MPTCP Threat analysis draft-bagnulo-mptcp-threat-00

marcelo bagnulo

IETF76 – MPTCP WG

Scope

- Understand what ***additional*** vulnerabilities are added by the MPTCP extensions
- i.e. What attacks are possible in MPTCP that are ***NOT*** possible in current TCP
- It is out of the scope of the current analysis to identify and attempt to prevent threats that already exist in current TCP

Scope: Types of attackers

- On-path vs. Off-path
- On-path attackers
  - Full time on the path
  - Passive (man on the side)
  - Active:
    - Blocking packets
    - Changing packets

Previous work

- Threat analysis for:
  - MIPv6 RO (see RFC 4225)
  - Shim6 threat analysis (see RFC4218)
  - SCTP security analysis (see RFC5062)
- Relevant differences
  - In MIPv6/Shim6 what is a stake is the whole identity of the host, while in MPTCP/SCTP, only one connection is at stake…
    - So, we may want to be less conservative
    - **It is very important to keep in mind that re-using ID/loc/crypto state can change what it is at stake**

Basic MPTCP

- In order to understnad all possible threats, we will use a very basic MPTCP
- MPTCP will use the TCP 3 way handshake for the first flow
  - Will have a data seq number that is synch in that exchange
  - The address pair used in this 3 way handshake are the application identifiers (maybe learnt through DNS or DNSSEC, passed through a referral)
    - May have some level of trust embedded

Basic MPTCP (cont)

- Once the first flow is established, MPTCP will use extensions for adding addresses
  - Implict: the address is convey in the source IP address field
  - Explicit: an option carrying an address list.
- We assume that MPTCP will distribute the load accross all address pairs, based on congestion.

Scenario

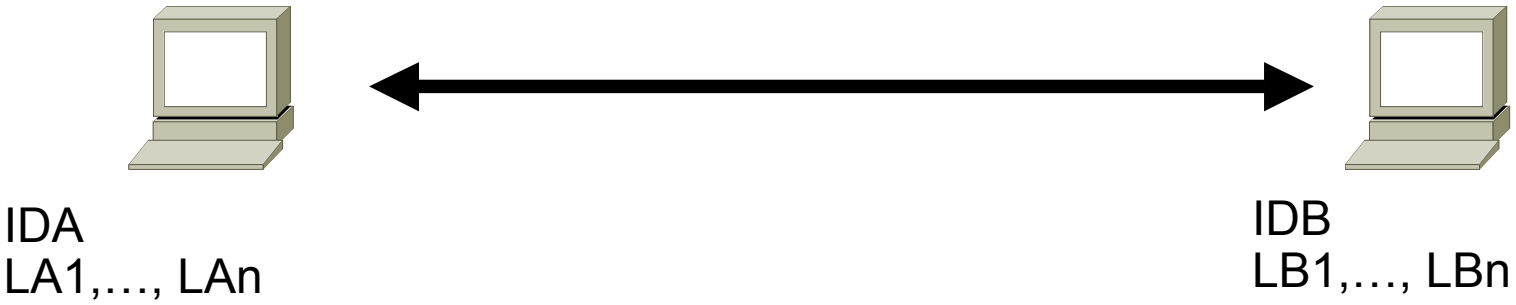

Scenario

IDA
LA1,…, LAn

IDB
LB1,…, LBn

IDB
LB1,…, LBn

IDA
LA1,…, LAn

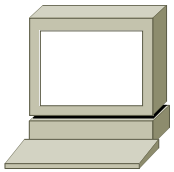

IDX
LX1,…, LXn

Redirection attacks

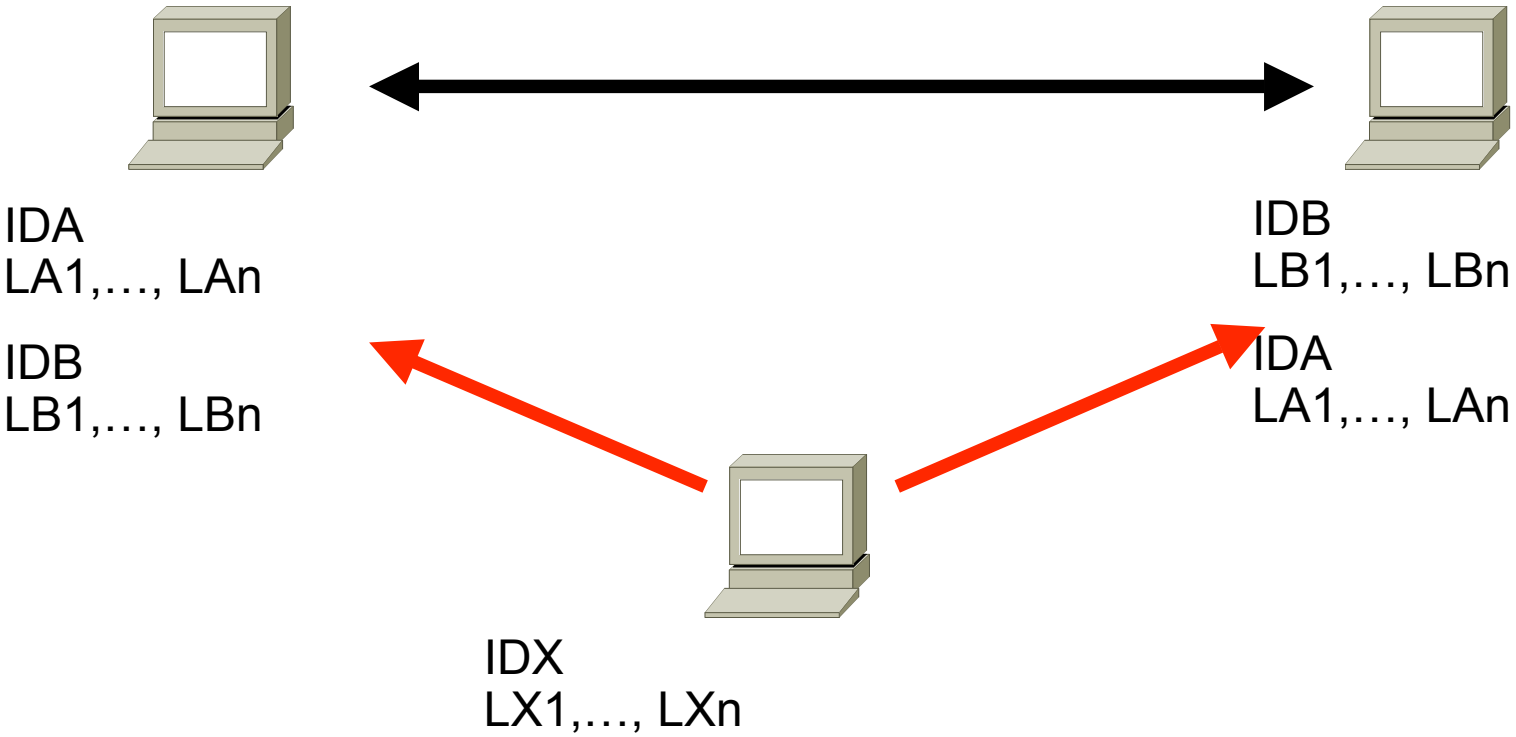

Flooding

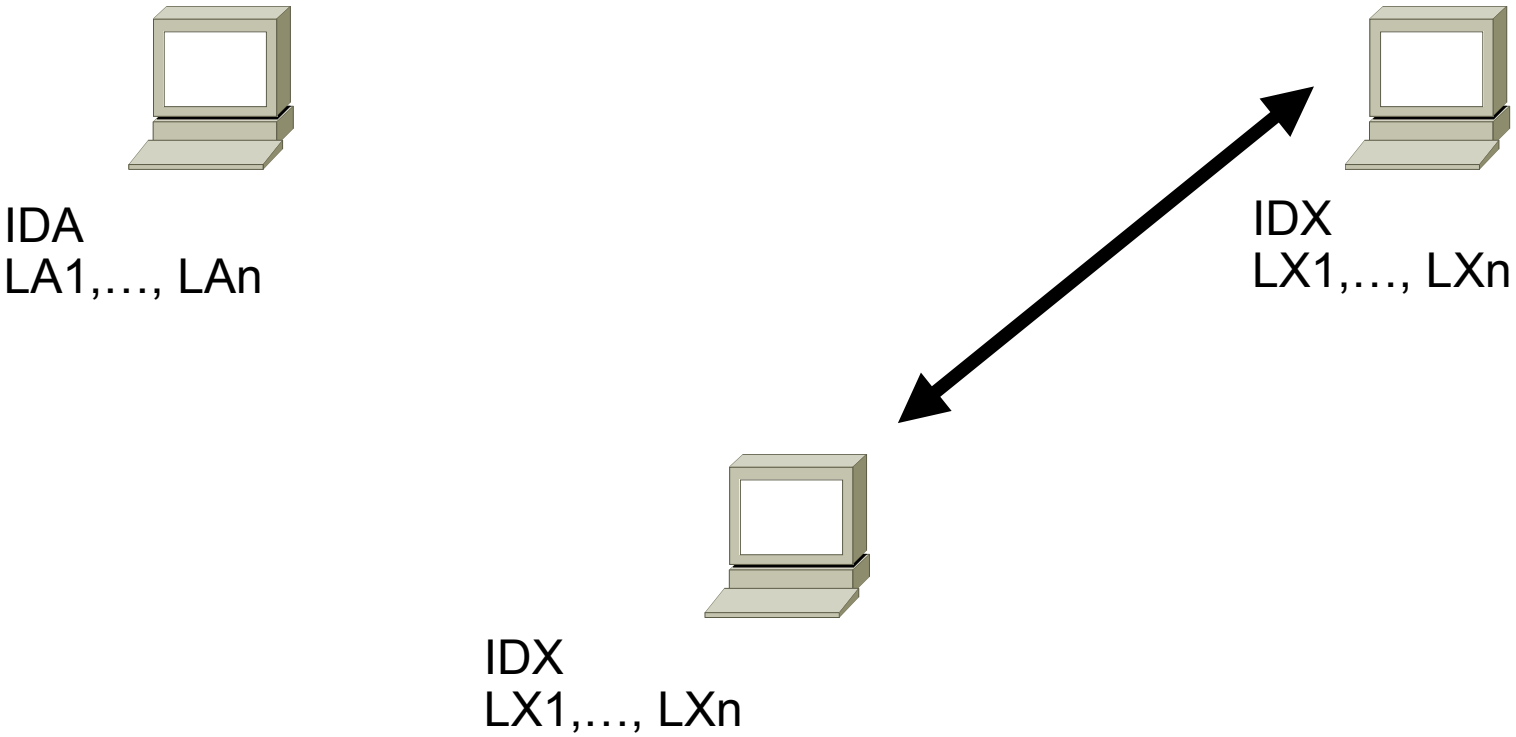

Flooding

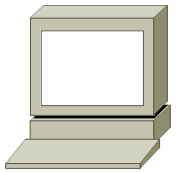

IDA
LA1,…, LAn

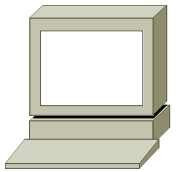

IDX
LX1,…LAi,…, LXn

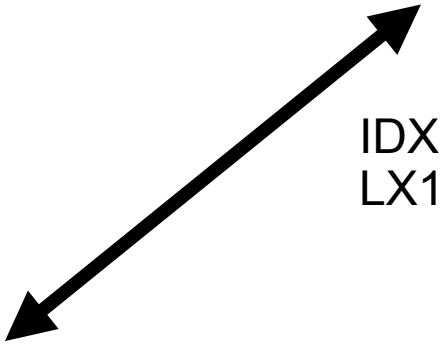

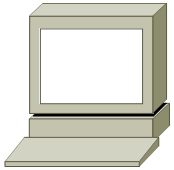

IDX
LX1,…, LXn

Attacker adds Lai
Implicit mode: only works if no ingress fitlering
Explicit mode always works

Flooding

IDA
LA1,…, LAn

IDX
LX1,…LAi,…, LXn

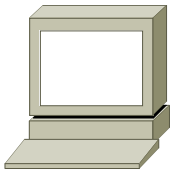

IDX
LX1,…, LXn

Server sends traffic to target
Attacker closes other addresses or pretends there is congestion
Attacker send ACKs for data sent to target
How is data acked? Only from acks containing the crrect src and dst?
Target will issue a RST
How does the server reacts upon reception of both ACKs and RST for the same data? Is the flight size enough to

Additional threat

- In TCP, an on path attacker can launch a flooding attack to the infrastructure along the path, but off path attackers can’t.
- MPTCP security must prevent off path attackers to launch flooding attacks.

Flooding

- Standard protection against flooding attack is a reachability test.
  - Before start sending packets to a new locator, a reachability test is perform, exchanging some connection identifier.
  - If the identifier does not correspond to any existing connection, the victim/target/receiver will reject the connection and the attack will be prevented

Flooding and MPTCP

- If MPTCP performs a 3-way handshake per new flow and they identify the connection
- This provides the reachability check required to prevent flooding attacks
- It is very important to NOT send data without a prior reachability check

Connection Hijacking

IDA
LA1,…, LAn

IDB
LB1,…, LBn

IDB
LB1,…, LBn

IDA
LA1,…,…, LAn

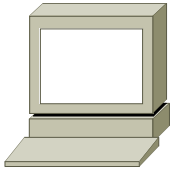

IDX
LX1,…, LXn

Connection Hijacking

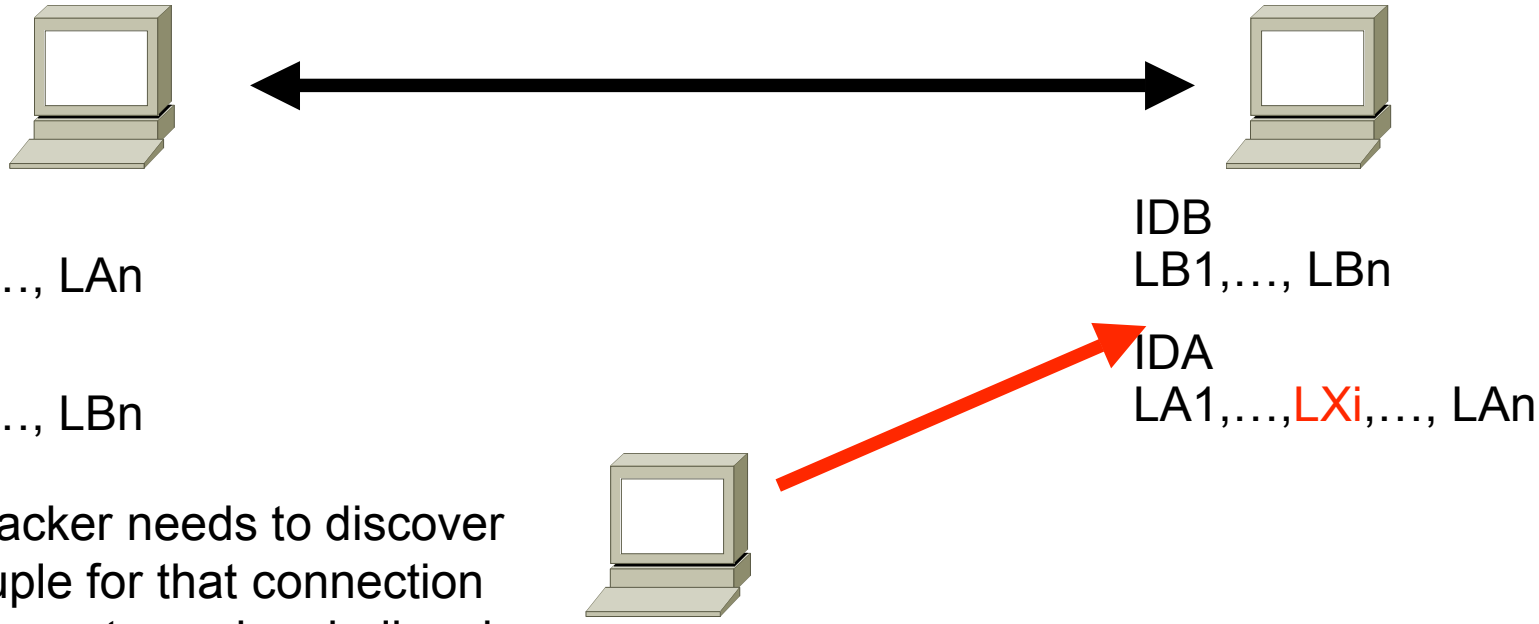

The attacker needs to discover the 4-tuple for that connection
-The src port may be challenging

Connection Hijacking

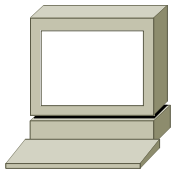

IDA
LA1,…, LAn

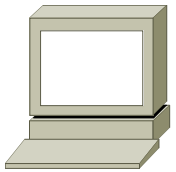

IDB
LB1,…, LBn

IDA
LA1,…,LXi,…, LAn

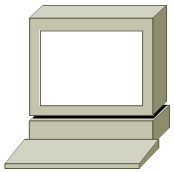

IDX
LX1,…, LXn

Packets flowing from B to the attacker.
Node B will distribute the load based on congestion, so some packet will flow to A
-Node A won't receive all packets
In order to do a full hijacking, the attacker should remove LAi from the connection
-- change on implicit or explicit mode

Connection Hijacking

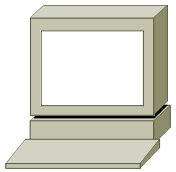

IDA
LA1,…, LAn

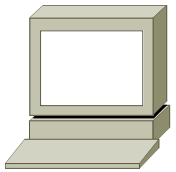

IDB
LB1,…, LBn

IDA
LA1,…,LXi,…, LAn

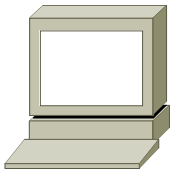

IDX
LX1,…, LXn

Packets flowing from attacker to node B
The attacker will be able to inject data (as long as the seq# is valid
Node A will still be able to inject data
Attacker can try to remove IP address es from A

Additional Threat

- In current TCP, an on-path attacker can launch a hijacking attack, but an off-path attacker can’t.
  - It may be able to inject some packets (depending on the seq# and ingress filtering), but certainly cannot receive packets
- So, MPTCP security must prevent off path atackers to perform hijacking attacks

Hijacking and MPTCP with cookie based security

- MPTCP can use a combination of seq# and cookie for security. (as in draft-ford-mptcp-multiaddressed)
  - By Seq# i refer to the data seq# (not the one per flow, but the one of the data)
  - They are both exchanged in the first 3 way exchange, when the ULID pair is defined for the connection.
- So what residual hijacking attacks can be performed with this protection?

Time-shifted/future attacks

- A time-shifted attack is an attack where:
  - The attacker is on-path during a period of time and obtains information (e.g. The cookie and the seq#) or even installs state if needed.
  - Then the attacker leaves the on path location
  - The attakcs continues even after the attacker left the on path position
- Current TCP is not vulnerable to time-shifted attacks
  - i.e. When the attacker leaves the position, it no longer receives the packets of the TCP connection

Time shifted attack in MPTCP Flavour 1

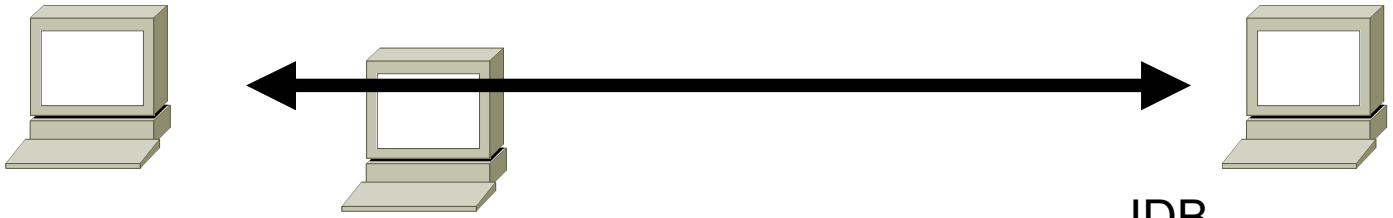

IDA
LA1,…, LAn

IDB
LB1,…, LBn

Attacker on path learns cookie and seq#

IDB
LB1,…, LBn

IDA
LA1,…, LAn

Any side initiates the connection

Time shifted attack in MPTCP Flavour 1

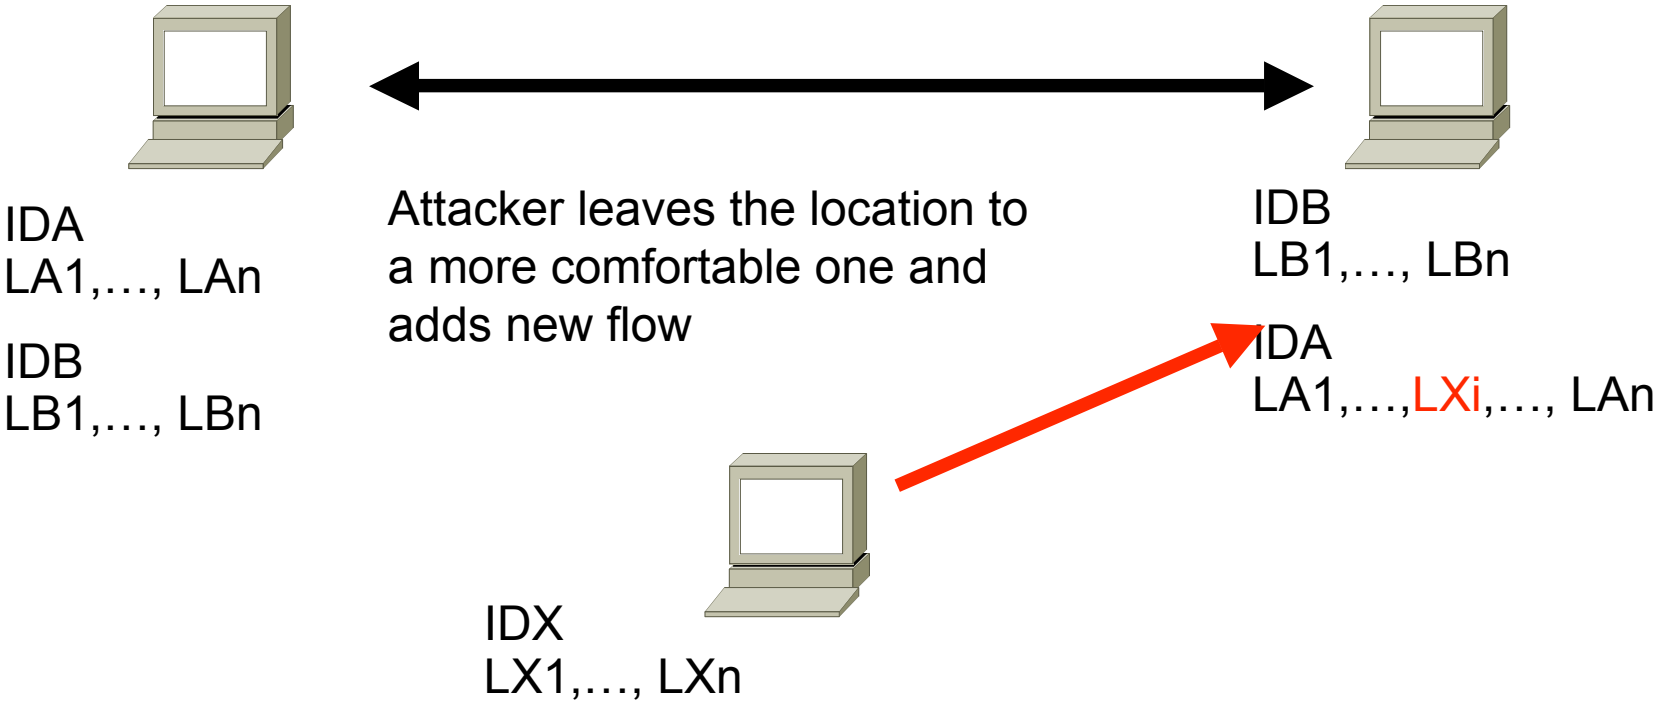

Time shifted attack in MPTCP Flavour 2

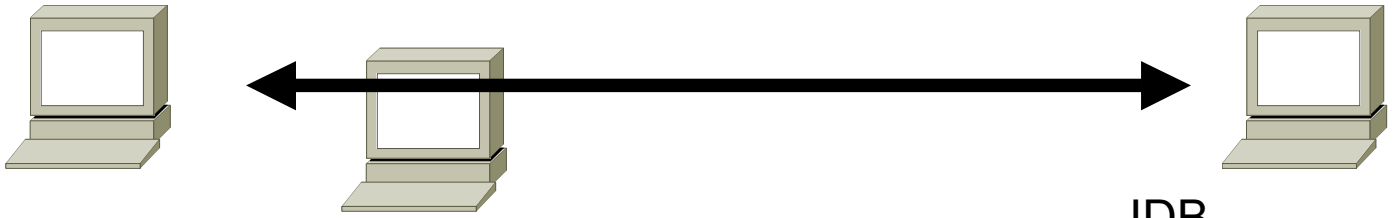

IDA
LA1,…, LAn

IDB
LB1,…, LBn

Attacker on path learns
cookie and seq#
Then it closes this flow ☺

IDB
LB1,…, LBn

IDA
LA1,…, LAn

Any side initiates the connection

Time shifted attack in MPTCP Flavour 2

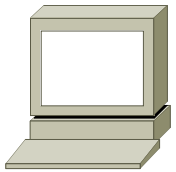

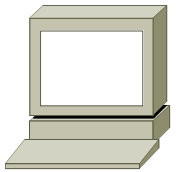

IDA
LA1,…, LAn

IDB
LB1,…, LBn

Attacker leaves the location to a more comfortable one and adds new flow

IDB
LB1,…, LBn

IDA
LA1,…,LXi,…, LAn

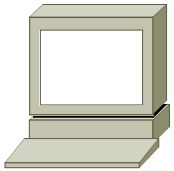

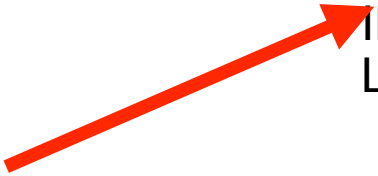

Impact:
Complete hijacking of the connection

IDX
LX1,…, LXn

Time shifted attack in MPTCP Flavour 3

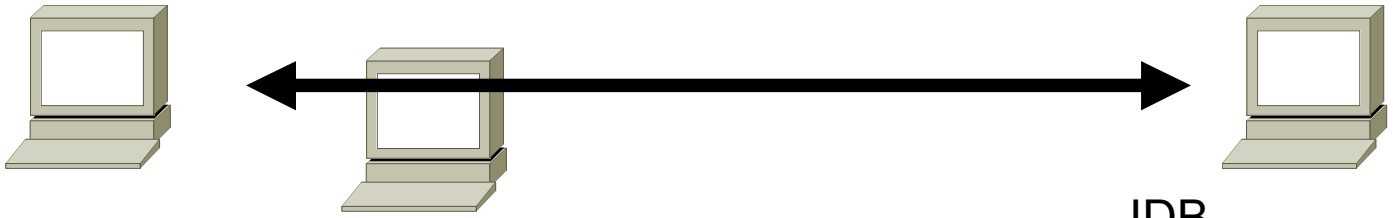

IDA
LA1,…, LAn

IDB
LB1,…, LBn

Attacker initiates the connection from the topologically correct location

IDB
LB1,…, LBn

IDA
LA1,…, LAn

Time shifted attack in MPTCP Flavour 3

IDA
LA1,…, LAn

IDB
LB1,…, LBn

Attacker leaves the location to a more comfortable one and adds new flow

IDB
LB1,…, LBn

IDA
LA1,…,LXi,…, LAn

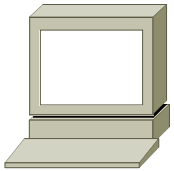

IDX
LX1,…, LXn

Impact:
Complete impersonation of the victim