Host Identity Protocol

Prof. Sasu Tarkoma

23.02.2009

Part of the material is based on lecture slides by Dr. Pekka Nikander (HIP) and Dmitrij Lagutin (PLA)

Contents

- Introduction
- Current state
- Host Identity Protocol (HIP)
- Overlays (i3 and Hi3)
- Summary

Introduction

- Current Internet is increasingly data and content centric
- The protocol stack may not offer best support for this
- End-to-end principle is no longer followed
  - Firewalls and NAT boxes
  - Peer-to-peer and intermediaries
- Ultimately, hosts are interested in receiving valid and relevant information and do not care about IP addresses or host names
- This motivate the design and development of new data and content centric networking architectures
  - Related work includes ROFL, DONA, TRIAD, FARA, AIP, ..

The Internet has Changed

- A lot of the assumptions of the early Internet has changed
  - Trusted end-points
  - Stationary, publicly addressable addresses
  - End-to-End
- We will have a look at these in the light of recent developments
- End-to-end broken by NATs and firewalls

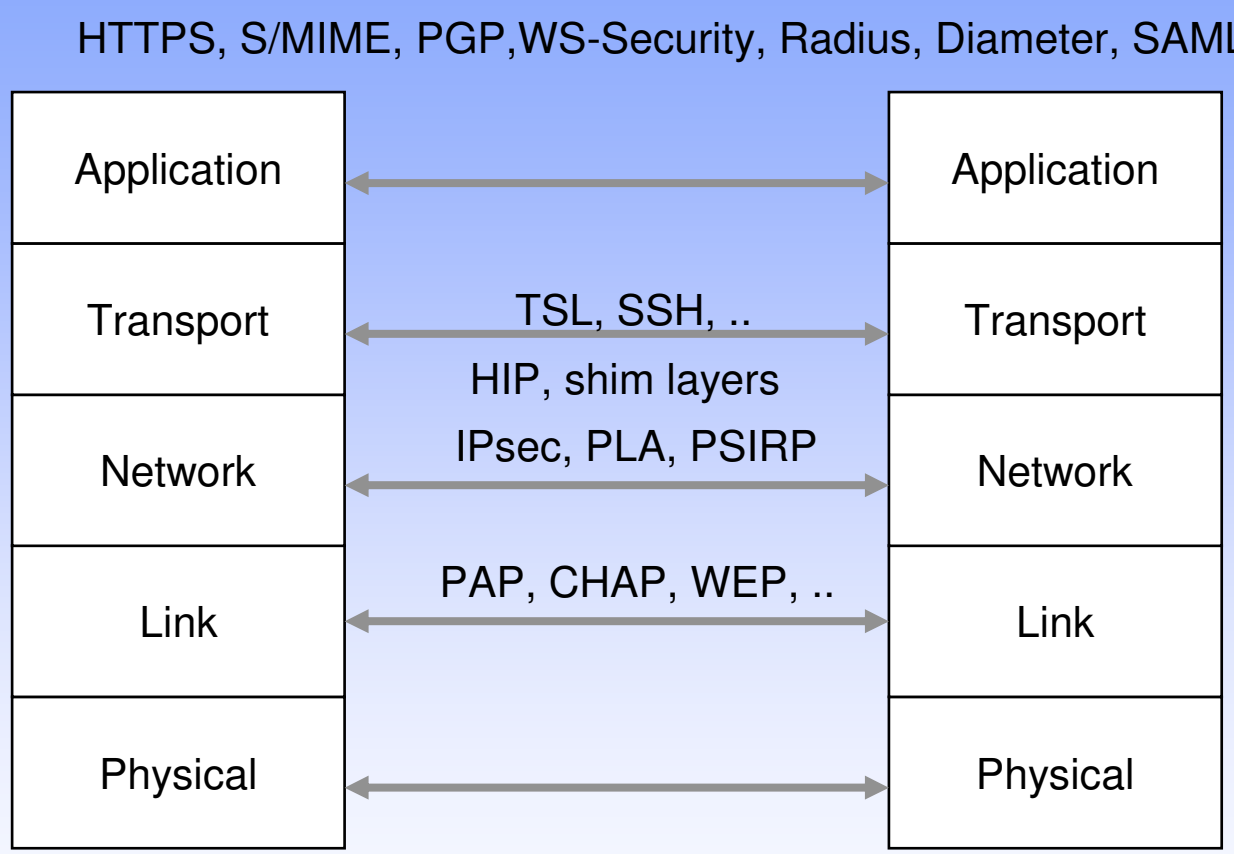
HTTPS, S/MIME, PGP,WS-Security, Radius, Diameter, SAML 2.0 ..
Application
Application
Transport
TSL, SSH, ..
Transport
HIP, shim layers
Network
IPsec, PLA, PSIRP
Network
Link
PAP, CHAP, WEP, ..
Link
Physical
Physical

Current State

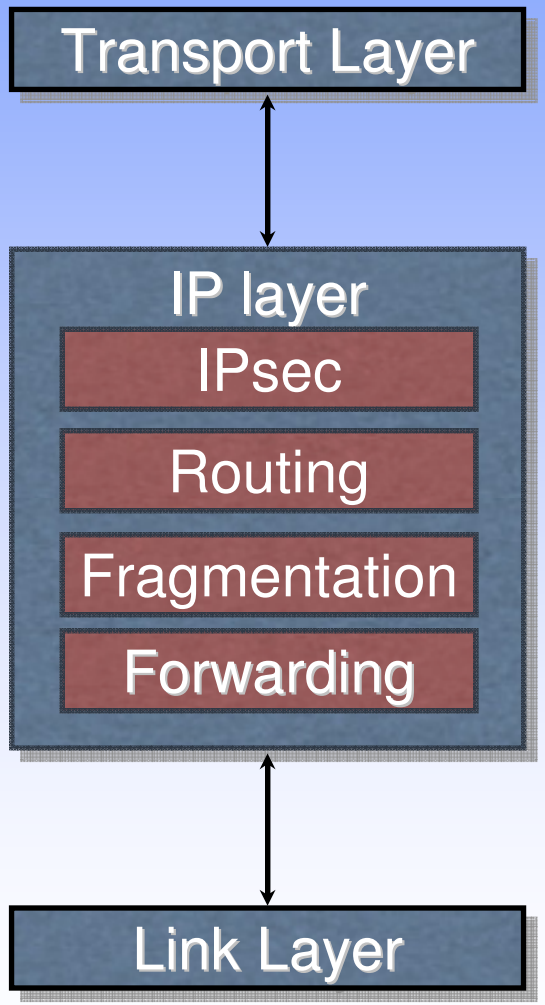

Observations

End-to-end reachability is broken

Unwanted traffic is a problem

Mobility and multi-homing are challenging

Multicast is difficult (does not scale)

Security is difficult

Not optimal fit for broadcast and all-optical networking

HIP

What is HIP?

- HIP = Host Identity Protocol
- A proposal to separate identifier from locator at the network layer of the TCP/IP stack
  - A new name space of public keys
  - A protocol for discovering and authenticating bindings between public keys and IP addresses
- Secured using signatures and keyed hashes (hash in combination with a secret key)

Motivation

- Not to standardise a solution to a problem
  - No explicit problem statement
- Exploring the consequences of the id / loc split
  - Try it out in real life, in the live Internet
- A different look at many problems
  - Mobility, multi-homing, end-to-end security, signalling, control/data plane separation, rendezvous, NAT traversal, firewall security, ...

HIP in a Nutshell

- Architectural change to TCP/IP structure
- Integrates security, mobility, and multi-homing
  - Opportunistic host-to-host IPsec ESP
  - End-host mobility, across IPv4 and IPv6
  - End-host multi-address multi-homing, IPv4/v6
  - IPv4 / v6 interoperability for apps
- A new layer between IP and transport
  - Introduces cryptographic Host Identifiers

The Idea

- A new Name Space of Host Identifiers (HI)
  - Public crypto keys!
  - Presented as 128-bit long hash values, Host ID Tags (HIT)
- Sockets bound to HIs, not to IP addresses
- HIs translated to IP addresses in the kernel

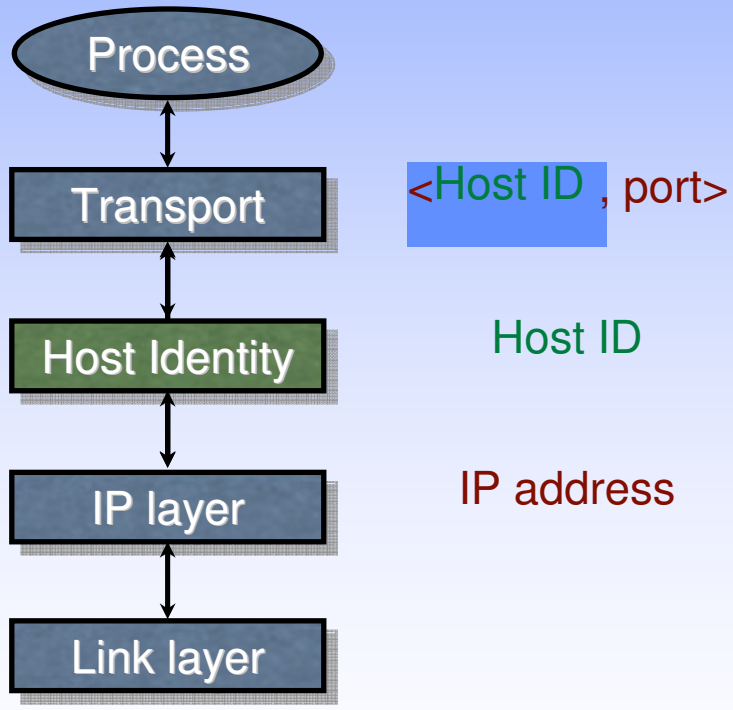

Protocol overview

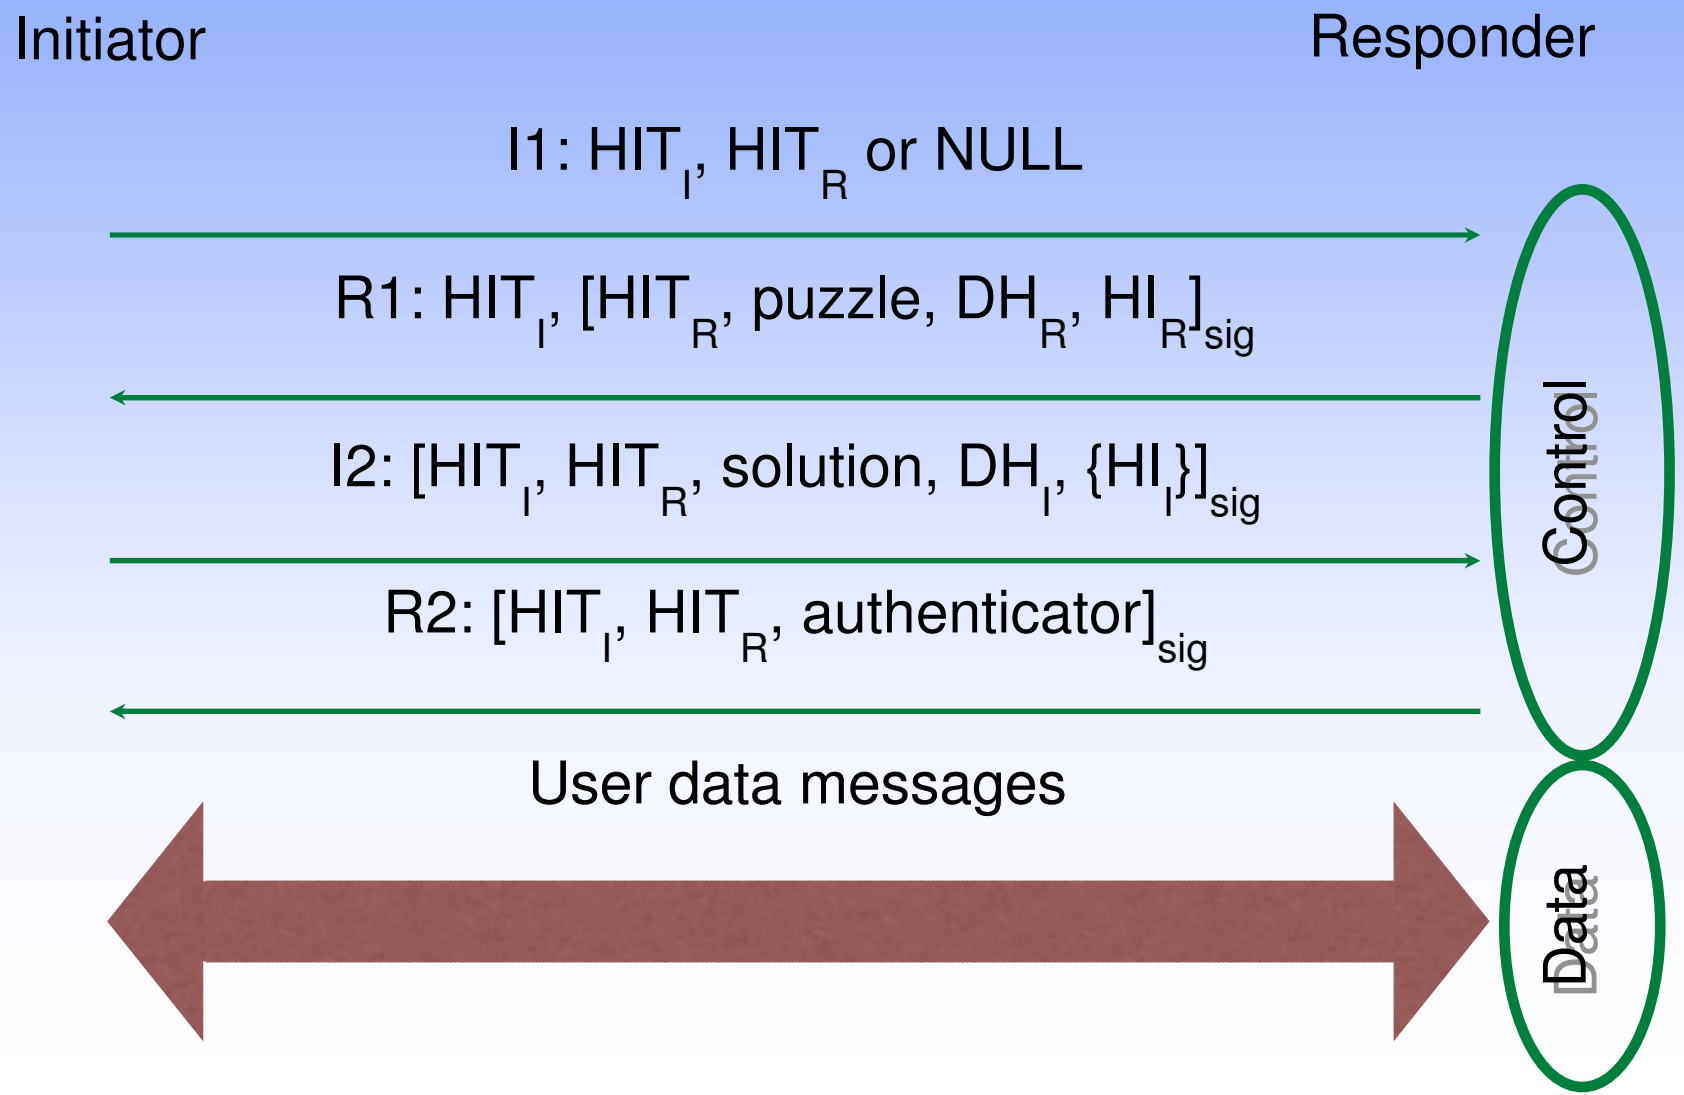

Base exchange

- Based on SIGMA family of key exchange protocols

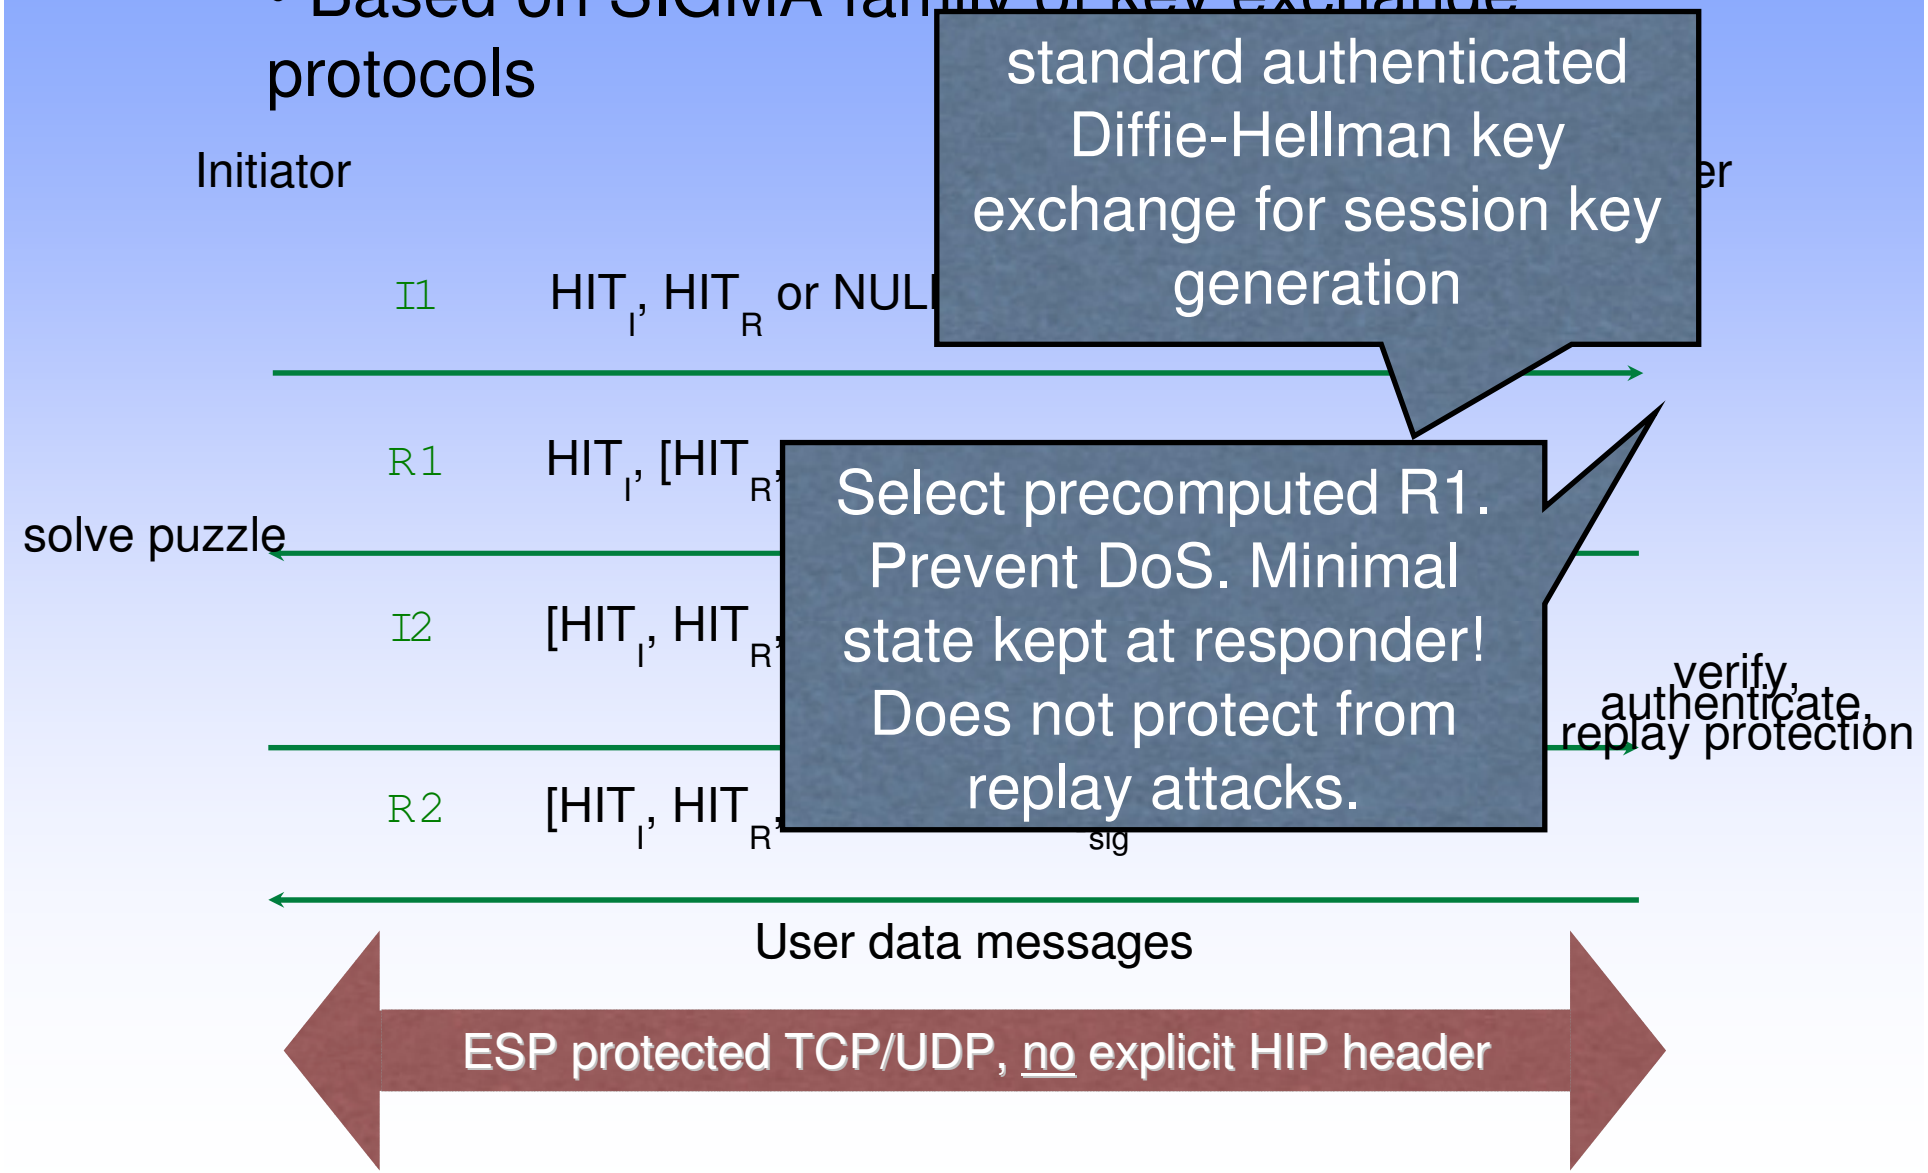

Other Core Components

- Per-packet identity context
  - Indirectly, through SPI if ESP is used
  - Directly, e.g., through an explicit shim header
- A mechanism for resolving identities to addresses
  - DNS-based, if FQDNs used by applications
  - Or distributed hash tables (DHTs) based

Using HIP with ESP

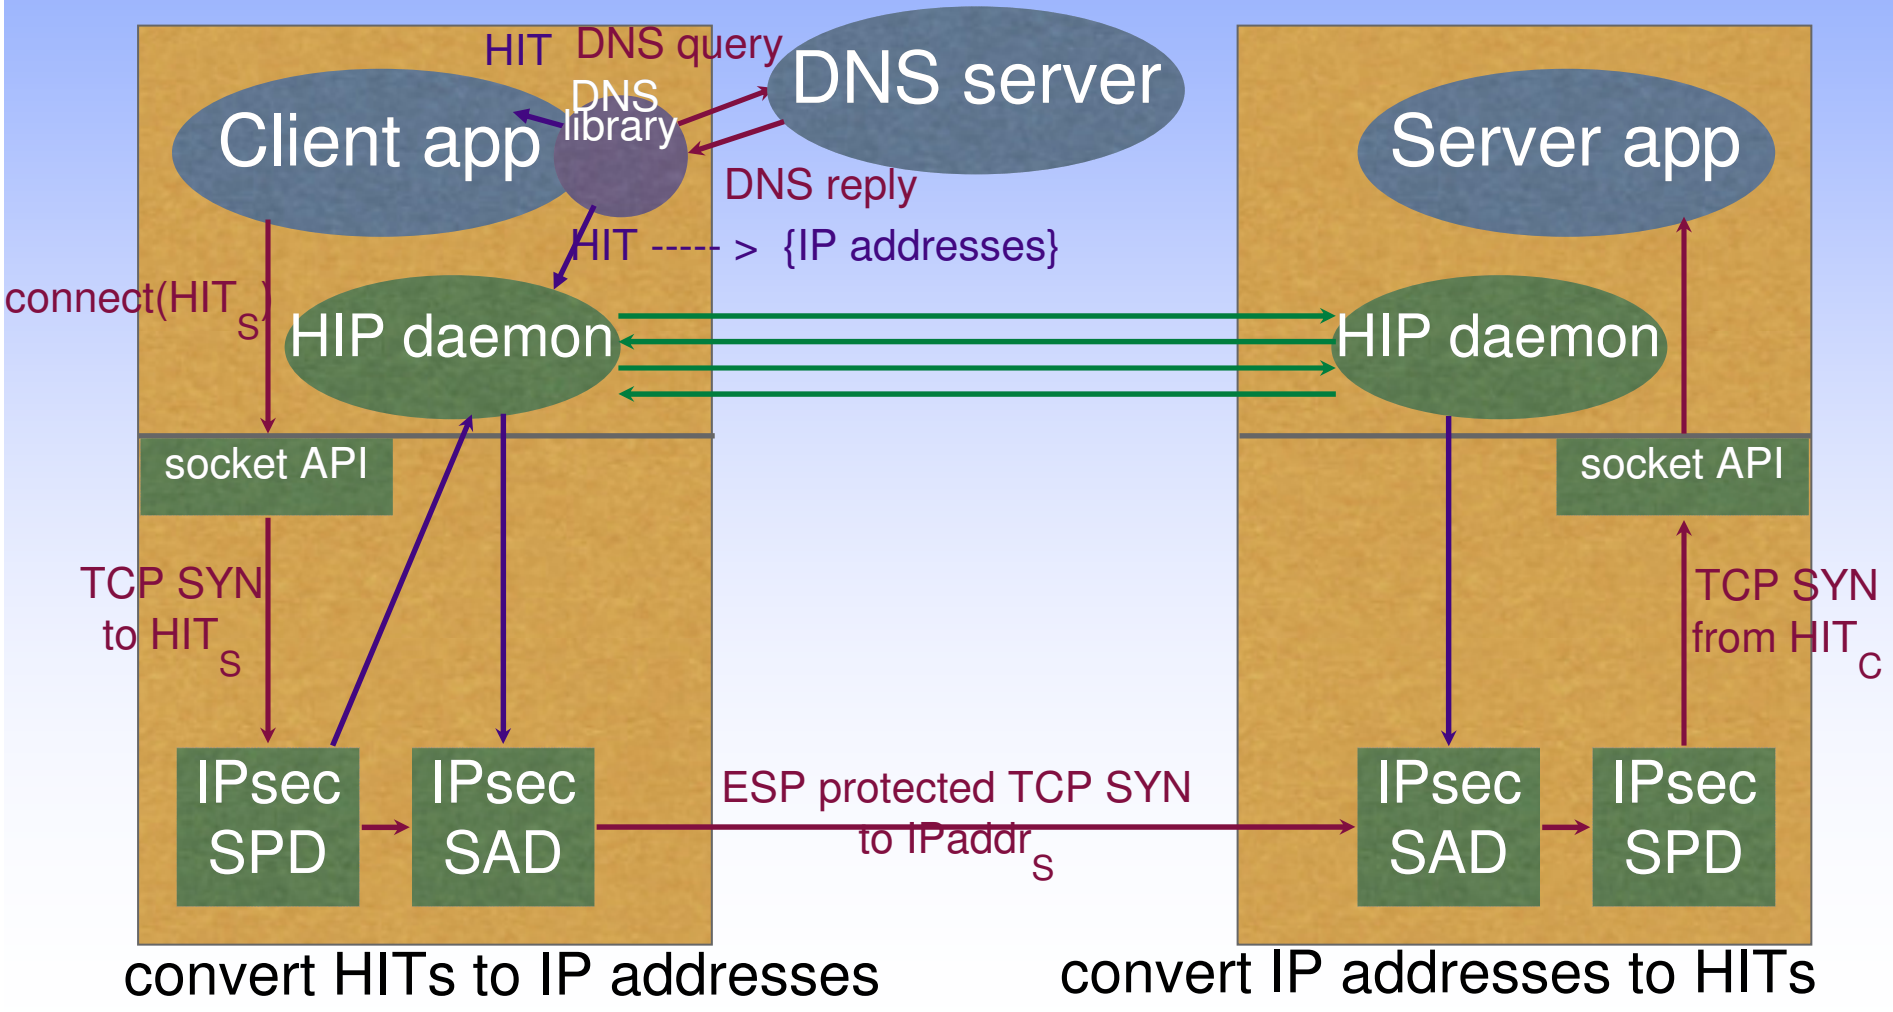

Many Faces

- More established views:
  - A different IKE for simplified end-to-end ESP
  - Super Mobile IP with v4/v6 interoperability and dynamic home agents
  - A host multi-homing solution
- Newer views:
  - New waist of IP stack; universal connectivity
  - Secure carrier for signalling protocols

HIP as the new waist of TCP/IP

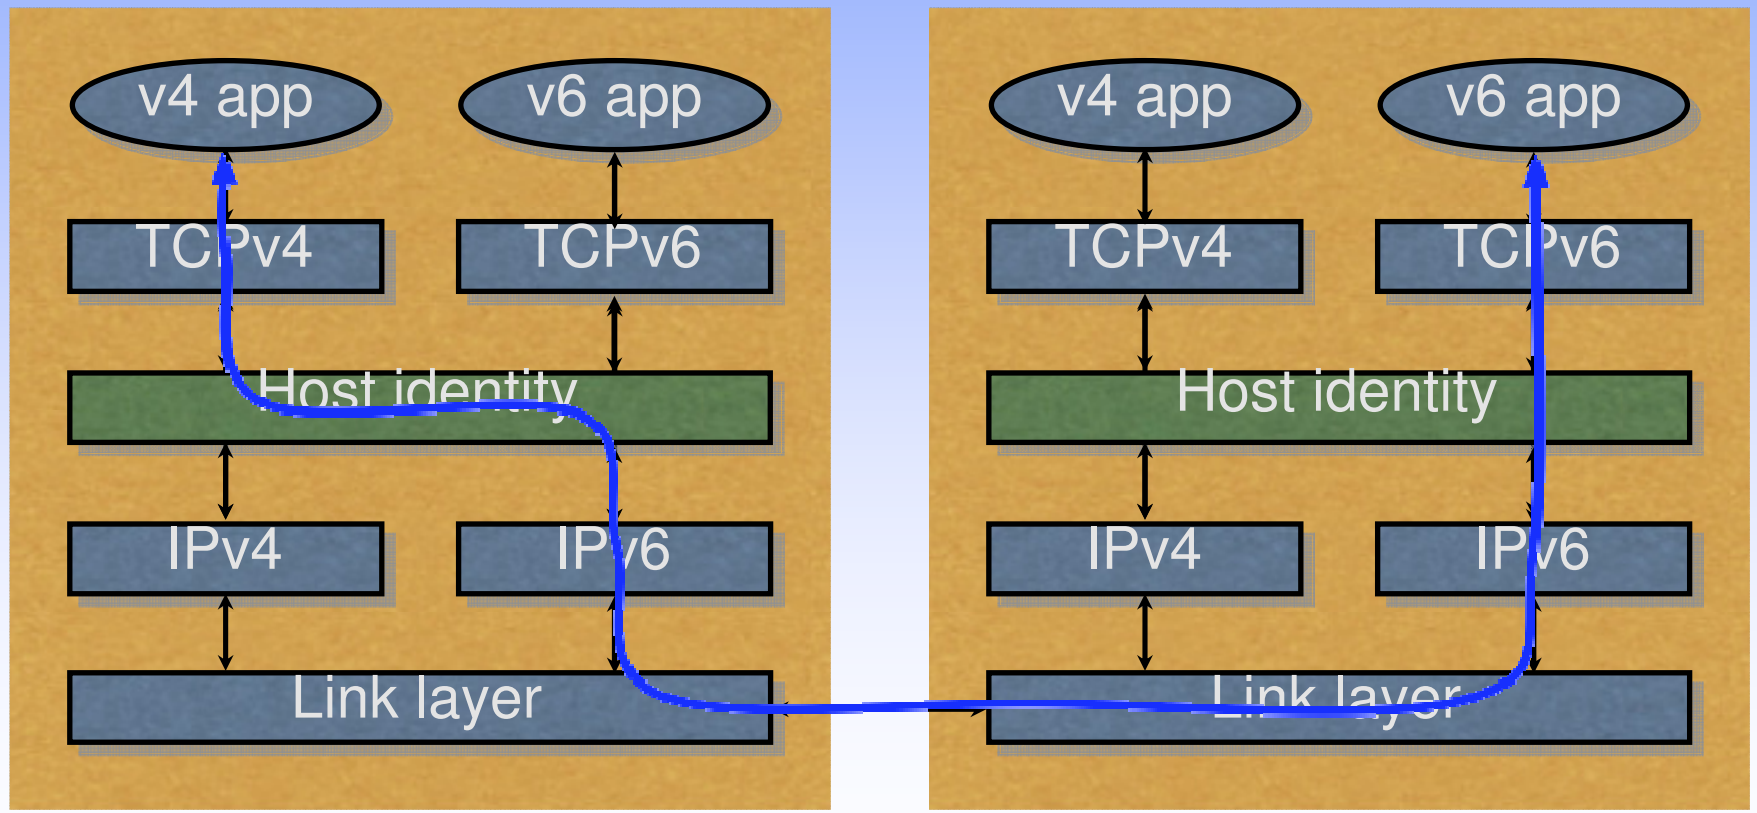

Rendezvous

- Initial rendezvous
  - How to find a moving end-point?
  - Can be based on directories
  - Requires fast directory updates
    - Bad match for DNS
- Tackling double-jump
  - What if both hosts move at same time?
  - Requires rendezvous point

Mobile IP

- Home Agent (HA)
  - Serves a Home Address
  - Initial reachability
  - Triangular routing
- Route optimization
  - Tunnels to bypass HA
  - HA as rendezvous point

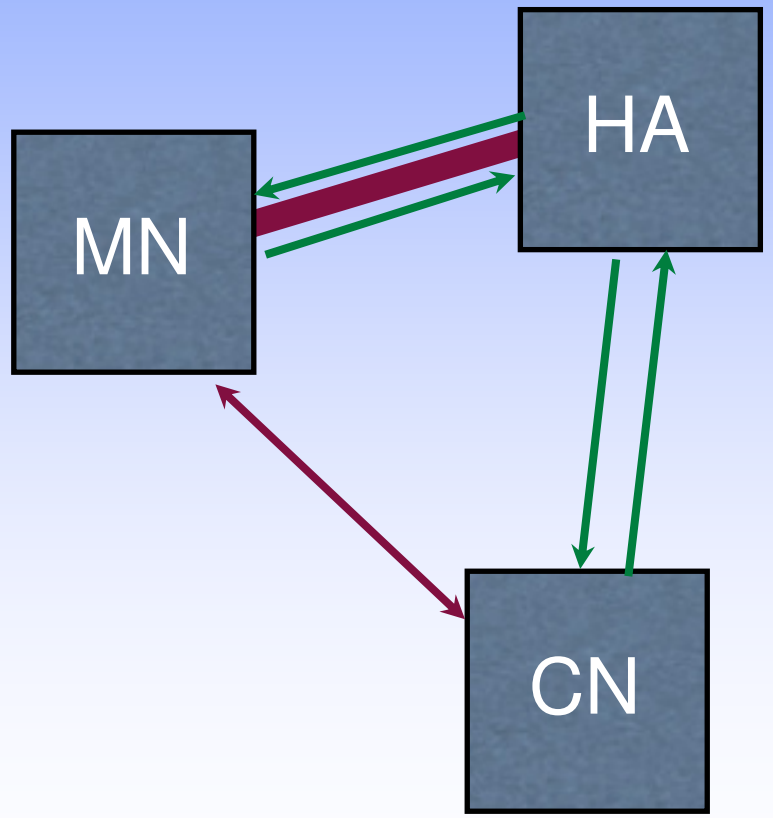

Mobility protocol

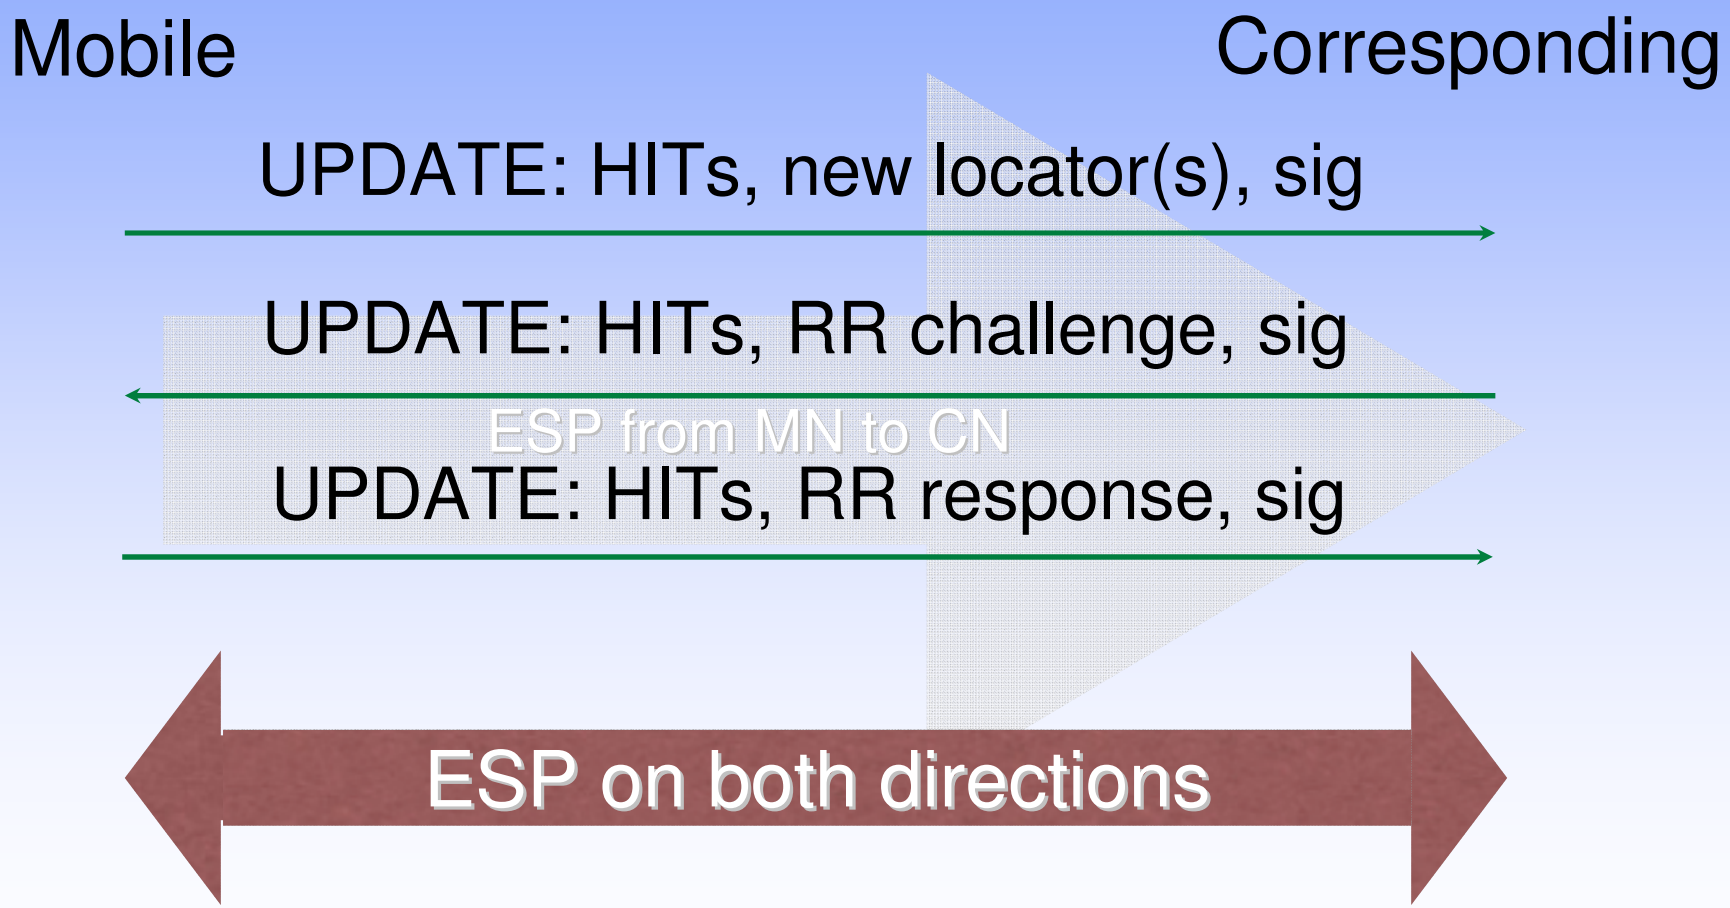

Key distribution for HIP

- Depends on application
- For multi-addressing, self-generated keys
- Usually keys in the DNS
- Can use PKI if needed
- Opportunistic mode supported
  - SSH-like leap-of-faith
  - Accept a new key if it matches a fingerprint

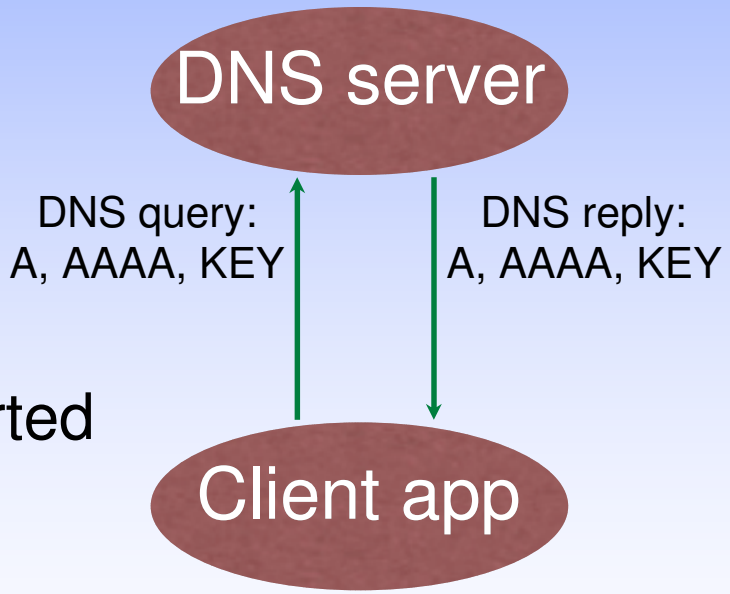

Basic HIP rendezvous

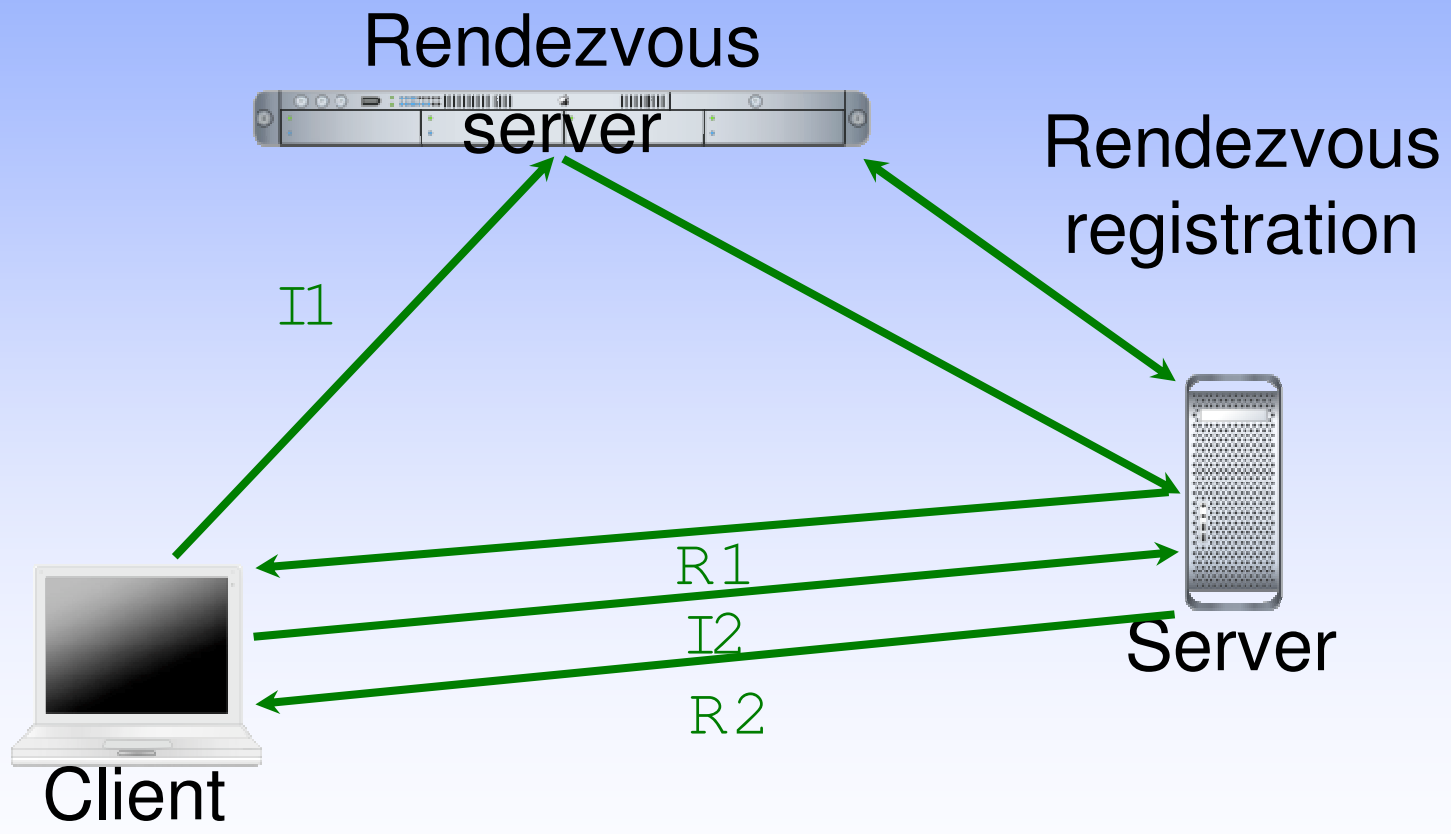

The infrastructure question

- HIs originally planned to be stored in the DNS
  - Retrieved simultaneously with IP addresses
  - Does not work if you have only a HIT
- Question: How to get data based on HIT only?
  - HITs look like 128-bit random numbers
- Possible answer: DHT based overlay like i^3

Distributed Hash Tables

- Distributed directory for flat data
- Several different ways to implement
- Each server maintains a partial map
- Overlay addresses to direct to the right server
- Resilience through parallel, unrelated mappings
- Used to create overlay networks

i^3 rendezvous abstraction

- Trigger inserted by receiver(s)
- Packets addressed to identifiers
- i^3 routes packet to the receiver(s)

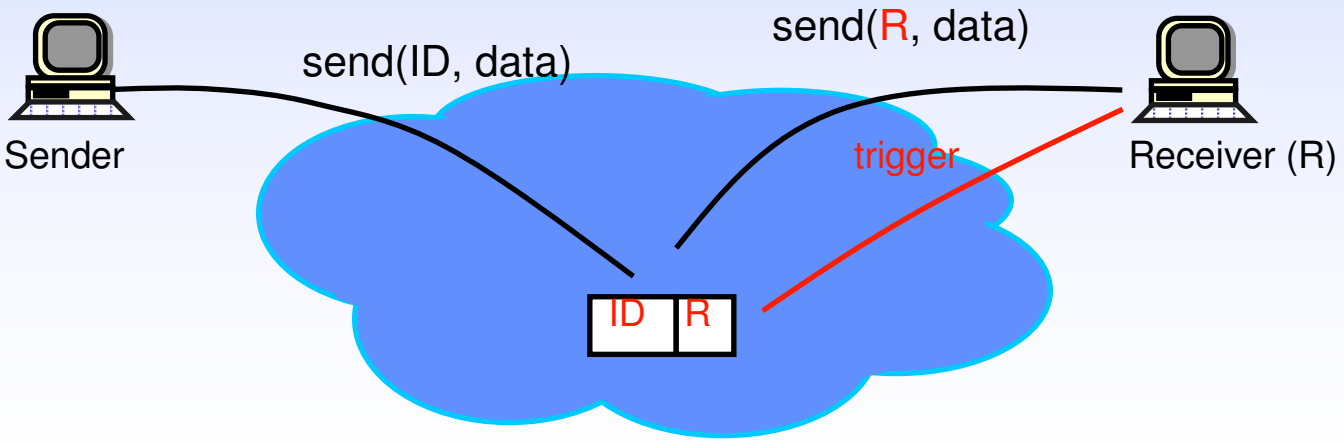

Hi^3: combining HIP and i3

- Developed at Ericsson Research IP Networks
- Uses i^3 overlay for HIP *control* packets
  - Provides rendezvous for HIP
- *Data* packets use plain old IP
  - Cryptographically protected with ESP
- Only soft or optional state in the network

Hi^3 and DHT-based rendezvous

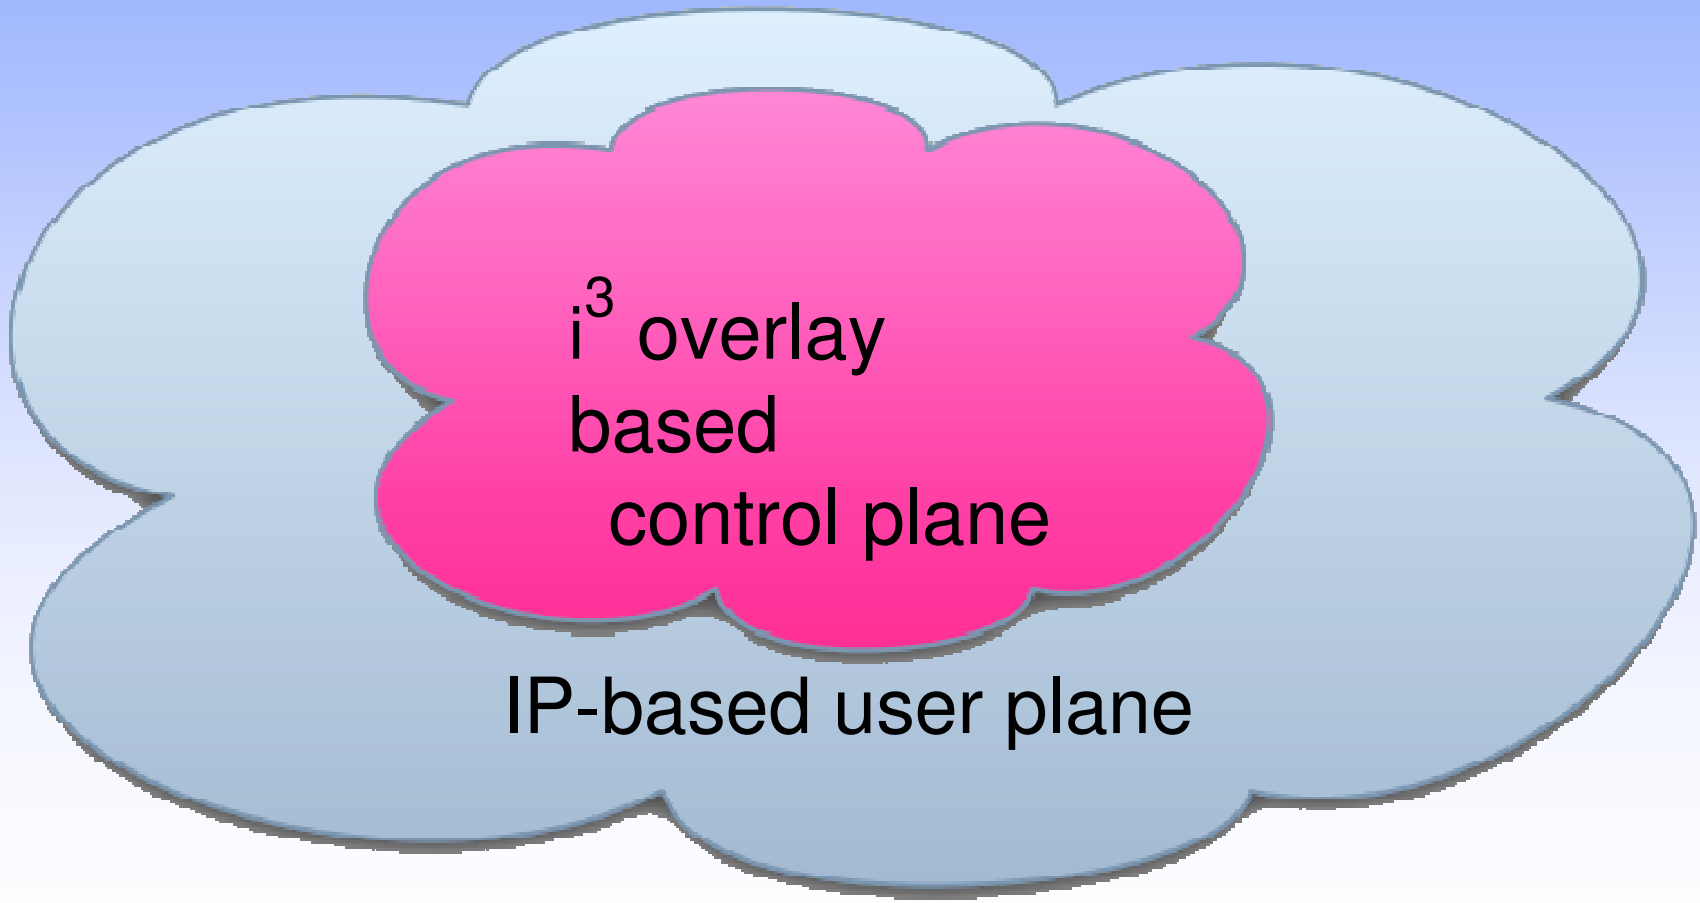

Control/data separation

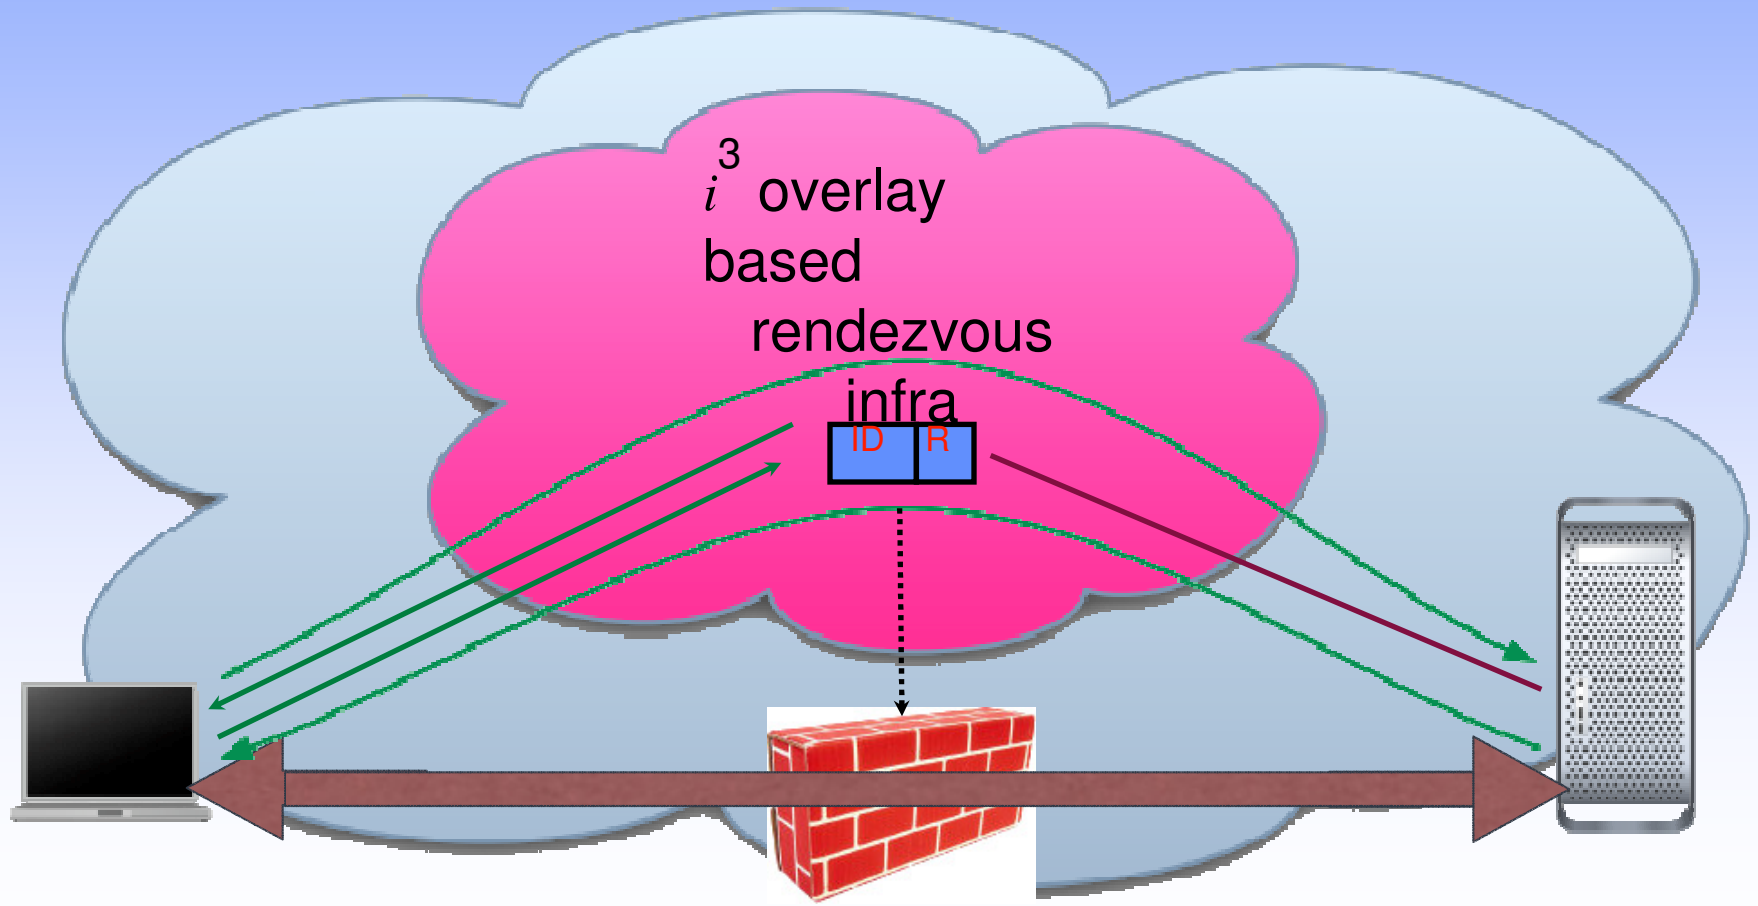

An Internet control plane?

- HIP separates control and data traffic
- Hi^3 routes control traffic through overlay
  - Control and data packets take potentially very different paths
- *Allows* telecom-like control …
  - … but does not *require* it

Current status

- RFCs
  - Host Identity Protocol (HIP) Architecture (RFC 4423) (60977 bytes)
  - Host Identity Protocol (RFC 5201) (240492 bytes)
  - Host Identity Protocol (HIP) Domain Name System (DNS) Extensions (RFC 5205) (34799 bytes)
  - Host Identity Protocol (HIP) Registration Extension (RFC 5203) (26620 bytes)
  - Using the Encapsulating Security Payload (ESP) Transport Format with the Host Identity Protocol (HIP) (RFC 5202) (68195 bytes)
  - Host Identity Protocol (HIP) Rendezvous Extension (RFC 5204) (30233 bytes)
  - End-Host Mobility and Multihoming with the Host Identity Protocol (RFC 5206) (99430 bytes)
  - Using the Host Identity Protocol with Legacy Applications (RFC 5338) (34882 bytes)
- Internet-Drafts
  - Basic HIP Extensions for Traversal of Network Address Translators (75933 bytes)
  - Basic Socket Interface Extensions for Host Identity Protocol (HIP) (42500 bytes)
  - HIP Certificates (19638 bytes)
  - HIP BONE: Host Identity Protocol (HIP) Based Overlay Networking Environment (42692 bytes)

Implementations

- HIP for Linux (HIPL) infrahip.hiit.fi
- HIP for NetBSD
- OpenHIP