Document Clarifications

- All NAT mappings are bidirectional
  - MAP mappings apply for both inbound and outbound packets
- MAP mappings are ***by definition*** EIM
  - Purpose of MAP mappings is to receive traffic from ***any*** peer
- PEER mappings may be EIM or EDM

PCP, 80th IETF, Prague
0900-1130 Friday 1st April 2011

Stuart Cheshire
Apple Inc.

Document Clarification

THIRD_PARTY Option and Liveness

- Clients using THIRD_PARTY option (e.g. IGD IWF) MUST verify ongoing liveness of the third party
  - e.g. periodic test connections to the service, etc.
- Purpose of lifetime & renewals is to clean stale state
- If device goes away, its mappings should clean up too
- Clients using THIRD_PARTY option MUST NOT defeat this mechanism by renewing unwanted mappings forever

Stuart Cheshire
Apple Inc.

Document Enhancement

Sample Code Improvements

- Improve sample code to illustrate event-driven operation
- PCP mappings are ***by necessity*** dynamic
  - May move your laptop to a new network and get a new external address and port
  - NAT gateway may be rebooted and give you a new external address and port
- Current sample code suggests:
  - Client asks for mapping
  - Client gets it
  - Client never has to think about it again.
- This is **not** what we intend to suggest

Protocol Question

PCP Port Number 5351 or 44323?

- PCP Packet format is based on NAT-PMP
- Initial fields of header are the same
  - Version 0 ⇒ NAT-PMP
  - Version 1 ⇒ PCP

- Using 5351 for both NAT-PMP and PCP eases transition
  - Dual-mode server listens on only one port and handles both kinds of requests
  - Dual-mode client sends PCP-format request to 5351; from NAT-PMP-only server gets immediate "bad version" error so client can re-issue request as NAT-PMP-format

Protocol Clarification

PCP Lifetime Extension with Active Traffic

- Uniform treatment of all mapping types:
- Outbound packet & PCP request
  - Creates mapping if necessary
  - Extends expiration timer if necessary
- Inbound packet
  - Does neither
  - (Remote peer is not necessarily trusted)

Stuart Cheshire
Apple Inc.

Protocol Enhancement

Notification of State Changes

- Needed for capability parity with NAT-PMP
- State Loss:
  - On reboot, NAT MAY multicast announcement
  - Clients MAY listen for multicast announcements
- Reconfiguration:
  - NAT MUST send new unicast replies to clients
  - Clients MUST handle unsolicited responses

Design Philosophy Comment

- Purpose of PCP server is to ***serve*** PCP clients
- Reasonable for server to reject:
  - Malformed client requests (client software mistake)
  - Excessive client requests (user mistake)
    - Resource limits should be scoped so this rarely happens
- ***Unreasonable*** for server to reject:
  - Well-formed requests
  - This is why we'd like to eliminate a couple of bogus error codes

Protocol Simplification (1)

Eliminate NONEXIST_PEER error

- Allow PEER opcode to ***create*** a mapping?
  - Consistent handling of TCP SYN & PCP PEER opcode
  - Avoids race condition between which is received first
- Allow PEER opcode to ***recreate*** a mapping?
  - With addition of *suggested port* field
  - Allows connection recovery after reboot
  - No NAT is obliged to respect *suggested port* field

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS error

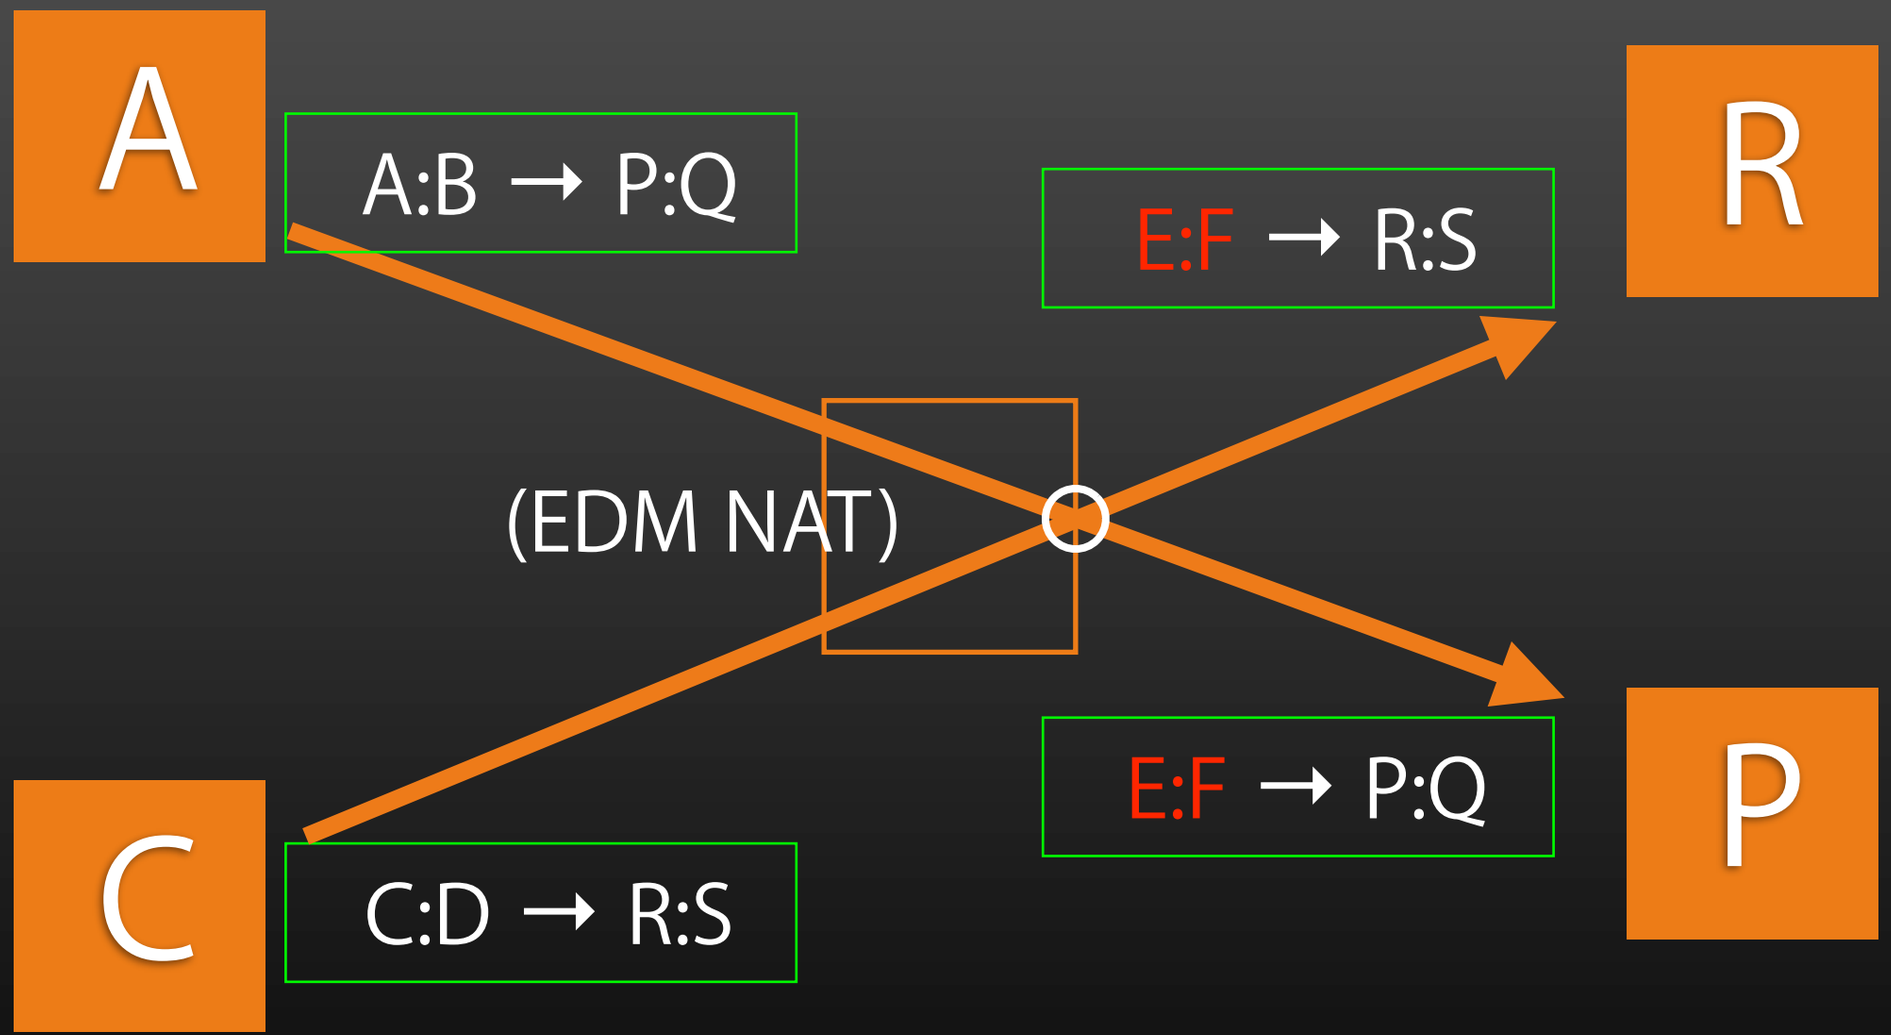

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS error

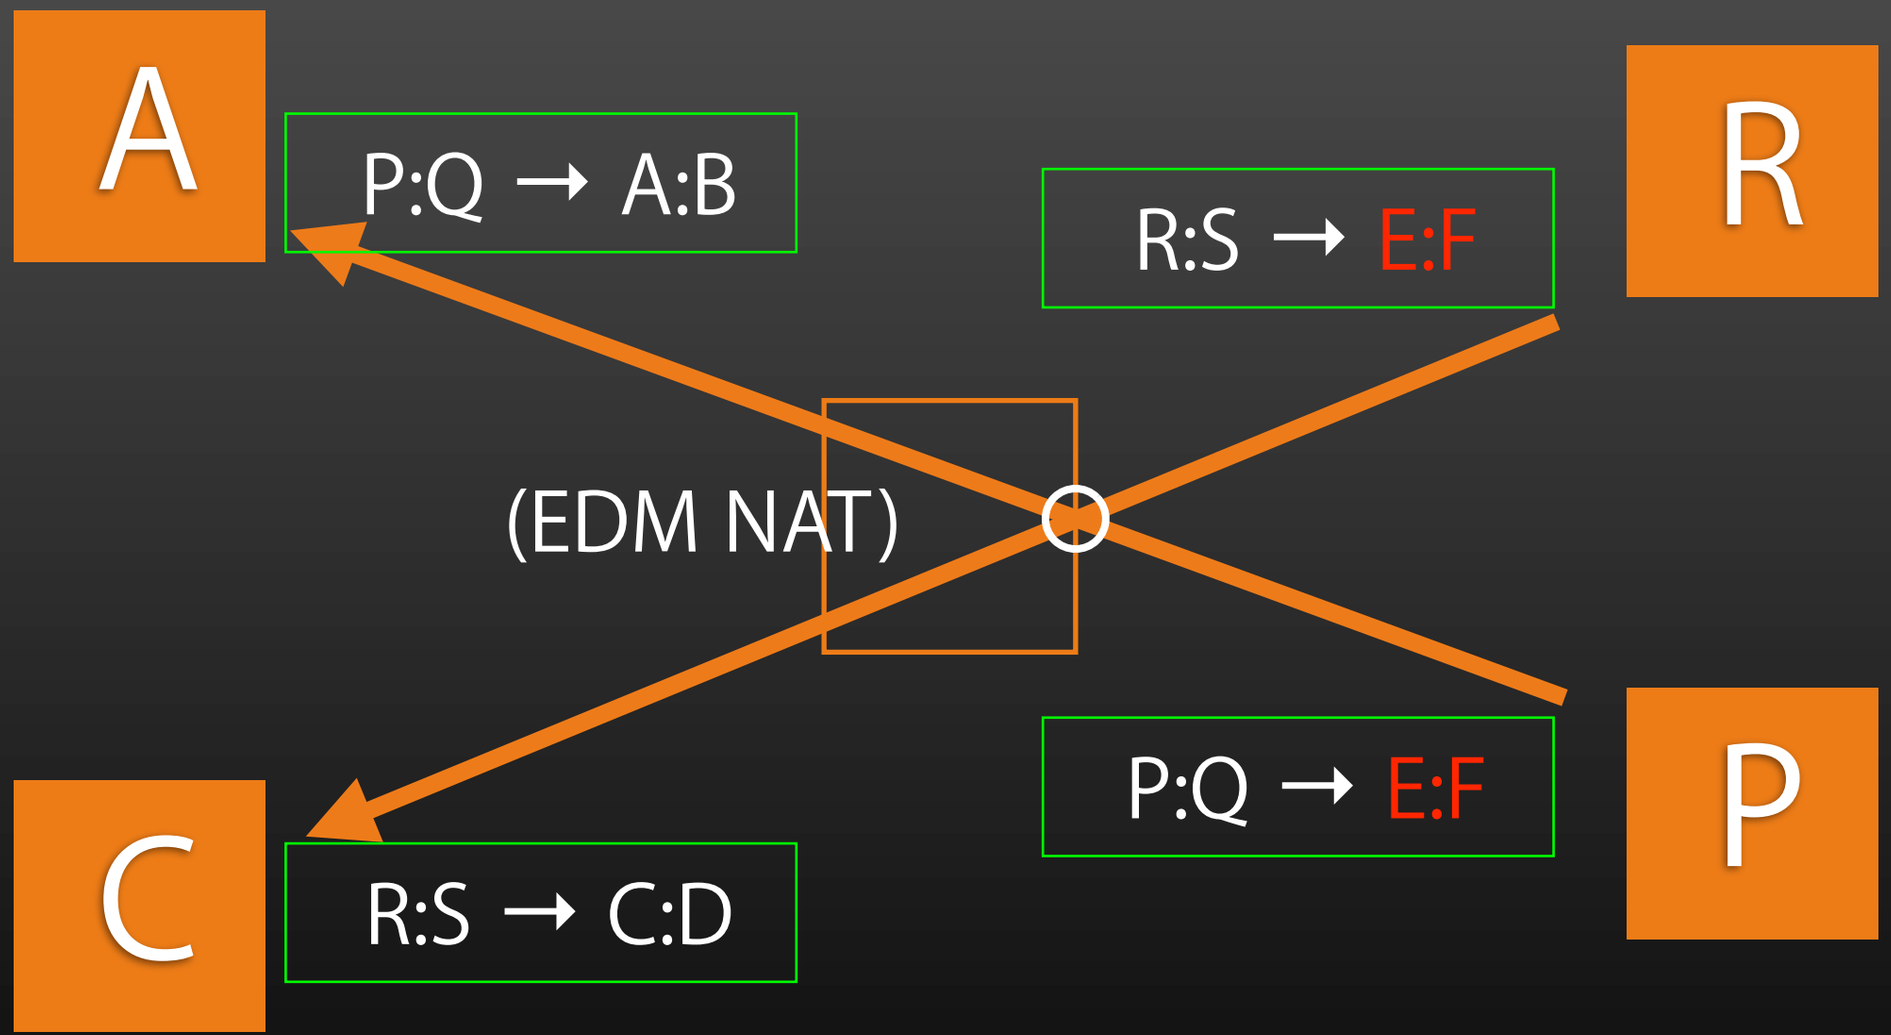

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS error

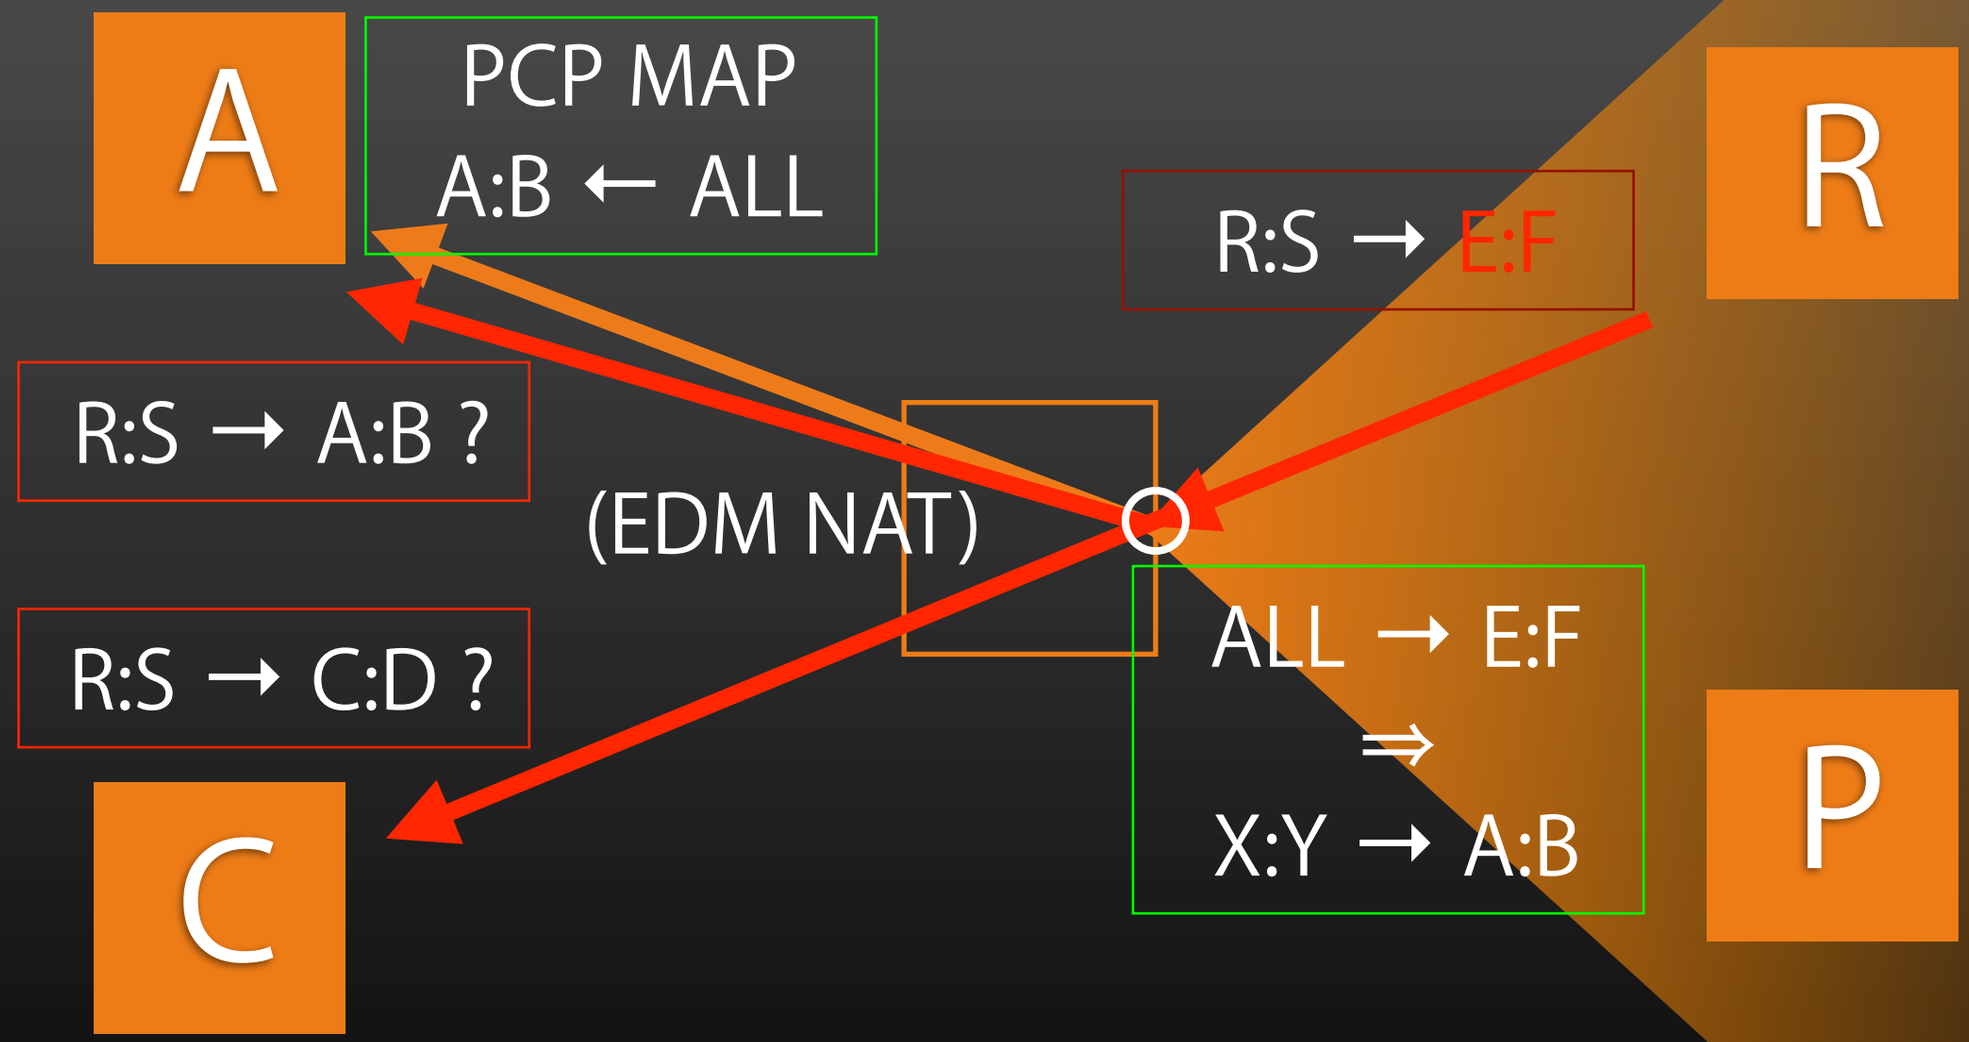

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS error

- What to do on reception of PCP MAP A:B request?
  - Give (A:B ← ALL) mapping external addr:port E:F and kill (R:S → E:F) ⇒ (R:S → C:D) mapping? No!
  - Give (A:B ← ALL) mapping ***different*** ext addr:port G:H?
    - (P:Q → E:F) ⇒ (P:Q → A:B)
    - (R:S → E:F) ⇒ (R:S → C:D)
    - (X:Y → G:H) ⇒ (X:Y → A:B)
    - (P:Q → G:H) ⇒ (P:Q → A:B)

These are the same!
What source to use for (A:B → P:Q) replies?

  - Give (A:B ← ALL) mapping ***different*** ext addr:port G:H and kill (P:Q → E:F) ⇒ (P:Q → A:B) mapping?

Stuart Cheshire
Apple Inc.

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS: Solution (1)

- Keep EDM (P:Q → E:F) ⇒ (P:Q → A:B) mapping
- Keep EDM (R:S → E:F) ⇒ (R:S → C:D) mapping
- Give EIM (A:B ← ALL) mapping same ext addr:port E:F but "subbordinate" to any existing EDMs
- If a packet matches both an EDM mapping and an EIM mapping, then the EDM mapping is used
- If an EIM mapping exists, no new EDM mappings are made using the same external addr:port
- If an outbound packet matches only an EIM mapping, but a ***reply*** to the translated packet would match an existing EDM mapping and go to wrong internal host then a new EDM mapping needs to be made

Stuart Cheshire
Apple Inc.

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS: Solution (2)

- If an outbound packet matches only an EIM mapping, but a ***reply*** to the translated packet would match an existing EDM mapping and go to wrong internal host then a new EDM mapping needs to be made
- If we translate A:B → R:S ⇒ E:F → R:S
  then reply will translate R:S → E:F ⇒ R:S → C:D
- In this case, only solution is that outgoing A:B → R:S packet has to make its own new EDM using different external addr:port not in use by any EIM
- Can be mitigated if NAT partitions its port space into ports for EDM use and ports for EIM use

Protocol Simplification (2)

Eliminate IMPLICIT_MAPPING_EXISTS

Remember:

This only applies to EDM NAT!