Routing Architecture for the Next Generation Internet (RANGI)

Rangi = Sky father
Maori Mythology

Xiaohu Xu, Dayong Guo,
Raj Jain, Jianli Pan, Subharthi Paul

Presented to Routing Research Group (RRG),
Internet Research Task Force Meeting
Minneapolis, November 21, 2008

These slides are available at:
http://www.cse.wustl.edu/~jain/ietf/rangi.htm

- Part I: Long Term View – Internet 3.0
  - Internet 3.0: Next Generation Internet
  - User- Host- and Data Centric Models
  - Triple Tier Virtualization
- Part II: Short Term View – RANGI
  - A proposal to meet RRG Design Goals and More

Internet 3.0: Next Generation Internet

- Internet 3.0 is the name of the Washington University project on the next generation Internet
- Named by me along the lines of "Web 2.0"
- Internet 3.0 is more intuitive then GENI/FIND
- Goal 1: Develop a *clean slate architecture* to overcome limitations of the current internet
- Goal 2: Develop an *incremental approach* to implement the architecture

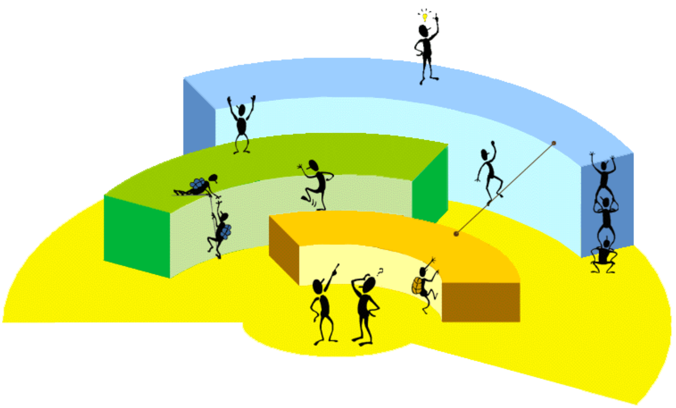

Washington University in St.Louis

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Internet 3.0: Next Generation Internet

- **Internet 1.0** (1969 – 1989) – Research project
  - RFC1 is dated April 1969.
  - ARPA project started a few years earlier
  - IP, TCP, UDP
  - Mostly researchers
  - Industry was busy with proprietary protocols: SNA, DECnet, AppleTalk, XNS

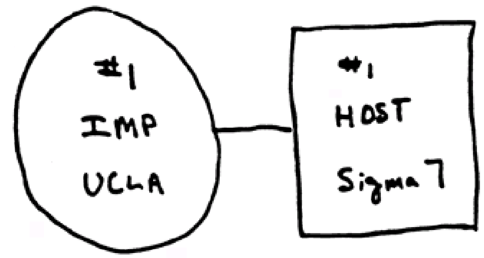

- **Internet 2.0** (1989 – Present) – Commerce ⇒ new requirements
  - Security RFC1108 in 1989
  - NSFnet became commercial
  - Inter-domain routing: OSPF, BGP,
  - IP Multicasting
  - Address Shortage IPv6
  - Congestion Control, Quality of Service,…

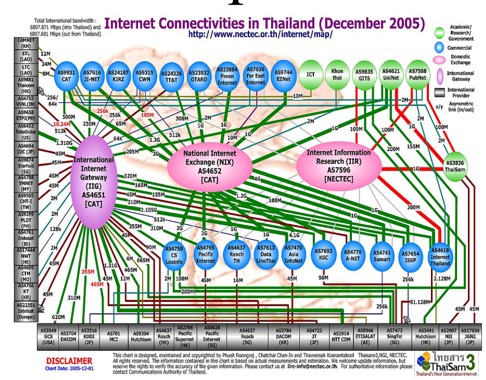

Washington University in St.Louis

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Key Problems with Current Internet

1. Security:
   - Inability to enforce policies related to Authorization, authentication, privacy, resource utilizations
   - Perimeter based representation of organization is not sufficient

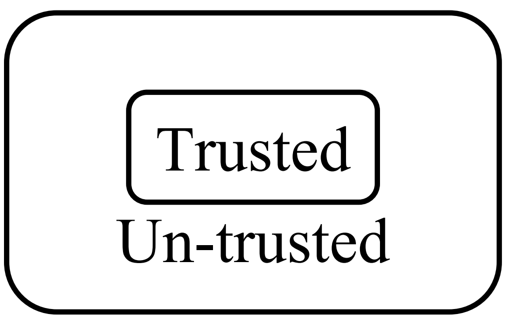

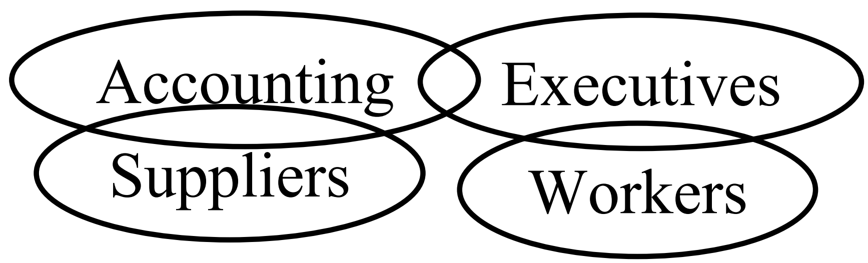

Problems (cont)

2. No representation for real end systems: the human.

3. Identity and location in one (IP Address) Makes mobility complex. [Well known]

Ref: For a bigger list see our Milcom 2006 paper [1]

Realms

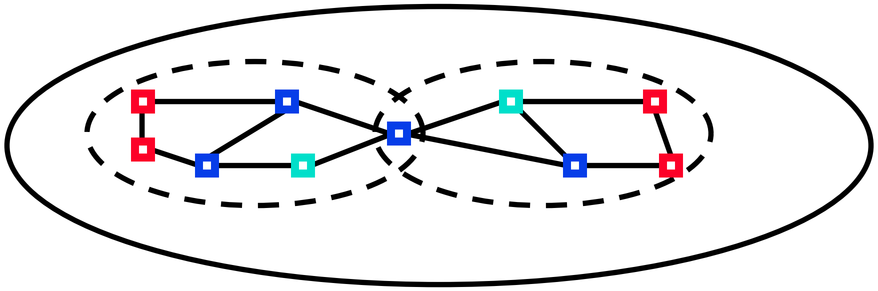

- Object names and Ids are defined within a realm
- A realm is a **logical** grouping of objects under an administrative domain
- The Administrative domain may be based on Trust Relationships
- A realm represents an organization
  - Realm managers set policies for communications
  - Realm members can share services.
  - Objects are generally members of multiple realms
- Realm Boundaries: Organizational, Governmental, ISP, P2P,…

Realm = Administrative Group

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Physical vs. Logical Connectivity

- Physically and logically connected:
  All computers in my lab
  = Private Network,
  Firewalled Network
- Physically disconnected but logically connected:
  My home and office computers

- Physically connected but logically disconnected: Passengers on a plane, Neighbors, Conference attendees sharing a wireless network, A visitor

Physical connectivity ≠ Trust

Id-Locator Split Architecture (MILSA)

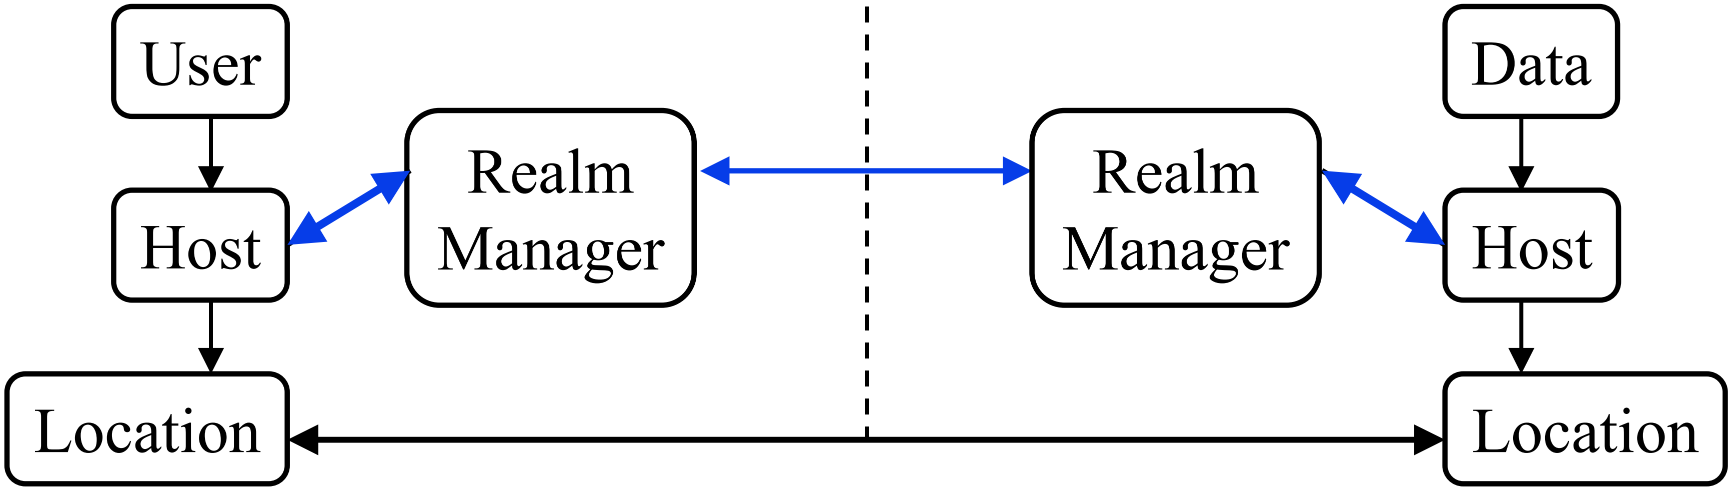

- Realm managers:
  - Resolve current location for a given host-ID
  - Enforce policies related to authentication, authorization, privacy
  - Allow mobility, multi-homing, location privacy
  - Similar to several other proposals
- Ref: Our Globecom 2008 paper [2]

User- Host- and Data Centric Models

- All discussion so far assumed host-centric communication
  - Host mobility and multihoming
  - Policies, services, and trust are related to hosts
- User Centric View:
  - Bob wants to watch a movie
  - Starts it on his media server
  - Continues on his iPod during commute to work
  - Movie exists on many servers
  - Bob may get it from different servers at different times or multiple servers at the same time
- Can we just give addresses to users and treat them as hosts? No! ⇒ Policy Oriented Naming Architecture (PONA)

Policy Oriented Naming/Routing

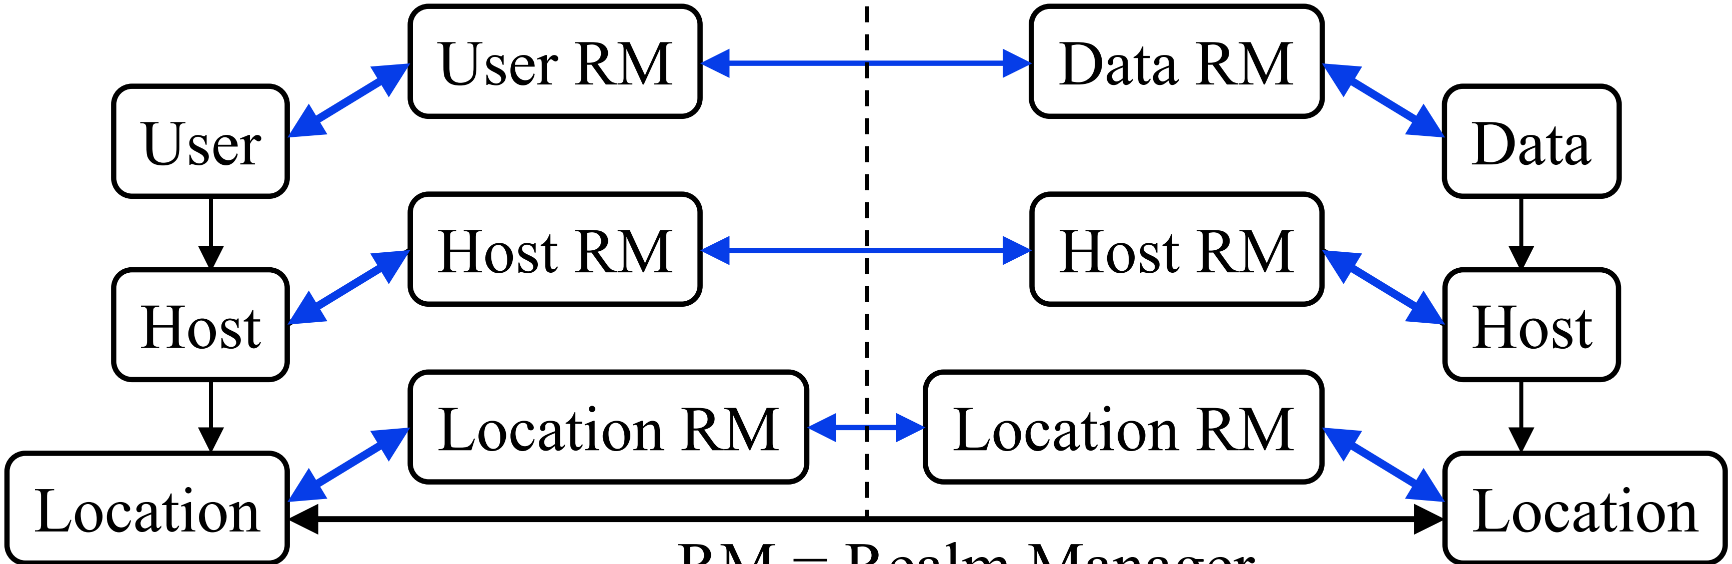

RM = Realm Manager

- Both Users and data need hosts for communication
- Data is easily replicable. All copies are equally good.
- Users, Hosts, Infrastructure, Data belong to different realms (organizations).
- Each object has to follow its organizational policies.

Virtualizable Network Concept

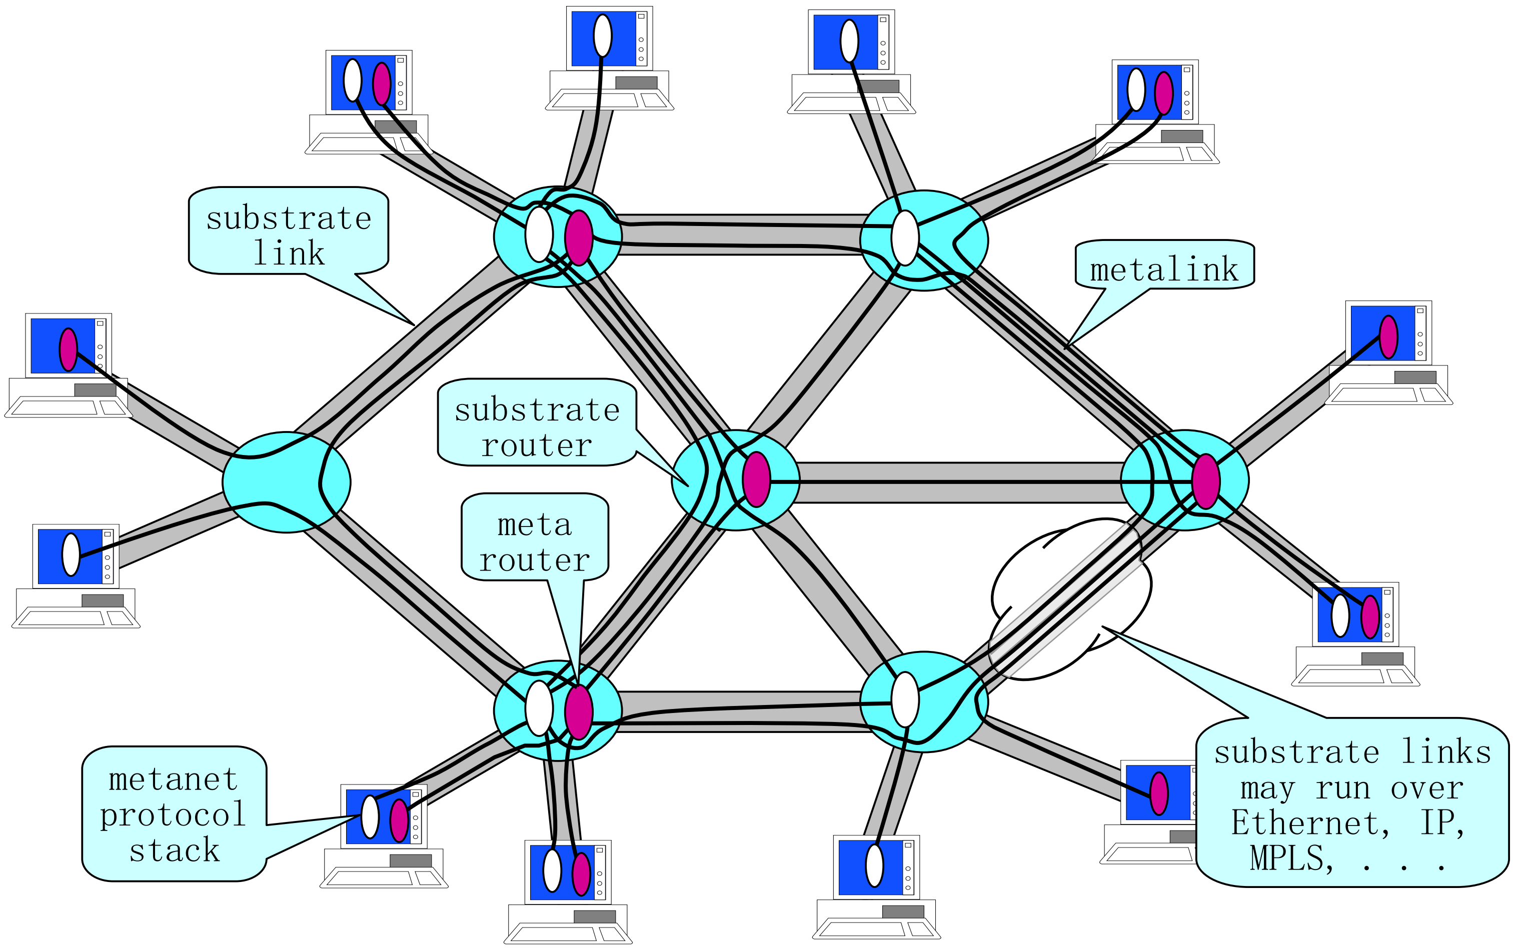

Ref: T. Anderson, L. Peterson, S. Shenker, J. Turner, "Overcoming the Internet Impasse through Virtualization," Computer, April 2005, pp. 34 – 41.

Slide taken from Jon Turner's presentation at Cisco Routing Research Symposium

Realm Virtualization

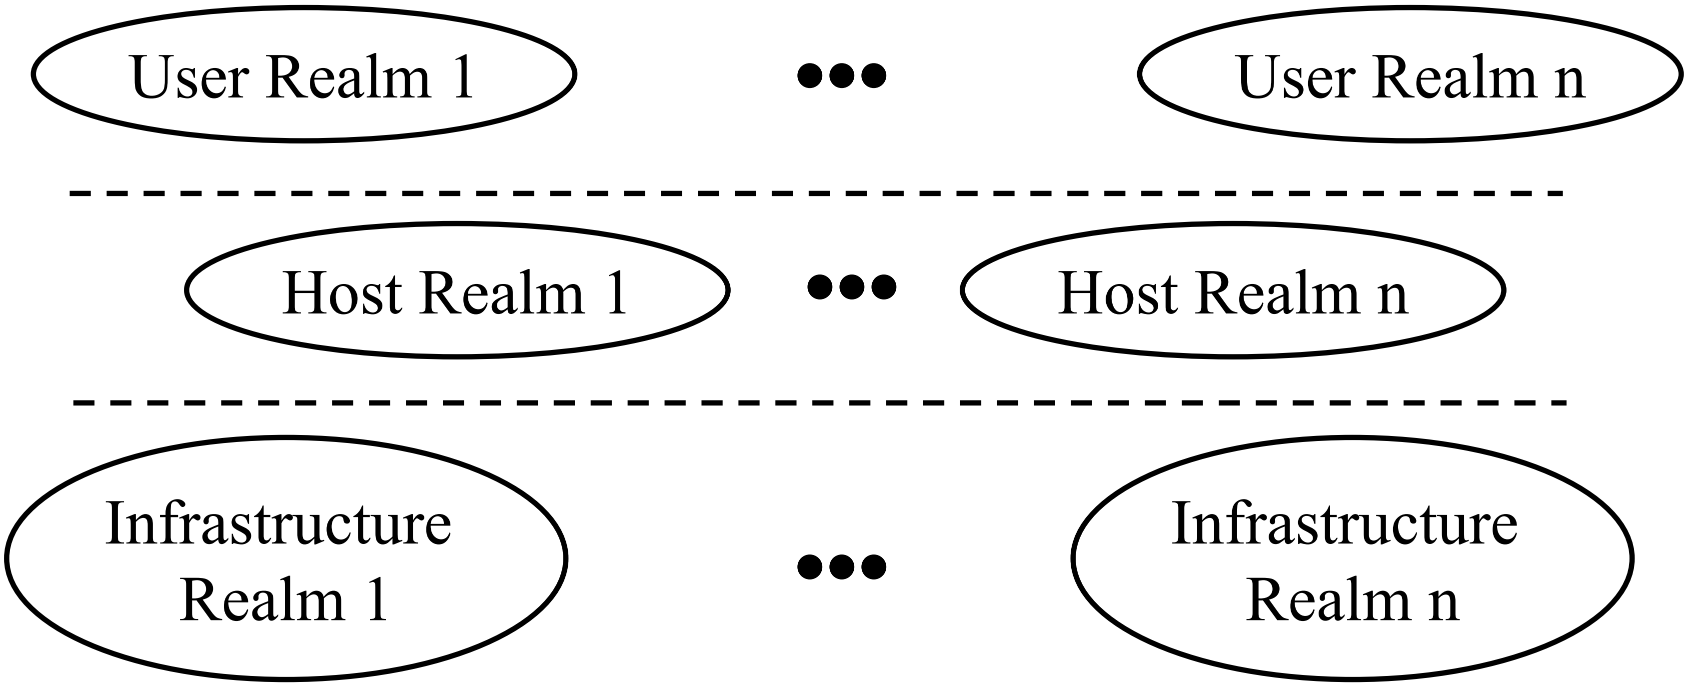

- Old: Virtual networks on a common infrastructure
- New: Virtual user realms on virtual host realms on a group of infrastructure realms. 3-level hierarchy not 2-level. Multiple organizations at each level.
- Ref: Our PONA paper [3]

Summary: Part I

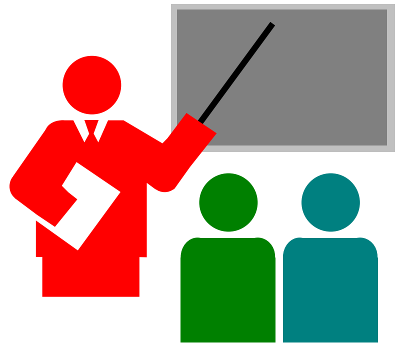

1. Internet 3.0 is the next generation of Internet.
2. It must be secure, allow mobility, and be energy efficient.
3. Must be designed for commerce
   ⇒ Must represent multi-organizational structure and policies
4. Moving from host centric view to user-data centric view
   ⇒ Important to represent users and data objects
5. Users, Hosts, and infrastructures belong to different realms (organizations). Users/data/hosts should be able to move freely without interrupting a network connection.

Part II: Immediate Goals for the Next Generation Routing

1. Routing Scalability
2. Traffic Engineering
3. Mobility and Multihoming
4. Simplified Renumbering
5. Decoupling Location and Identification
6. Routing Quality
7. Routing Security
8. Incremental Deployability

Ref: RRG Workshop

Current State of the Internet

- IPv4 is ubiquitous among hosts and routers
- IPv6 has been implemented in hosts (Windows) But most routers are still IPv4
- Inter-Domain routing is complex
  - Renumbering ⇒ Customers want PI addresses
  - Service providers have difficulty supporting PI addresses
- Need a solution for the current state ⇒ Routing Architecture for the Next Generation Internet (RANGI) [Under development]

RANGI Design Goals

1. Routing Scalability
2. Traffic Engineering
3. Mobility and Multihoming
4. Simplified Renumbering
5. Decoupling Location and Identification
6. Routing Quality
7. Routing Security: Also avoids ID theft
8. Incremental Deployability
9. **Business friendly realm and domain boundaries**

Ref: HRA paper [4]

RANGI Assumptions

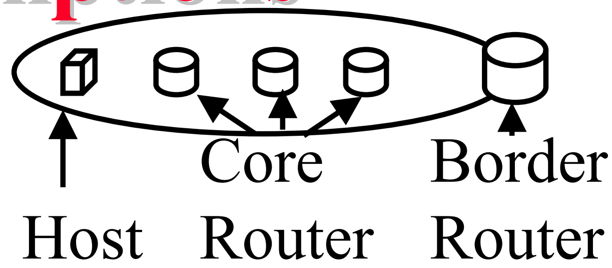

- **Hosts**:
  - Have IPv4 *local* addresses
    (Local = assigned by the organization network manager)
  - Have IPv6 128-bit global addresses
  - Have 128-bit *global* IDs (Hierarchical)
  - Support IPv6 over IPv4 tunnel
  - Have IPv6 aware higher layer protocols: TCP, UDP, FTP,...
- **Border Routers**:
  - Support all requirements of the hosts (Routers = n hosts)
  - Can establish BGP session using IPv6 global address
- **Core Routers** (non-border):
  - Have IPv4 local or IPv6 address. Understand IPv4 or IPv6.

RANGI Mechanisms

1. ID/Locator split ⇒ Mobility
2. Hierarchical ID ⇒ Administrative Scalability
3. Cryptographic ID ⇒ Security (like HIP)
4. 128-bit ID = IPv6 Addresses (like CGA)
   ⇒ Easy Application Transition
5. Local IPv4 embedded in IPv6
   ⇒ Simplify renumbering (like ISATAP)
6. IPv6 tunnel over IPv4 (ISATAP tunnel)
   ⇒ Easy transition (allow IPv4 intra-domain routers)
7. Address overwriting at border routing (Six/One or GSE)
   ⇒ Traffic engineering
8. Policy control (during ID to locator translation)

RANGI Overview

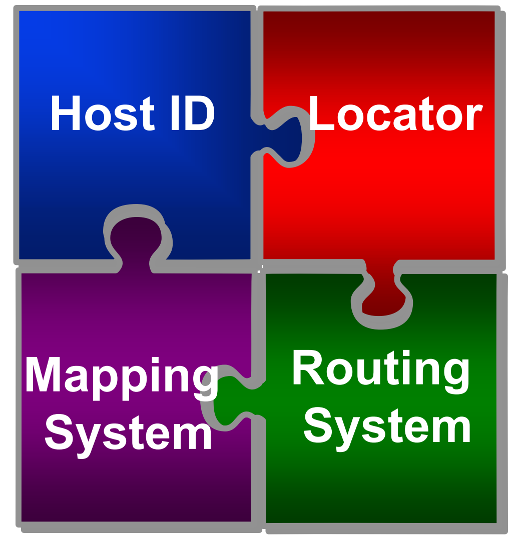

LD= Locator Domain
LDID = Locator Domain ID
LDBR = LD Border Router

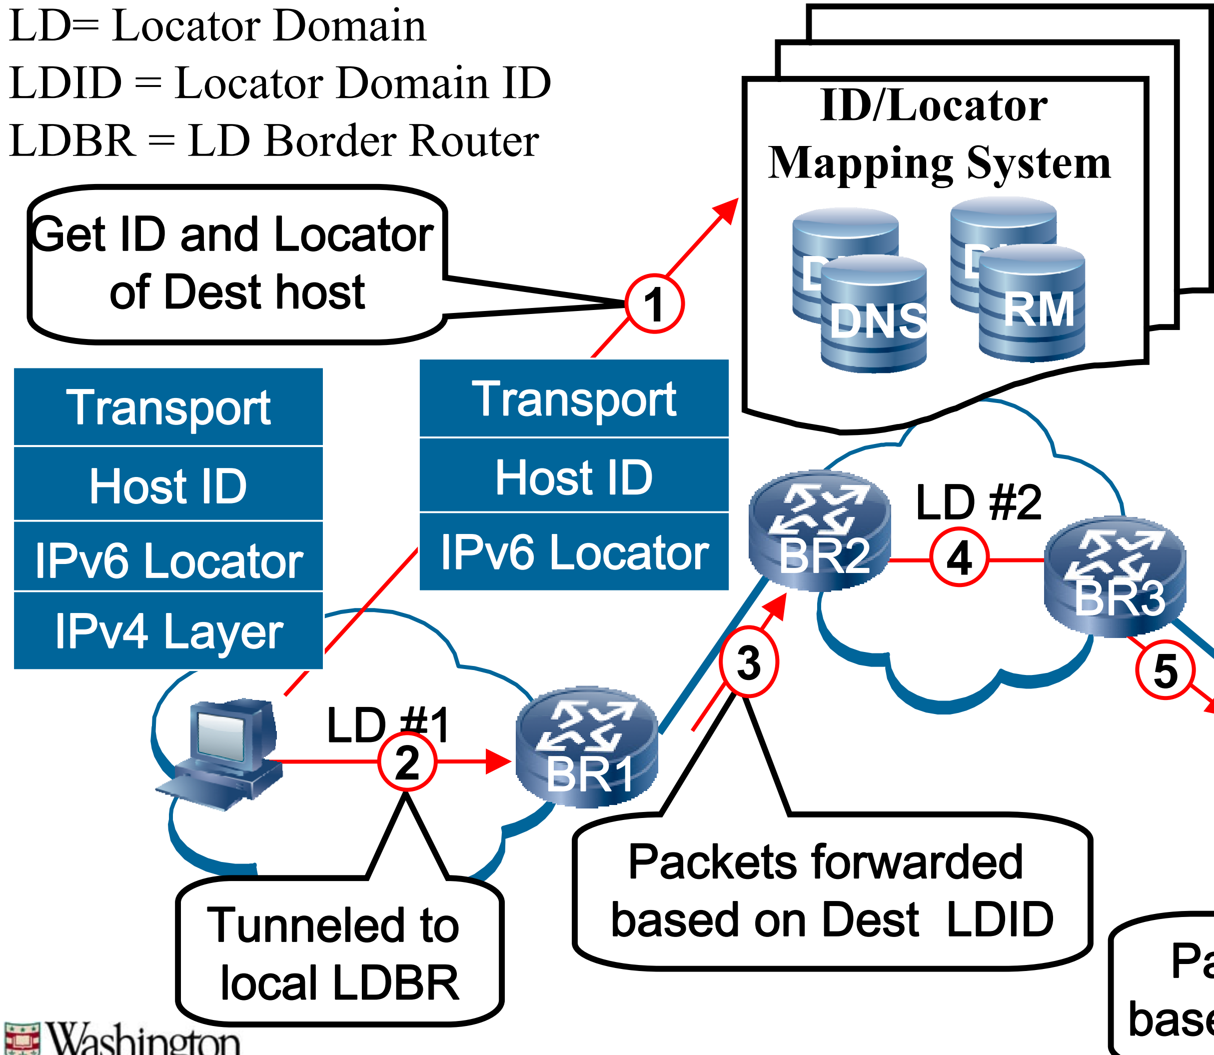

Washington University in St.Louis

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Hierarchical Host ID

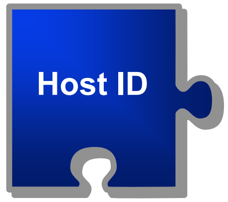

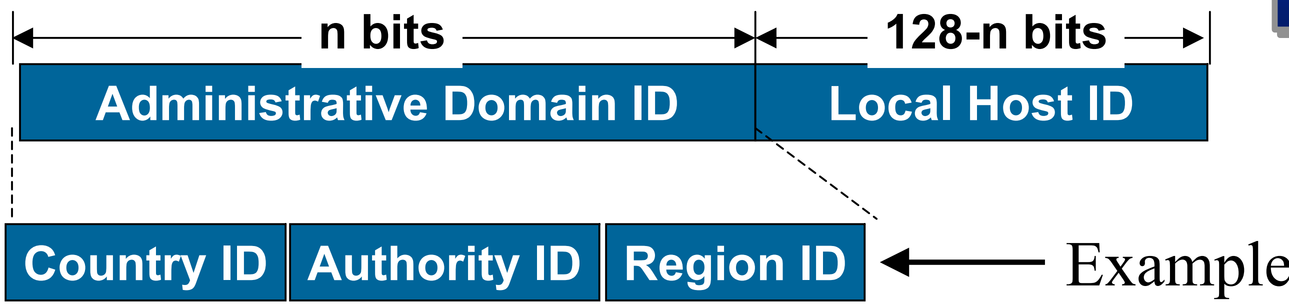

- Administrative Domain ID
  - Organizational semantics
  - Easy to deploy filtering policy based on organization boundary
- Local Host ID
  - The Hash of the public key and the AD ID

Scalability with security

Hierarchical Locator

- LD (Locator Domain) ID
  - To globally identify each LD, that is a /96 IPv6 prefix
  - Has a hierarchical structure
- LL (Local Locator) = IPv4
  - Each LD adopts independent (local) IPv4 address space
- GL (Global Locator)=LD ID + Local Locator
  - Special IPv6 address with IPv4 address embedded

Local IPv4 address ⇒ Easy renumbering

Hierarchical Routing

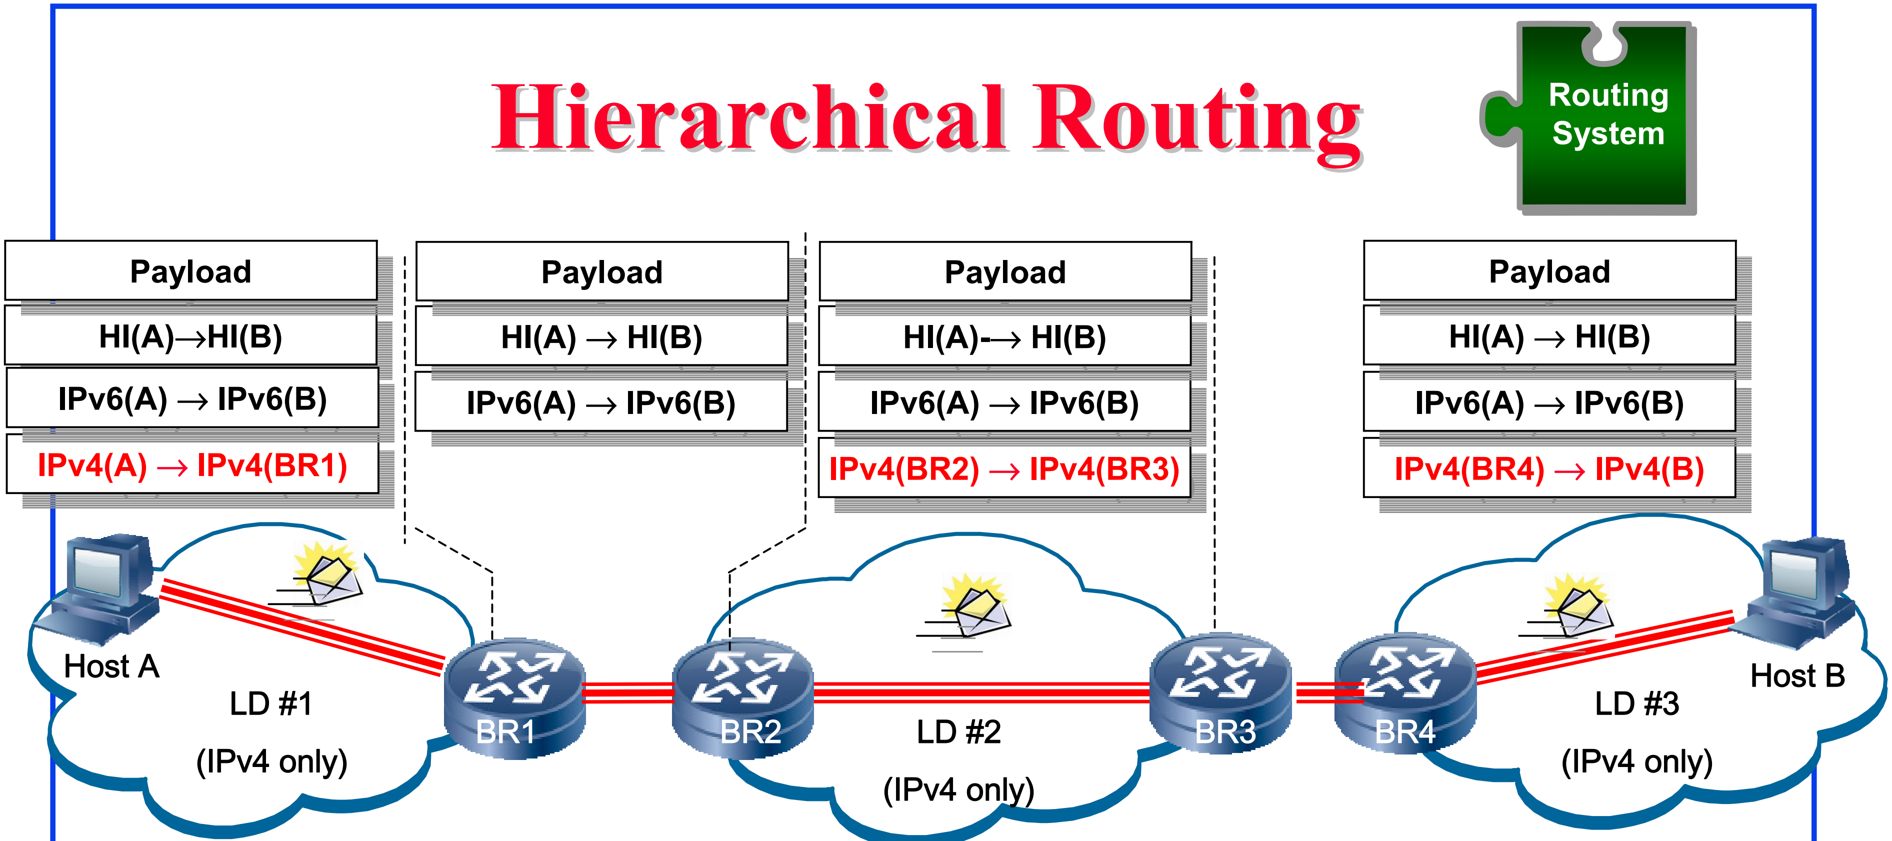

- LD ID(/96 IPv6 Prefix) based routing by LDBR
- IPv4 based routing by internal router within each LD
  - IPv6 over IPv4 tunnel between LDBRs

IPv4 Internal routers ⇒ Quick transition

Overlay View

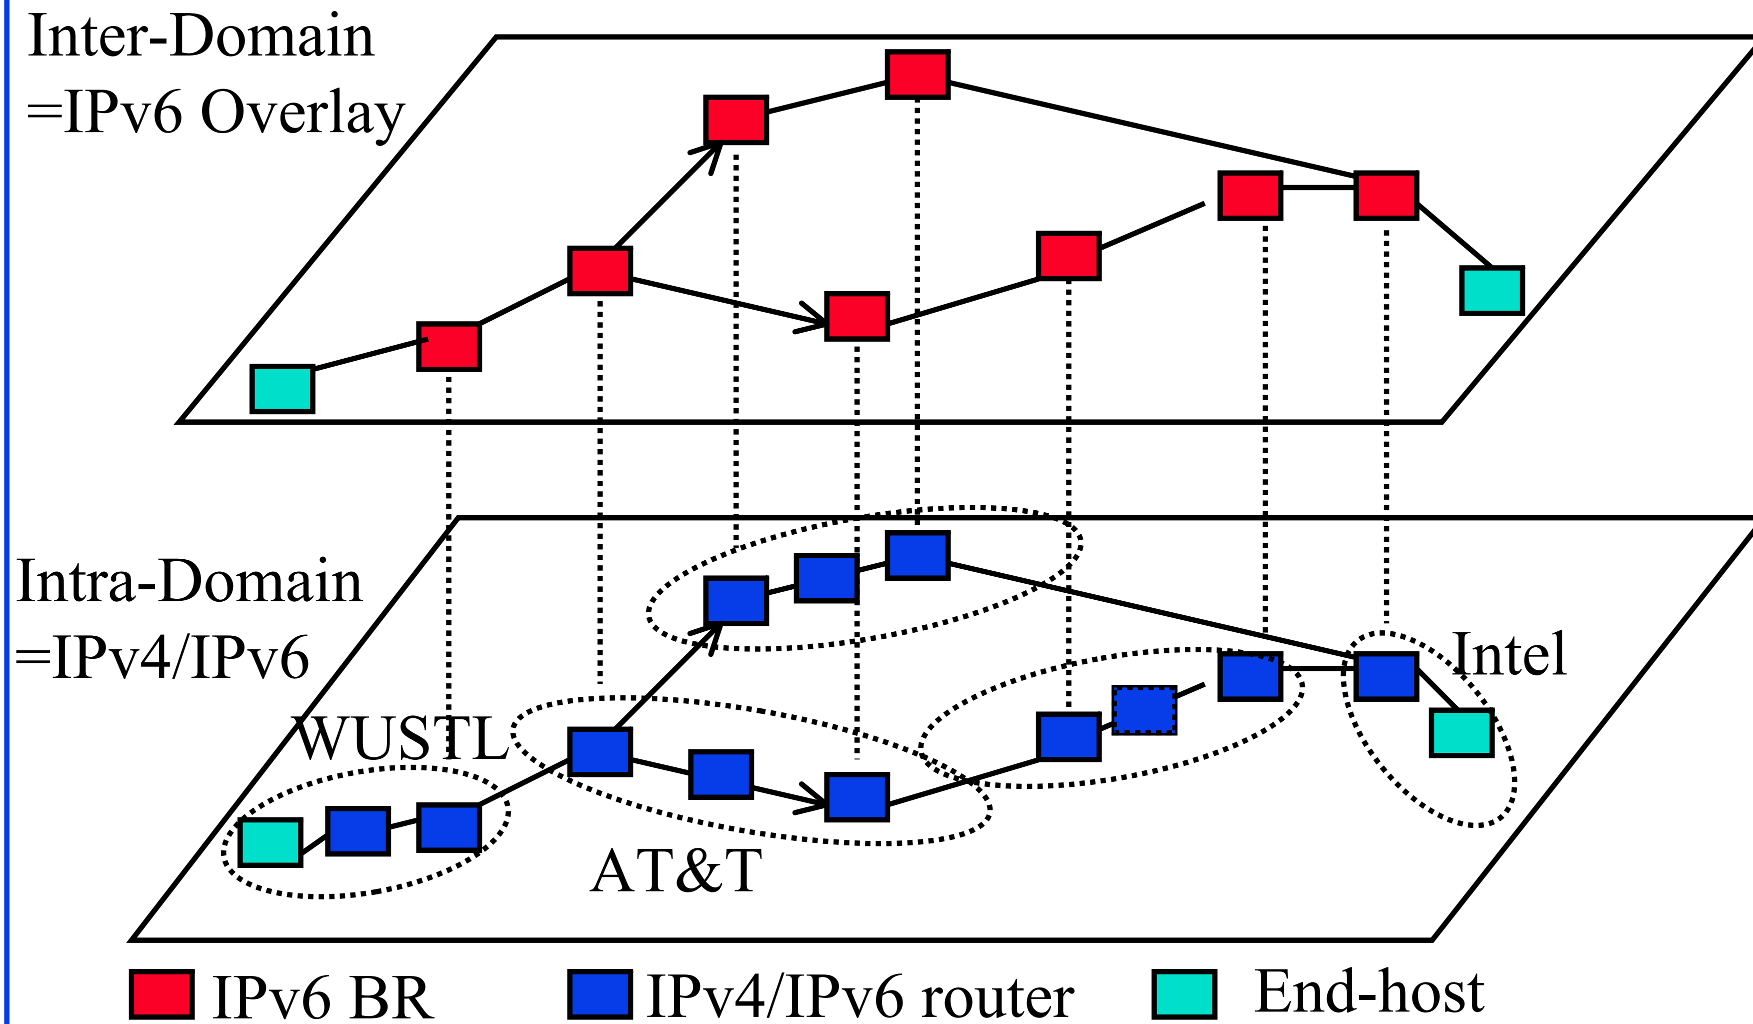

Washington University in St.Louis

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Key RANGI Features

1. Allows easy transition from IPv4 to IPv6
2. Allows site multi-homing
3. Allows site traffic engineering
4. Allows network mobility
5. And more …

Transition from IPv4 to IPv6

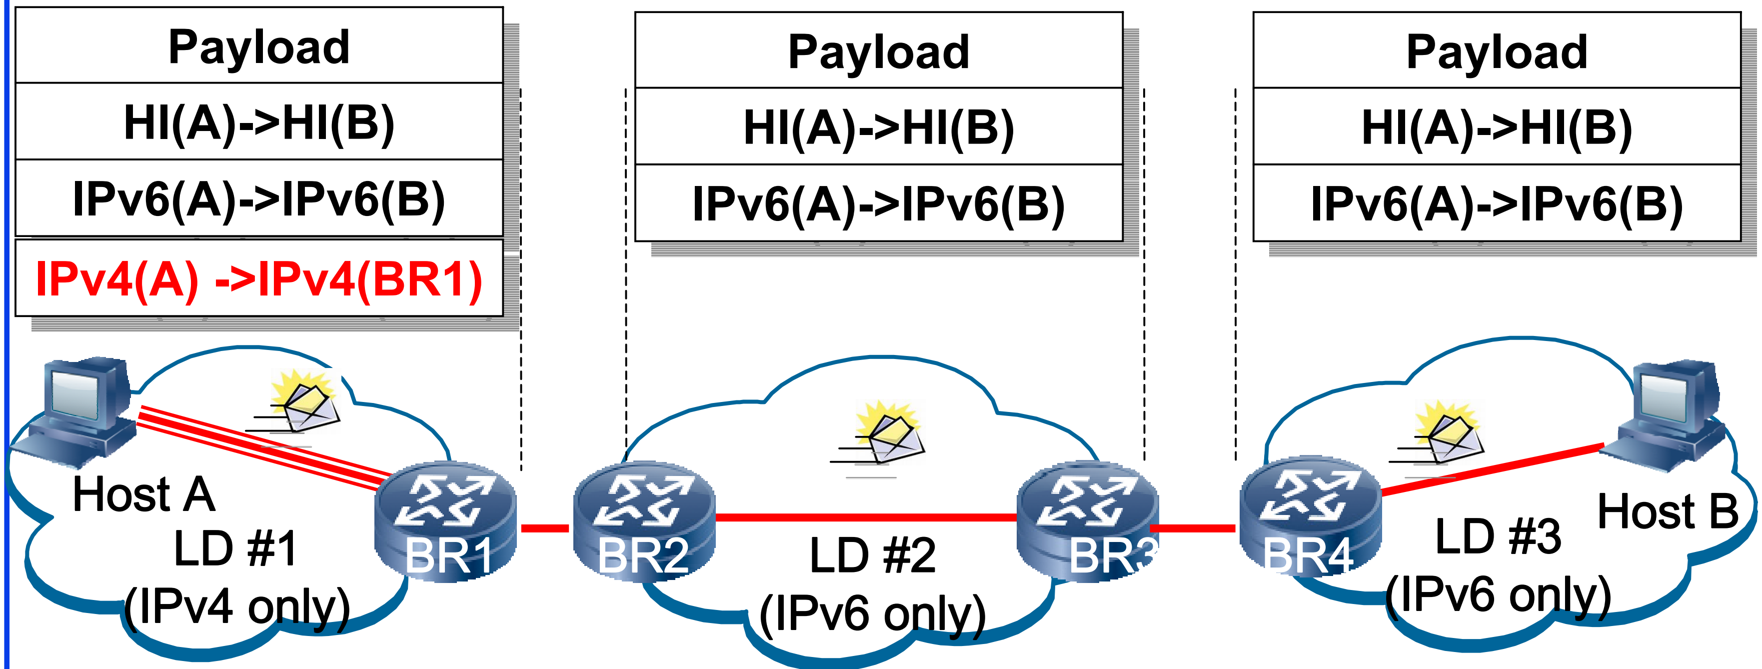

- Eliminate the IPv6 over IPv4 tunnel layer between LDBRs once the internal routers within LD are upgraded to IPv6

Smooth the transition from IPv4 to IPv6

Site Multi-homing

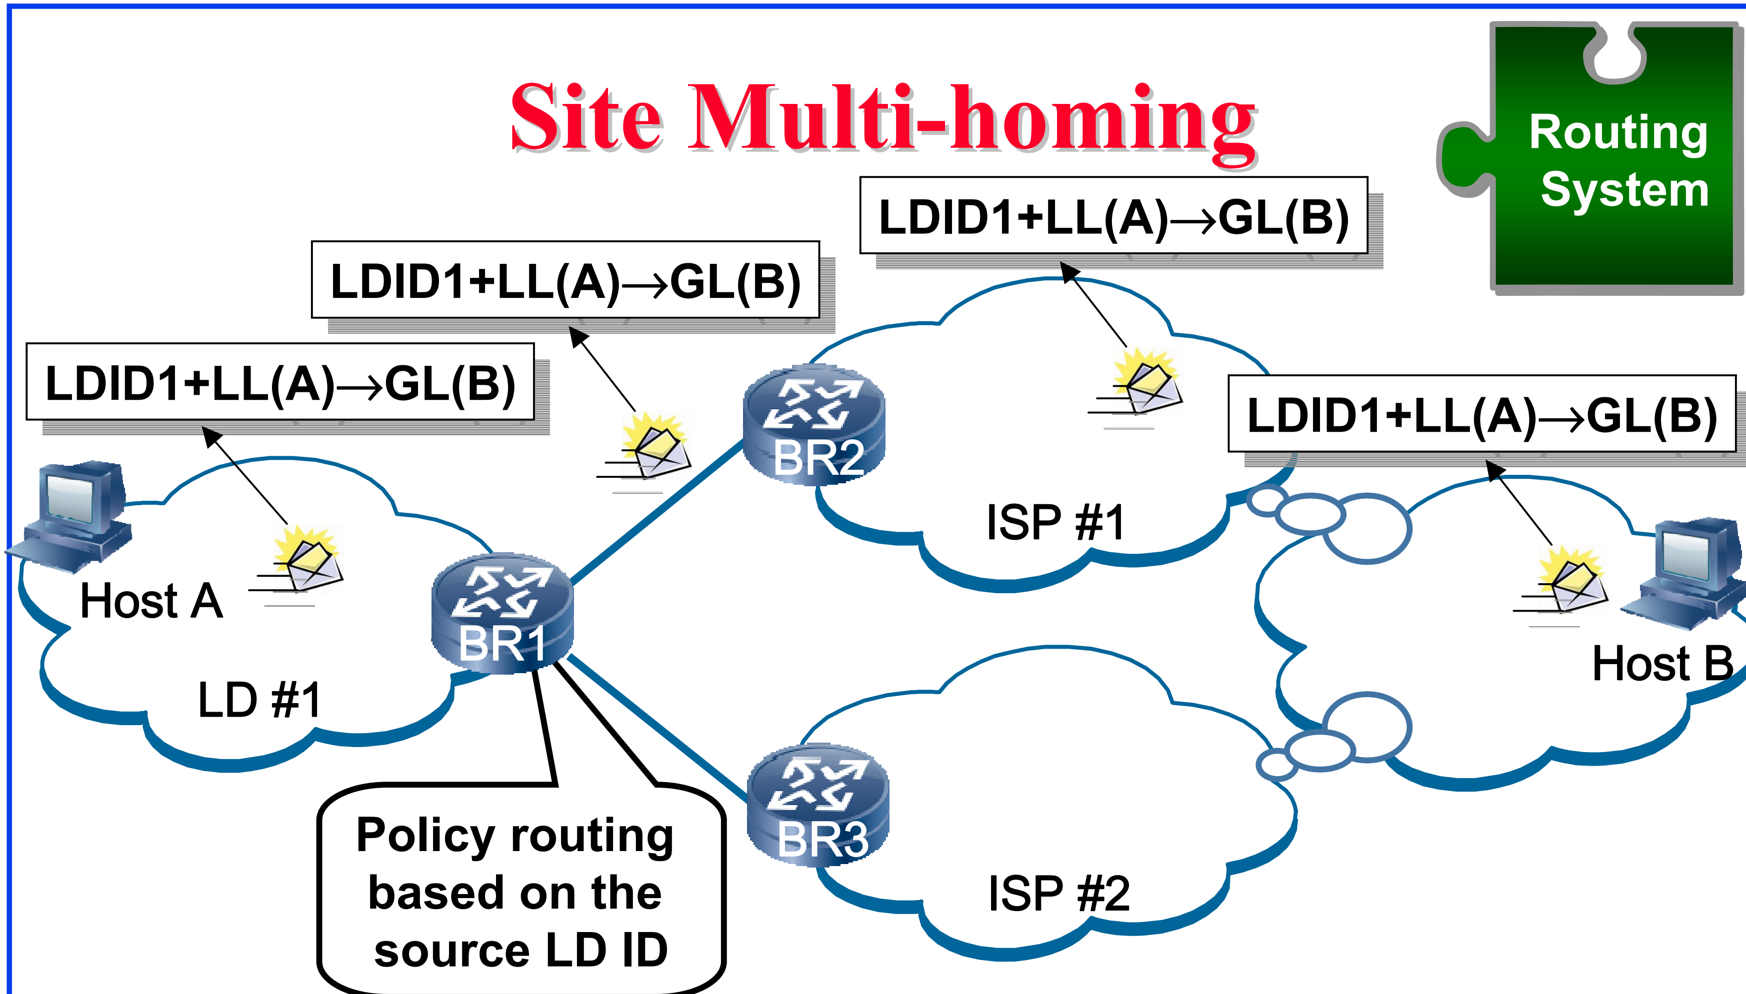

- Multiple PA LDID assigned to the multi-homed site network
  - Routing system scales well due to the usage of the PA LDID

Site Traffic Engineering

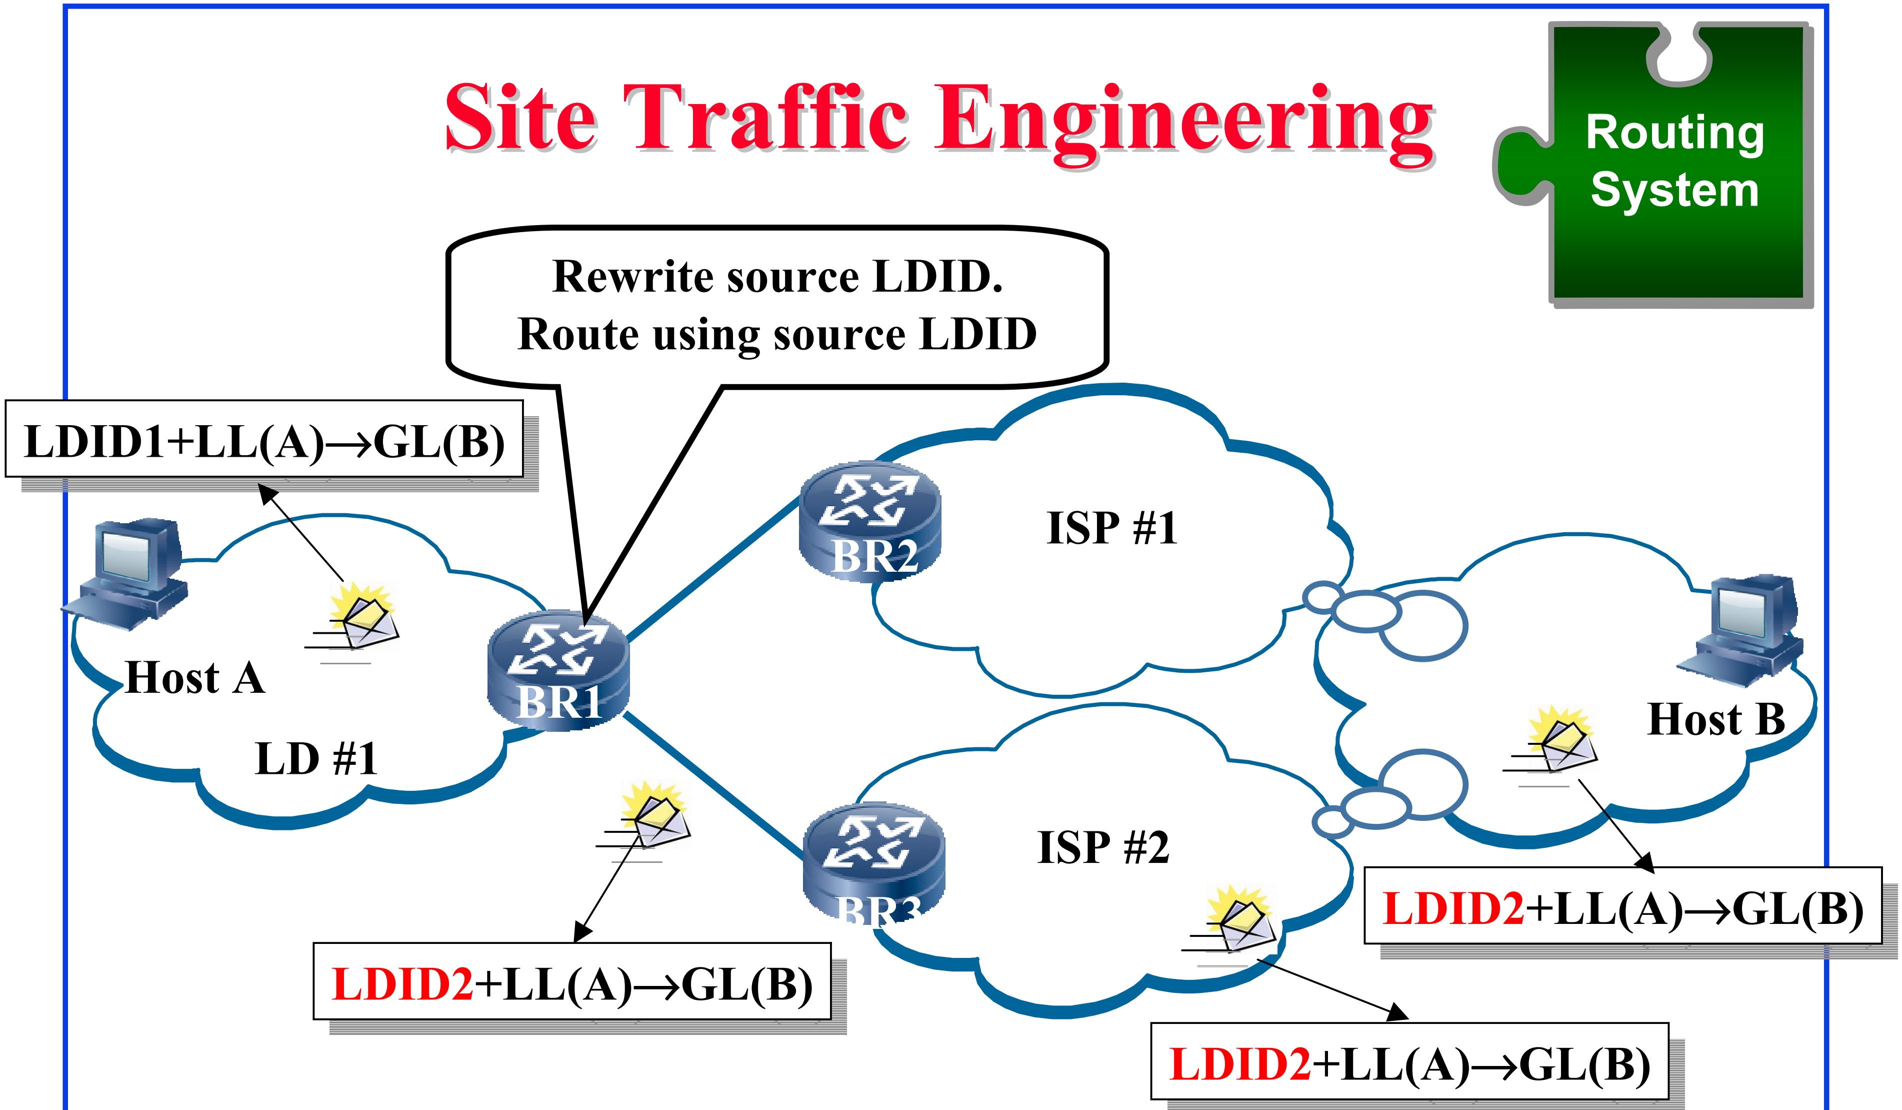

Site BR rewrites source LD of the outgoing packets

Washington University in St.Louis

http://www.cse.wustl.edu/~jain/ietf/rangi.htm

Site Traffic Engineering (Cont)

Routing System

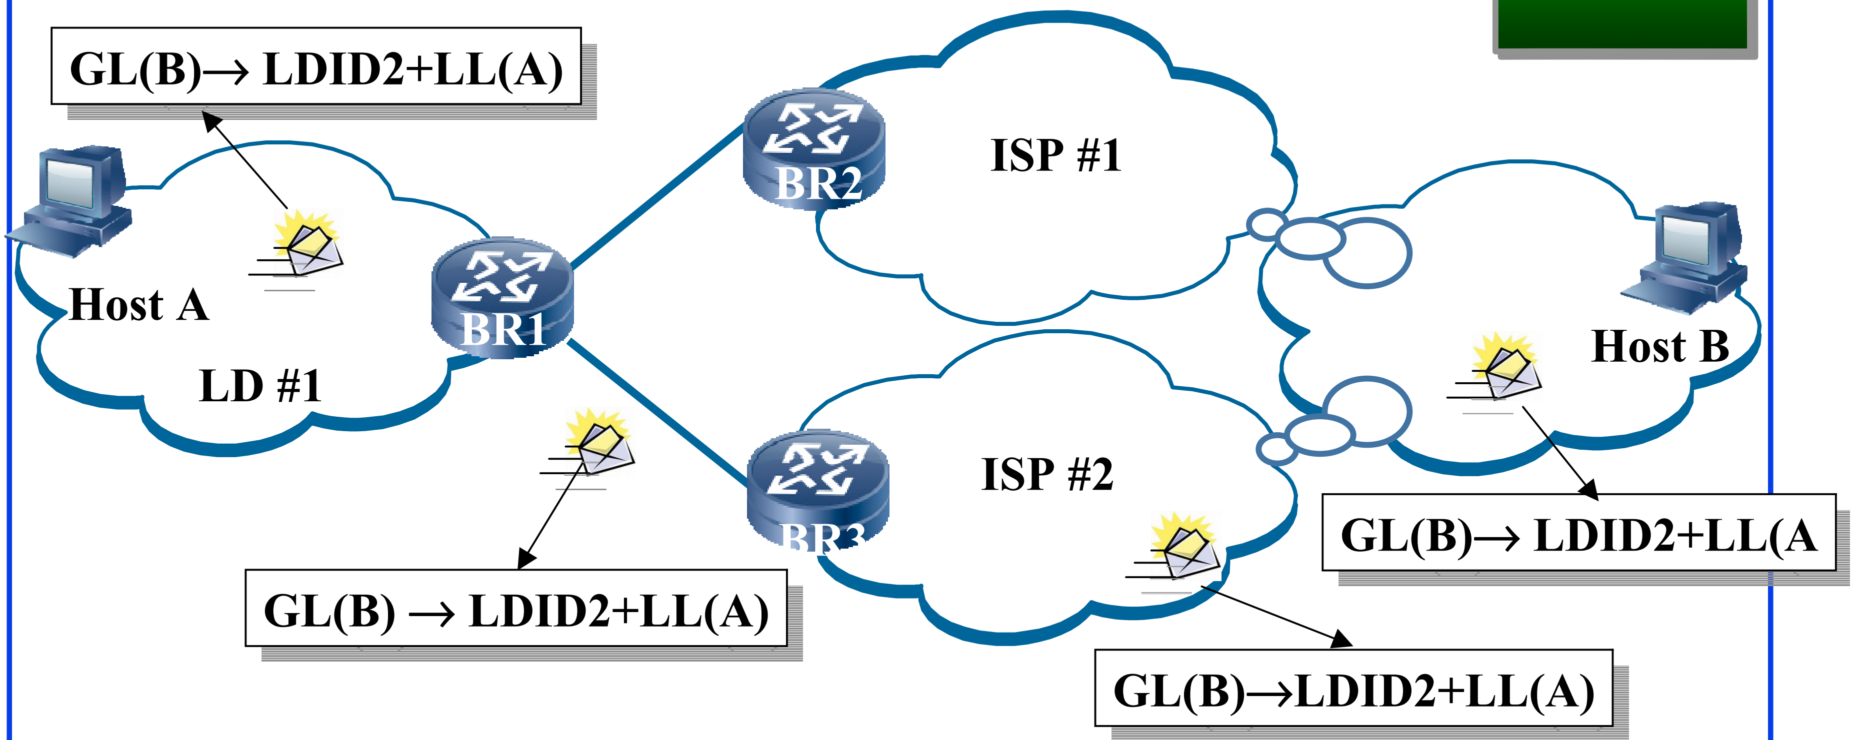

- Return packets follow the same path
- Possible to load balance also
- Idea similar to GSE, 8+8, Six/One

RANGI and RRG Design Goals

1. Routing Scalability
   - Solved by keeping separate local and global locators
   - Provider assigned locator domain ID
2. Traffic Engineering
   - Realm managers and border routers can select locator and path
3. Mobility and Multihoming
   - Identifier locator split $\Rightarrow$ Session portability
4. Simplified Renumbering
   - Local IPv4 addresses do not change
   - Global ID does not change

RANGI and RRG Design Goals (Cont)

5. Decoupling Location and Identification
6. Routing Quality
   - Allows BRs to select the paths with shorter delay or better performance
   - Size of global routing table and update frequency reduced significantly
7. Routing Security
   - RM enforce policies including security
   - Local addresses and paths are not disclosed outside
8. Incremental Deployability
   - Allow step by step deployment and long-term evolution

Summary

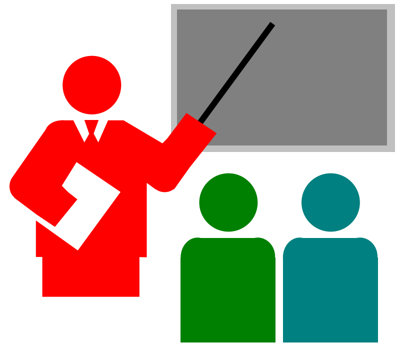

1. RANGI
   = Routing Architecture for the next generation Internet
   Solves scalability, mobility, multihoming, …, policy
2. RANGI-awareness required only in the hosts and in the border routers
3. Non-border routers can remain IPv4 or IPv6
4. Organizations have complete control over naming, addressing inside their organization (Local addressing) and resolution
5. Incremental deployment of RANGI and IPv6

Future Work

- Incremental deployment of RANGI border routers: Some clouds may not have RANGI border routers.
- Incremental deployment of RANGI in the domain Some hosts may and some may not have RANGI
- Policy enforcement of end-to-end trust
- Policy enforcement of path

References

1. Jain, R., "Internet 3.0: Ten Problems with Current Internet Architecture and Solutions for the Next Generation," in Proceedings of Military Communications Conference (MILCOM 2006), Washington, DC, October 23-25, 2006, http://www.cse.wustl.edu/~jain/papers/gina.htm
2. Subharthi Paul, Raj Jain, Jianli Pan, and Mic Bowman, "A Vision of the Next Generation Internet: A Policy Oriented View," British Computer Society Conference on Visions of Computer Science, Sep 2008, http://www.cse.wustl.edu/~jain/papers/pona.htm
3. Jianli Pan, Subharthi Paul, Raj Jain, and Mic Bowman, "MILSA: A Mobility and Multihoming Supporting Identifier-Locator Split Architecture for Naming in the Next Generation Internet,," Globecom 2008, Nov 2008, http://www.cse.wustl.edu/~jain/papers/milsa.htm

References (Cont)

4. Xiaohu Xu and Dayong Guo, "Hierarchical Routing Architecture," Proc. 4th Euro-NGI Conference on Next Generation Internetworks, Krakow, Poland, 28-30 April 2008, 7 pp., http://www.cse.wustl.edu/~jain/papers/hra.htm