Proxy Mobile IPv6 Extensions to Support Flow Mobility
draft-bernardos-netext-pmipv6-flowmob-01

Carlos J. Bernardos (Ed.) – Universidad Carlos III de Madrid

Beijing, NETEXT WG, 2010-11-09

Changes from version -00

- Added text on the prefix deployment models
- Added text about the partial handoff scenario
- Corrections here and there

Goals and requirements

- Define PMIPv6 extensions to allow to move flows among different simultaneously attached MN interfaces
- Analyze the different prefix deployment scenarios to be supported
- LMA is the controlling entity
  - The solution defines the signaling between MAG and LMA
  - The specifics on how the network nodes obtain the policies are out of scope
- The MN is equipped with one logical interface as described in draft-ietf-netext-logical-interface-01
  - We don't support flow mobility across different logical interfaces

Prefix (deployment) models (I)

- Multiple prefix models under a flow mobility solution may work:
  1. At the time of a new attachment, the MN obtains a new prefix or a new set of prefixes. This is the default behavior with RFC 5213
  2. At the time of a new attachment, the MN obtains the same prefix or the same set of prefixes as already assigned to an existing session
  3. At the time of a new attachment, the MN obtains a combination of prefix(es) in use and new prefix(es). This is a hybrid of the above two scenarios

Prefix (deployment) models (II)

- Scenario 2 needs extensions to RFC 5213 signaling at the time of a new attachment
  - No further signaling required between LMA and MAG
- Scenario 1 requires flow mobility signaling for relocating flows between the different attachments
- MAGs should be aware of the prefixes for which the MN is going to receive traffic
  - Signaling is required if involved prefixes are not limited to those delegated to the MAG upon attachment of the new interface

Flow mobility scenarios

- Flow mobility signaling takes place whenever the LMA decides to move a flow from one access to another. At this point, either the prefix corresponding to the flow is already valid on the target MAG, or it needs to be signaled
  - If already valid, LMA just moves the flow: “shared prefix” scenario
  - If not valid, LMA informs the MAG: “unique prefix” scenario
    - By default at prefix level, granularity MAY include detailed flow descriptors

Unique prefix per physical interface

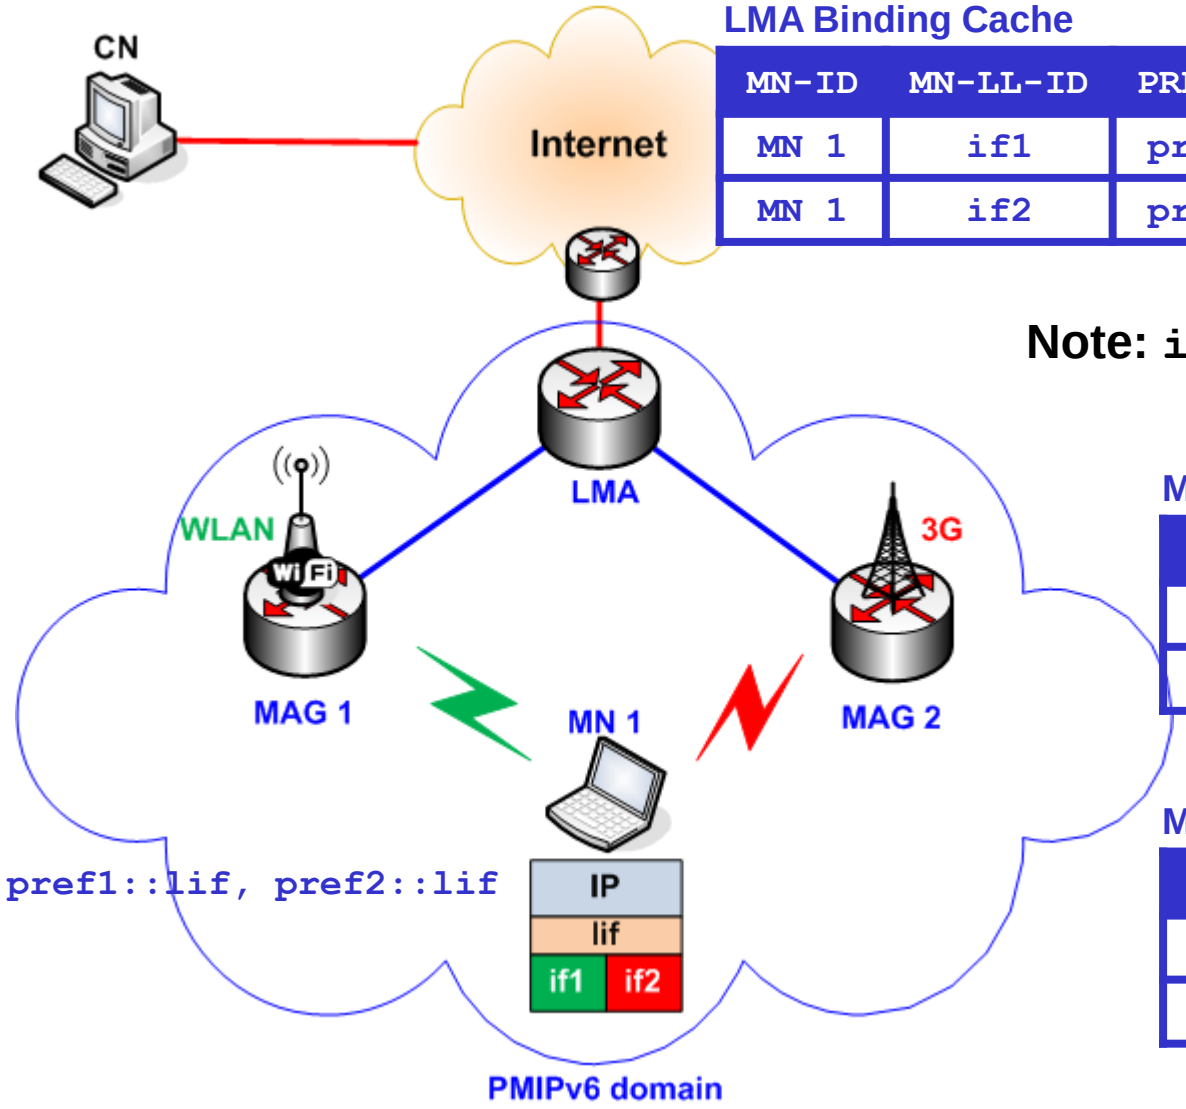

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref2	MAG 2

Note: `if1` and `if2` may be the same (`lif`)

MAG1 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
::/0	LMA

MAG2 routing state

dest	next hop
pref2::/64	p2p-iface-with-MN1
::/0	LMA

Unique prefix per physical interface

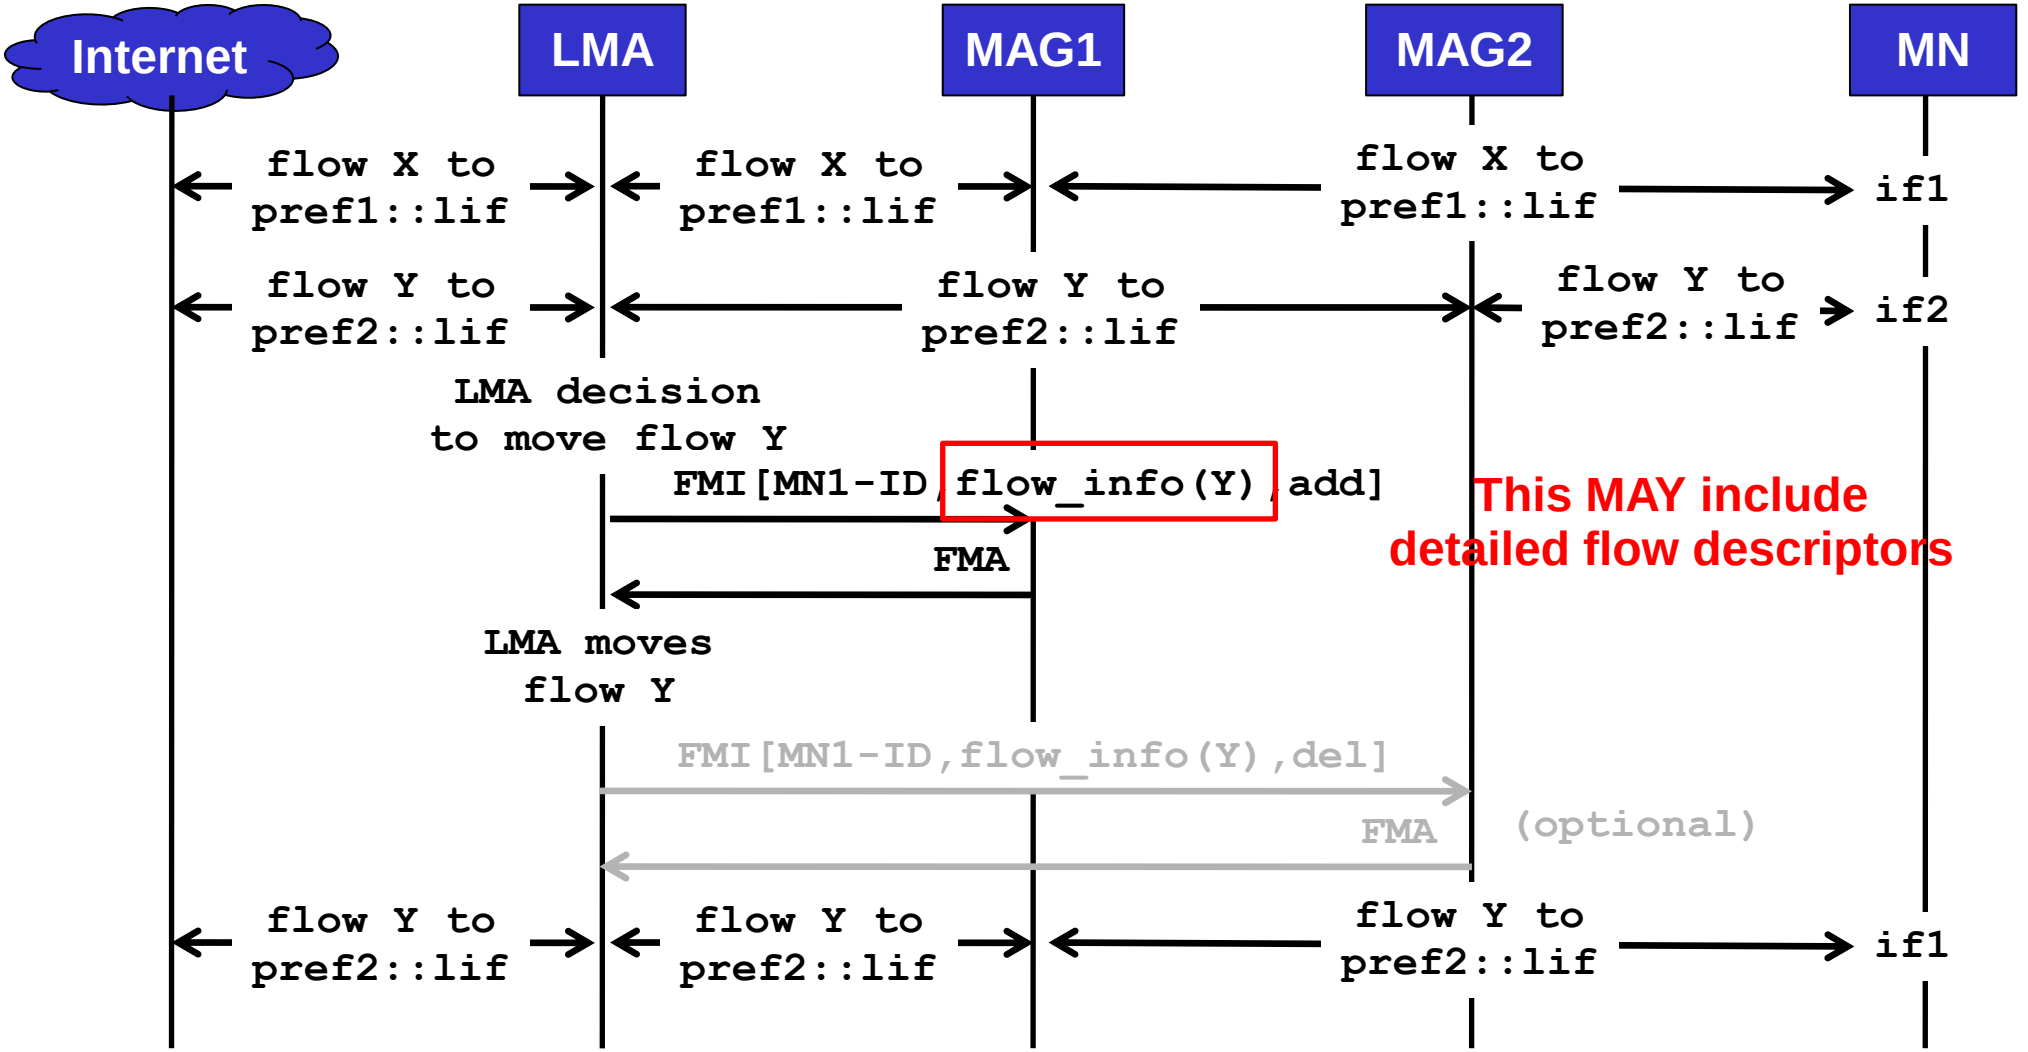

Unique prefix per physical interface

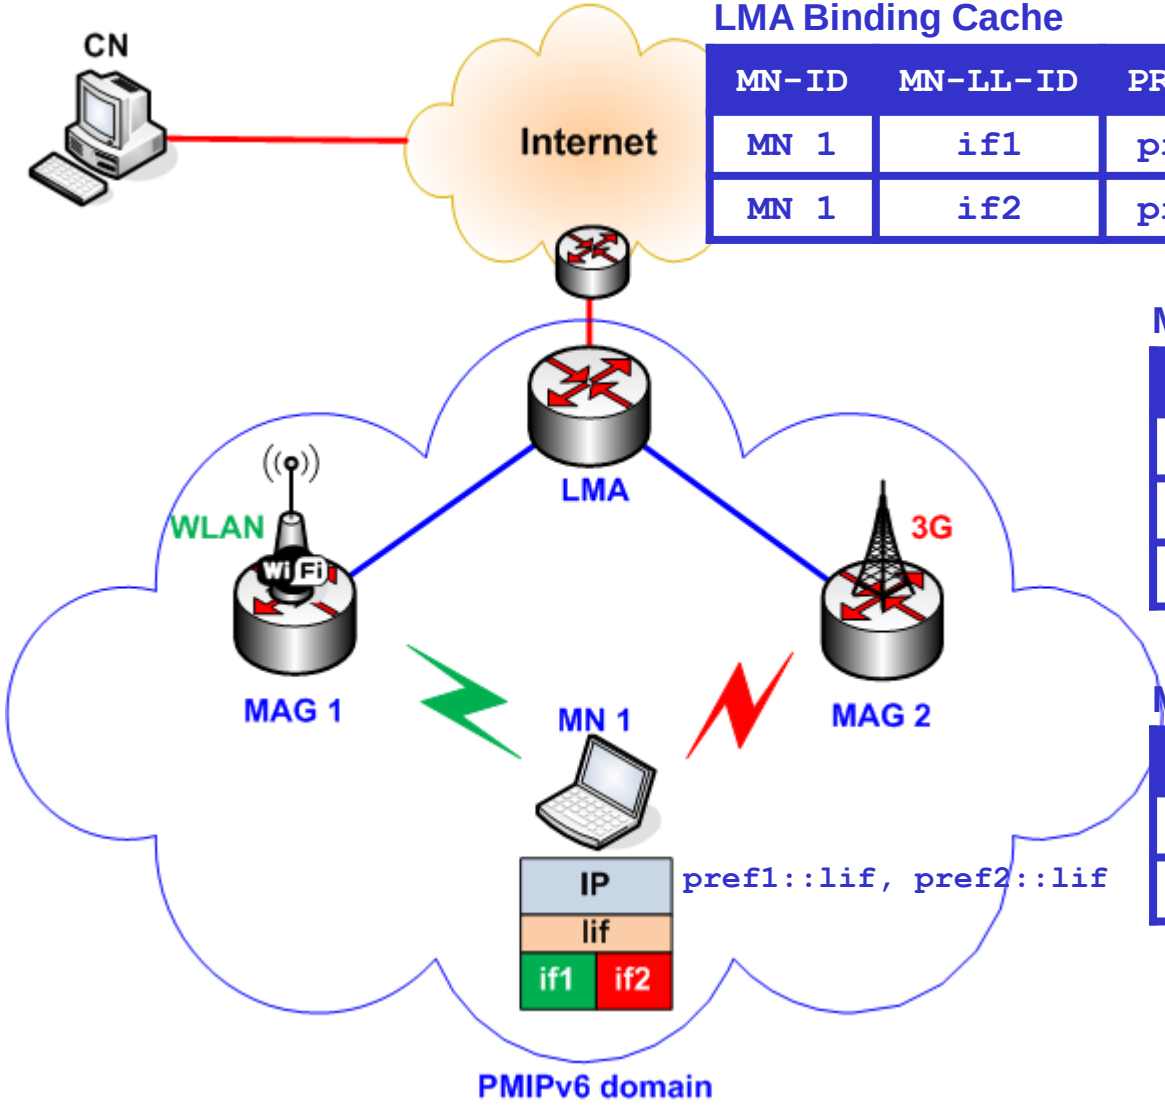

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref2	MAG 2

LMA flowmob state

MN-ID	flow	MAG
MN 1	flow X	MAG 1
MN 1	flow Y	MAG 1

MAG1 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
::/0	LMA
pref2::/64	p2p-iface-with-MN1

MAG2 routing state

dest	next hop
pref2::/64	p2p-iface-with-MN1
::/0	LMA

Unique prefix per physical interface
(partial handoff to a new interface)

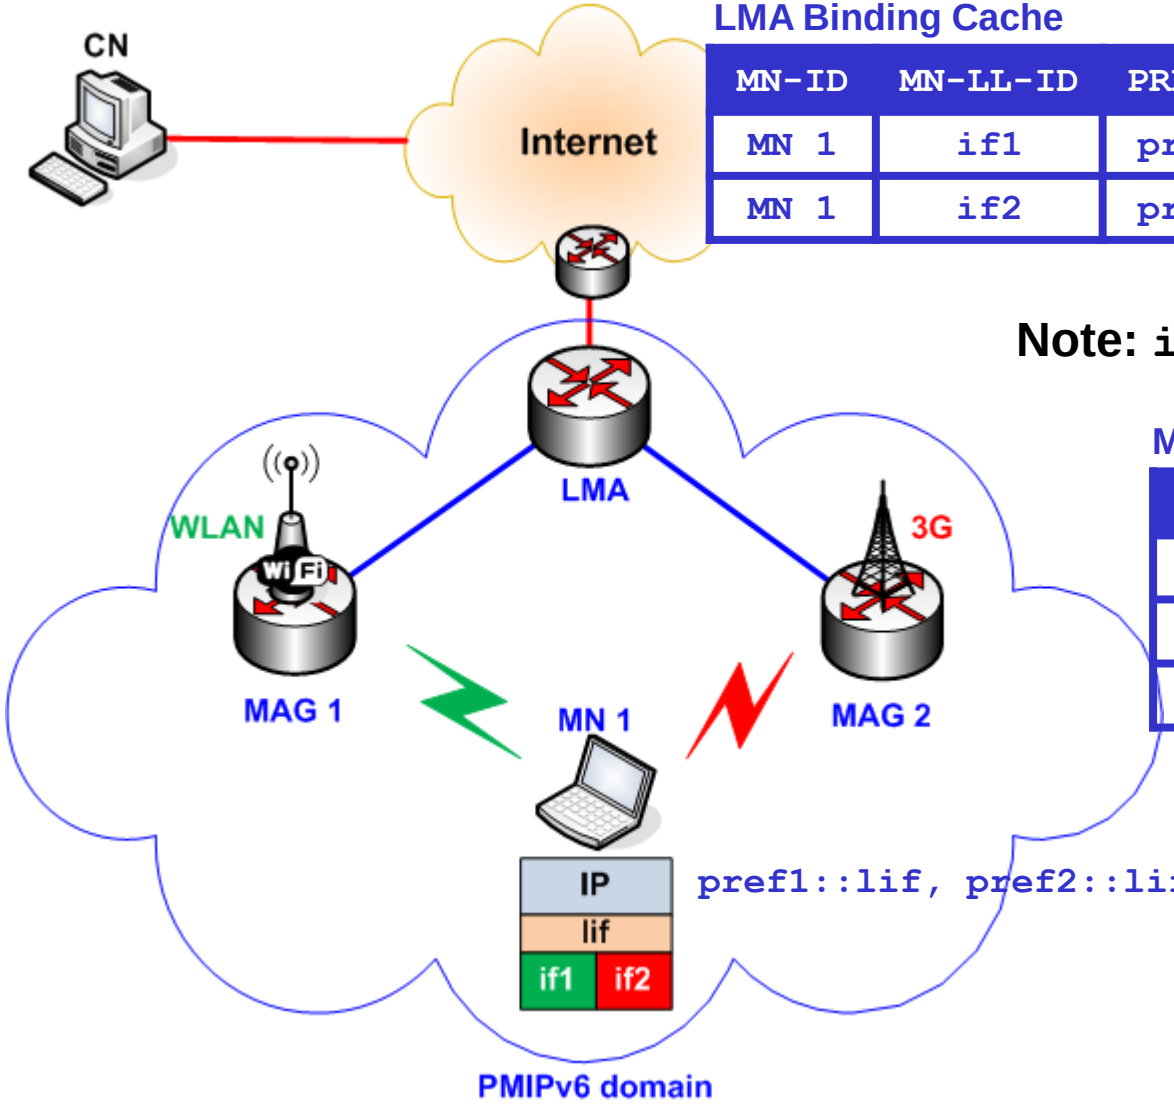

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref2	MAG 1

Note: `if1` and `if2` may be the same (`lif`)

MAG1 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
pref2::/64	p2p-iface-with-MN1
::/0	LMA

Unique prefix per physical interface
(partial handoff to a new interface)

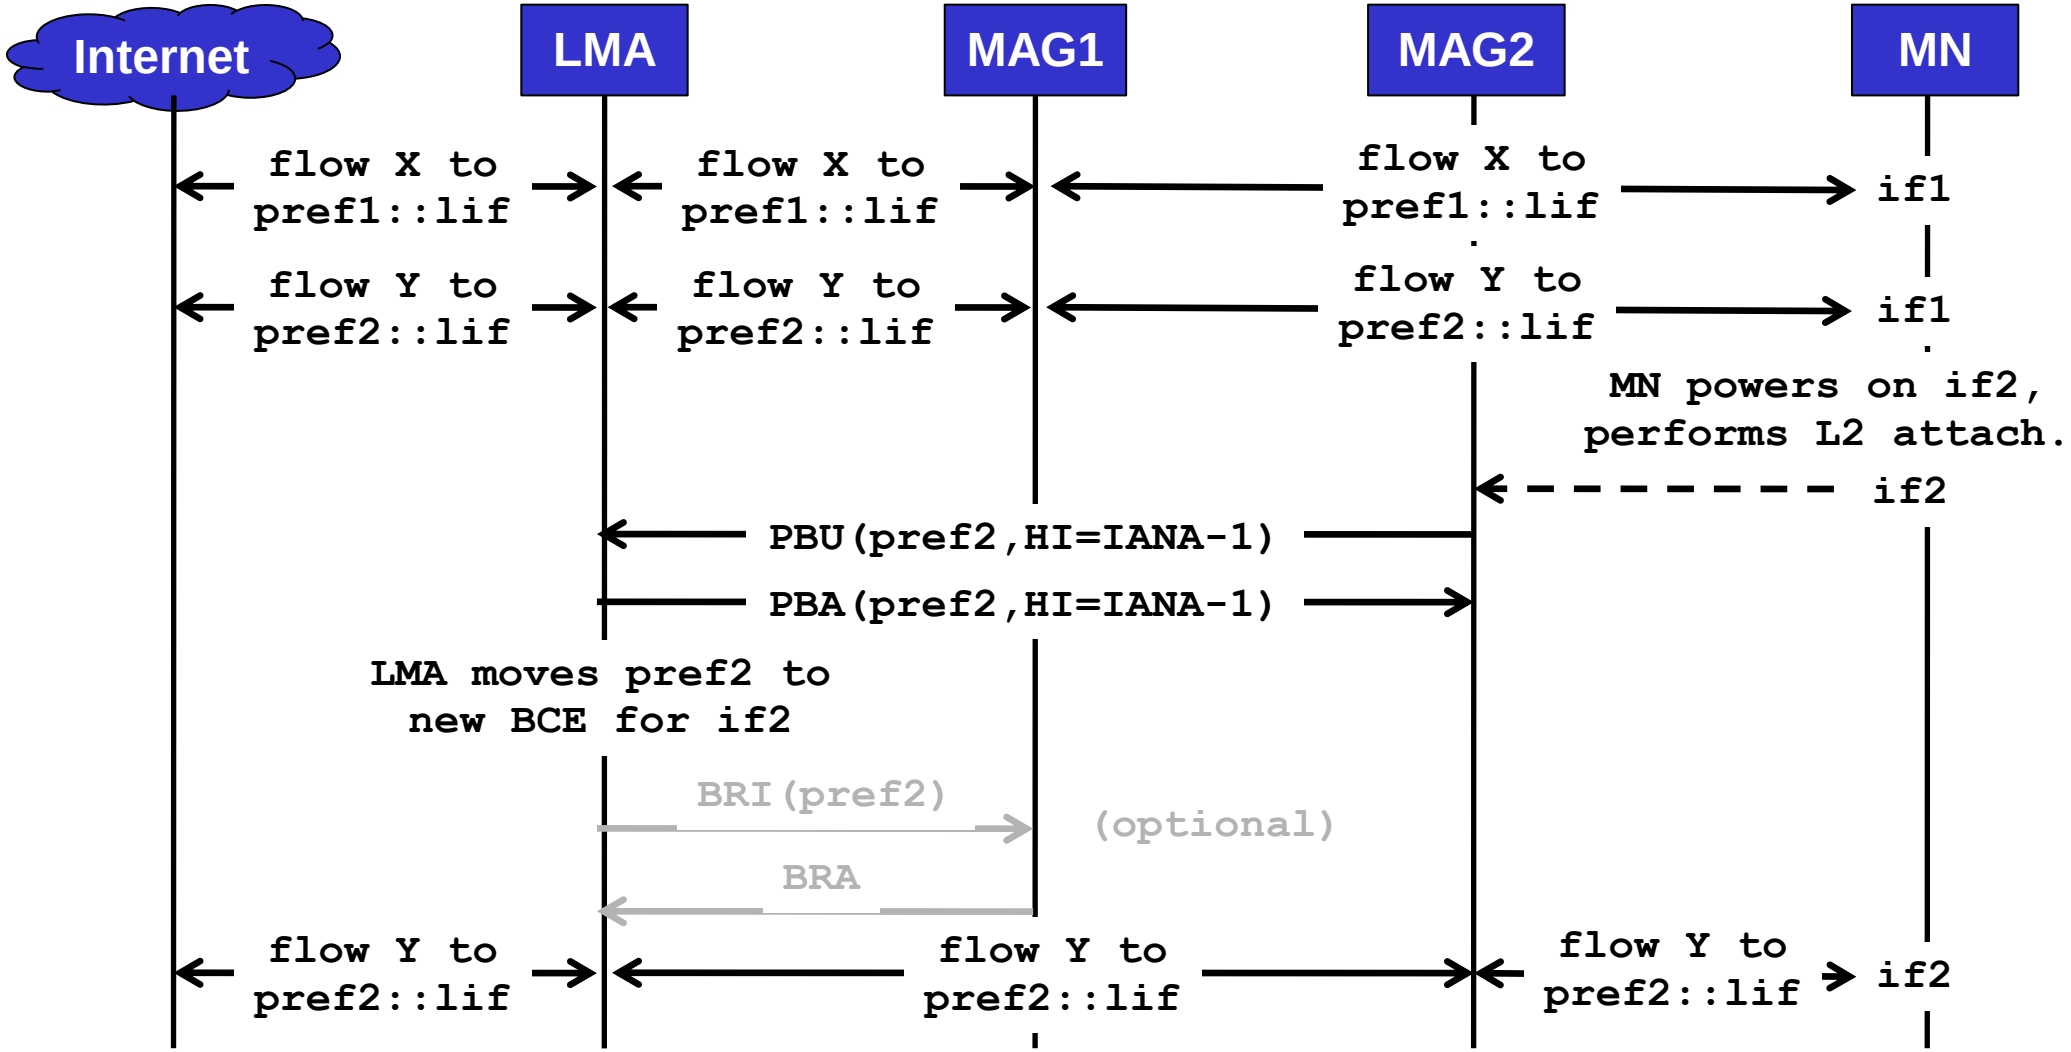

Unique prefix per physical interface
(partial handoff to a new interface)

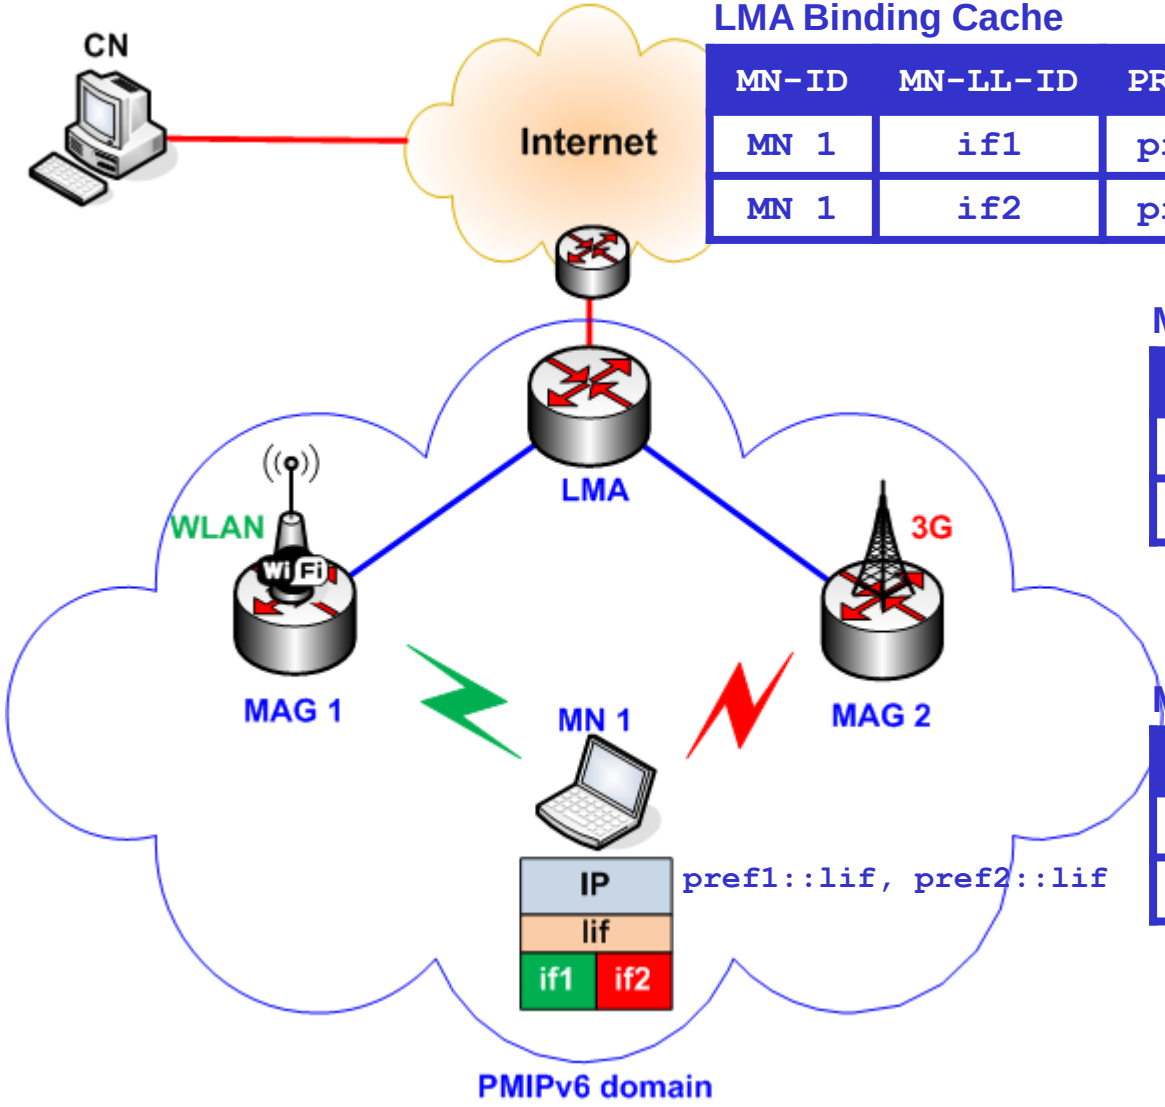

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref2	MAG 2

MAG1 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
::/0	LMA

MAG2 routing state

dest	next hop
pref2::/64	p2p-iface-with-MN1
::/0	LMA

Shared prefix across physical interfaces

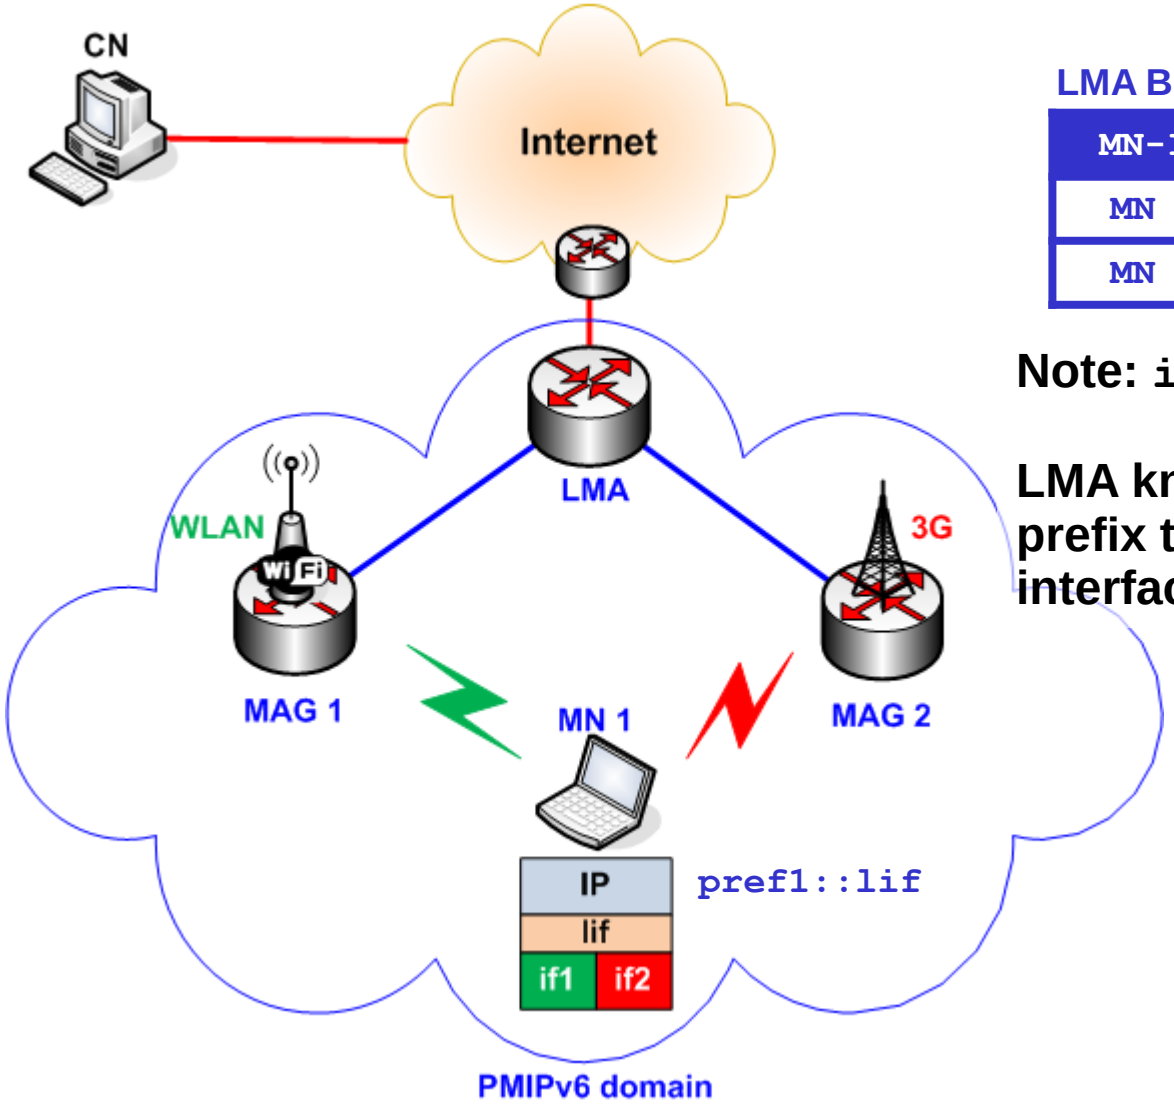

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref1	MAG 2

Note: `if1` and `if2` may be the same (`lif`)

LMA knows that it has to assign the same prefix to upon attachement of different interfaces (TBD)

Shared prefix across physical interfaces

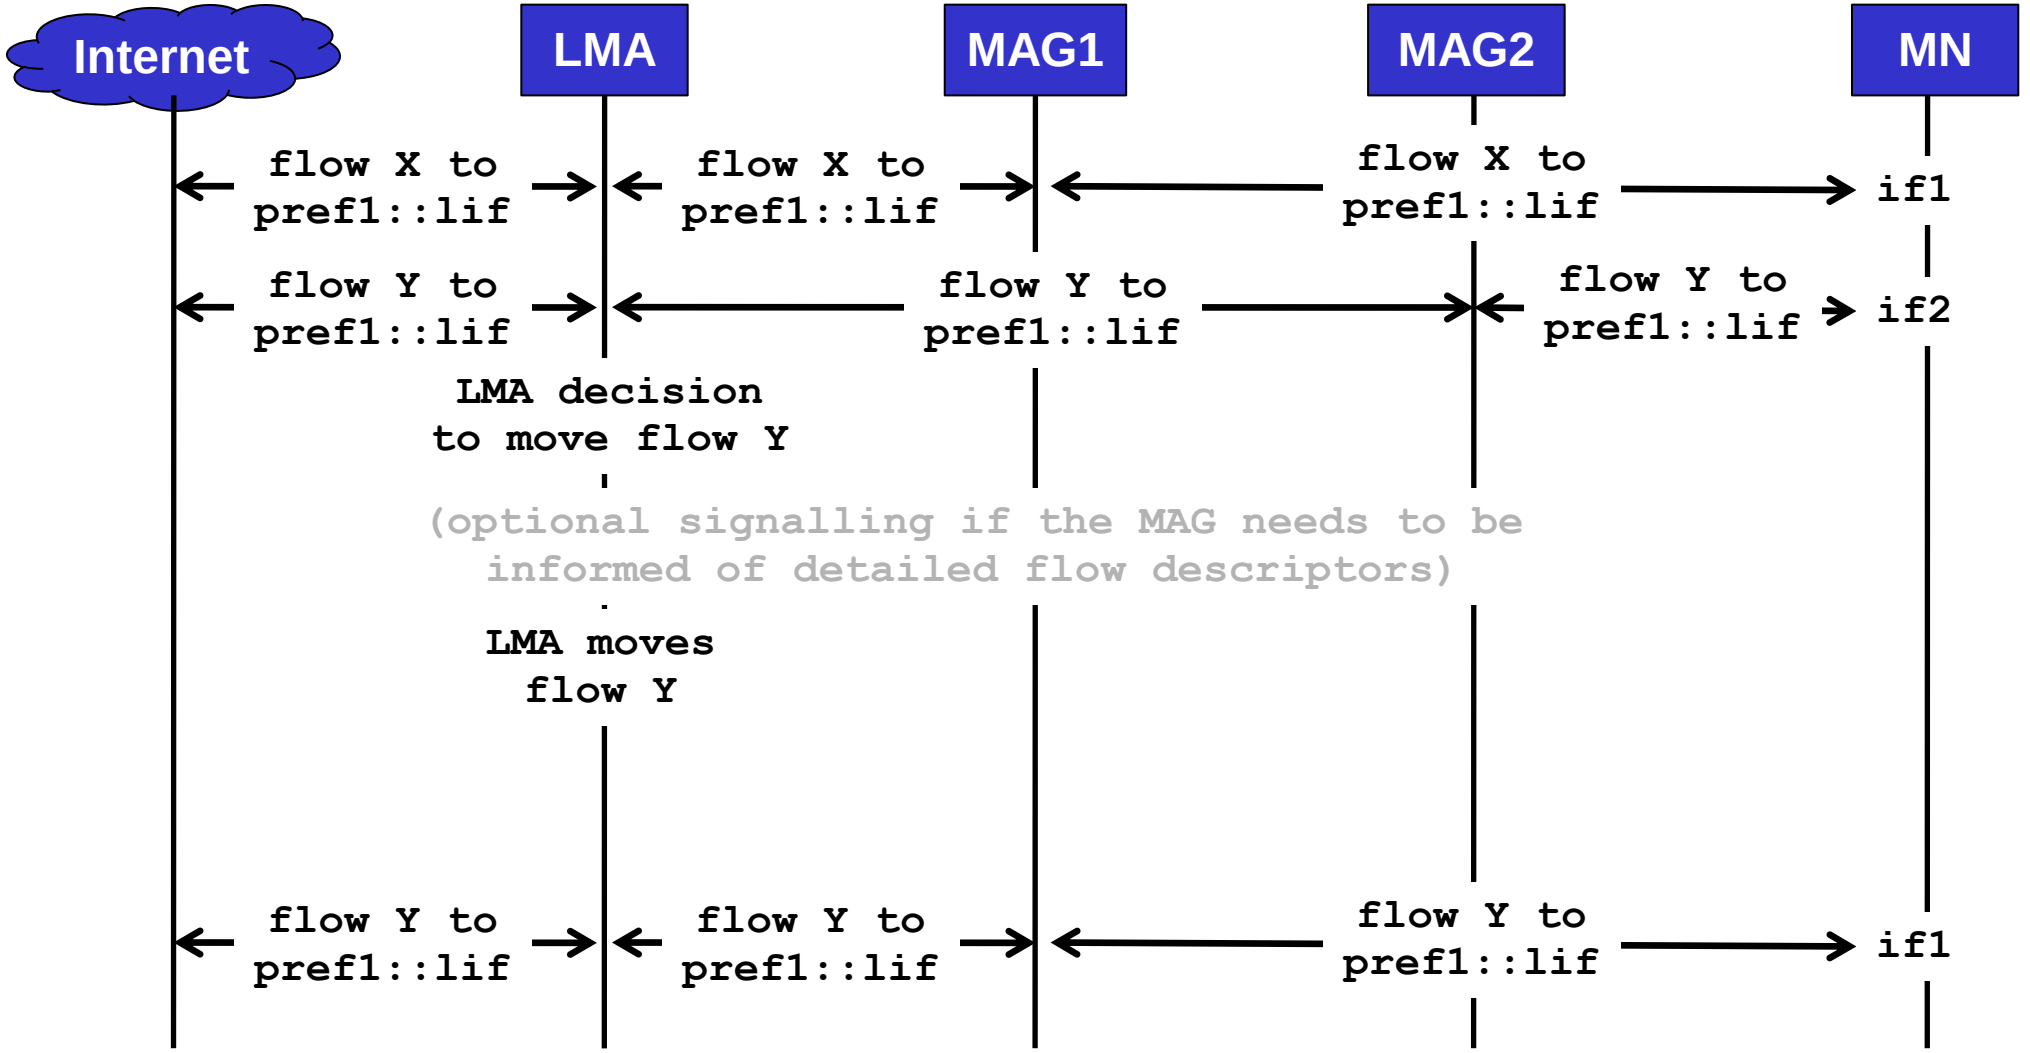

Shared prefix across physical interfaces

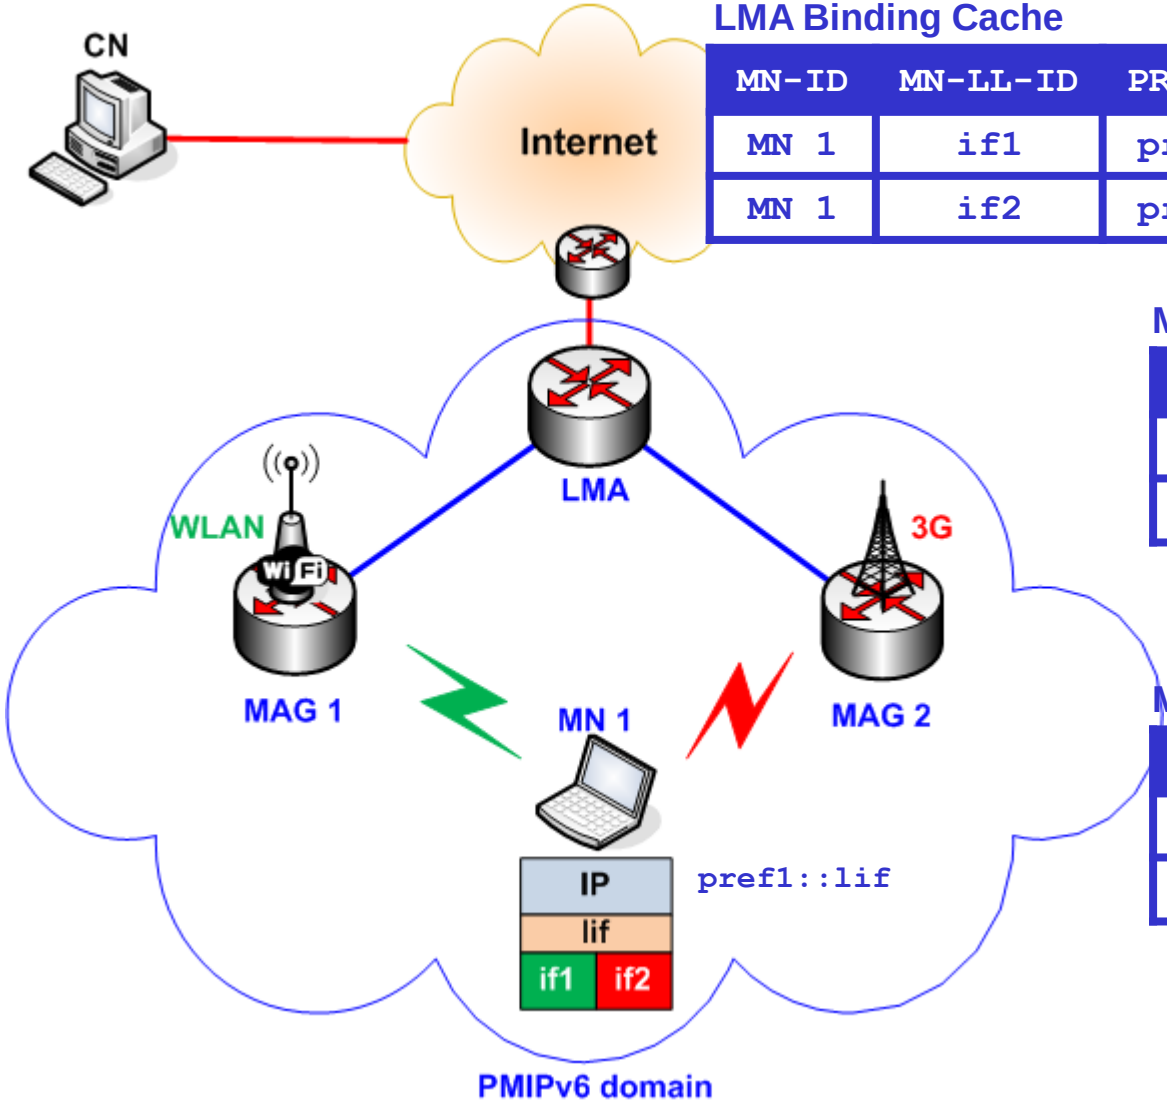

LMA Binding Cache

MN-ID	MN-LL-ID	PREFIX	MAG
MN 1	if1	pref1	MAG 1
MN 1	if2	pref1	MAG 2

LMA flowmob state

MN-ID	flow	MAG
MN 1	flow X	MAG 1
MN 1	flow Y	MAG 1

MAG1 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
::/0	LMA

MAG2 routing state

dest	next hop
pref1::/64	p2p-iface-with-MN1
::/0	LMA

Next steps

- Some design choices are still open for discussion
  - Feedback from the WG very welcome
- Ask for WG adoption