

Update on Future Internet Research

Prof. Sasu Tarkoma

13.10.2009

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

Contents

- Introduction
- Current state
- Research Activities
- Host Identity Protocol (HIP)
- Packet Level Authentication (PLA)
- Overlays (i3 and Hi3)
- DONA
- Clean-slate design: PSIRP
- 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

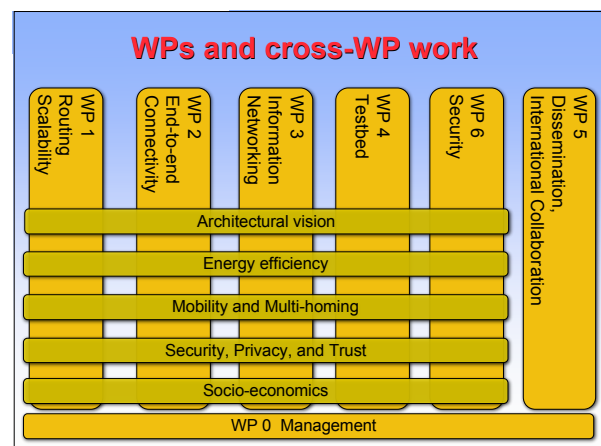
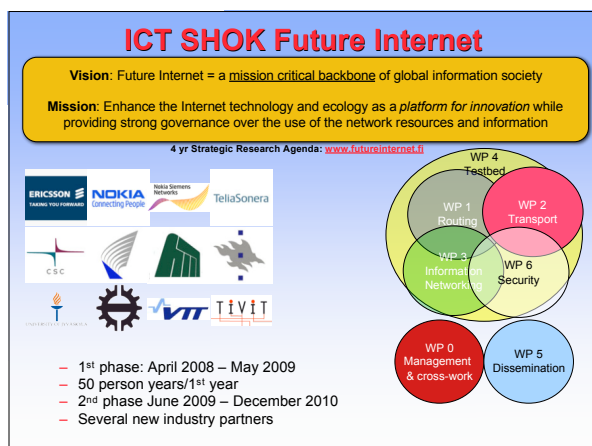
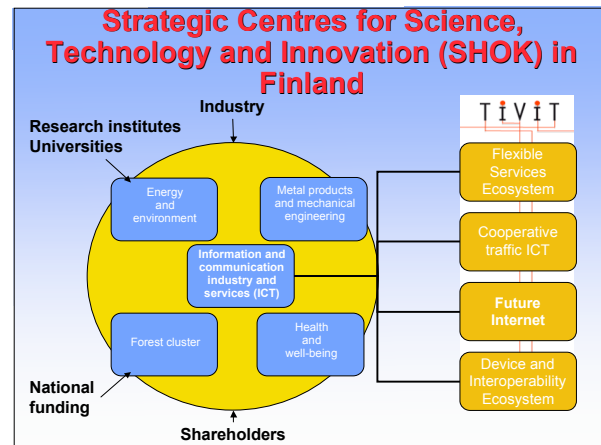
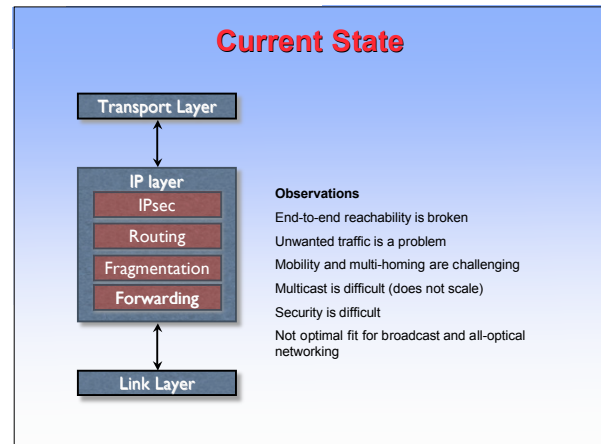
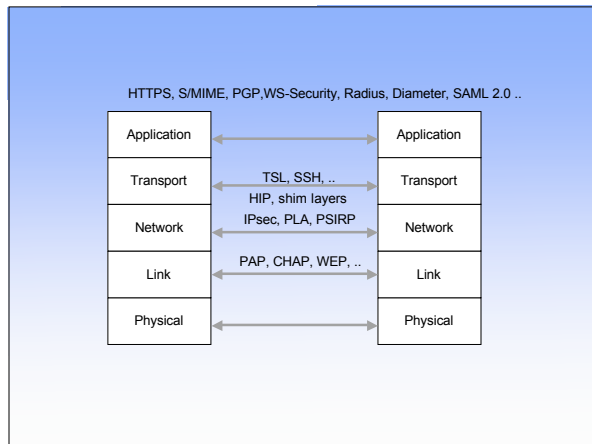
- A lot of the assumptions of the early Internet has changed
 - Trusted end-points
 - Stationary, publicly addressable addresses
 - End-to-End
- We will have a look at these in the light of recent developments
- End-to-end broken by NATs and firewalls

Convergence and Divergence



Current State

- Internet is growing fast (40%+ annual growth)
- Much of the growth comes from P2P and video delivery
- There are circa 1 billion Internet users and 3.3 billion mobile phone users
- Mobile Internet is anticipated to grow rapidly
 - Many problems with applications and services
- It is very difficult to change the Internet backbone and large access networks
 - Overlay solutions
 - Middleboxes
- A lot of discussion on Internet architecture
 - Clean-slate vs. incremental



FI SHOK Summary

- WP 1 Routing: short- & long-term: **active contribution in IRTF → IETF**
- WP 2 End-to-end connectivity: **medium term – simulation & modeling**
- WP 3 Information networking: **long-term – exploration**
- WP 4 Testbed: **preparing test platform**
- WP 5 Dissemination & cooperation
- WP 6 Security: unwanted traffic prevention and malware detection
- WP 0 Management

ICT Research in Europe

- In the FP7 Work Programme 2009-2010
 - *ICT remains central for sustainable economic growth and adjusting to the changing social realities*
 - *Lower carbon emission economy, globalisation, new value chains, higher quality health and social care, inclusion, security,...*
 - *3 major technology and socio-economic transformations*
 - **Future Internet**
 - Alternative paths to ICT components and systems
 - ICT for sustainable development

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

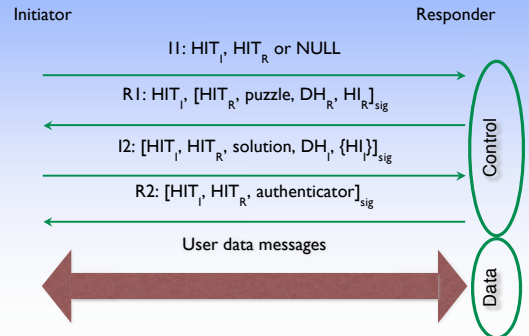
- Architectural change to TCP/IP structure
- Integrates security, mobility, and multi-homing
 - Opportunistic host-to-host IPsec ESP
 - End-host mobility, across IPv4 and IPv6
 - End-host multi-address multi-homing, IPv4/v6
 - IPv4 / v6 interoperability for apps
- A new layer between IP and transport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

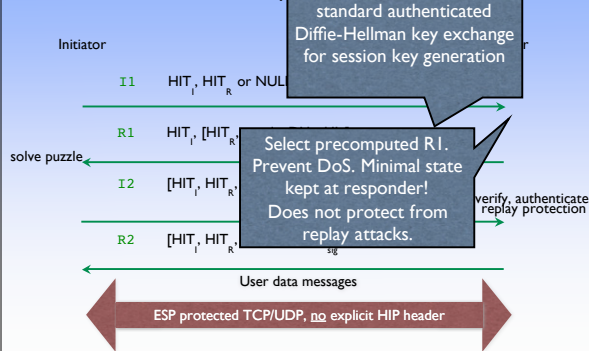


Protocol overview



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

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

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

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)




Hi³: combining HIP and i^3

- 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

Hi³ and DHT-based rendezvous



Control/data separation



An Internet control plane?

- HIP separates control and data traffic
- Hi³ routes control traffic through overlay
 - Control and data packets take potentially very different paths
- Allows telecom-like control ...
 - ... but does not *require* it

PLA

Packet Level Authentication (PLA)

- We assume that per packet public key cryptography operations are feasible in Internet's scale because of new digital signature algorithms and advances in semiconductor technology
- PLA is a novel solution for protecting the network infrastructure against various attacks (e.g., DoS) by providing availability
- The network should be able to fulfill its basic goal: to deliver valid packets of valid users in reliable and timely manner in all situations

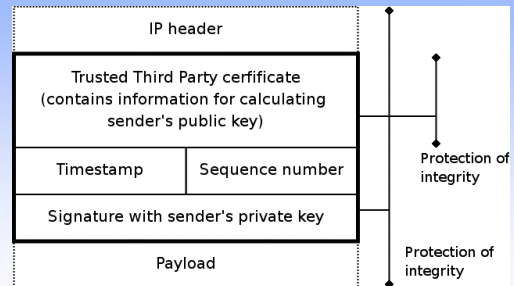
PLA continued

- The main aim of PLA is to make it possible for any node to verify authenticity of every packet without having previously established trust relation with the sender of the packet
 - Malicious packets can be detected and discarded quickly before they can cause damage or consume resources in the rest of the network
 - Good analogy for PLA is a paper currency: anyone can verify the authenticity of the bill by using built-in security measures like watermark and hologram, there is no need to contact the bank that has issued the bill

PLA continued

- PLA accomplishes its goals by using public key digital signature techniques. PLA adds an own header to the packet using standard header extension technique
 - The PLA header contains all necessary information for detecting modified, duplicated and delayed packets
 - PLA complements existing security solutions instead of replacing them. PLA can work together with other security solutions such as Host Identity Protocol (HIP) and IPSec
- Initial PLA implementation has been built on top of IPv6, however PLA is not dependent on the network layer protocol used and it can be also be positioned on top of layer 2 protocols

PLA Header




PLA Performance

- With the help of dedicated hardware acceleration, per packet public key cryptography is scalable to high speed core networks and mobile devices
 - Simulation results show that an FPGA based accelerator developed for PLA is capable of performing 166,000 verifications per second
 - Transferring the design into a 90nm ASIC using Altera's Hardcopy technology would improve performance to 850,000 verifications per second with power consumption of 26μJ per verification
 - Such performance would be enough to verify 50Gbps of traffic with jumbo frames (60kbits of payload per frame)

DONA

Anycast Routing of Fetches in DONA



- If there's an entry for a data item, follow next-hop
- Otherwise, send to parent
- Standard routing behavior, but at DONA-layer

DONA

- Naming makes it easy to authenticate data
- DONA-layer provides easy access to data:
 - name-based "resolution through routing"
 - caching and replication infrastructure
- DONA makes it easier to build transport, applications

PSIRP: Project Overview

Project Coordinator Arto Karila Helsinki University of Technology, HIIT Tel: +358 50 384 1549 Fax: +358 9 694 9768 Email: arto.karila@hiit.fi	WP1 Management (TKK-HIIT)
Partners <ul style="list-style-type: none"> Helsinki University of Technology Helsinki Institute for Information Technology (FI) RWTH Aachen University (DE) British Telecommunications Plc (GB) Oy L M Ericsson Ab (FI) Nokia Siemens Networks Oy (FI) Institute for Parallel Processing of the Bulgarian Academy of Science (BG) Athens University of Economics and Business (GR) Ericsson Magyarorszag Kommunikacios Rendszeres K.F.T. (HU) 	WP2 Architecture Design (TKK-HIIT)
Duration: January 2008 – June 2010 Total Cost: €4.1m EC Contribution: €2.5m Contract Number: INFSO-ICT-216173	WP3 Implementation, Prototyping & Testing (LMF)
	WP4 Validation and Tools (BT)
	WP5 Dissemination and Exploitation (NSNF)
	Project website: www.psirp.org

Observation: It's All About Information

Internet Today:

- In 2006, the amount of digital information created was 1.288×10^{18} bits
- 99% of Internet traffic is information dissemination & retrieval (Van Jacobson)
- HTTP proxying, CDNs, video streaming...
- Akamai's CDN accounts for 15% of traffic
- Between 2001 and 2010, information will increase 1million times from 1 petabyte (10^{15}) to 1 zettabyte (10^{21})
- Social networking is information-centric
- Most solutions exist in silos
- overlays over IP map information networks onto endpoint networks

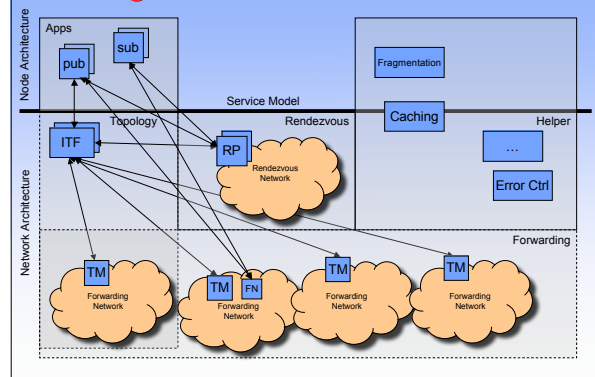
Internet Tomorrow:

- Proliferation of dissemination & retrieval services, e.g., context-aware services & sensors
- aggregated news delivery
- augmented real life
- Personal information tenfold in the next ten years (IBM, 2008)
- Increase of personalized video services
- e.g., YouTube, BBC iPlayer
- Vision recognized by different initiatives & individuals
- Internet of Things, Van Jacobson, D. Reed
- lack of interworking of silo solutions will slow innovation and development speed

Publish/Subscribe Internet Routing


- We propose a future network design that
 - gives more **trust** and more **anonymity** to Internet
 - ensures network and data **availability**
 - ensures **rapid** and **accurate** dissemination of crucial information
- The **publish/subscribe** model
 - Subscribers and publishers
 - Many-to-many communication
 - End-points described in terms of data and local links
 - Incorporating support for end-point identification
 - Flat self-certifying labels
 - Data-centric** routing, forwarding, rendezvous

High-Level Architecture




Node Architecture: Component Wheel

- Components may be decoupled in space, time, and context
 - Layerless protocol suite
- Applications may insert or request new components to the wheel at runtime
 - Implemented as **helper functions**
- The components are attached to the local blackboard (**BB**)
 - Components are attached to the local blackboard, sharing publications, state
 - Pub/sub is used to signal changes to blackboard state



PSIRP



Observations

- No topological addresses, only labels
- Security enhanced using self-certification
- End-to-end reachability, control in the network
- Natural support for multicast, it is the norm
- Support for broadcast and all-optical label-switching technologies

Dynamic state is introduced into the network

How do we make it scale?

Security and Trust

- We are going towards identity-based service access
 - A number of identities per host
 - Pseudonyms, privacy issues
 - Delegation and federation are needed
- Decentralization: the user has the freedom of choosing who manages identity and data
- Solutions for authentication
 - Below applications: HIP, PLA
 - ID-FF
 - Web-based standard (top-down)
 - OpenID and OAuth
 - Web-based practice (bottom-up)
 - SAML 2.0

Summary of Future Internet Developments

- Incremental using overlays and middleboxes
 - Short term solutions
 - HIP
 - Difficult to introduce new protocols
 - Connectivity and reachability problems
 - A lot of issues are solved in application layer
- Radical with clean-slate
 - Impossible to deploy?
 - Long haul development
 - PLA, PSIRP

Thank You