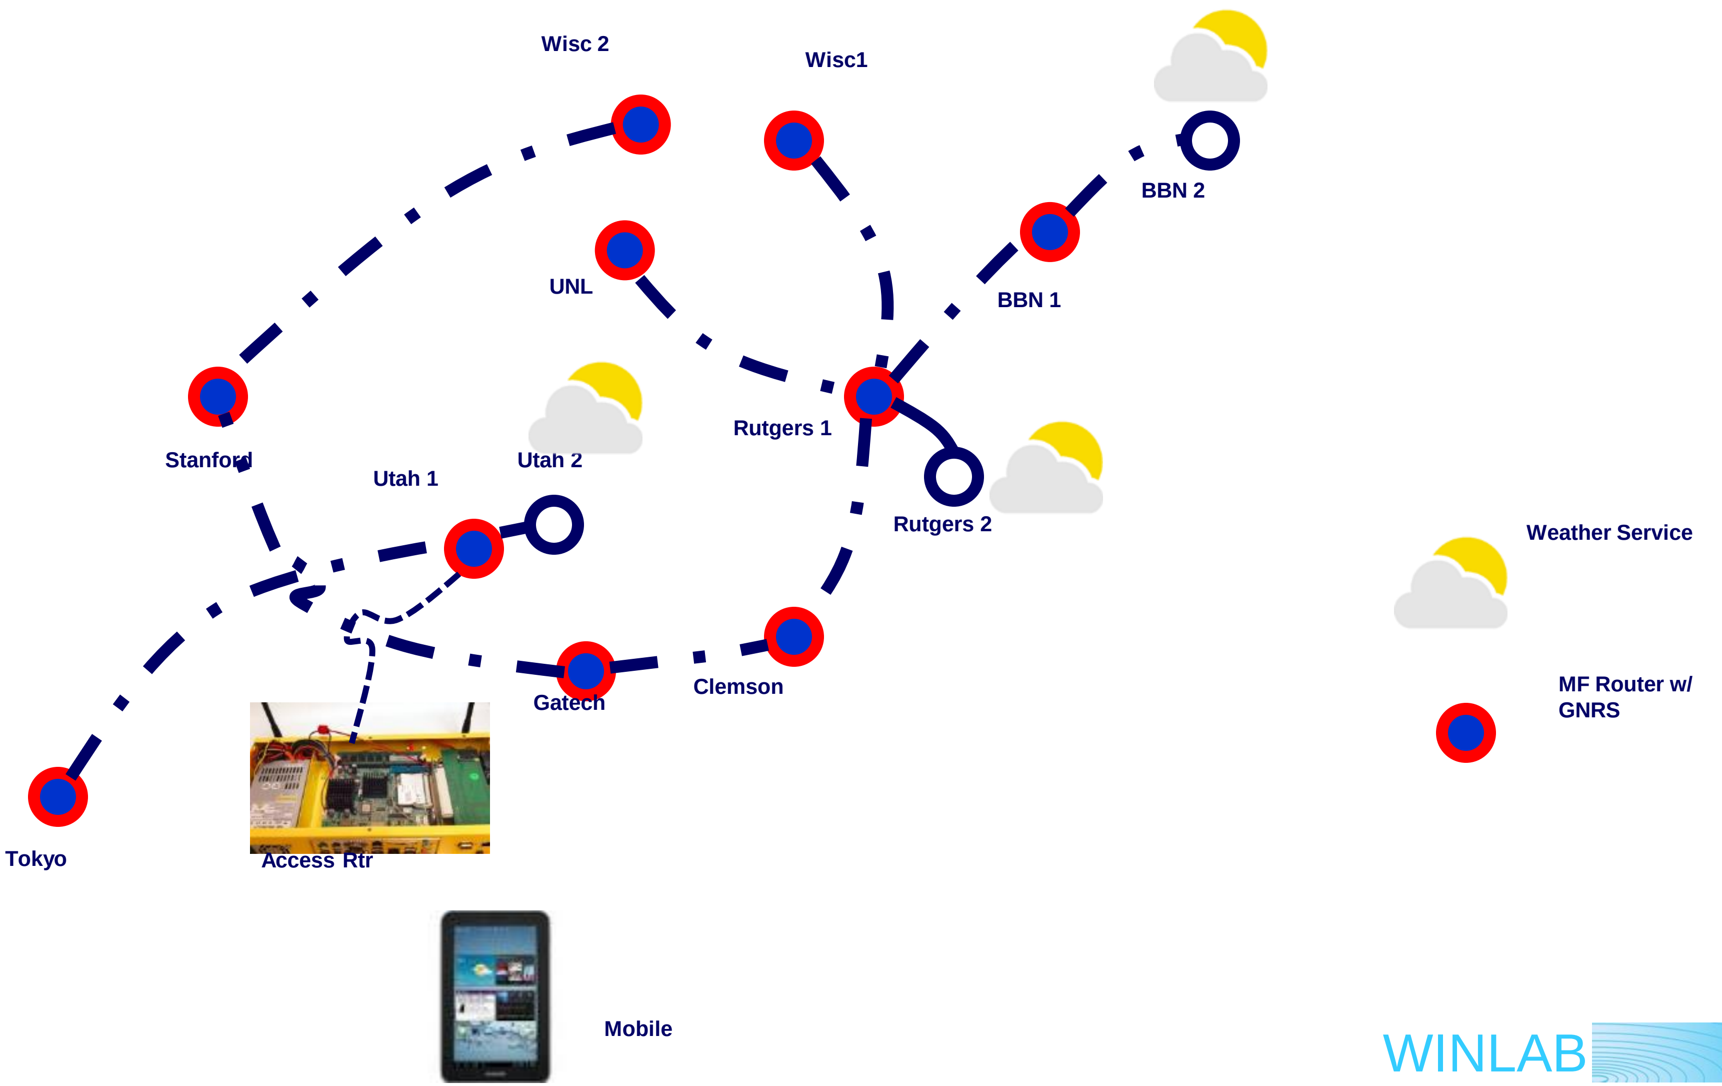

MobilityFirst GENI Demo – GUID Network & Service Identifiers

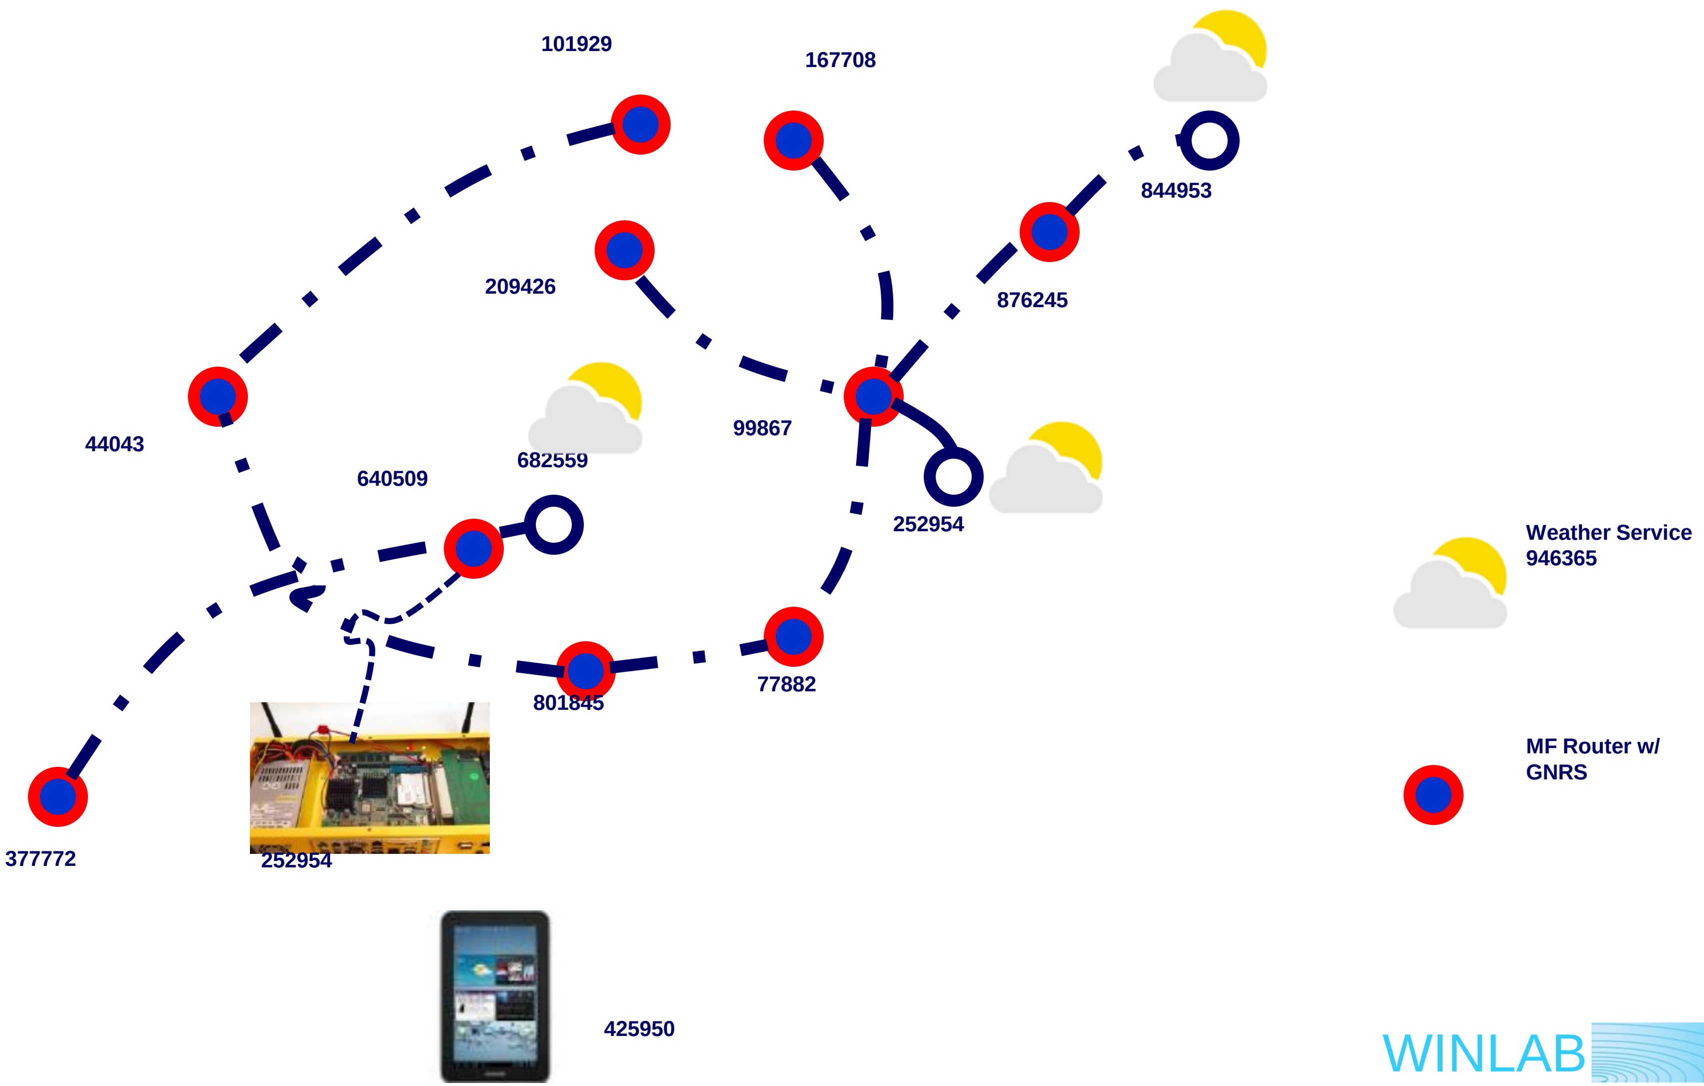

WINLAB

Resources

- Project website: http://mobilityfirst.winlab.rutgers.edu
- GENI website: www.geni.net
- ORBIT website: www.orbit-lab.org