End-Middle-End Architecture for the Internet

Olli-Pekka Lamminen
TKK Networking Laboratory

Outline

- End-Middle-End Architecture
- IRTF EME Research Group
  - Requirements
- NUTSS
  - Architecture
  - Feasibility
- Considerations for EME Architectures
- Future Work
- References

End-Middle-End Architecture

- Middleboxes have changed Internet
  - End-to-end traffic model has been broken
  - Firewalls, NATs, etc.
- Middleboxes need to be included in connection establishment
  - Endpoints should be aware of middle
  - Endpoints need means to request services

IRTF EME Research Group

- IRTF Established EME RG in 2006
  - Set of requirements for EME architectures
    - Common naming
      - globally-unique
      - long-term stable
      - user-friendly
    - Access control policies
    - 2-way authentication
      - endpoint <> middlebox
    - Information protection
      - anonymity
      - encryption
    - Middlebox discovery
      - when allowed
    - Flow redirection
      - endpoint > endpoint
      - middlebox > endpoint
      - mobile rerouting
    - Protocol negotiation
    - Multi-homing
    - Multicast
    - Fast flow initiation
      - optimally 1st packet contains data
    - Incremental deployment

NUTSS Architecture

- EME compatible architecture
  - Created by people behind IRTF EME RG
- Policy-aware edge networks connected to policy-free core
- Edge nets have *P-Boxes* and *M-Boxes*
  - *P-Box:* controls network policies
    - form a tree-like hierarchy
    - used during name-routed signalling
  - *M-Box:* 'regular' middlebox
    - just like middleboxes today (NAT, firewall, ...)
    - handles data flows
- Endpoints register to P-Boxes

NUTSS Architecture
[naming]

- Stable, user-friendly naming
  - location-independent
  - 3-tuple *(user, domain, service)*
    - *user:* not globally-unique, identifies user
    - *domain:* globally-unique, hierarchical DNS name
    - *service:* globally-unique service identifier
  - Mapped to 5-tuple address during connection establishment
- Assumes the presence of DNS

NUTSS Architecture
[name-routed signalling]

- Used to create path for data flow
- Signalling traverses through P-Box tree
  - Up until core
  - Down to endpoint
- P-Boxes add next-hop tokens
  - Tokens used for address routing

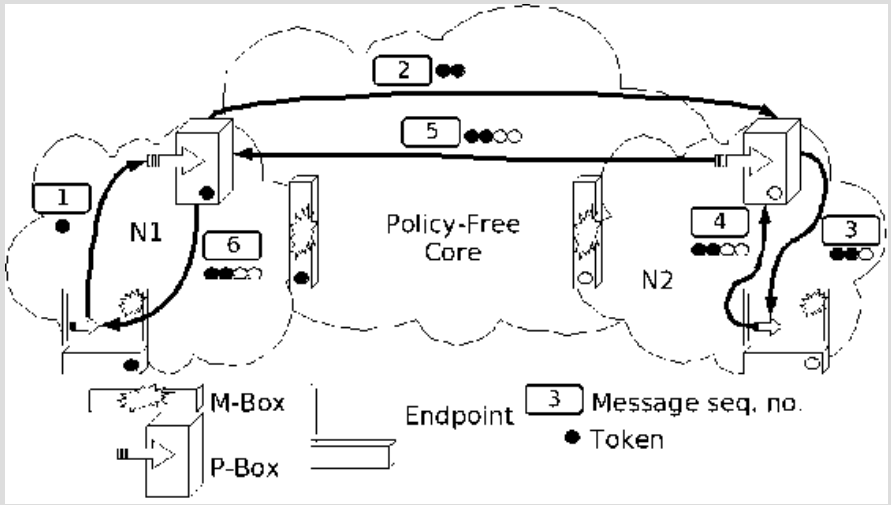

NUTSS Architecture
[address-routed messages]

- Data flows through M-Boxes
- Routing by next-hop tokens
- Endpoint addresses given during name-routing

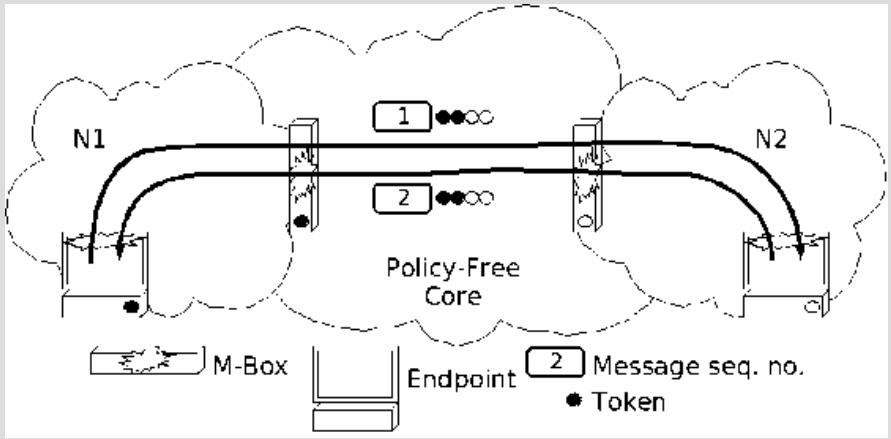

NUTSS Feasibility

- Fulfills many IRTF EME requirements
  - Mobility by short-lived addresses and rapid renegotiation
  - Multi-homing from location independence
  - Multicast with extended naming
    - 3-tuples changed 4-tuples
- Performance may be an issue
  - Lots of signalling overhead
  - Payload in 1st packet not possible
- Deployment strategy at draft-stage
  - Which is OK

Considerations
[naming]

- Users want user-friendly names
- Names should be if not globally-unique at least scope-unique
  - Uniqueness requires coordination
  - Coordination requires authority (NS)
- Mobile endpoints will be commonplace
  - Changing location requires updating NS
  - Network registration helps mobility
  - Registration and updates require authentication

Considerations
[middleboxes]

- Middleboxes are already commonplace
  - Home routers, web proxies, ...
- Endpoints should be aware of them
  - Awareness enables flexibility
  - NAT traversal, firewall control, ...
- Middleboxes need means to advertise
- Endpoints need means to request
- Endpoints should be able to trust middleboxes and vice versa
  - Service authentication is required

Considerations
[trust]

- Trust needs to be provided
  - Joining network
  - Name and location updates
  - Middlebox services
- Who trusts whom?
  - Endpoints vs. Endpoints
  - Endpoints vs. Networks, Middleboxes
  - Between the networks
  - Inside a network
- Who provides the trust?
  - Trust relationships between operators
  - 3rd party trust authorities

References

- IRTF EME Research Group
  - [http://www3.tools.ietf.org/group/irtf/trac/wiki/EME]
- NUTSS
  - [http://www3.tools.ietf.org/group/irtf/trac/wiki/EME_NUTSS]
  - An End-Middle-End Approach to Connection Establishment
    Guha & Francis: In Proceedings of SIGCOMM'07
- Middleboxes
  - Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World
    Blumenthal & Clark: ACM TOIT, vol.1, no. 1, Aug. 2001
  - Middleboxes No Longer Considered Harmful
    Walafish et al.: In Proceedings of OSDI'04

End-Middle-End Architecture for the Internet

Questions?