PATTERNS IN NETWORK ARCHITECTURE:

VERTICAL COMPOSITION

OR

LAYERING

PURPOSE:

TO BUILD A NETWORK WITH A LARGER SPAN OUT OF SMALLER, HETEROGENEOUS NETWORKS

The Internet has its own name space, protocols, headers, routing, etc.

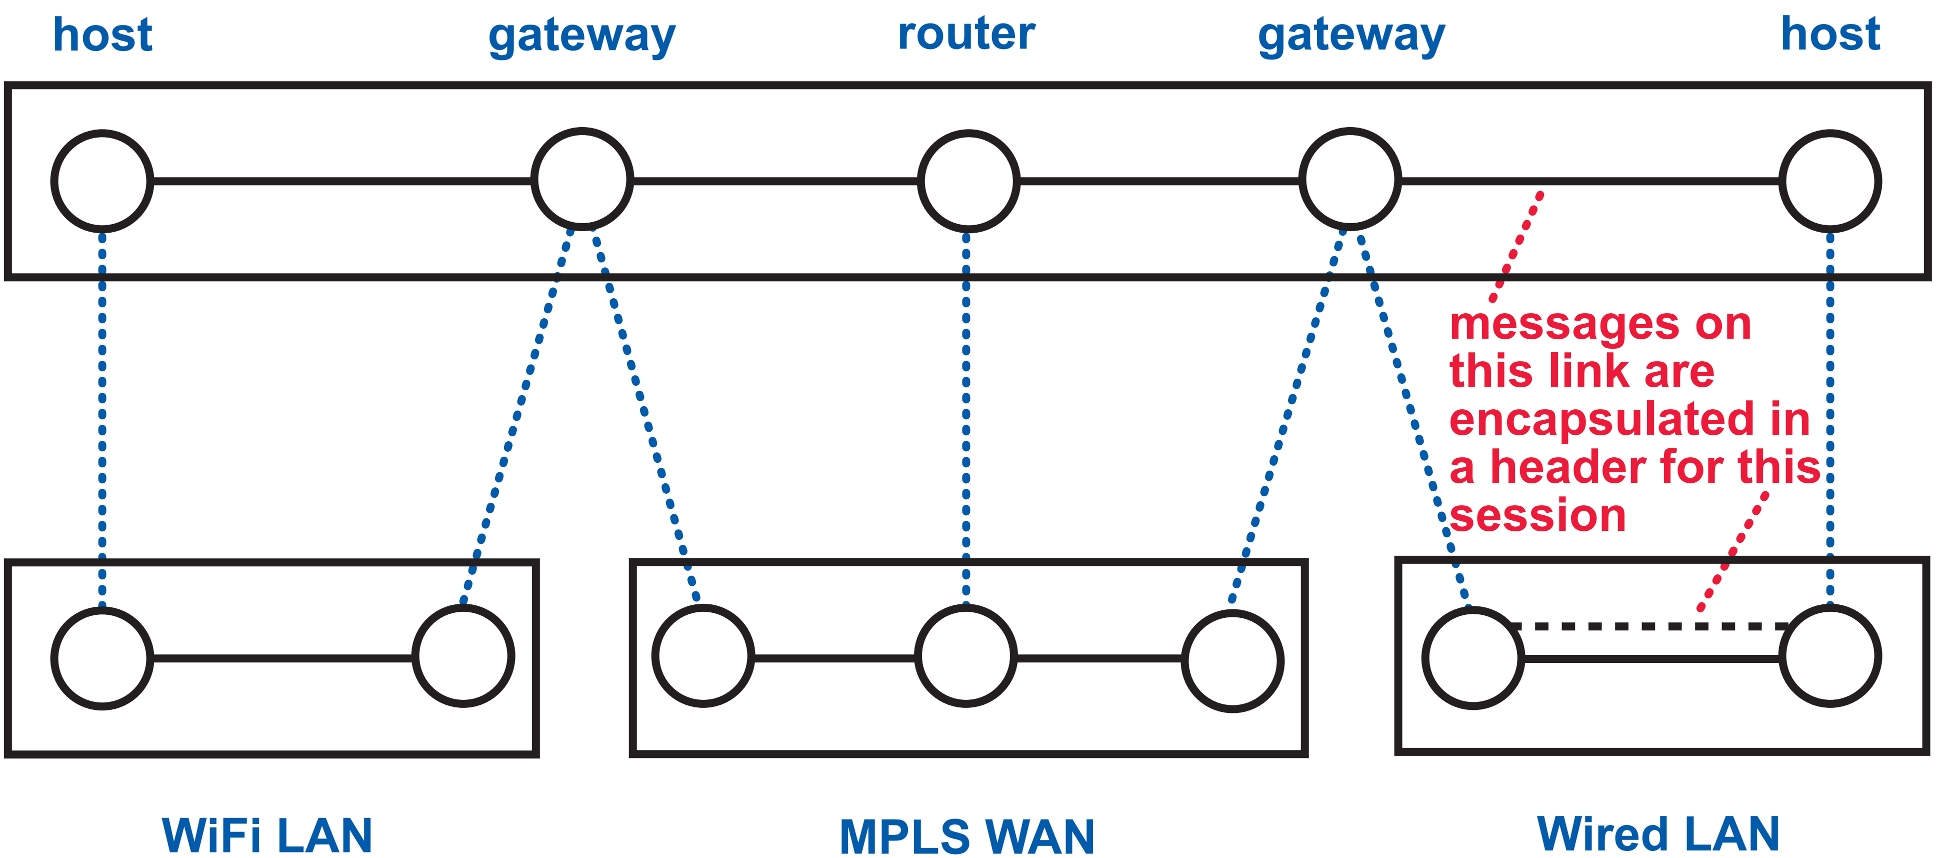

Each underlay network has its own name space, protocols, headers, routing, etc.

OTHER PURPOSES

TO BUILD A NETWORK WITH BETTER PERFORMANCE OR RELIABILITY ON TOP OF AN EXISTING NETWORK

for example, Resilient Overlay Networks

TO SHARE THE RESOURCES OF AN EXISTING NETWORK IN A DISCIPLINED WAY

for example, Virtual LANs

TO BUILD A NETWORK WITH LINKS THAT OFFER A SUPERIOR COMMUNICATION SERVICE

for example, Virtual Private Networks

during the semester, be on the lookout for other purposes!

ATTACHED NODES AND LOCATIONS

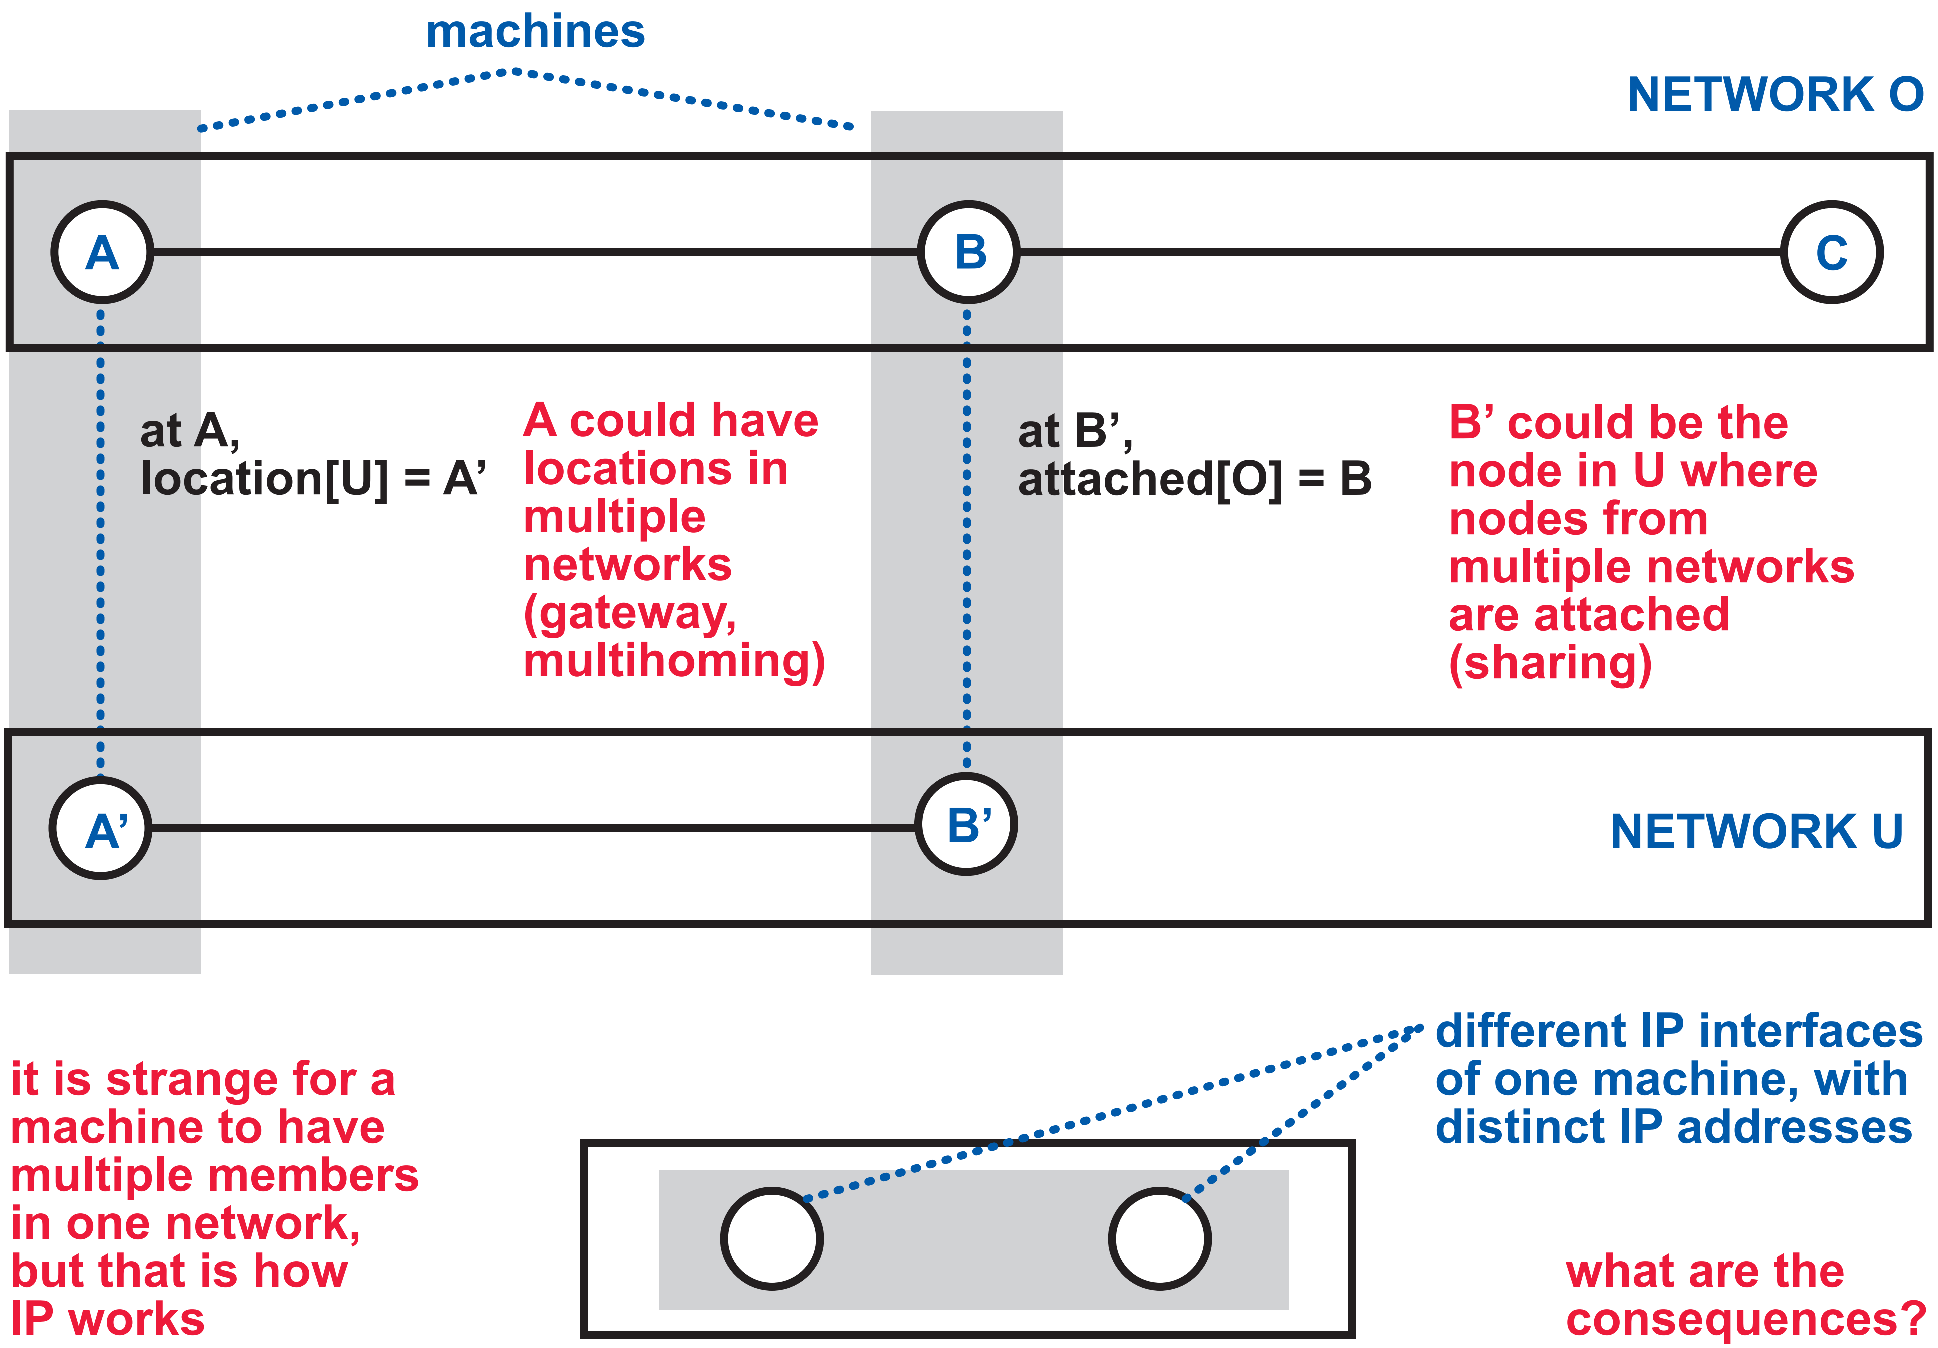

THE DETAILS OF LAYERING 1

header H

source = A
destination = C
sessionIdent = s

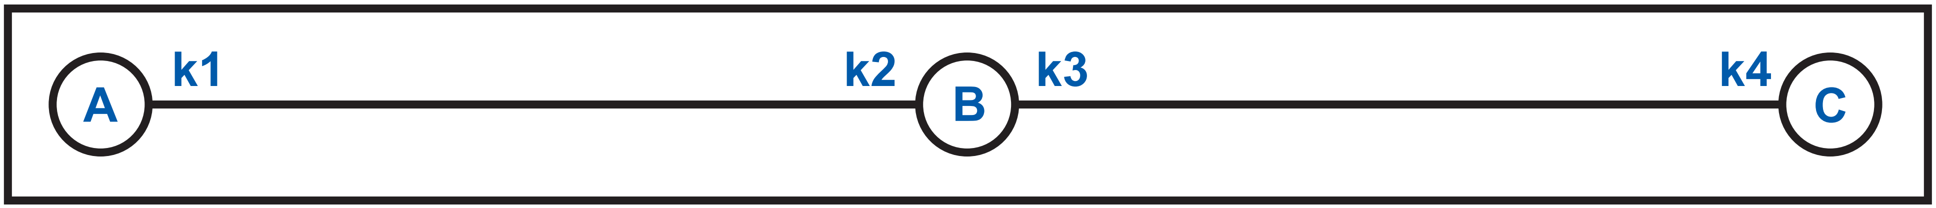

1 A sends packet in session s encapsulating it in header H

2 at A, forwarding[H, Self] = k1

3 A transmits packet on link with local linkIdent k1

4 B acquires packet on link with local linkIdent k2

5 at B, forwarding[H,k2] = k3

6 B transmits packet on link with local linkIdent k3

7 C acquires packet on link with local linkIdent k4

8 H.destination = C, so C receives packet in session s, decapsulating it from header H

THE DETAILS OF LAYERING 2

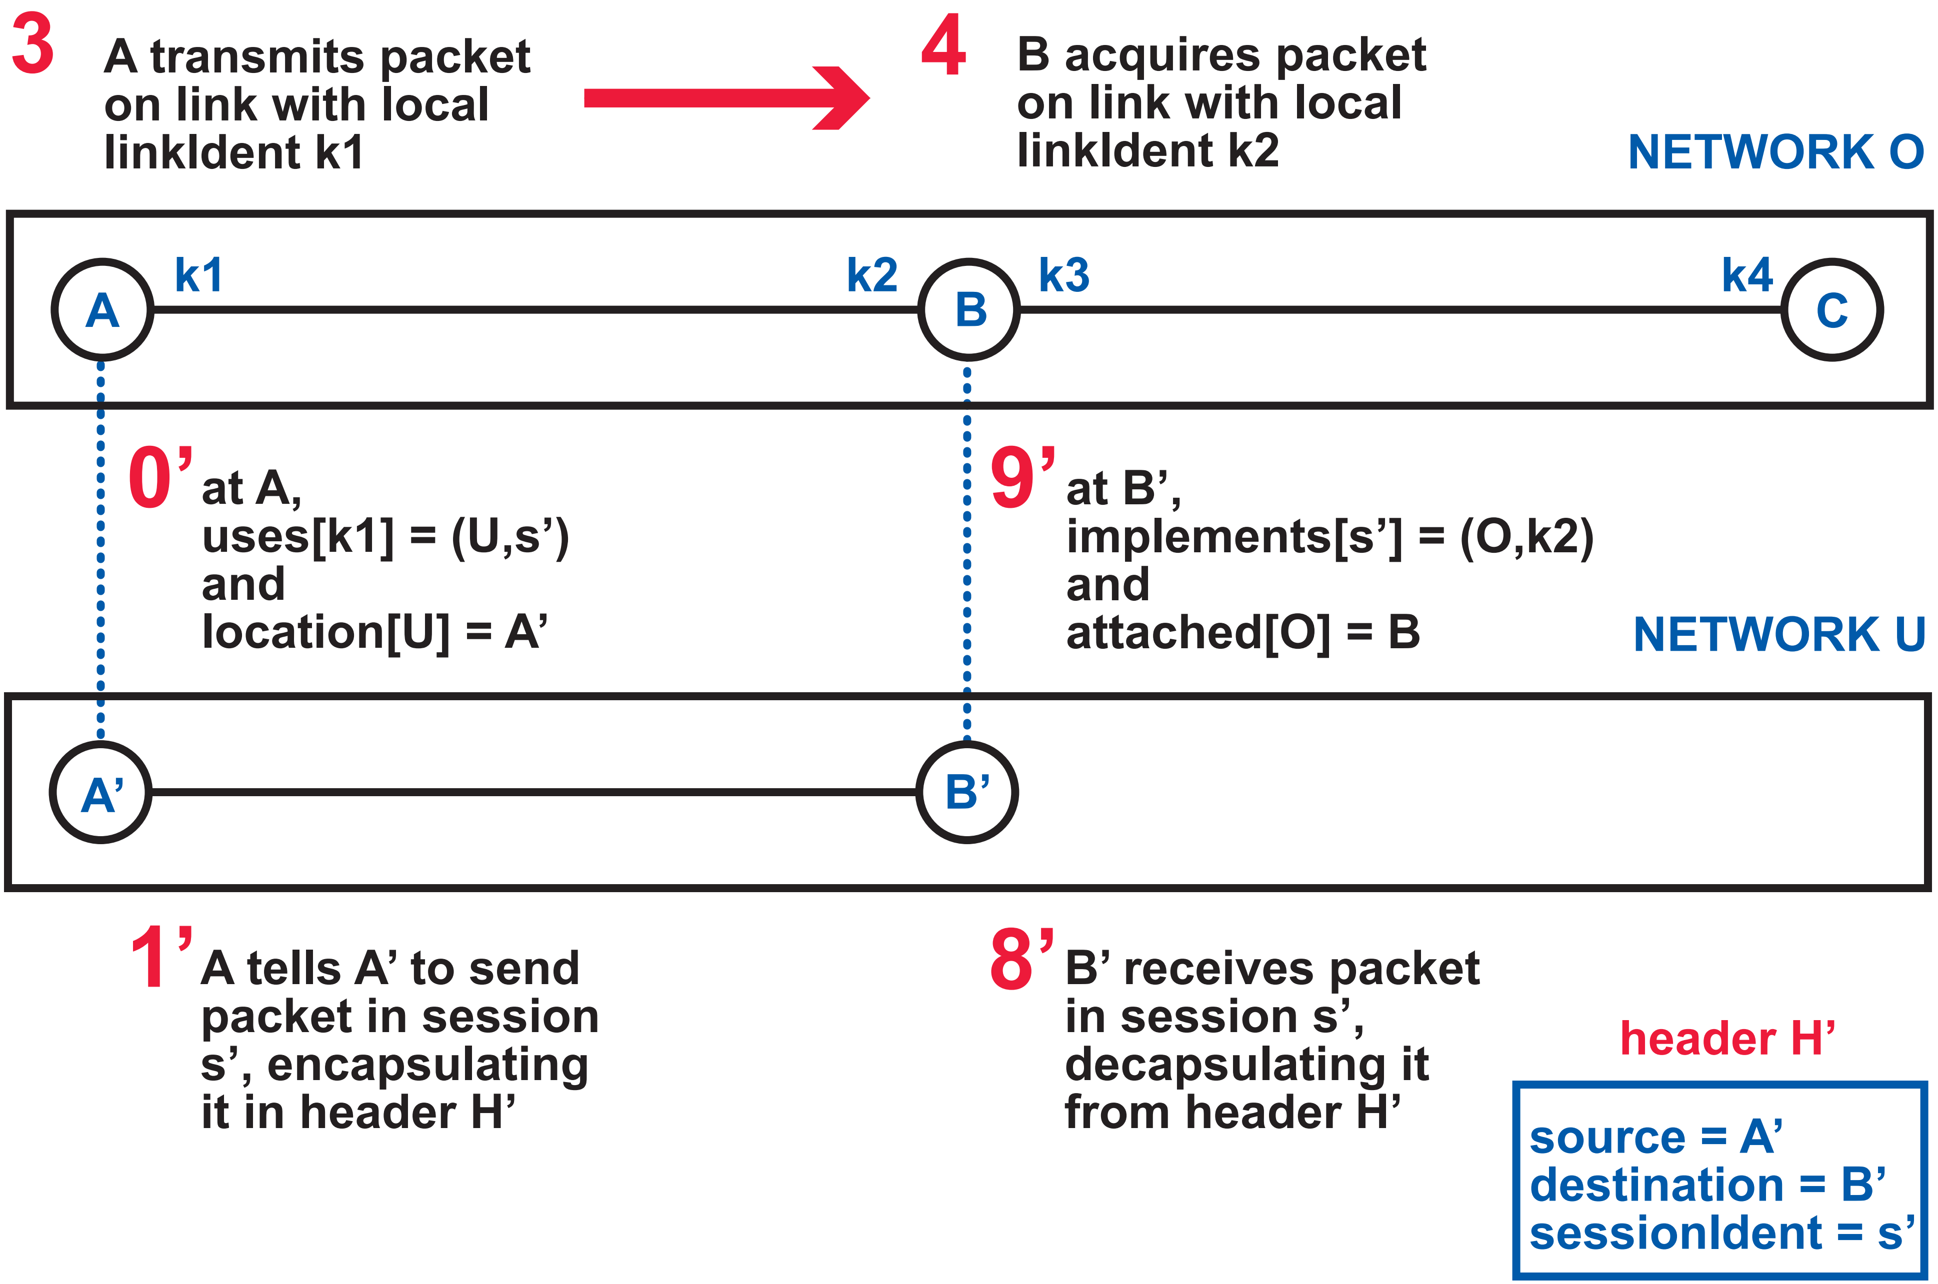

THE DETAILS OF LAYERING 3

NETWORK STATE CAN BE SET UP DYNAMICALLY

3 A transmits packet on link with local linkIdent k1

4 B acquires packet on link with local linkIdent k2

NETWORK O

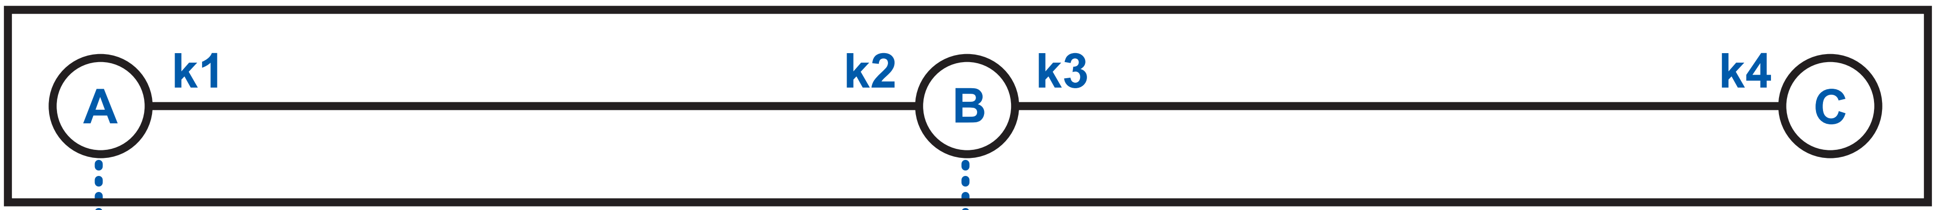

0’ at A, location[U] = A’

A asks A’ for an implementation s’ of k1 to B; location[B,U] = B’

STORE uses[k1] = (U,s’)

9’ at B’, attached[O] = B

B’ asks B for a linkIdent k2 implemented by s’

STORE implements[s’] = (O,k2)

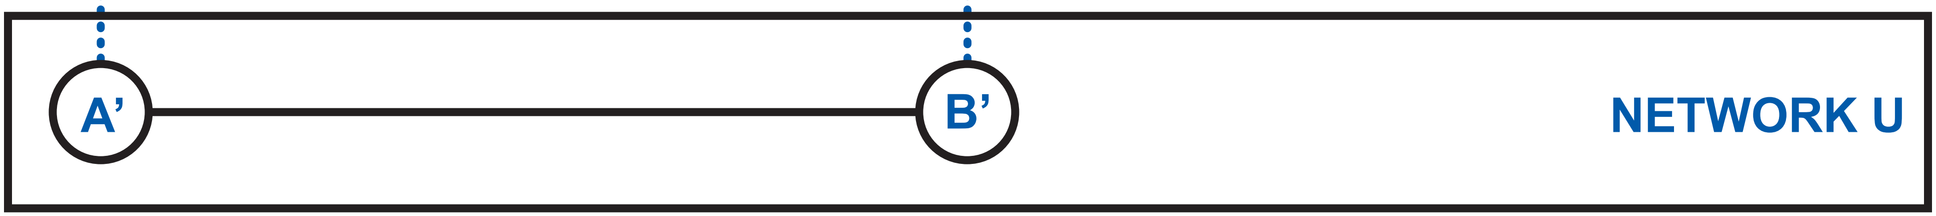

NETWORK U

1’ A tells A’ to send packet in session s’, encapsulating it in header H’

8’ B’ receives packet in session s’, decapsulating it from header H’ and seeing overlay = O

header H’

source = A’
destination = B’
sessionIdent = s’
overlay = O

IP NETWORK LAYERED ON AN ETHERNET LAN

3 A transmits packet on link with local linkIdent k1

4 B acquires packet on link with local linkIdent k2

IP NETWORK O

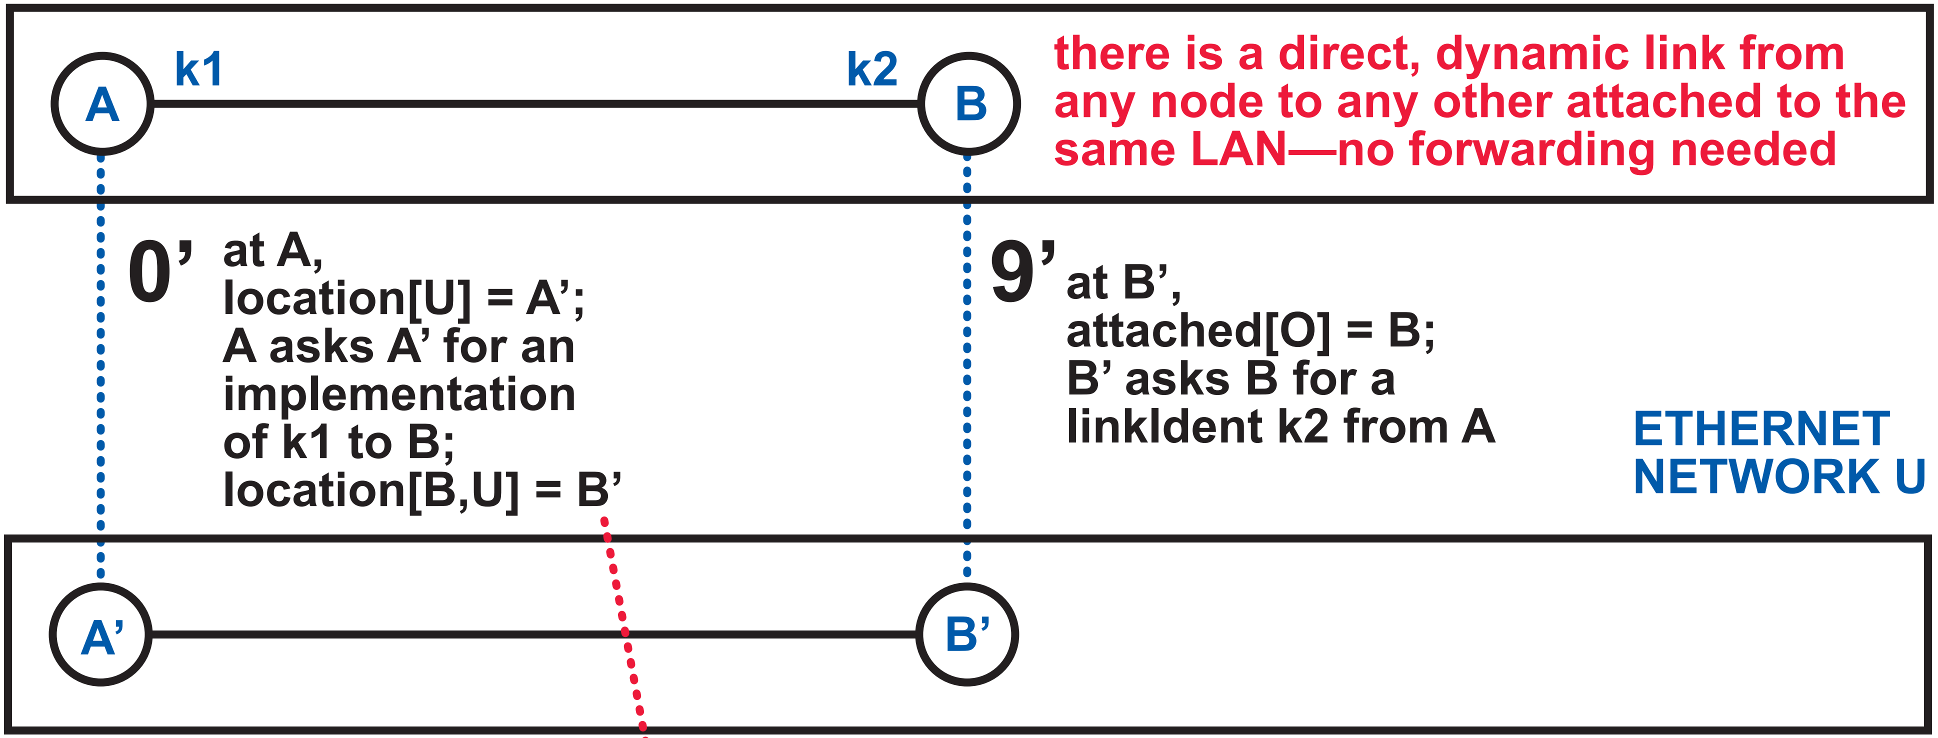

1’ encapsulate packet in header H’ and send

A’ gets the location of B by broadcasting an ARP request

8’ B’ receives packet, decapsulating it from header H’

header H’

source = A’
destination = B’
~~sessionIdent = s’~~
~~overlay = O~~

H’ does not need s’ and O because they are constants

OTHER PURPOSES

TO BUILD A NETWORK WITH BETTER PERFORMANCE OR RELIABILITY ON TOP OF AN EXISTING NETWORK

for example, Resilient Overlay Networks

TO SHARE THE RESOURCES OF AN EXISTING NETWORK IN A DISCIPLINED WAY

for example, Virtual LANs

TO BUILD A NETWORK WITH LINKS THAT OFFER A SUPERIOR COMMUNICATION SERVICE

for example, Virtual Private Networks

during the semester, be on the lookout for other purposes!

IP NETWORK LAYERED ON A VLAN ?

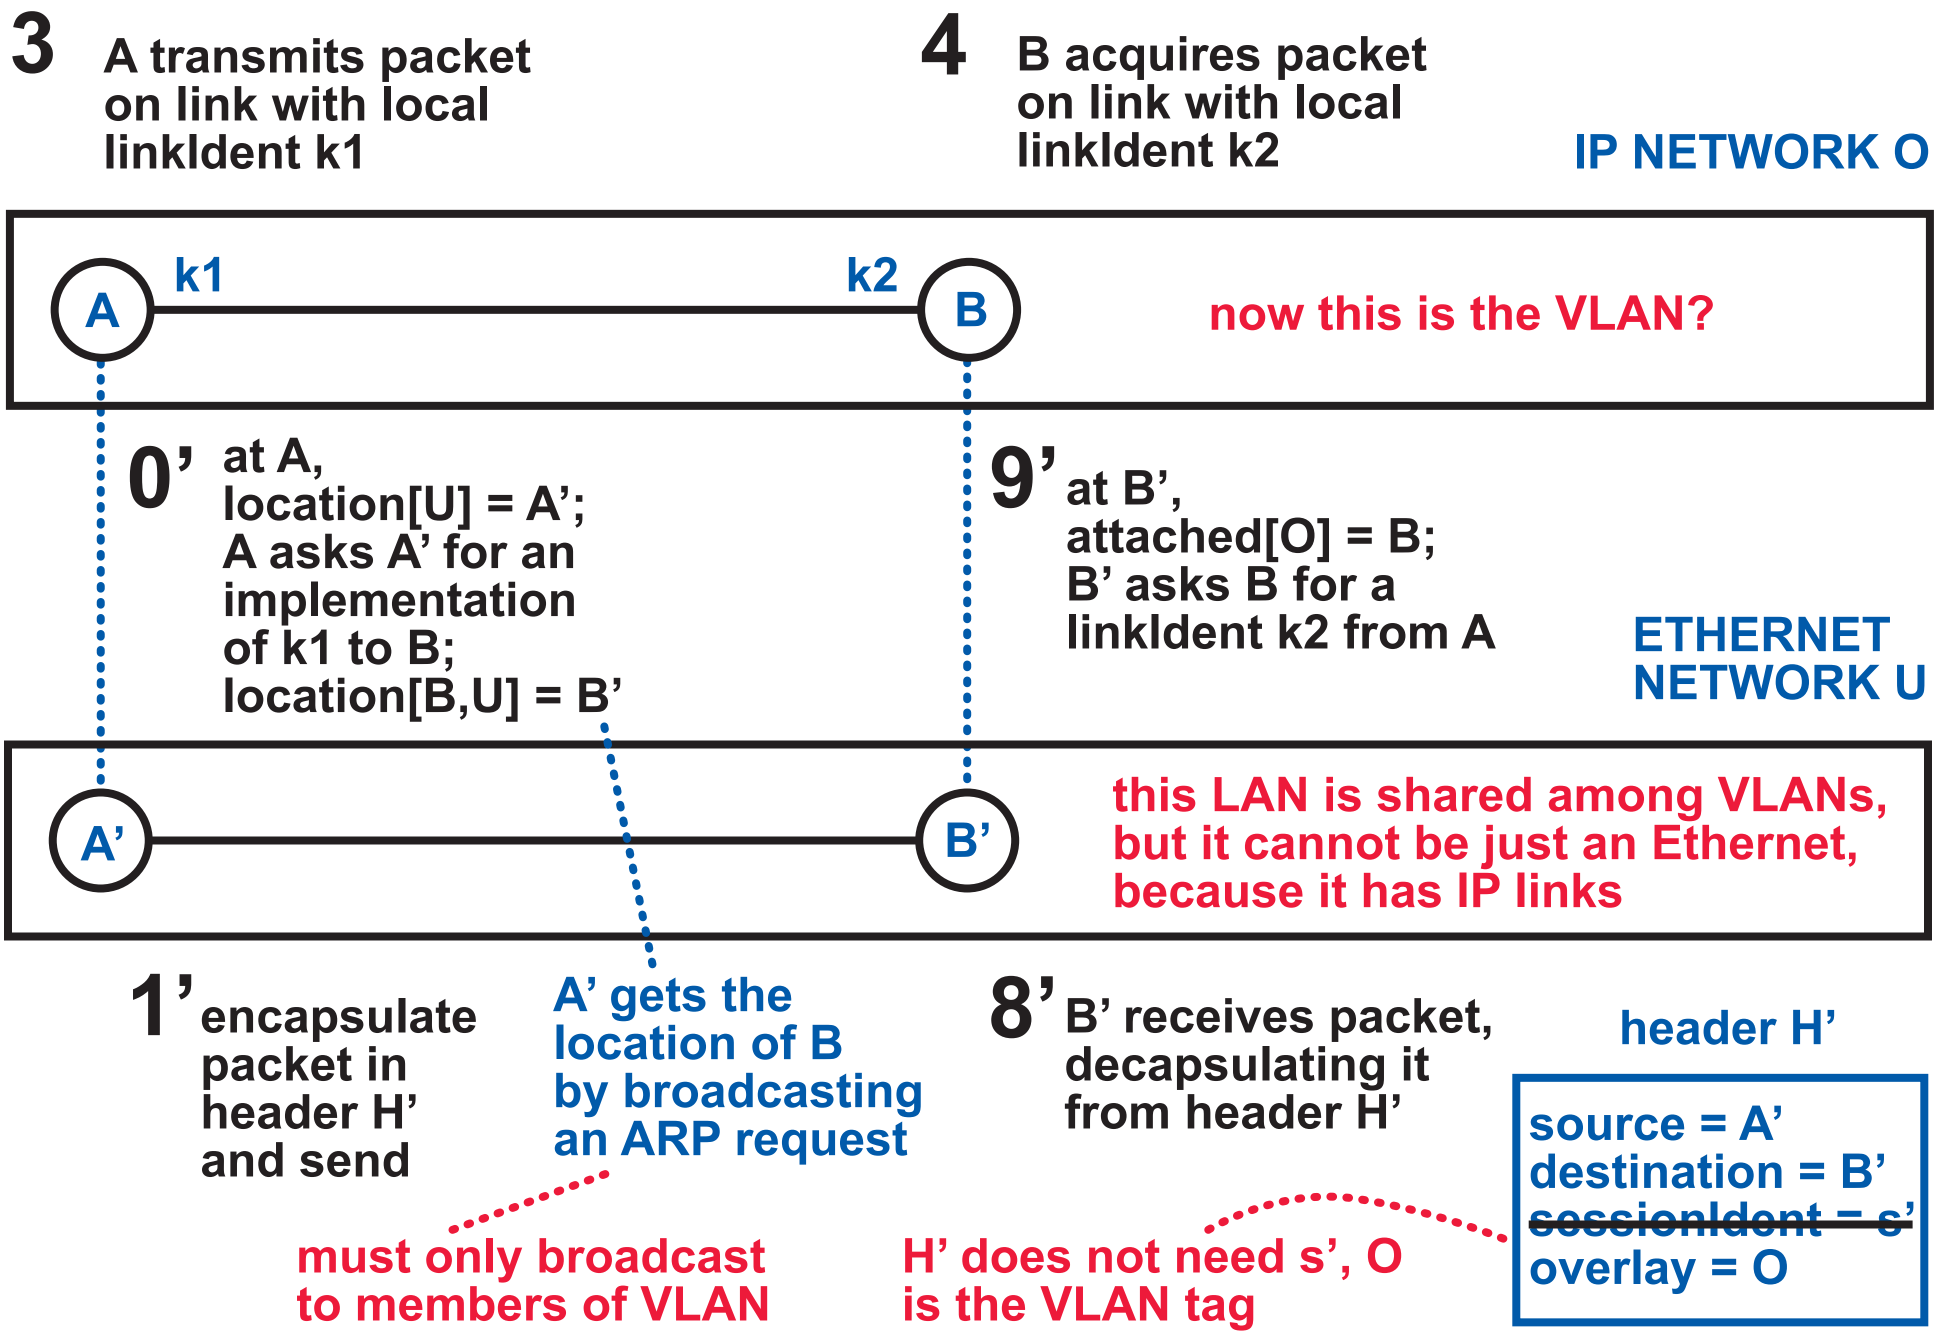

IP NETWORK LAYERED ON A VLAN

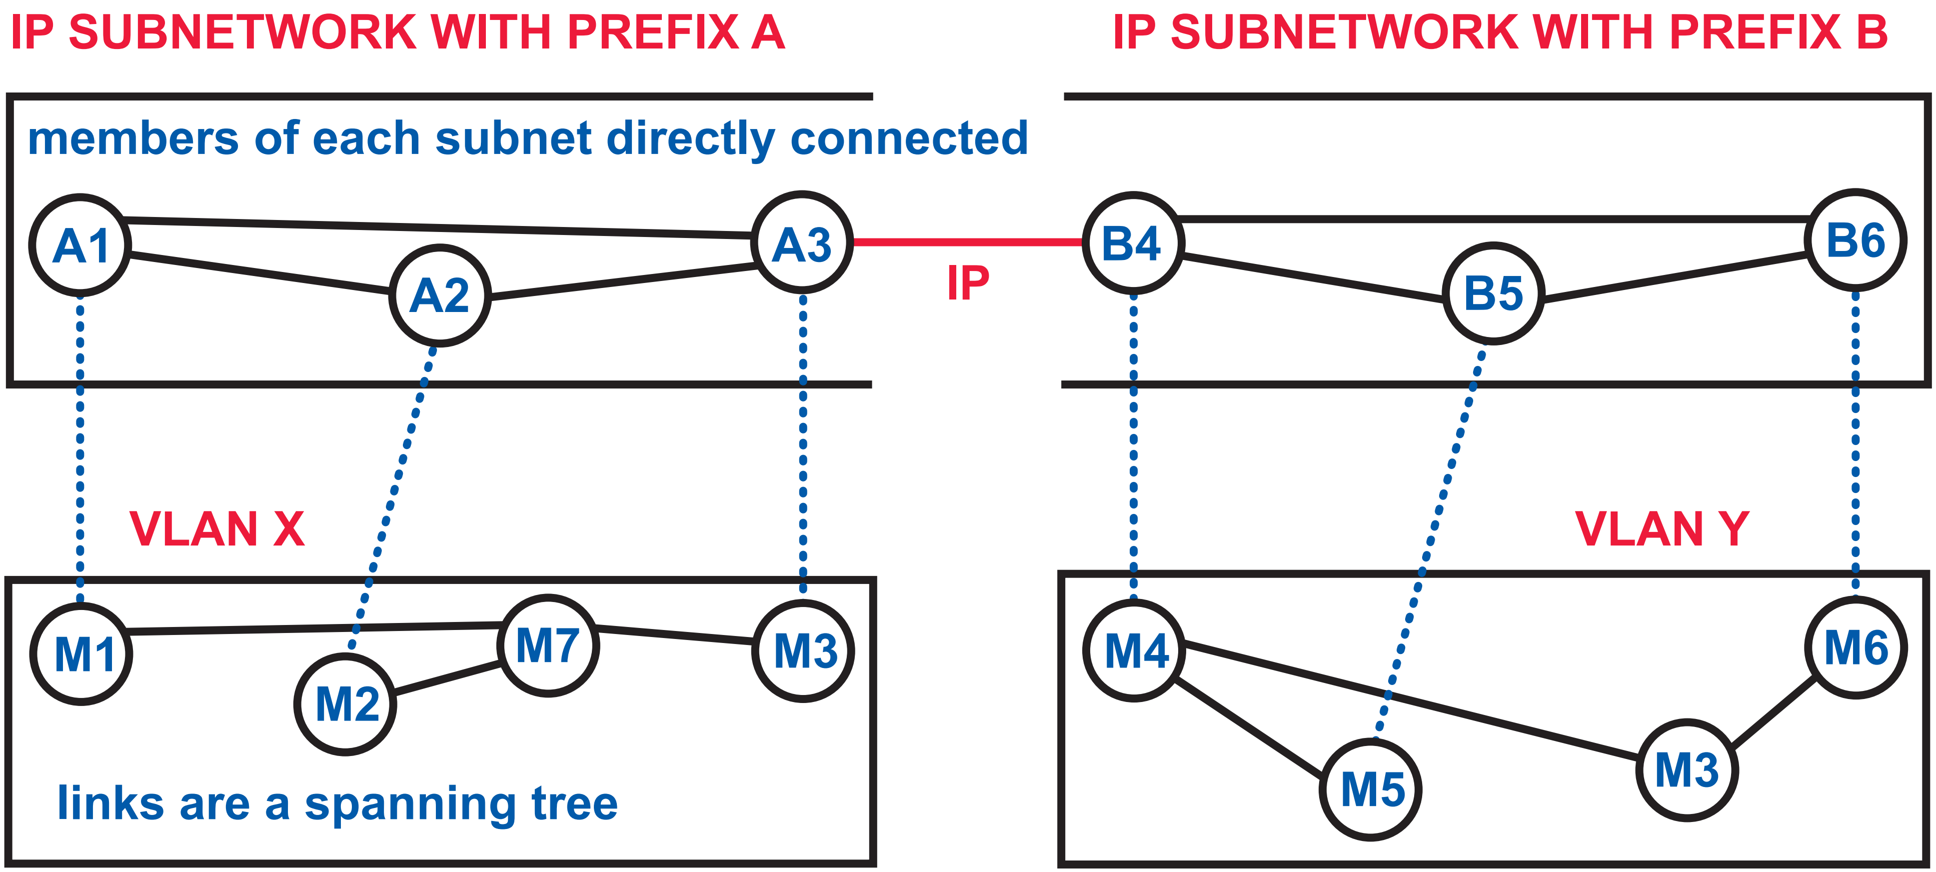

PHYSICAL LAN has
host members with
names M1, M2, M6
switches with names
M3, M7

IP BACKBONE has
routers with names
A3, B4

PHYSICAL LAN has
host members with
names M5, . . .
switches with names
M4

one machine does 3 completely different things (name, links, forwarding) in its roles as A3, M3 in X, M3 in Y

DISCUSSION OF

"A SURVEY OF VIRTUAL LAN USAGE IN CAMPUS NETWORKS"

OTHER PURPOSES

TO BUILD A NETWORK WITH BETTER PER-FORMANCE OR RELIABILITY ON TOP OF AN EXISTING NETWORK

for example, Resilient Overlay Networks

TO SHARE THE RE-SOURCES OF AN EXISTING NETWORK IN A DISCIPLINED WAY

for example, Virtual LANs

TO BUILD A NETWORK WITH LINKS THAT OFFER A SUPERIOR COMMUNICATION SERVICE

for example, Virtual Private Networks

during the semester, be on the lookout for other purposes!

WHY A VIRTUAL PRIVATE NETWORK?

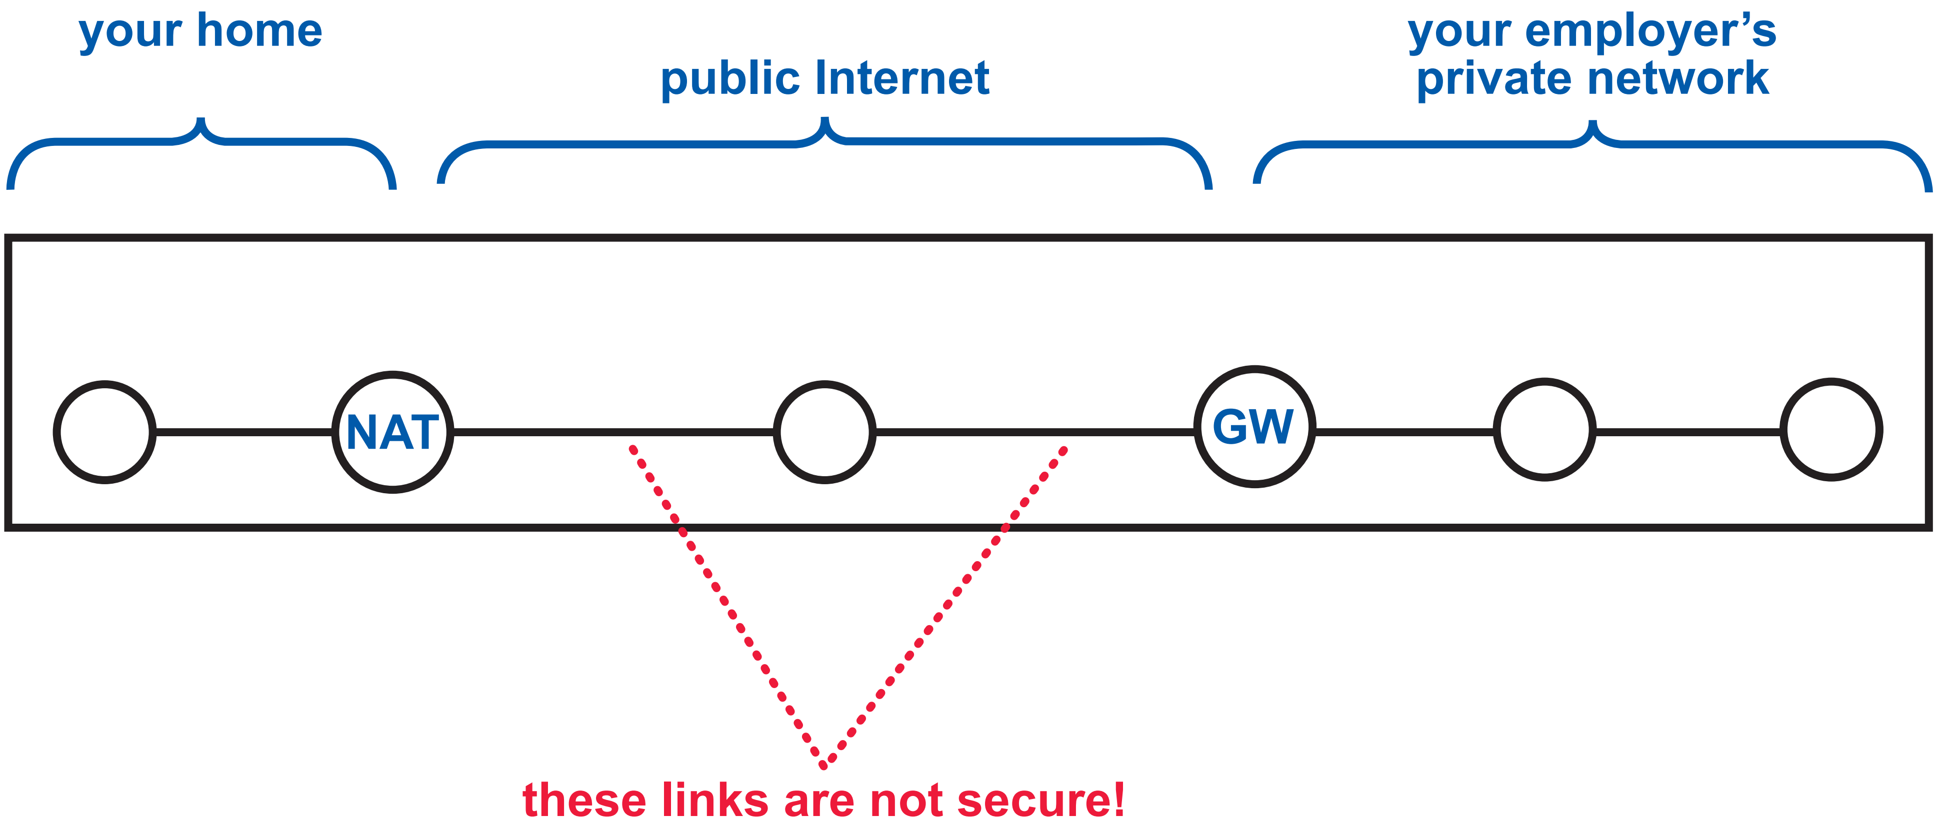

data could be read or tampered with

attackers could insert packets with false source addresses

there could be DoS attacks, replay attacks, etc.

VIRTUAL PRIVATE NETWORK

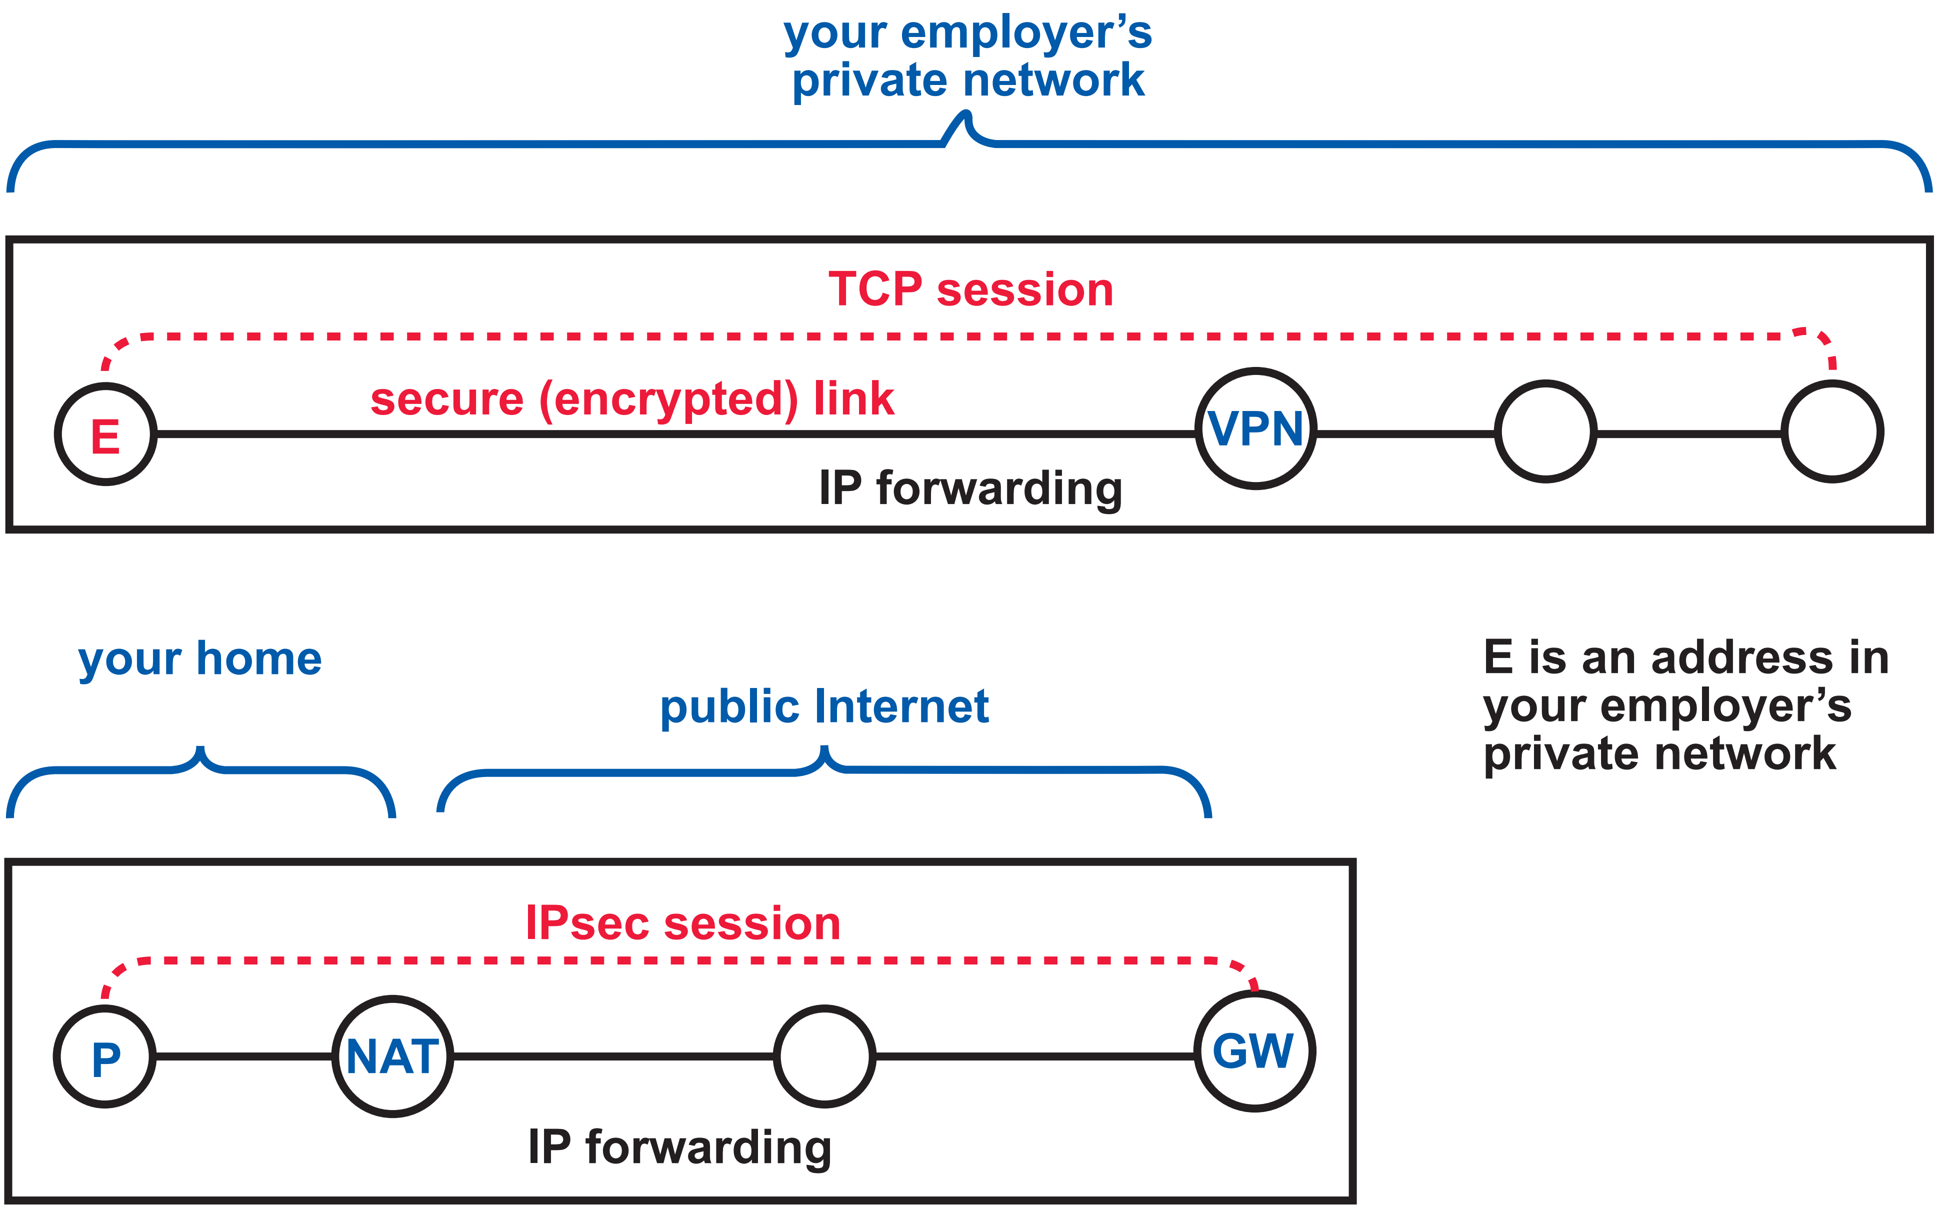

IPsec AND NAT TRAVERSAL

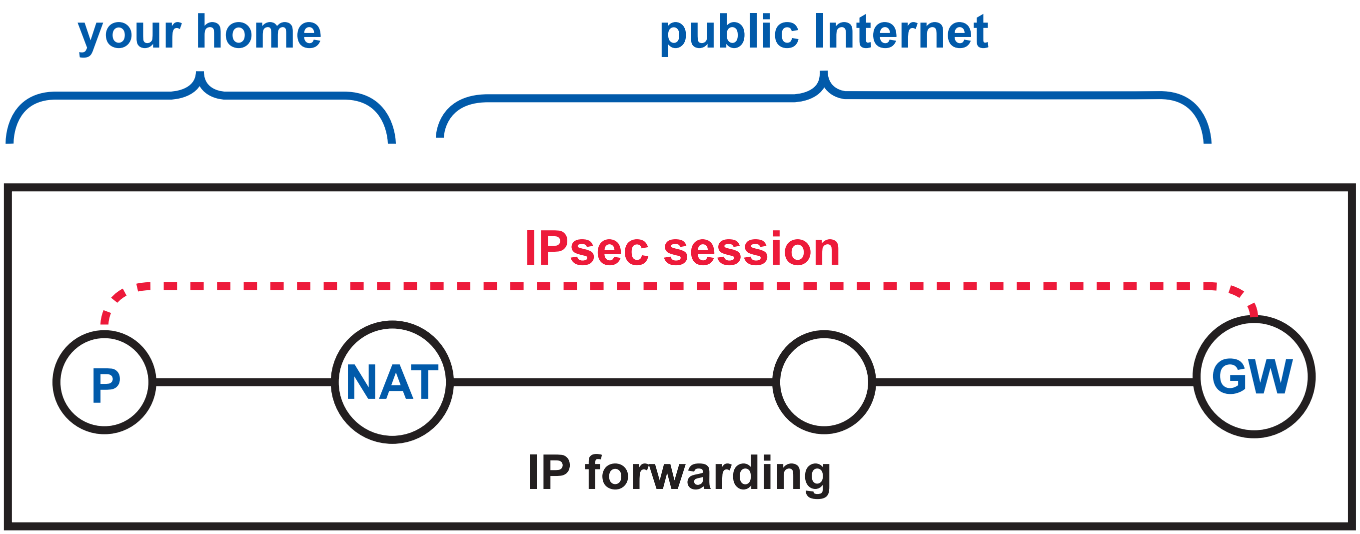

P is an IP address in your home network

NAT has an IP address in your ISP's network

the IP session must be able to traverse the NAT

the NAT is expecting to see TCP or UDP ports as the session identifier, but an IPsec packet does not have these

an IPsec packet has Security Parameters Indices (SPIs), but they are different in each direction

so the answer is some sort of ad hoc fix built into the NAT, helped by the fact that the IPsec session begins with the Internet Key Exchange (IKE) protocol using UDP port 50

PROPERTIES OF LAYERING 1

Usually, many sessions in a network share the resources of its links.

Usually, session packets must be routed over a path of multiple links to get to their destinations.

With layering, a network can create a dynamic, source-to-destination link for each session.

This is possible because an underlay will implement the link as a session, and routing in the underlay will do all the work.

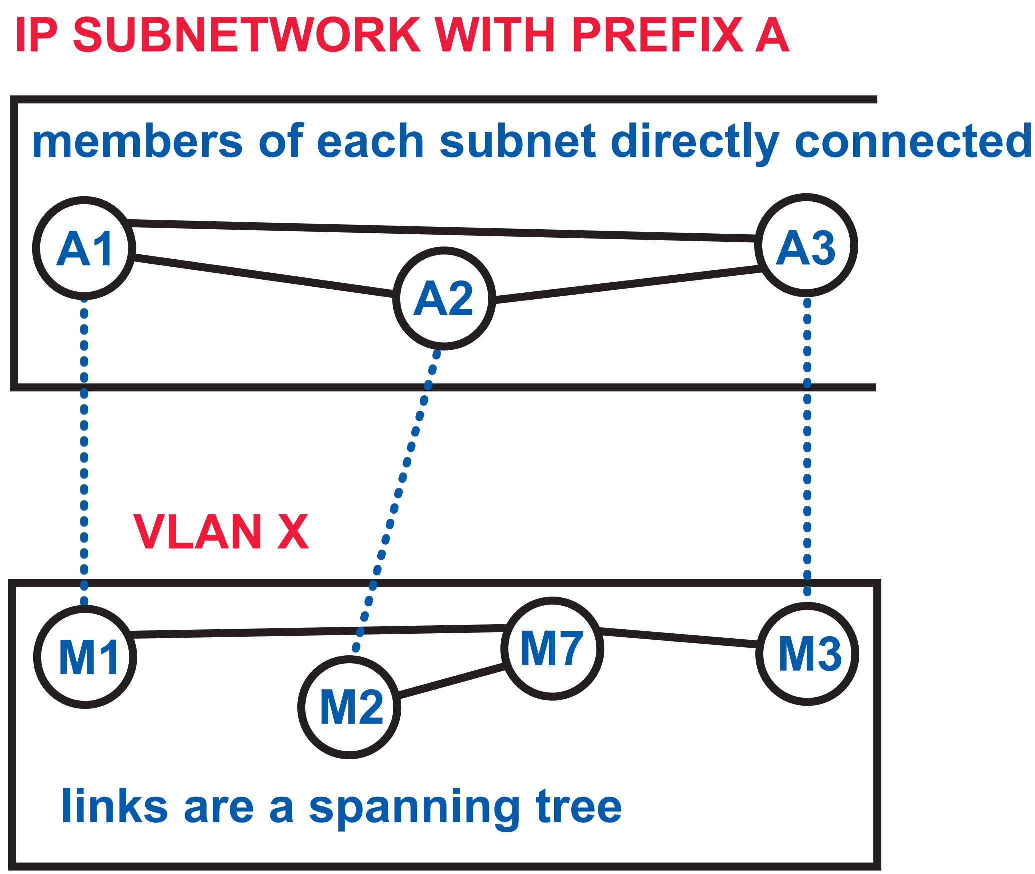

PROPERTIES OF LAYERING 2

FOR A PARTICULAR OVERLAY AND UNDERLAY . . .

. . . WHAT ARE THE RELATIONSHIPS BETWEEN LINKS AND THE SESSIONS THAT IMPLEMENT THEM?

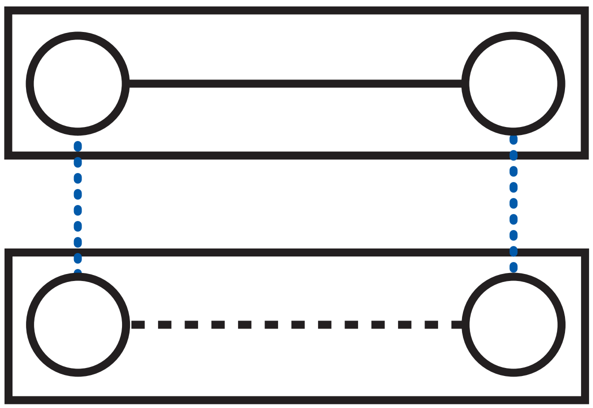

one-to-one is good

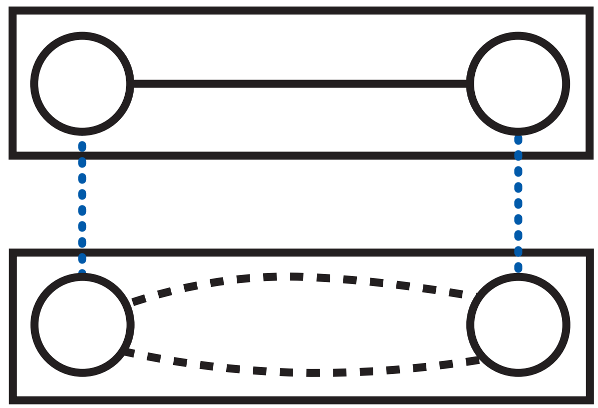

one-to-many is awkward

- when a packet is transmitted on the link, requires a mechanism to choose which session to use

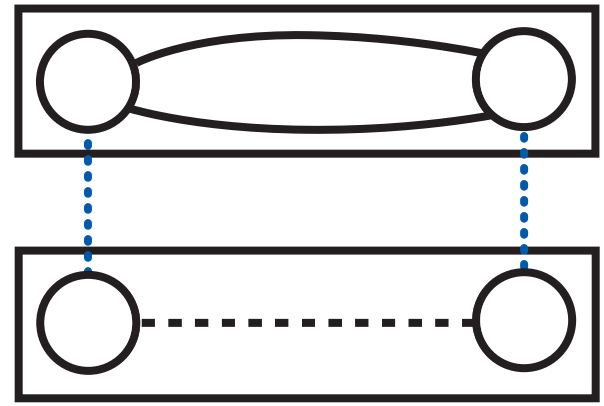

many-to-one is worse

- when a packet is received in the session, there is no way to know which link should acquire it

grouping of packets is done with sessions, and we have excluded this mechanism

there is no reason to do this—extra sessions are cheap

PROPERTIES OF LAYERING 3

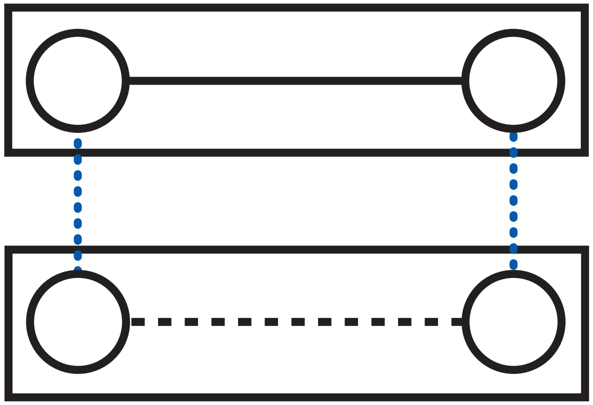

a network can have as many virtual links as are required . . .

. . . implemented one-to-one by as many sessions as required

what is wrong with the word "tunneling"?

it implies "link over path" layering

- ignores the possibility that the underlay session protocol could provide a service
- ignores the need (for general purposes) to send a session identifier and overlay

WHAT PROPERTIES ARE REQUIRED

OR DESIRABLE?

WHAT PROPERTIES SHOULD BE

VERIFIED, AND HOW?

PROPERTIES OF LAYERING 4

In the underlay, session destination must be reachable from session source.

note that this is a motivation for reachability requirements in the underlay

Two different links in the overlay are implemented using different resources in the underlay.

MORE?

Compute the load on a network, based on its overlays.

Compute the capacity of a network, based on its underlays.

CODE GENERATION AND OPTIMIZATION

The “Details of Layering” slide showed a step-by-step algorithm for packet processing, in networks that follow the model.

Imagine that we wrote a program to do this processing in any network.

Customizations:

- specific types for names, link identifiers, session identifiers
- extra, protocol specific information in headers
- functions like . . . . . . can become . . .

functions like . . .	. . . can become . . .
location [A,U] = A’	constants
attached[B’,O] = B	table lookups
	remote queries

Optimization:

- compile efficiently to run where it needs to (router, host, VM, NIC, etc.)

NOW WE CAN GENERATE VERIFIED SOFTWARE FOR THE DATAPLANE OF ANY NETWORK.

NEXT STEP IS TO BRING THE SAME BENEFITS TO CONTROL PLANES.

WHAT DOES THE COMPOSITION OF RON AND A VPN LOOK LIKE?

. . . RON AND A VLAN?

. . . A VLAN AND A VPN?

. . . ALL THREE?

RON LAYERED ON A VPN

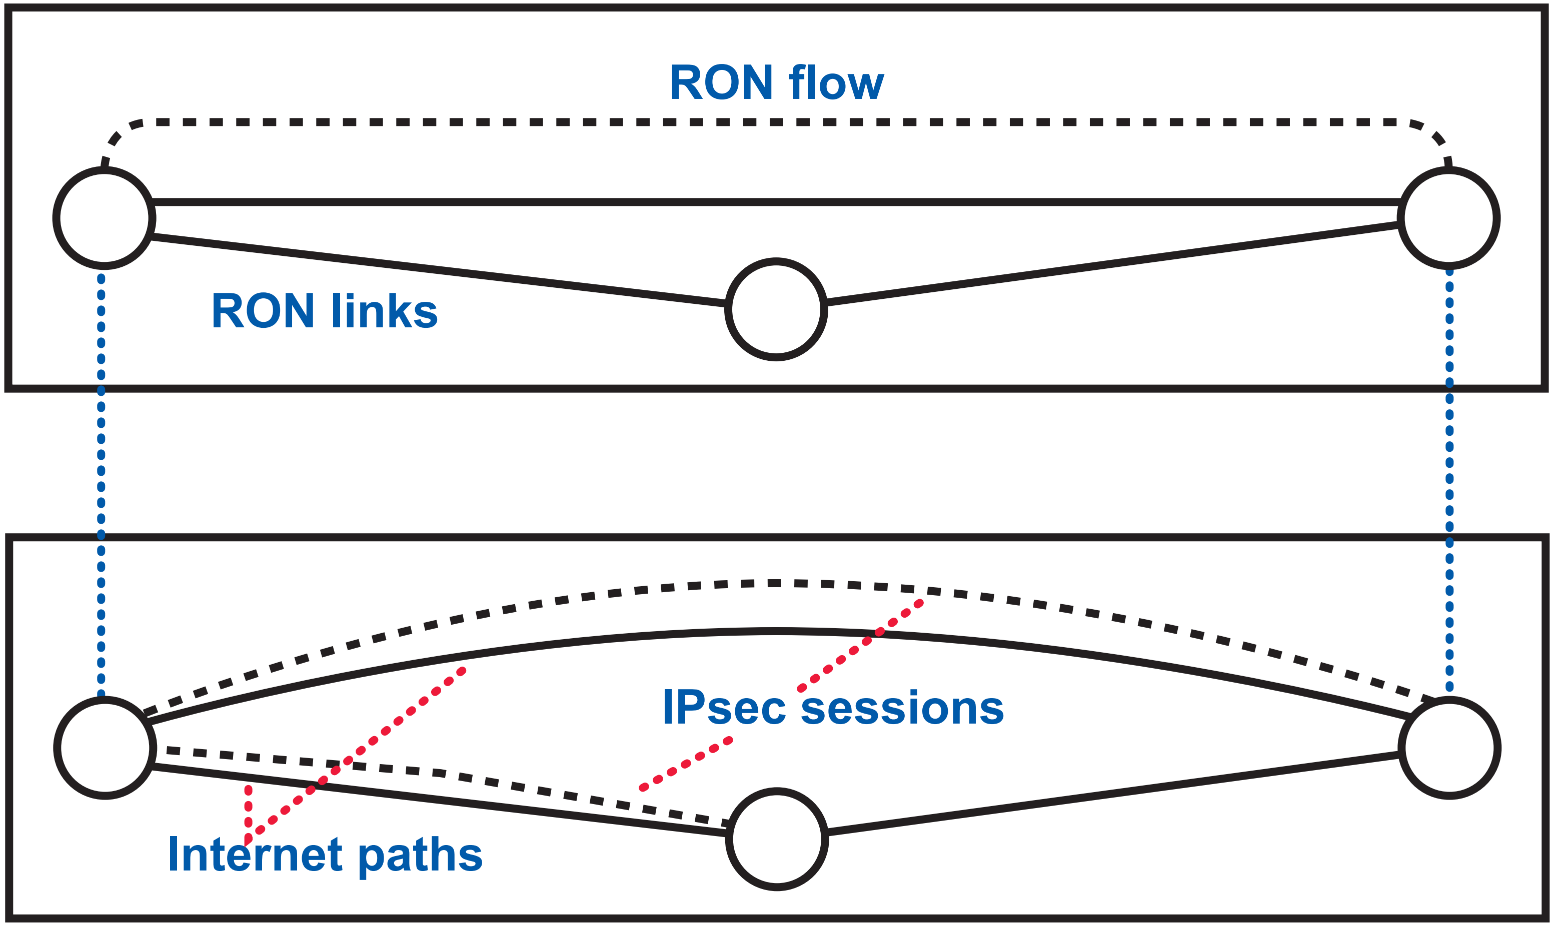

VPN LAYERED ON A RON

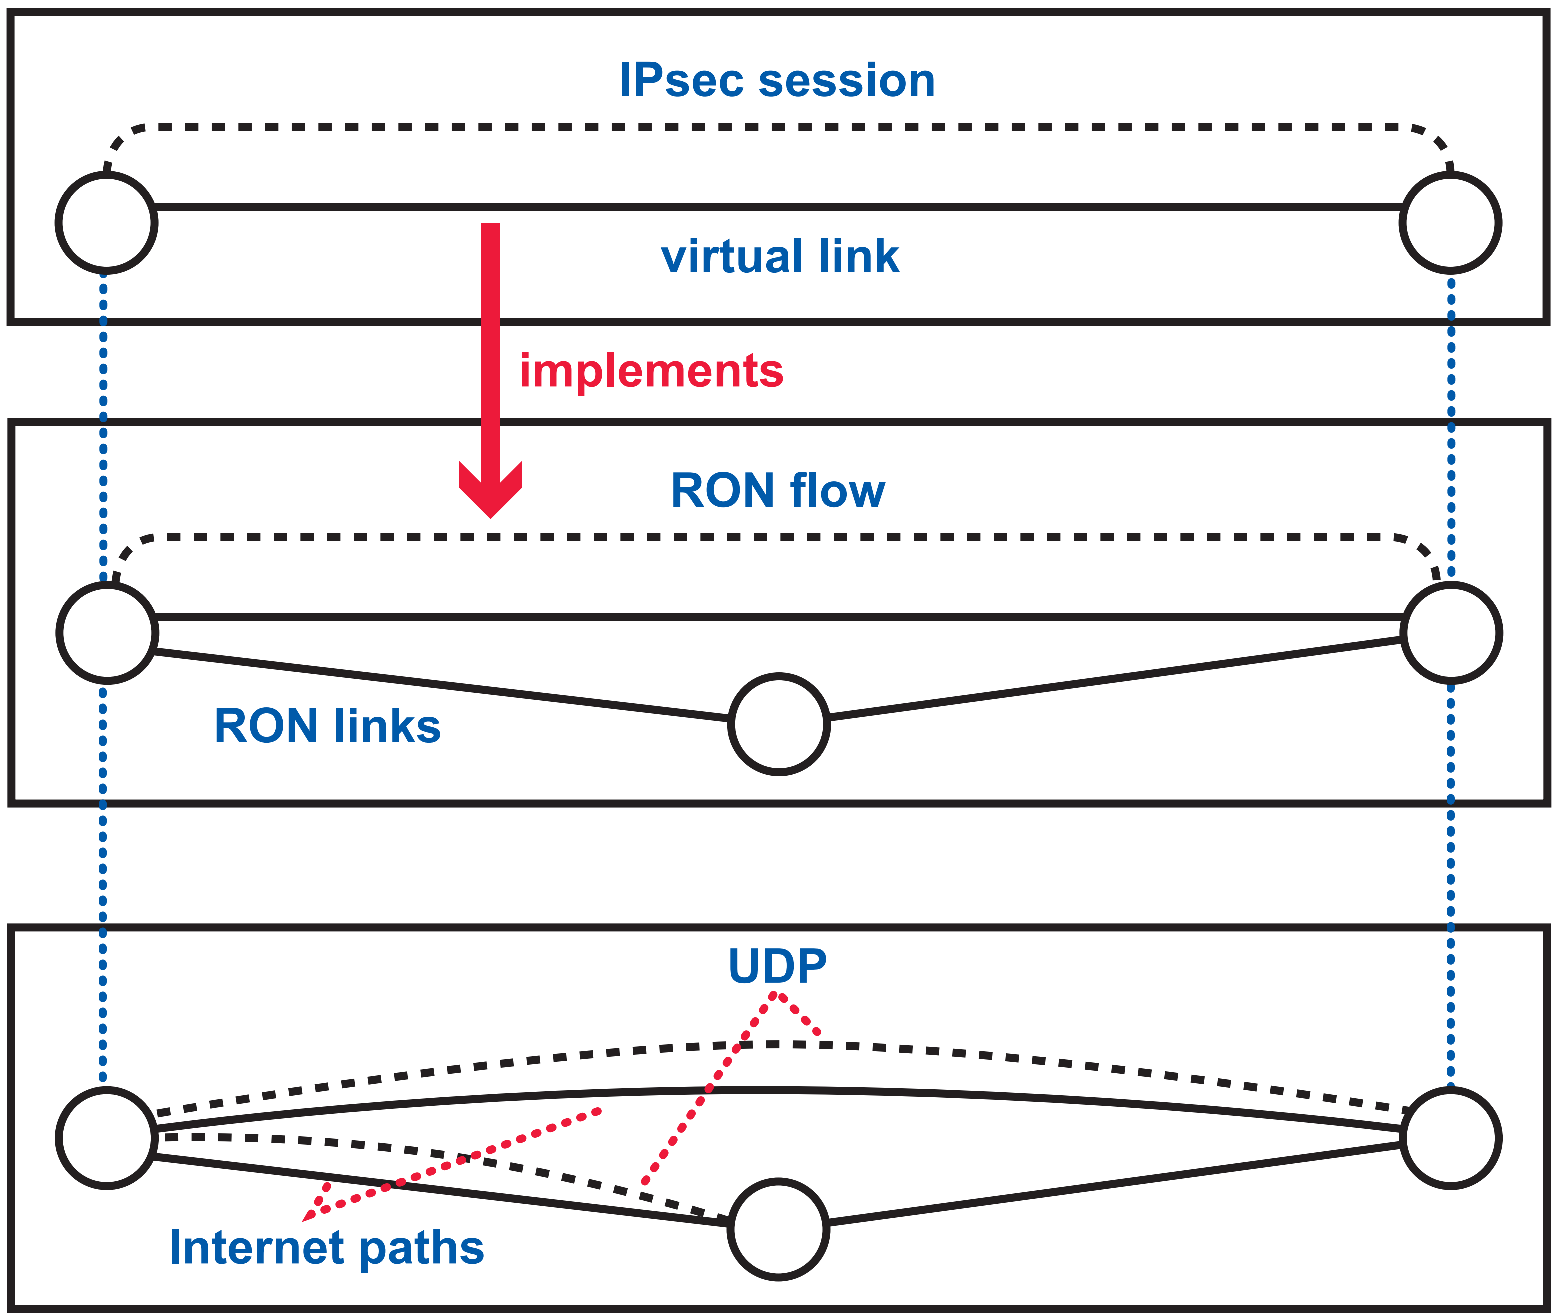

A BETTER IDEA

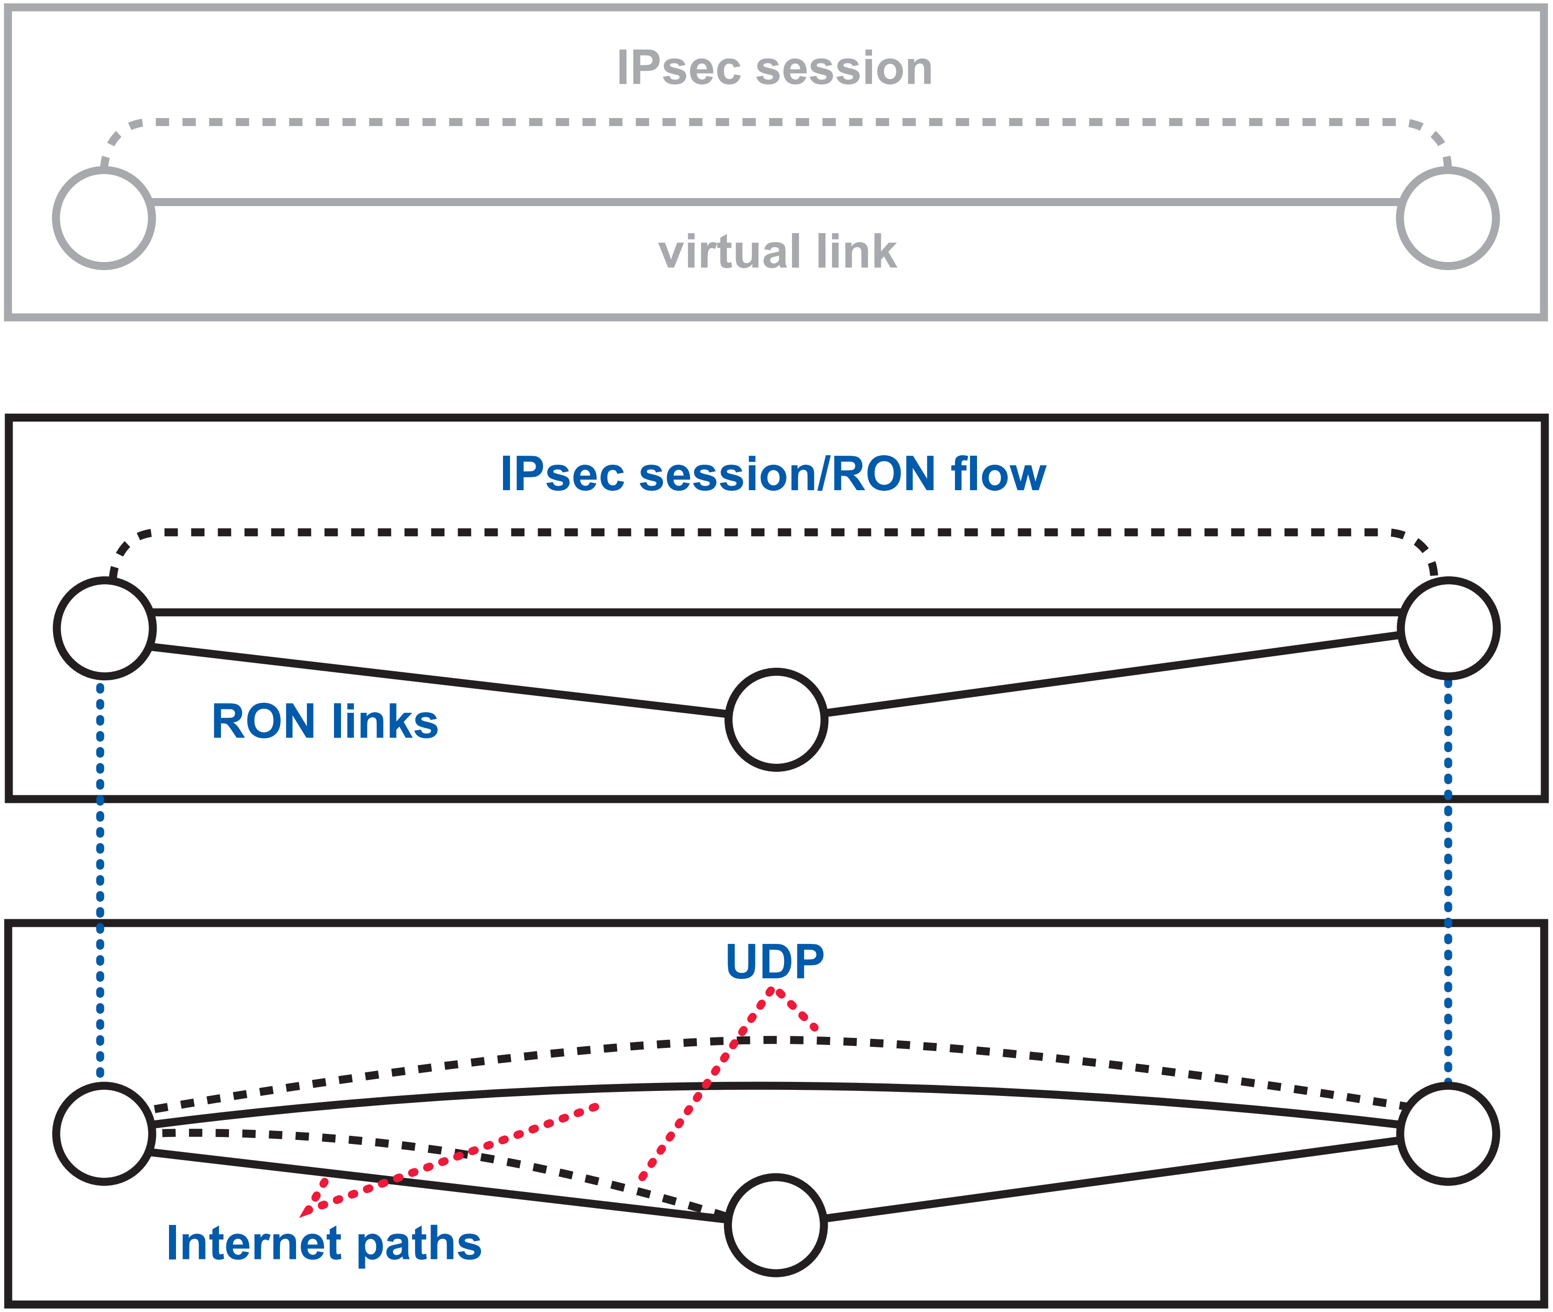

if both IPsec and RON had "normal" or standard session identifiers rather than something quirky, there would be no problem with running IPsec as a session protocol in a RON network

DISCUSSION OF

“RETHINKING THE DESIGN OF THE INTERNET: THE END-TO-END ARGUMENTS VS. THE BRAVE NEW WORLD