NAT Behavioral Requirements for TCP

Saikat Guha, Kaushik Biswas, Bryan Ford, Paul Francis, Senthil Sivakumar, Pyda Srisuresh

draft-ietf-behave-tcp-01

IETF 66

Changes Since -00

Now a standalone document

- Much easier to read
- (Re)defines terminology shared with UDP
- References UDP only for IP requirements

Handling Unsolicited SYN?

SYNs that . . .

- are inbound
- are NOT part of an in-progress TCP (S-O)
- are NOT allowed by filtering behavior

... basically the NAT cannot route

Unsolicited SYN: Option 1

Silent Drop

- Good for P2P
- Bad for erroneous SYNs
  - NATs do this today (92%)
  - Current WG consensus
  - Too rare a case?
  - Is it a problem today?

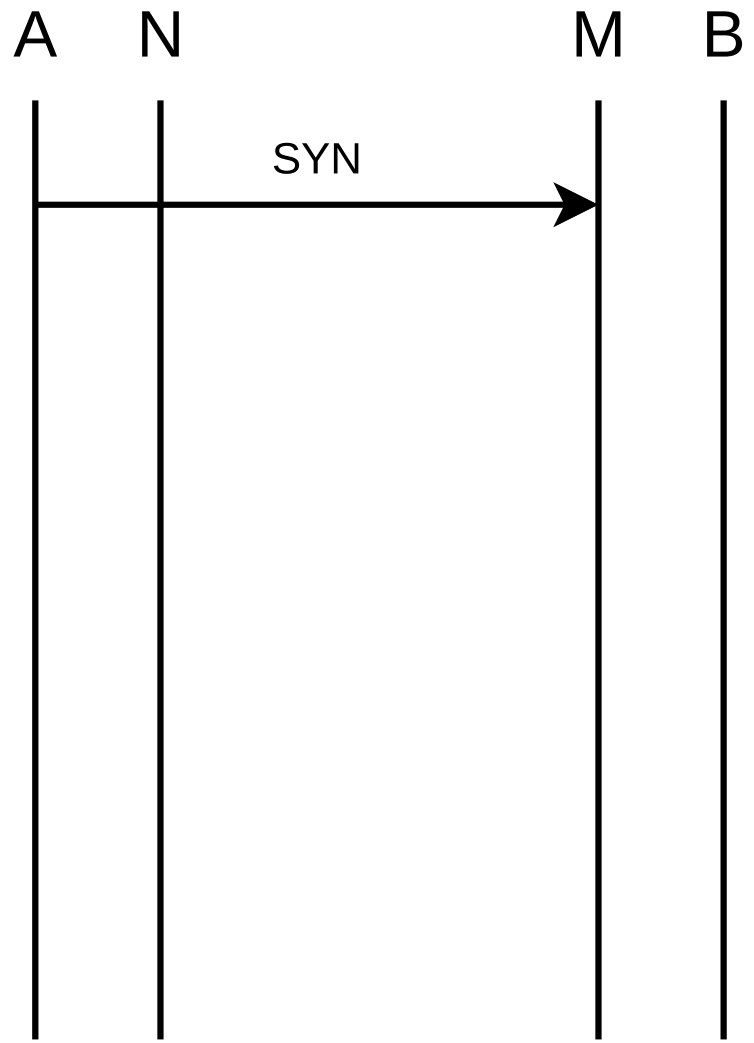

Unsolicited SYN: Option 1

Silent Drop

- Good for P2P
- Bad for erroneous SYNs
  - NATs do this today (92%)
  - Current WG consensus
  - Too rare a case?
  - Is it a problem today?

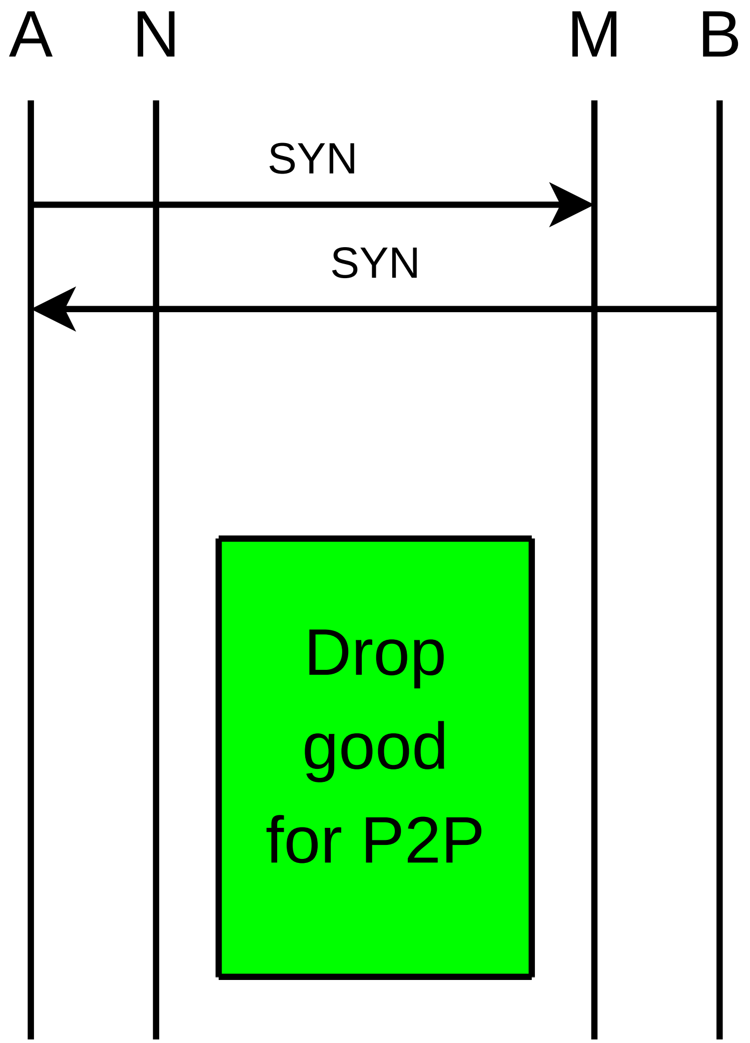

Unsolicited SYN: Option 1

Silent Drop

- Good for P2P
- Bad for erroneous SYNs
  - NATs do this today (92%)
  - Current WG consensus
  - Too rare a case?
  - Is it a problem today?

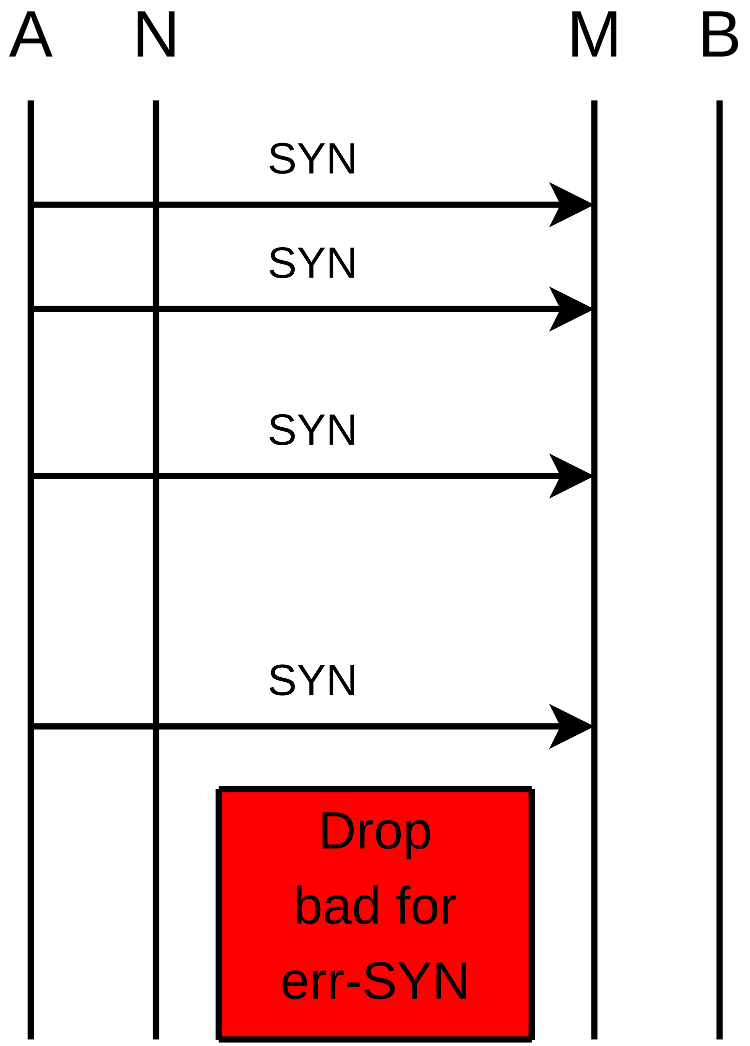

Unsolicited SYN: Option 2

ICMP Error

- Good for erroneous SYNs
- Good for P2P if . . .
  - error doesn't cause stack to abort[a]
- Otherwise, bad for P2P

[a] May need a new ICMP soft-error code proviso old stacks ignore undefined ICMPs, make sure Gont's TCPM draft (if it becomes a WG doc) retains this error as soft.

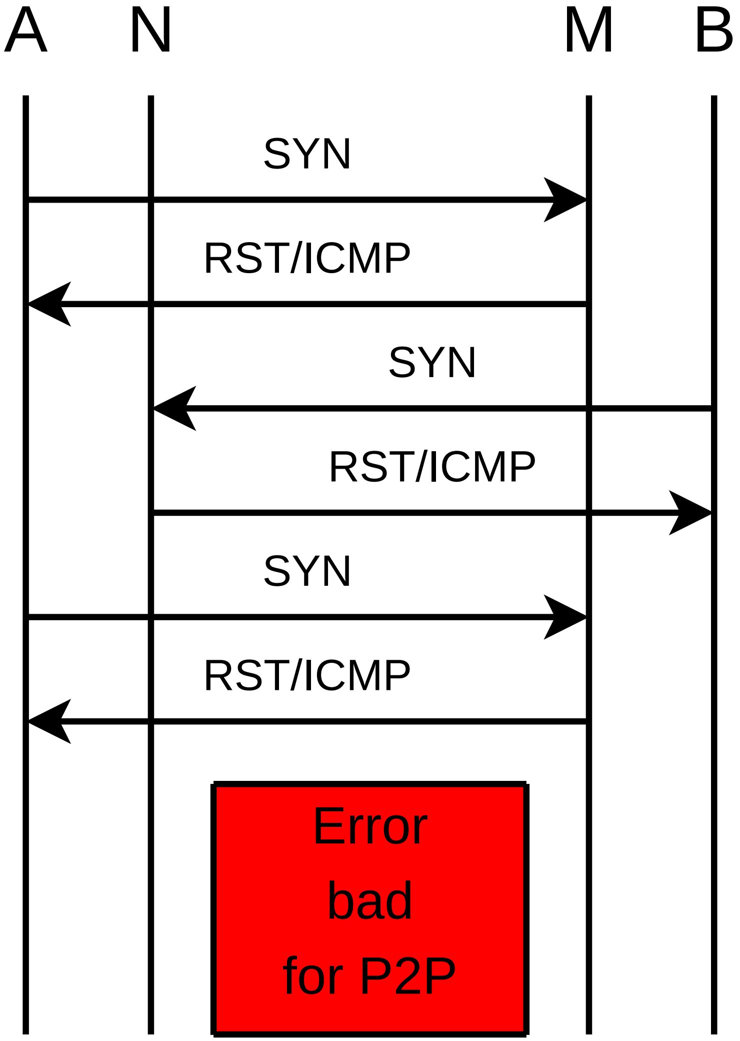

Unsolicited SYN: Option 2

ICMP Error

- Good for erroneous SYNs
- Good for P2P if . . .
  - error doesn't cause stack to abort[a]
- Otherwise, bad for P2P

[a] May need a new ICMP soft-error code proviso old stacks ignore undefined ICMPs, make sure Gont's TCPM draft (if it becomes a WG doc) retains this error as soft.

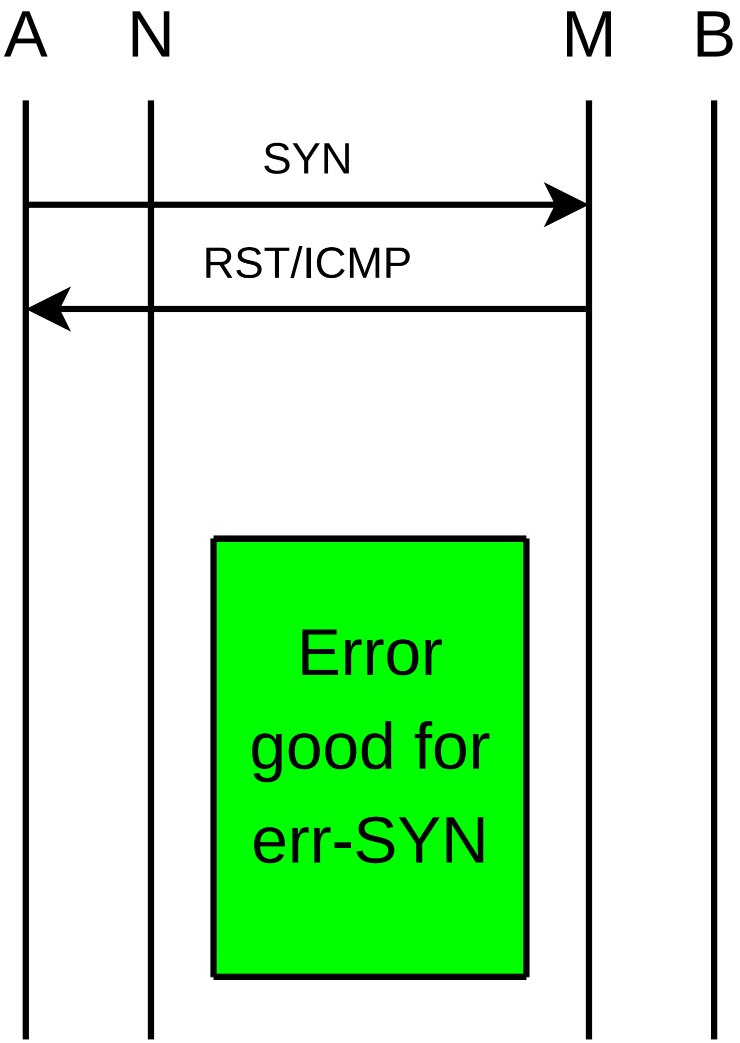

Unsolicited SYN: Option 3

Delayed Error

- Not bad for P2P
- Not bad for erroneous SYN
- Decide delay timeout
  - 6s too low for P2P?
  - 6s too high for err-SYN?

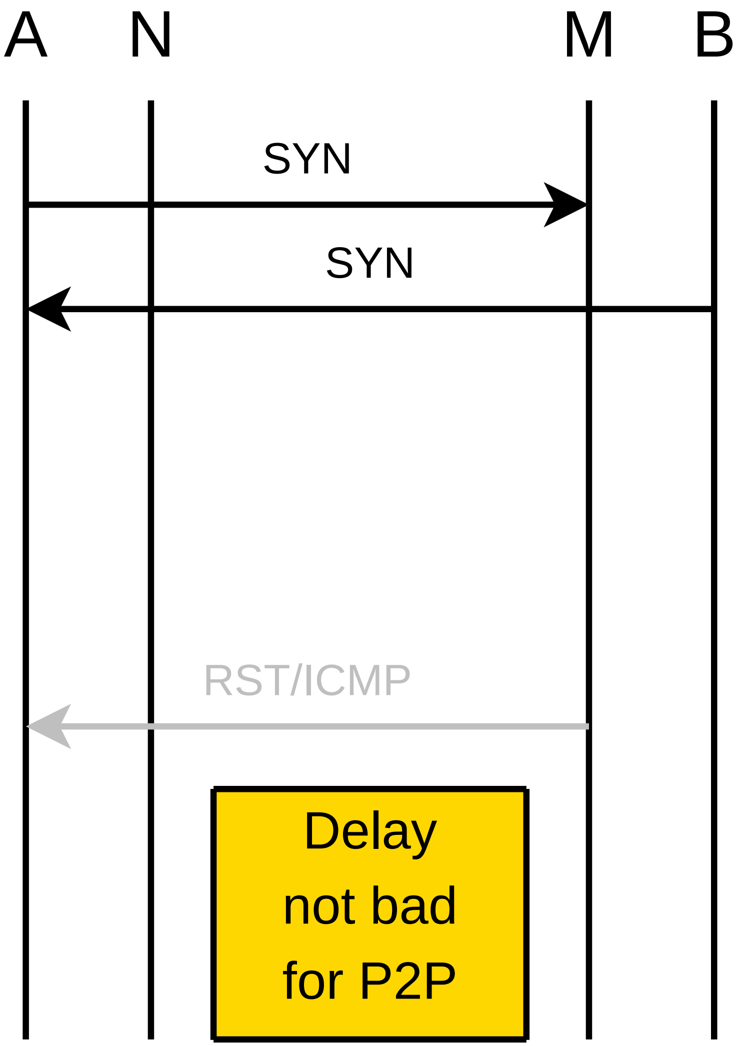

Unsolicited SYN: Option 3

Delayed Error

- Not bad for P2P
- Not bad for erroneous SYN
- Decide delay timeout
  - 6s too low for P2P?
  - 6s too high for err-SYN?

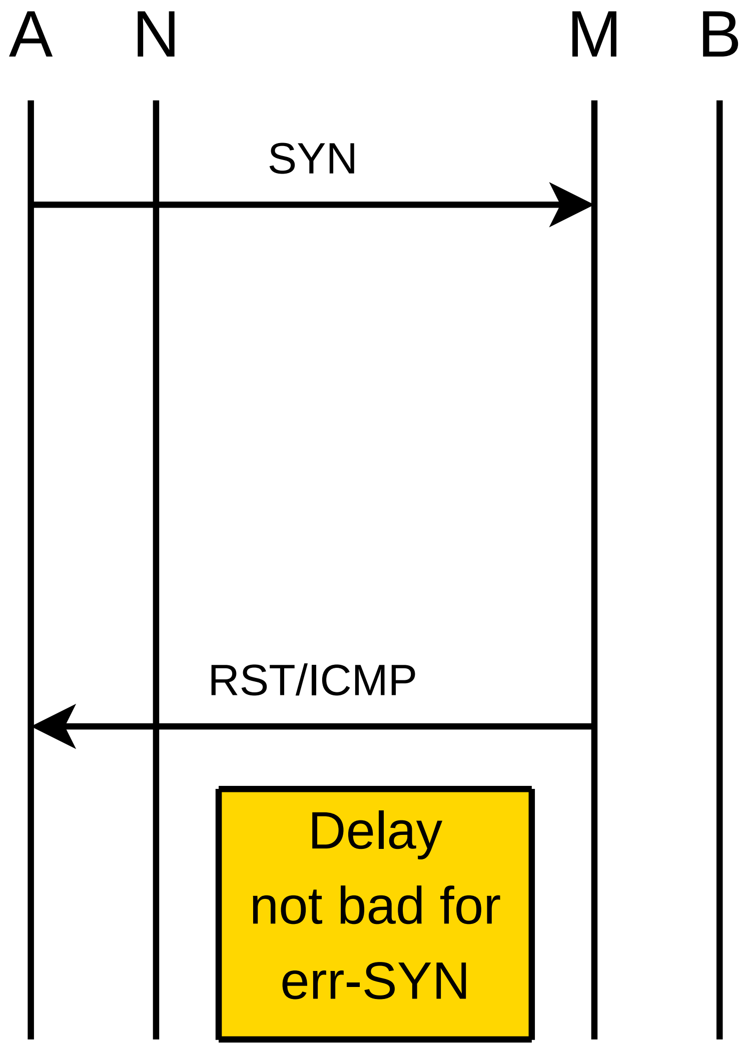

Unsolicited SYN

Opt. 1: Silently drop SYN (old WG consensus)

- What does TCPM think?

Opt. 2: Send ICMP, standardize new ICMP code

- Is this an option?

Opt. 3: Delay sending ICMP error

- Is 6s acceptable?[1]

[1] Variant allows for flexible timeouts if we can't decide on one

Unsolicited SYN: Option 4

Delayed Error 2

- Not bad for P2P
- Not bad for erroneous SYN
- Flexible timeouts
- Assumptions:
  - for P2P MUST do STUNT lookup first

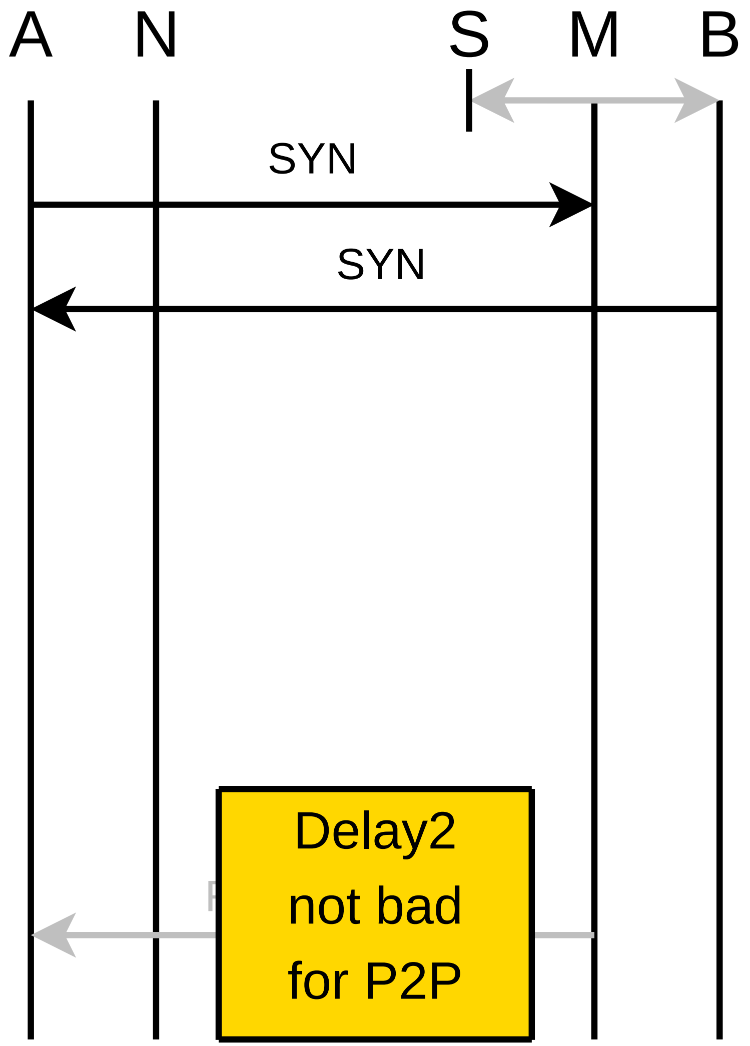

Delay2 not bad for P2P

Unsolicited SYN: Option 4

Delayed Error 2

- Not bad for P2P
- Not bad for erroneous SYN
- Flexible timeouts
- Assumptions:
  - for P2P MUST do STUNT lookup first

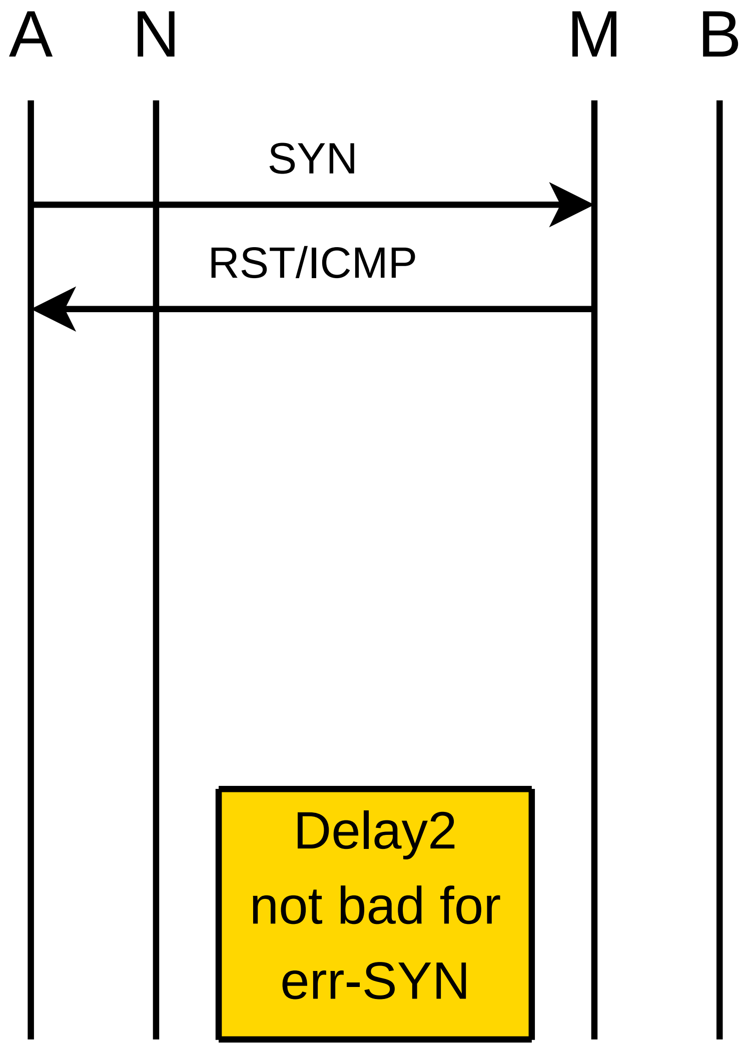

Open Issue: Port-range and ICMP

Port-Range Preservation

Does TCP need source port-range to be preserved ($<$1024, 1024–65535)?

ICMP Scope

Should ICMP handling of errors in response to TCP packets go in the ICMP draft or the TCP draft?
(to be discussed in ICMP slot)

Appendix

Extra slides

Appendix

Option 1

The NAT MUST silently drop unsolicited SYNs

Appendix

Option 2

If enabling P2P TCP apps is most important, a NAT MUST silently drop the SYN. If enabling quick diagnosis of network errors is most important, a NAT SHOULD signal an ICMP port unreachable. The behavior MAY be configurable by the administrator.

Appendix

Option 4

It is RECOMMENDED that a NAT respond to unsolicited SYN packets with an ICMP Port Unreachable error (Type 3, Code 3). If a NAT does so, it MUST delay the ICMP error by at least 6 seconds unless REQ-4a) applies. Furthermore, it MUST cancel this delayed ICMP if in that time it receives and translates an outbound SYN for the connection. If a NAT does not have resources to delay the ICMP error or chooses not to send it, the NAT MUST silently drop the unsolicited SYN.

a) If there is no active mapping that matches the unsolicited SYN, then the NAT SHOULD send the ICMP immediately.

Appendix

Option 3

It is RECOMMENDED that a NAT respond to unsolicited SYN packets with an ICMP Port Unreachable error (Type 3, Code 3). If a NAT does so, it MUST delay the ICMP error by at least 6 seconds. Furthermore, it MUST cancel this delayed ICMP if in that time it receives and translates an outbound SYN for the connection. If a NAT does not have resources to delay the ICMP error or chooses not to send it, the NAT MUST silently drop the unsolicited SYN.

Behave-App Recommendation

In order to establish TCP between two candidates[2],

- open 3 sockets (s1, s2, s3)
- bind() them all to the same local port
- listen(s1)
- connect(s2, peer.s1)
- connect(s3, peer.s3)

[2] think ICE