Mike Fisk
mfisk@lanl.gov
LANL/UCSD

draft-fisk-midcom-session-00

IETF 49
Midcom BOF

This presentation

- Is about:

  Interactions between middleboxes and untrusted, end-host applications

- Is not about: (but is complementary to)

  Middlebox control by trusted application-layer gateways

Man-in-the-street Definitions of a Middlebox

"A pragmatic device that transparently fixes packet flows between flawed endpoints."

— Engineer running a network

"A flawed device that breaks transparency by impeding the flow of packets between endpoints."

— End-to-end/transparency purist
— Frustrated user

Examples: Firewall, NAT, TCP PEP, etc...

Flawed? A question of agility:

Incapable of adequately supporting a necessary protocol or policy.

Incapable of adequately supporting a necessary protocol or policy.

Adequate = securely, fairly, ...

Protocol = IPv6, IPsec, window scaling, global addresses, session bundles, ...

Incapable = Not under your administrative domain of control

- Site's networking group can't manage all of site's desktops
- Desktop users can't control the network provider's firewall

End-to-end Nirvana

- Every end host is well-managed and supports everything necessary to work across every kind of network (IPv4, IPv6, Long Fat Networks, wireless, adversarial).

- Deployment cost: Deploy each new feature to each host

Reality

- Update end hosts once every few years

- All other changes made with middleboxes

Current Trend

- Network peers provide services that enable end nodes to work across every kind of network
  - e.g. SOCKS, RSIP, IPsec tunneling, HTTP Proxies
  - Clients are aware of middlebox and request functionality
- Deployment cost: Deploy each framework protocol to each host and then update/manage small number of servers.
  - Server deals with changing policies and protocols
  - Clients are less well managed and trusted
  - Clients submit themselves to the whims of servers

Observations

Application only sees end-to-end byte stream or message payloads.

Expectation is that packets are end-to-end, but frequently not so.

Middleboxes perform transport splicing/spoofing.

Why not have add a thin abstraction layer between the application and the transport protocols to provide end-to-end stream of messages over a series of transport connections?

Future? A Session Setup Protocol

Deploy a *single* framework protocol to each host and then update/manage small number of servers.

- Provide end-to-end byte stream or datagram payloads
- Relay data across a series of transport layer connections
- Middleboxes operate only at transport endpoints; no mucking with something that is supposed to be e2e.
- Applications agree (or not) to the requirements (policy) of the middlebox
- Is a single, flexible framework protocol feasible?
  Good question. Let's try.
  Encouraging thoughts:
  – SOCKS is used for many applications.
  – Who'd have thought that HTTP would be used for everything? (RPC, e-mail, etc.)

Requirements

- Middlebox discovery
- Mutual authentication (n-way)
- Encryption (e2e or between hops)
- Abstract view of e2e network connection without assuming e2e packets
  – Build (a series of) transport connections between middleboxes
- Compatibility with current protocols and middleboxes
  – Some new protocols (telephony, etc.) have more freedom
- Dynamic reconfiguration (mobility, topology changes)
  - New middlebox in path triggers renegotiation
  - One or more transport hops may change (TCP Connection Migration)
- Minimize need for Application Layer Gateways

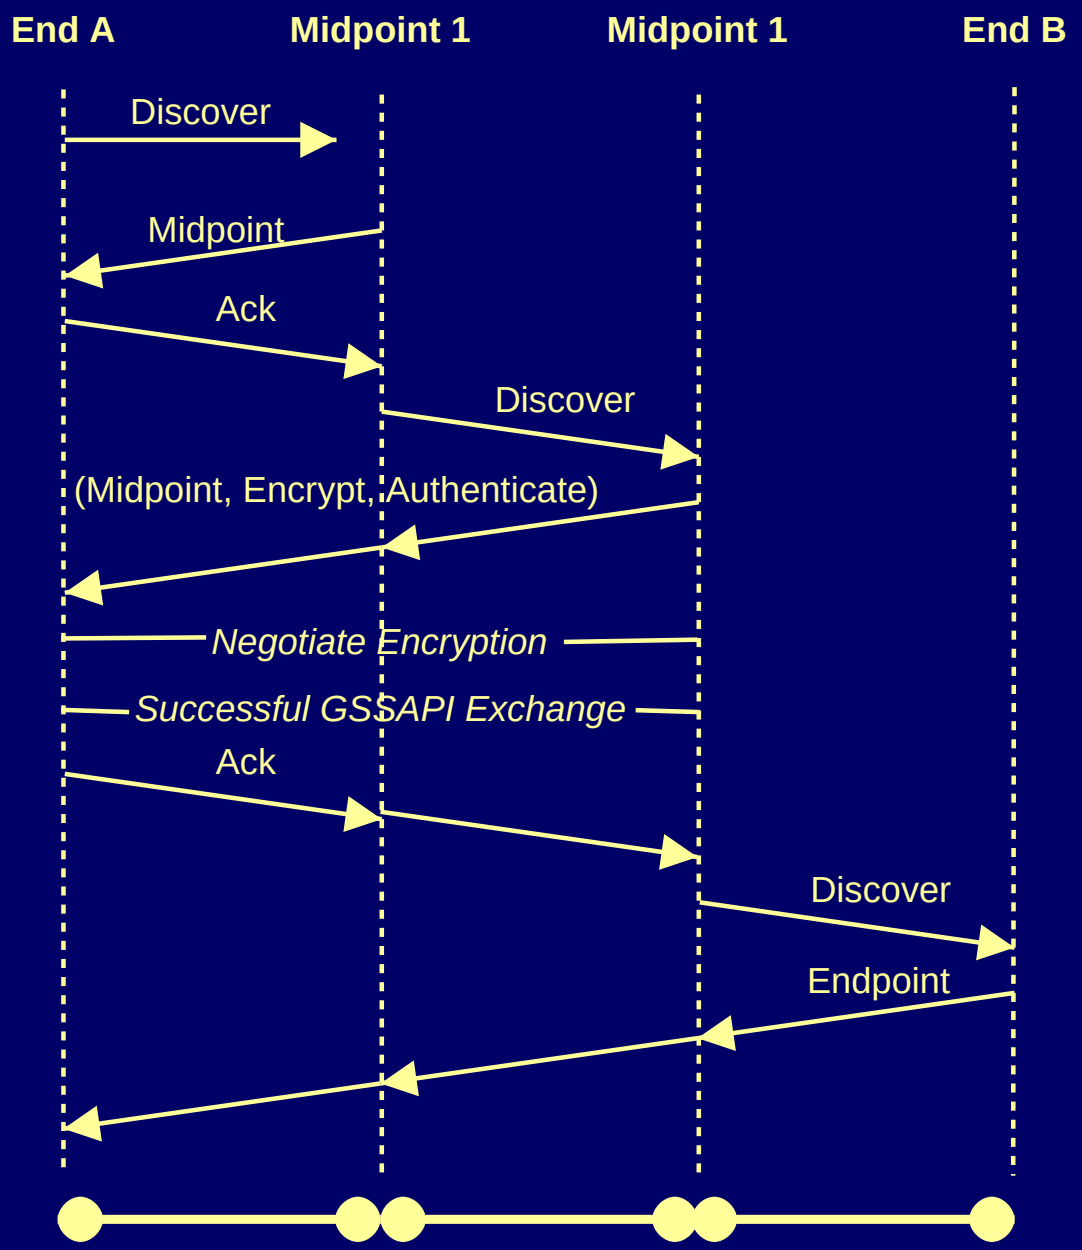
End A
Midpoint 1
Midpoint 1
End B
Discover
Midpoint
Ack
Discover
(Midpoint, Encrypt, Authenticate)
Negotiate Encryption
Successful GSSAPI Exchange
Ack
Discover
Endpoint

Design Decisions

How much of the needed functionality is already present in protocols like SOCKS and SIP?

New IP option for discovery.
Use SOCKS as base for setting-up each transport conn?

How much of this is just engineering and how much is still experimental?

We have experience with several point-solutions.
Now we just need to generalize.

Should this protocol be distinct from a protocol that allows ALG control?

Questions?
Comments?
Rants?

draft-fisk-midcom-session-00

mfisk@lanl.gov