UNIVERSITÄT KAISERSLAUTERN

NATURWISSENSCHAFTLICH-TECHNISCHE UNIVERSITÄT IM HERZEN DER PFALZ

ICSY Integrated Communication Systems

http://www.icsy.de

Site Multihoming and Provider-Independent Addressing using IPv6

Dipl.-Ing. (M.Sc.) Dirk Henrici
University of Kaiserslautern, Germany
Group Integrated Communication Systems

email: henrici@informatik.uni-kl.de

Outline

Renumbering and Multihoming in IPv6

- Motivation and Challenges
  - Renumbering
  - Multihoming
- Current Practices and Research
  - IPv4
  - IPv6
- Our Approach
  - ULIDs
  - NetMapping

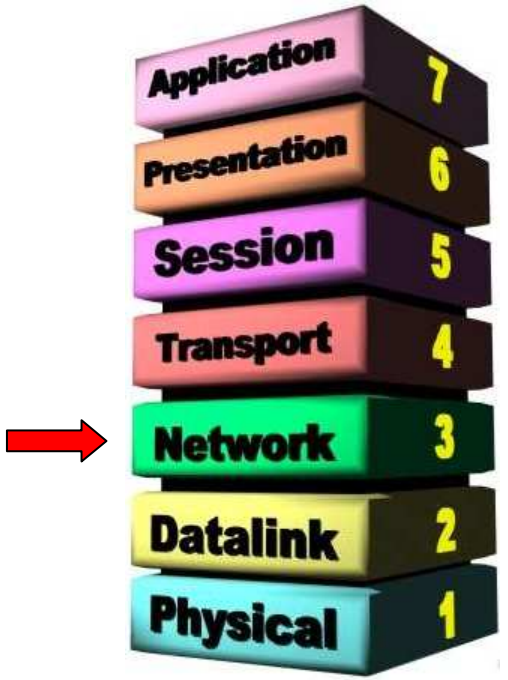

Picture source: How stuff works

Motivation

Renumbering

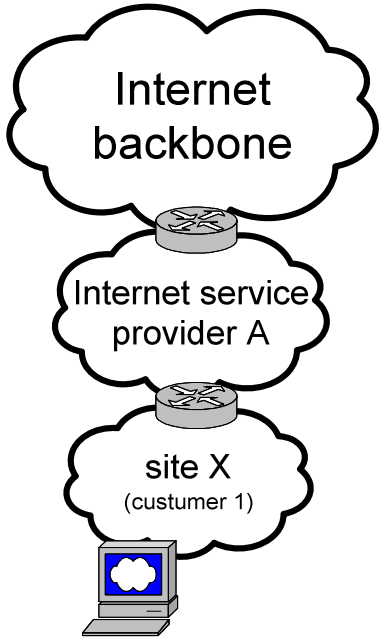

Addressing using IPv6:

INET::

INET:A::

INET:A:C1::

INET:A:C1: NETa:HOSTb

Hierarchical routing in IPv6

➔ A site needs to be renumbered if its ISP is changed

Motivation

Multihoming

➔ create failsafe Internet connection

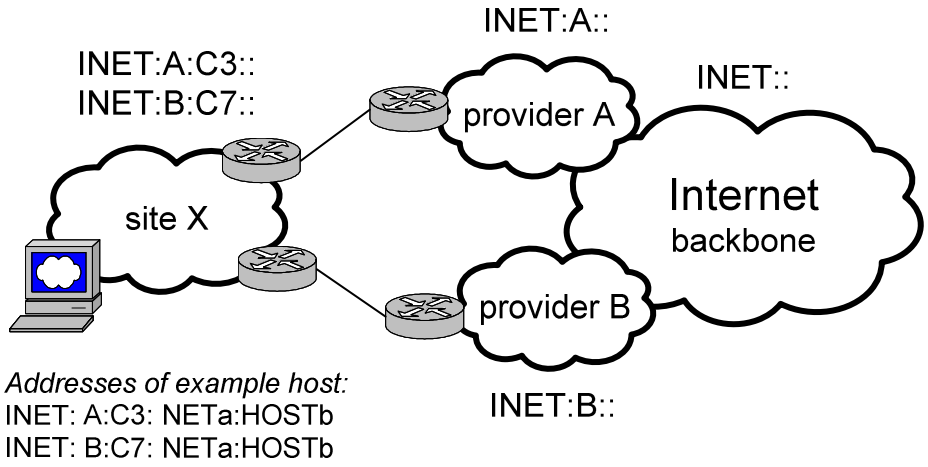

Hierarchical routing in IPv6

➔ Each host has more than a single IP address

Challenges

Renumbering

➔ not feasible in practice in large networks

(even if hostnames are used wherever possible)

because:

- much planning and many steps required
- renumbering without interruption of services difficult
- auto-configuration features not sufficient
- IP-based access control lists in routers, firewalls etc.
- IP-addresses present in configfiles of servers (e.g. resolv.conf)

➔ categorically avoid the need to renumber a network

Challenges

Multihoming

Functionality
- transport-layer survivability
- traffic engineering capability
- enforcement of administrative policies
- route selection by hosts optionally possible

Basic Requirements
- scalability
  (neither affect IPv6's hierarchical routing nor inject BGP routes)
- security

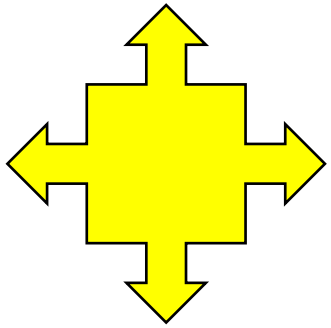

Manageability
- do not affect the end-to-end model
- simple setup and administration
- do not require new infrastructure

Feasibility
- do not require changes in hosts
- do not require cooperation between ISPs
- compatible to existing Internet standards
  (e.g. permit ingress and egress filtering)
- transition possible without a flag day

Current Practice: IPv4

IPv4 multihoming

Scenario 1:
public address space

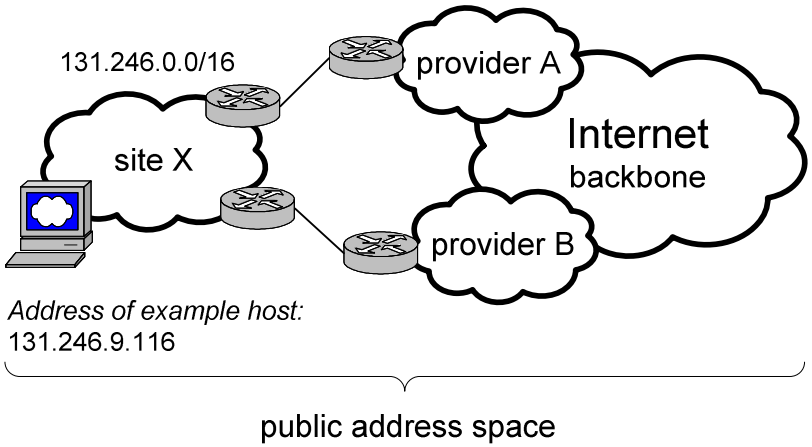

Scenario 2:
IP masquerading

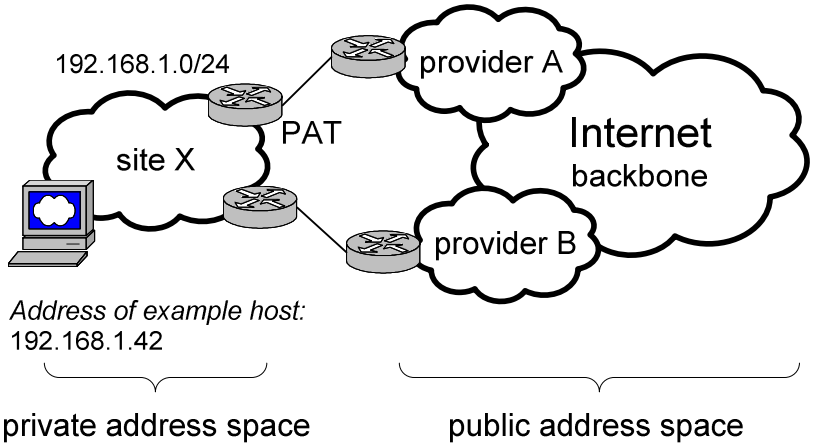

Evaluation

left | *right*

public address space | IP masquerading

IPv4 Multihoming

Functionality
- transport-layer survivability
- traffic engineering capability
- enforcement of administrative policies
- route selection by hosts optionally possible

Basic Requirements
- scalability
  (neither affect IPv6's hierarchical routing nor inject BGP routes)
- security

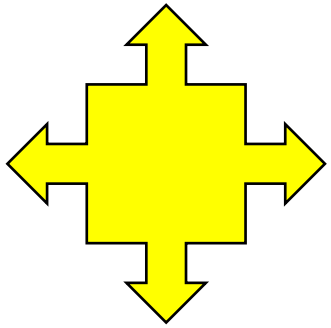

Manageability
- do not affect the end-to-end model
- simple setup and administration
- do not require new infrastructure

Feasibility
- do not require changes in hosts
- do not require cooperation between ISPs
- compatible to existing Internet standards
  (e.g. permit ingress and egress filtering)
- transition possible without a flag day

Current Practice: IPv6

IPv6 multihoming:

There is no standardized solution yet

RFC 4057: IPv6 Enterprise Network Scenarios; June 2005

> "4.9. Multihoming
>
> At this time, current IPv6 allocation policies are mandating the allocation of IPv6 address space from the upstream provider. If an enterprise is multihomed, the enterprise will have to determine how it wishes to support multihoming. This also is an area of study within the IETF and work in progress."

http://www.icsy.de

Current Practice: IPv6

Approaches to Multihoming
(categorization according to IETF draft)

- Routing
- Mobility
- Identity Considerations
  - Identity Protocol Element
  - Modified Protocol Element
- Modified Site-Exit and Host Behaviors

➔ many proposals exist

e.g. transport layer solution (SCTP), shim6

Source of picture:
http://adlib.blogs.com/andyblog/images/new-direction%20copy.jpg

Our solution: "SiMIA"

First step:
Use "Unique Local Addresses" within sites

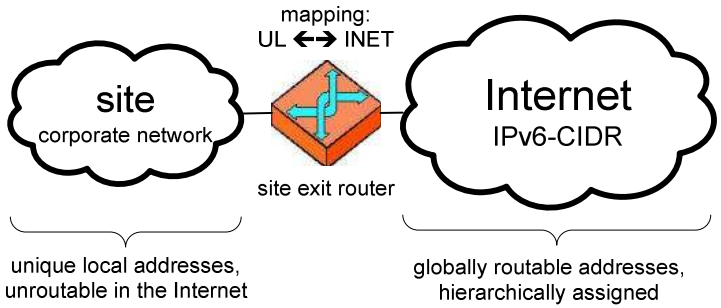

- Equivalent to private addresses in IPv4
  - ➔ not routable within the Internet
    but: globally unique
- 1:1-mapping at site-exit routers
  - ➔ between UL-addresses and globally routable addresses
    (network mapping: exchange of network prefix)
- Solely use UL-addresses within site
  - ➔ solves renumbering issue, eases access control lists etc.
- So far: Simple! But: End-to-end model not completely satisfied

Our solution: "SiMIA"

Second step:
Use UL-addresses for identifier/locator-split

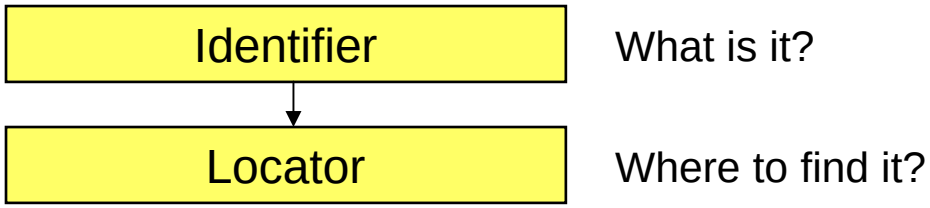

- Make UL-addresses available in the DNS
  - ➔ exploit longest prefix match (RFC3484) ➔ compatibility
- Use UL-addresses as identifiers wherever possible
  - ➔ always, excepting for Internet routing (locator) and non capable sites
- Network mapping at site-exit routers between address spaces

http://www.icsy.de

Our solution: "SiMIA"

Example

Communication between hosts in two sites, both employing "SiMIA"

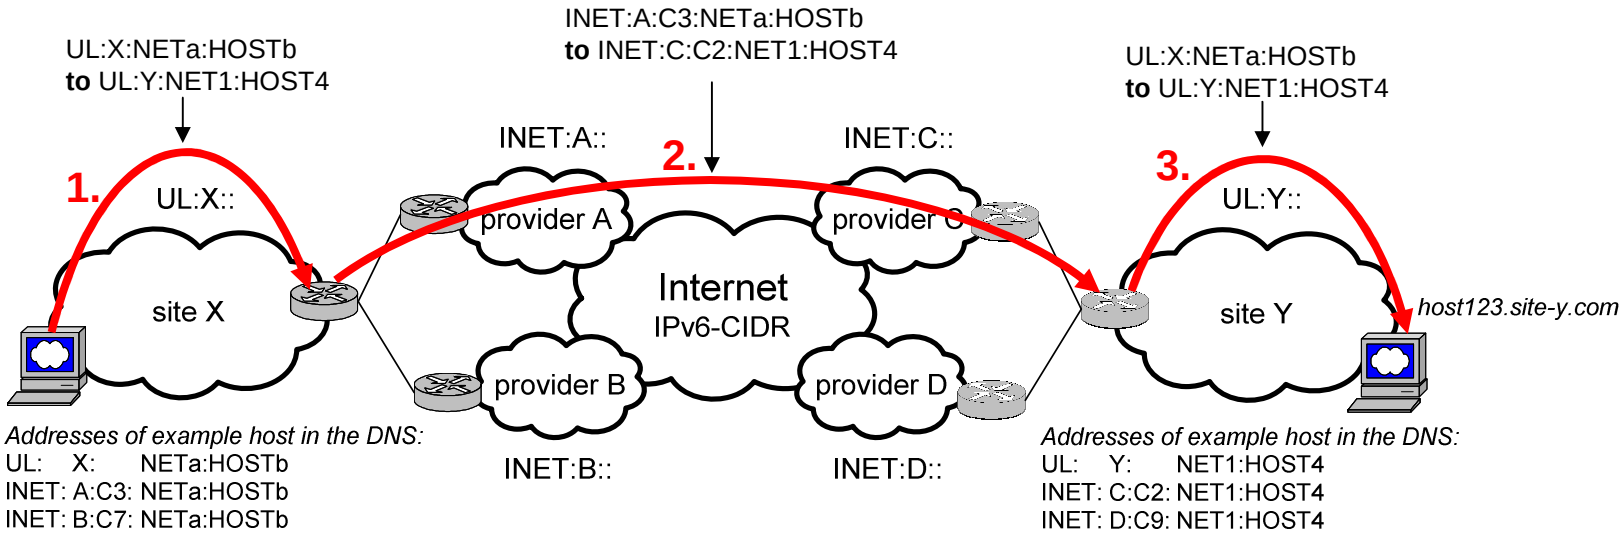

- UL-addresses are used as interface identifiers and as locators within sites
- INET-addresses are used as interface locators for ISPs and in the Internet backbone

Assessment of our solution

Multihoming
with “SiMIA”

Functionality
- transport-layer survivability
- traffic engineering capability
- enforcement of administrative policies
- route selection by hosts optionally possible

Basic Requirements
- scalability
  (neither affect IPv6’s hierarchical routing nor inject BGP routes)
- security (but we need to be careful!)

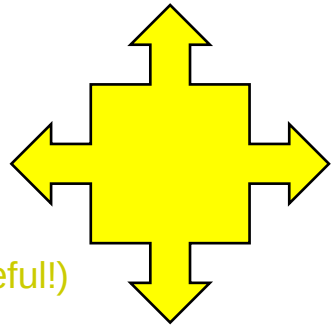

Manageability
- do not affect the end-to-end model
- simple setup and administration
- do not require new infrastructure

Feasibility
- do not require changes in hosts
- do not require cooperation between ISPs
- compatible to existing Internet standards
  (e.g. permit ingress and egress filtering)
- transition possible without a flag day

Integrated Communication Systems

http://www.icsy.de

Summary

- Introduction into challenges and current practices regarding renumbering and multihoming in IPv6
- Presentation of the idea behind our solution which is based on
  - Usage of “Unique Local Addresses” in LANs
  - Network mapping in site-exit routers
  - Use UL-addresses for identifier/locator-split

Thank you for your attention!

email: henrici@informatik.uni-kl.de

Publications can be found on our website “http://www.icsy.de”.