DCCP Mobility and Multihoming

Eddie Kohler
UCLA

IETF 60 DCCP Meeting
August 5, 2004

Thus began our quest for a mechanism that would support mobility and multihoming, at the DCCP level, with reasonable security and DoS- prevention, without using cryptography. (The DCCP working group's charter has been interpreted as preventing DCCP from including cryptography, even MD5 hashes.) DCCP's mobility support changed, often fundamentally, in every succeeding draft.

Unsurprisingly, we did not find a suitable mechanism, and I now believe no such mechanism exists. . . .

Unfortunately, mobility and multihoming support can't easily be implemented at a higher-level layer, and there are good arguments for supporting mobility and multihoming at the transport layer – not least required interactions with congestion control. This document, therefore, presents one potential design for DCCP mobility and multihoming support. It relaxes one of DCCP mobility's original requirements by using cryptography.

Requirements

- An endpoint does not need to announce a new address before moving to that address.
- Move requests must be safe against hijacking. Even attackers that can snoop on part or all of data traffic must not be able to move a connection. However, move requests need not be safe against man-in-the-middle attackers with control over which packets get delivered (such as routers).
- Mobility must not create new, large opportunities for denial-of-service attacks.
- Endpoints must be able to move freely between different NAT domains using the mobility mechanism.
- Simultaneous moves need not be supported.

Design overview

- Support for mobility is optional and defaults to off.
- Each endpoint of a mobility-capable connection has a public 128-bit Mobility ID.
- The endpoints share a Mobility Secret, a key communicated over a secure channel. The Secret is either transmitted out-of-band, or via public-key cryptography or Diffie-Hellman exchange. It is changed on every successful move.
- A Mobility Sequence number increases monotonically with moves, and identifies which Mobility Secret a packet is using.

Why two mobility identifiers?

- Mobility ID is public and static

  Lets the stationary endpoint (and any NATs) map a move announcement to an existing connection

  The original source address and port cannot be used for this purpose because of NATs

- Mobility Secret is private and dynamic

  Used during the move handshake to prevent hijackings

  Protected by cryptography

Packet exchange (1/2)

1. Mobile host sends DCCP-Move-Request from new address
   - Contains (1) stationary host's Mobility ID, (2) mobility token
   - Mobility token encrypted by Mobility Secret; contains Sequence Number, Acknowledgement Number, Mobility ID, and new half-Mobility Secret
   - Informs stationary endpoint of the move
2. Stationary host responds with DCCP-Move-Response
   - Includes a similar token, which completes the Mobility Secret
   - Stationary host remembers both new Mobility Secrets
   - Proves to mobile endpoint that true stationary endpoint received DCCP-Move-Request

Packet exchange (2/2)

3. Mobile host sends DCCP-Move-Confirm
   - Includes token encrypted by new Mobility Secret
   - Proves to stationary endpoint that true mobile endpoint received DCCP-Move-Response
4. Stationary host sends DCCP-Move-Complete
   - Removes old Mobility Secret(s)
   - Ends this mobility episode and informs NATs and middleboxes that the connection's endpoints have definitively changed

Remaining to be done

- Specify format of Mobility Secret and particular encryption algorithms

  Possibly using a feature to negotiate algorithms

- Security Considerations

Questions for the working group

- Is mobility and multihoming worthwhile?
- Should we continue with this draft?
- Is another approach preferred?