PATTERNS IN NETWORK ARCHITECTURE:

CLOUD COMPUTING

CLOUD COMPUTING

OUTLINE

1. Discussion of Alloy homework (net4.als)
2. Discussion of “VL2: A scalable and flexible data center network”
3. Models of VL2 and SEATTLE
4. Discussion of “Stratos: A network-aware orchestration layer for virtual middleboxes in clouds”
5. Model of a cloud design and comparisons to literature
6. The Nicira paper

VLAN TECHNOLOGY (ACCORDING TO VL2 PAPER)

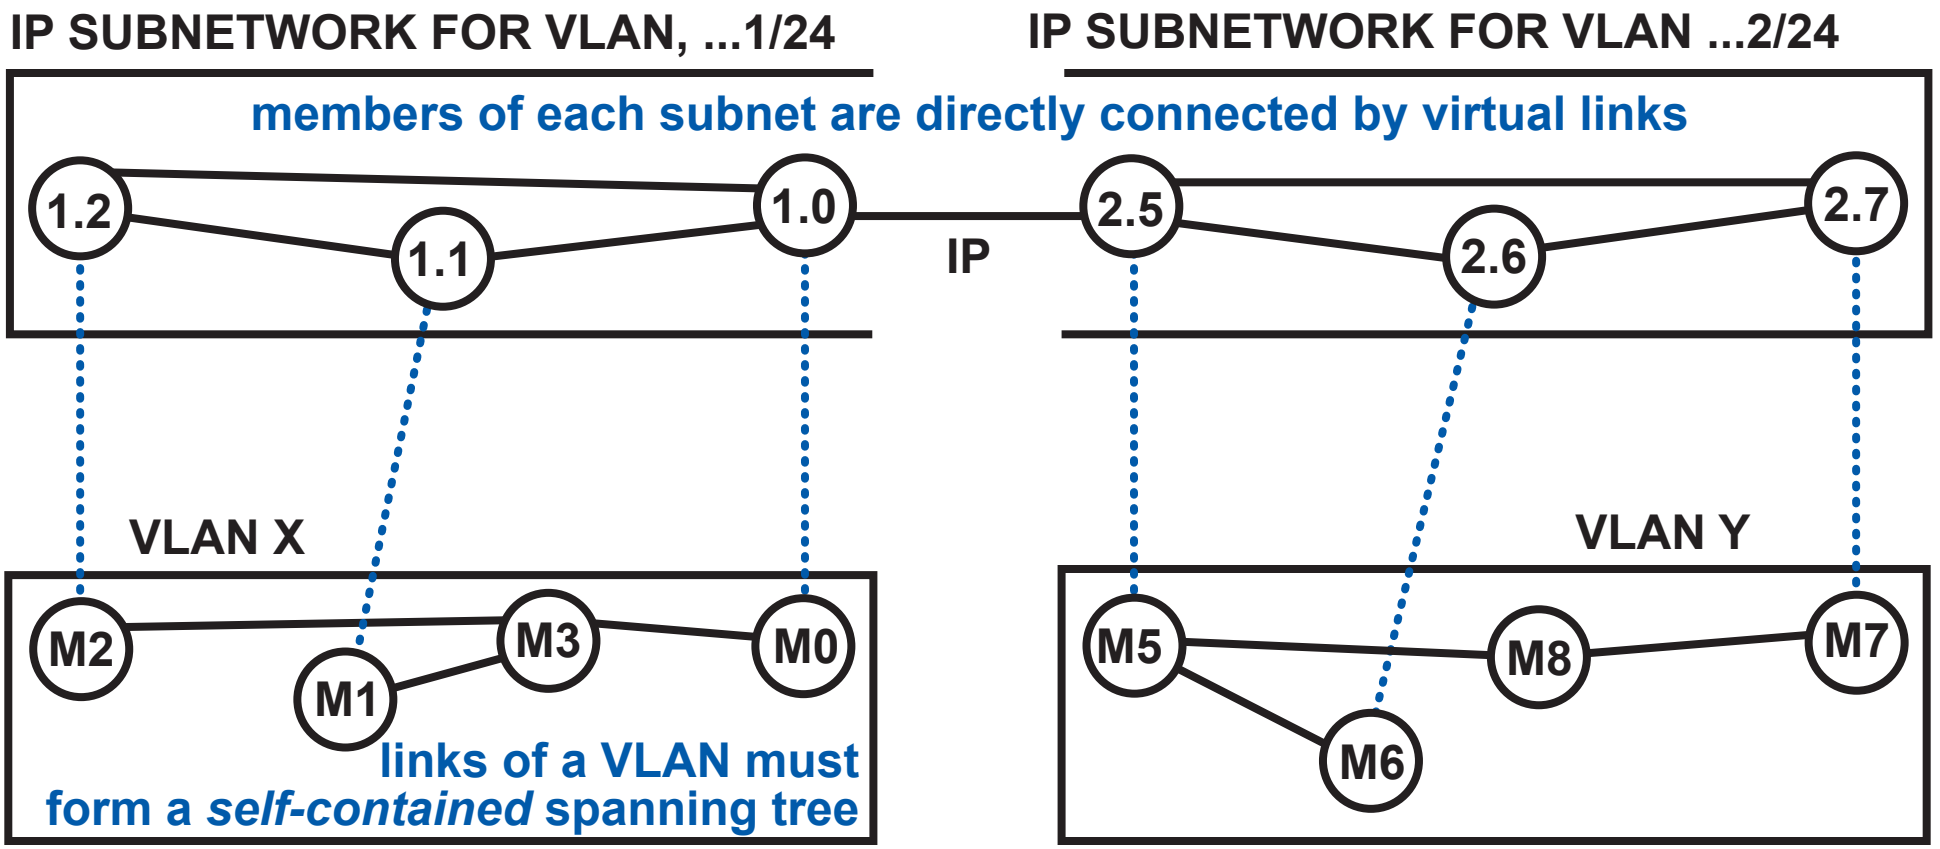

although each LAN has
2 access routers, inter-LAN
bandwidth is severely limited

LAN A

all members of VLAN X must be connected to this LAN

LAN B

all members of VLAN Y must be connected to this LAN

VL2 ARCHITECTURE

connections to the public Internet are not shown

VL2:

names are a random subset of the AA space

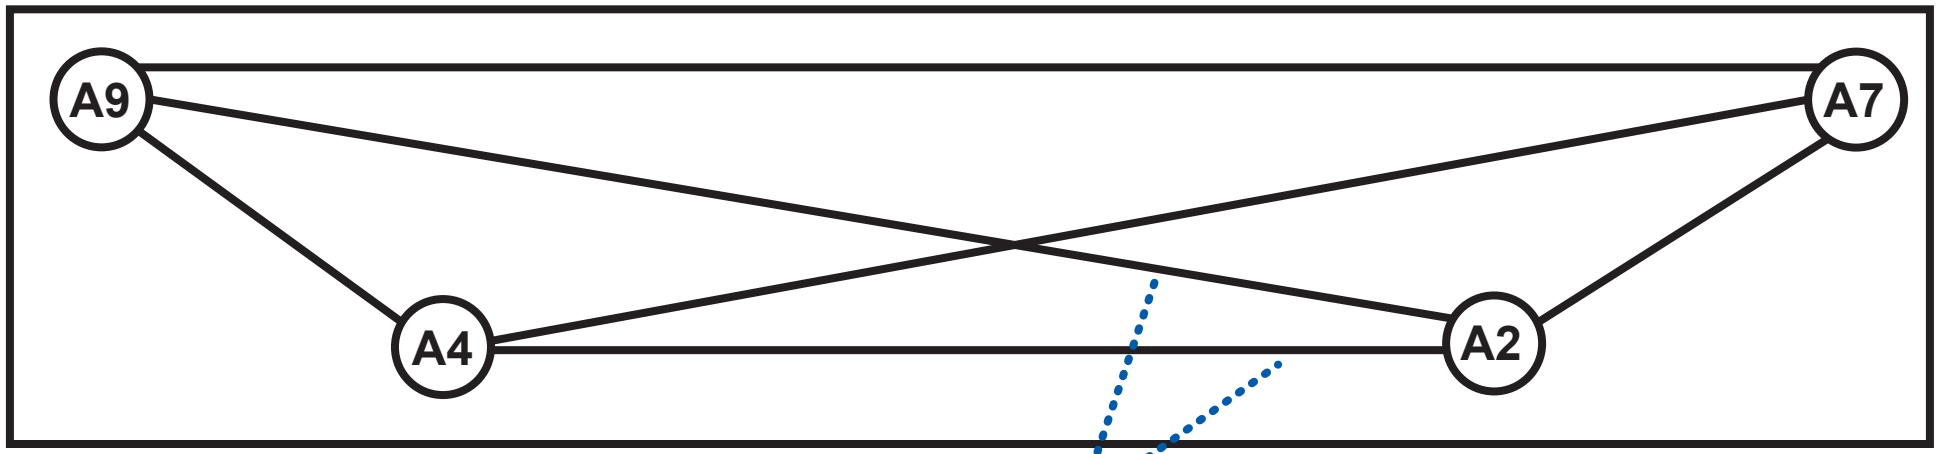

members are fully connected by dynamic links

a member is a virtual machine, meaning the data and processing state—what might be called the VM "image"

a VM can migrate from one location to another, without changing anything in the VL2

VL2 ARCHITECTURE

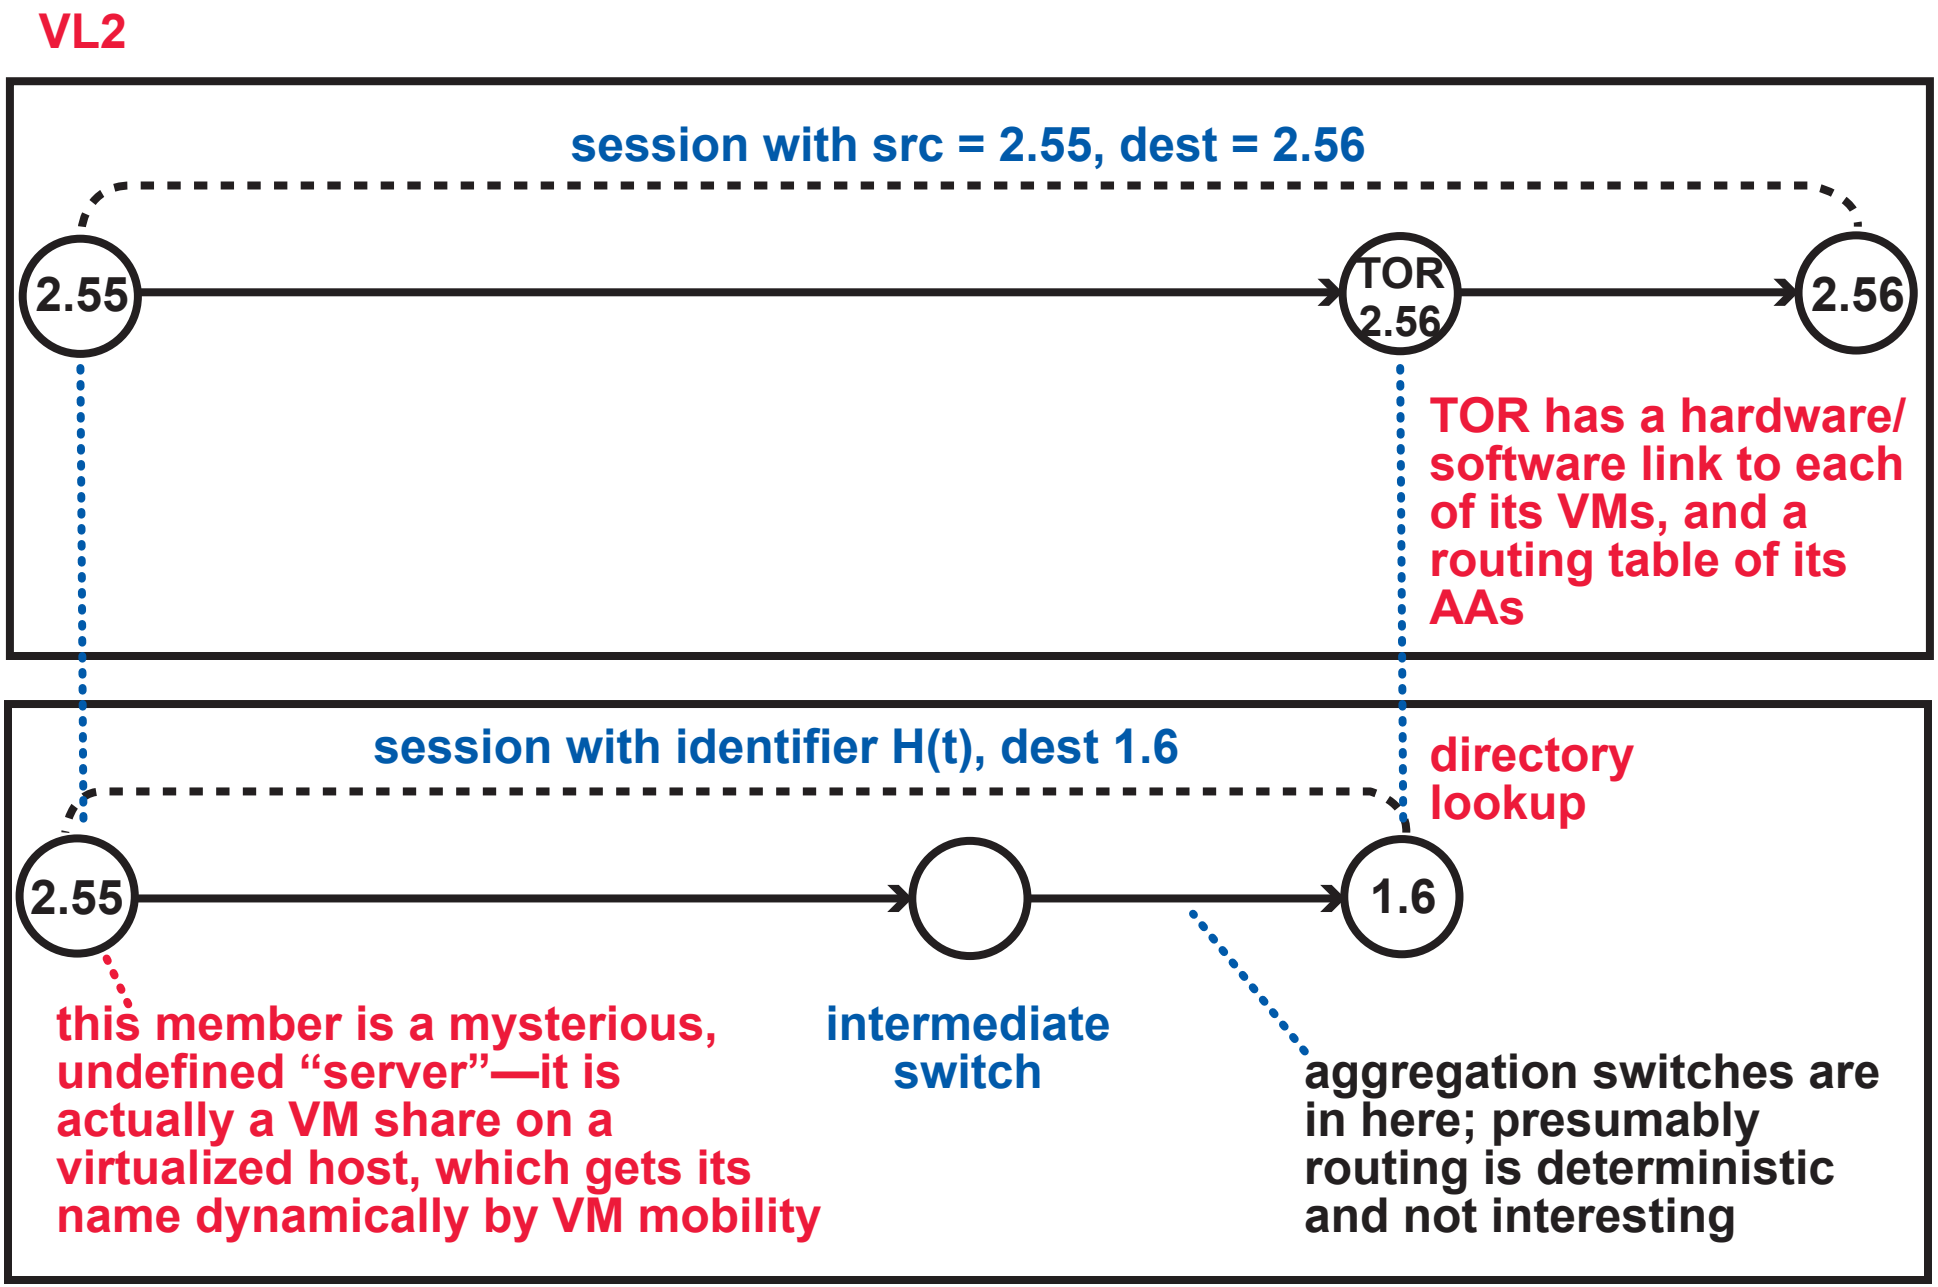

VL2 ARCHITECTURE

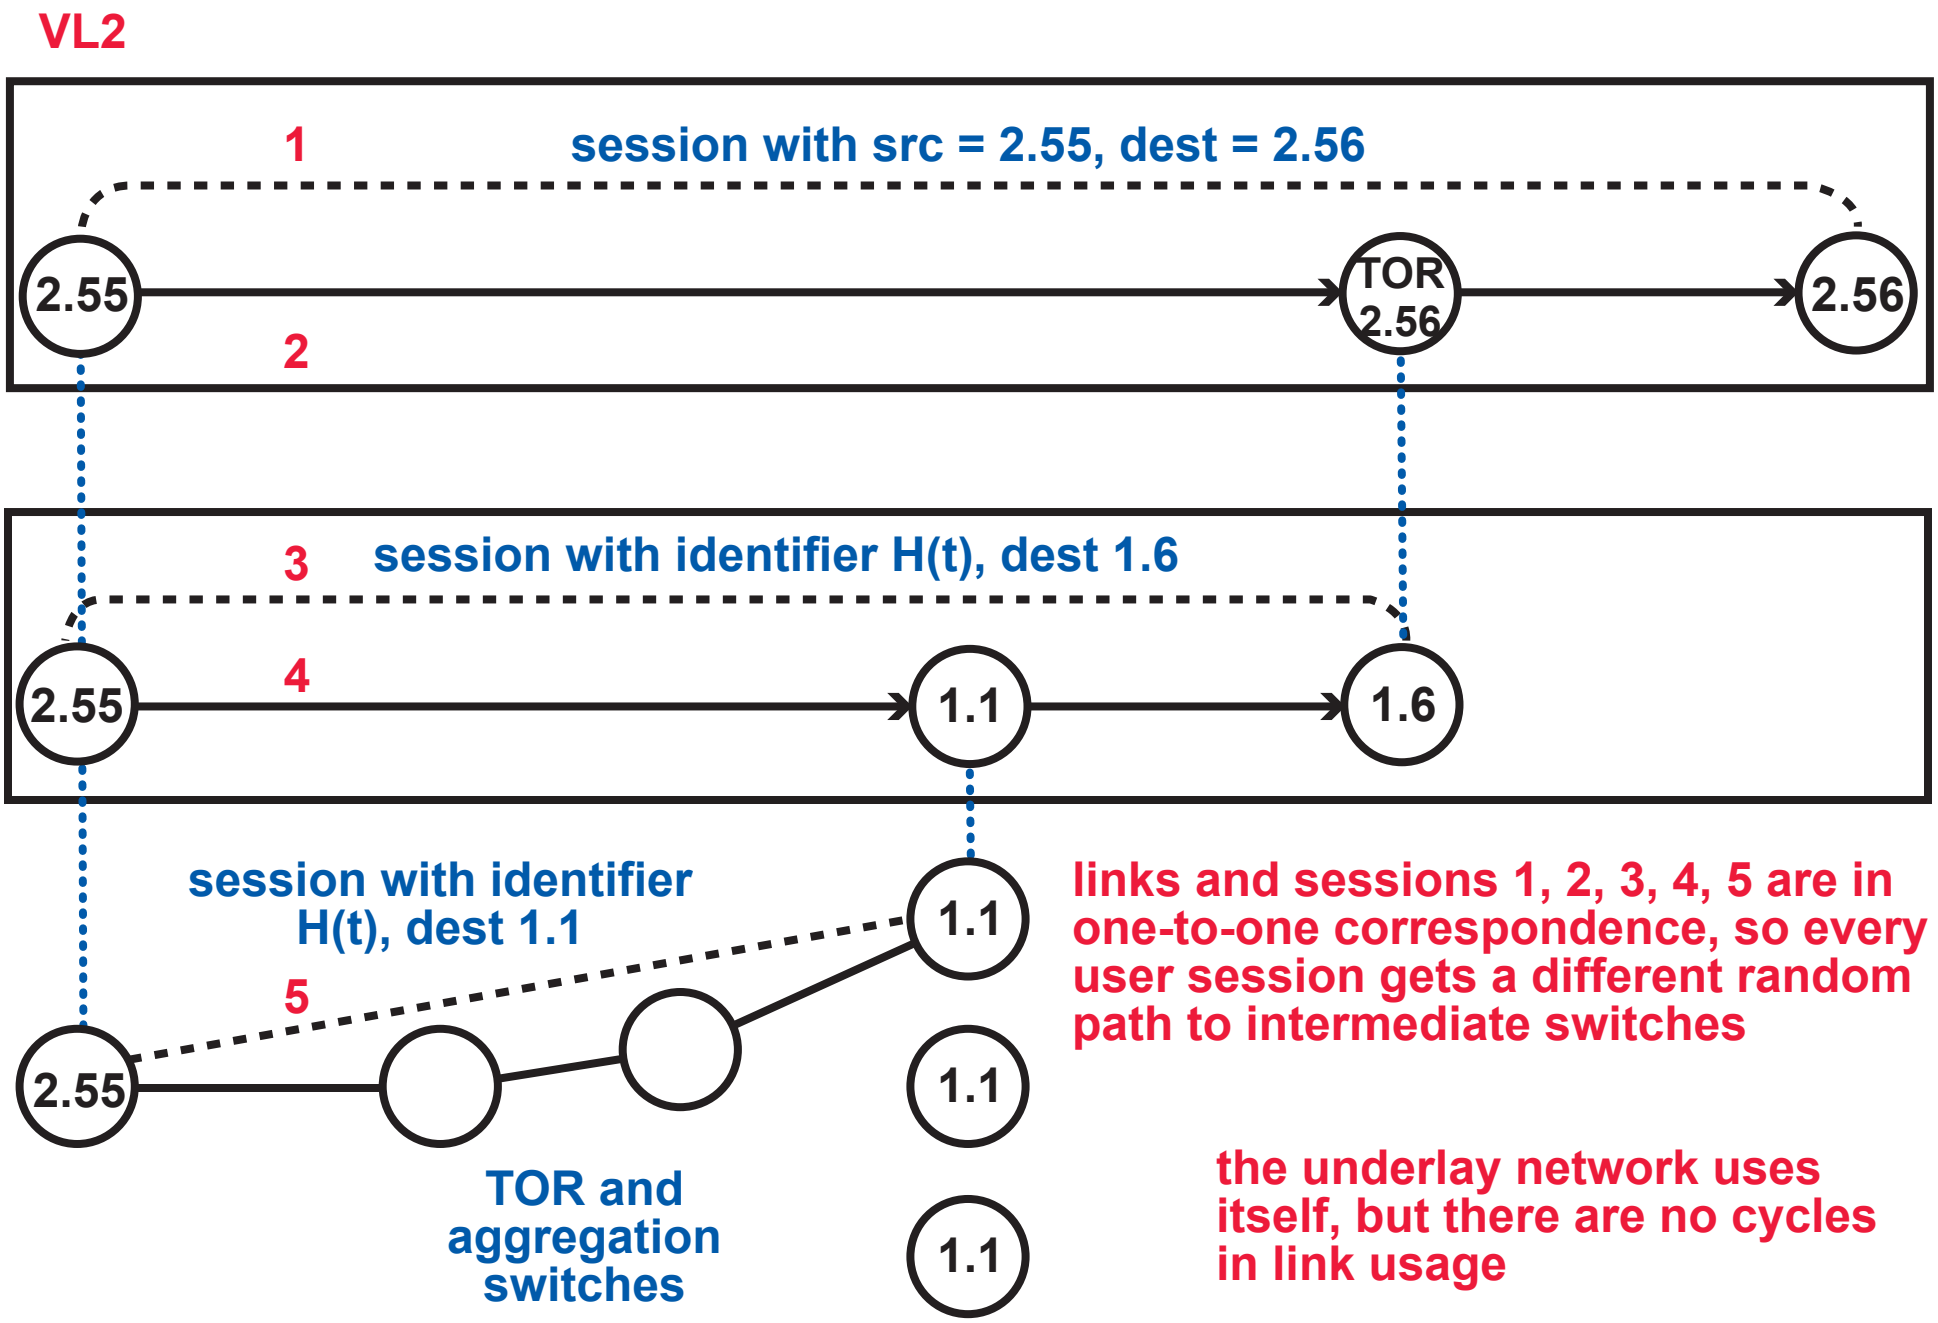

SEATTLE ARCHITECTURE

Ethernet hosts are connected to switches

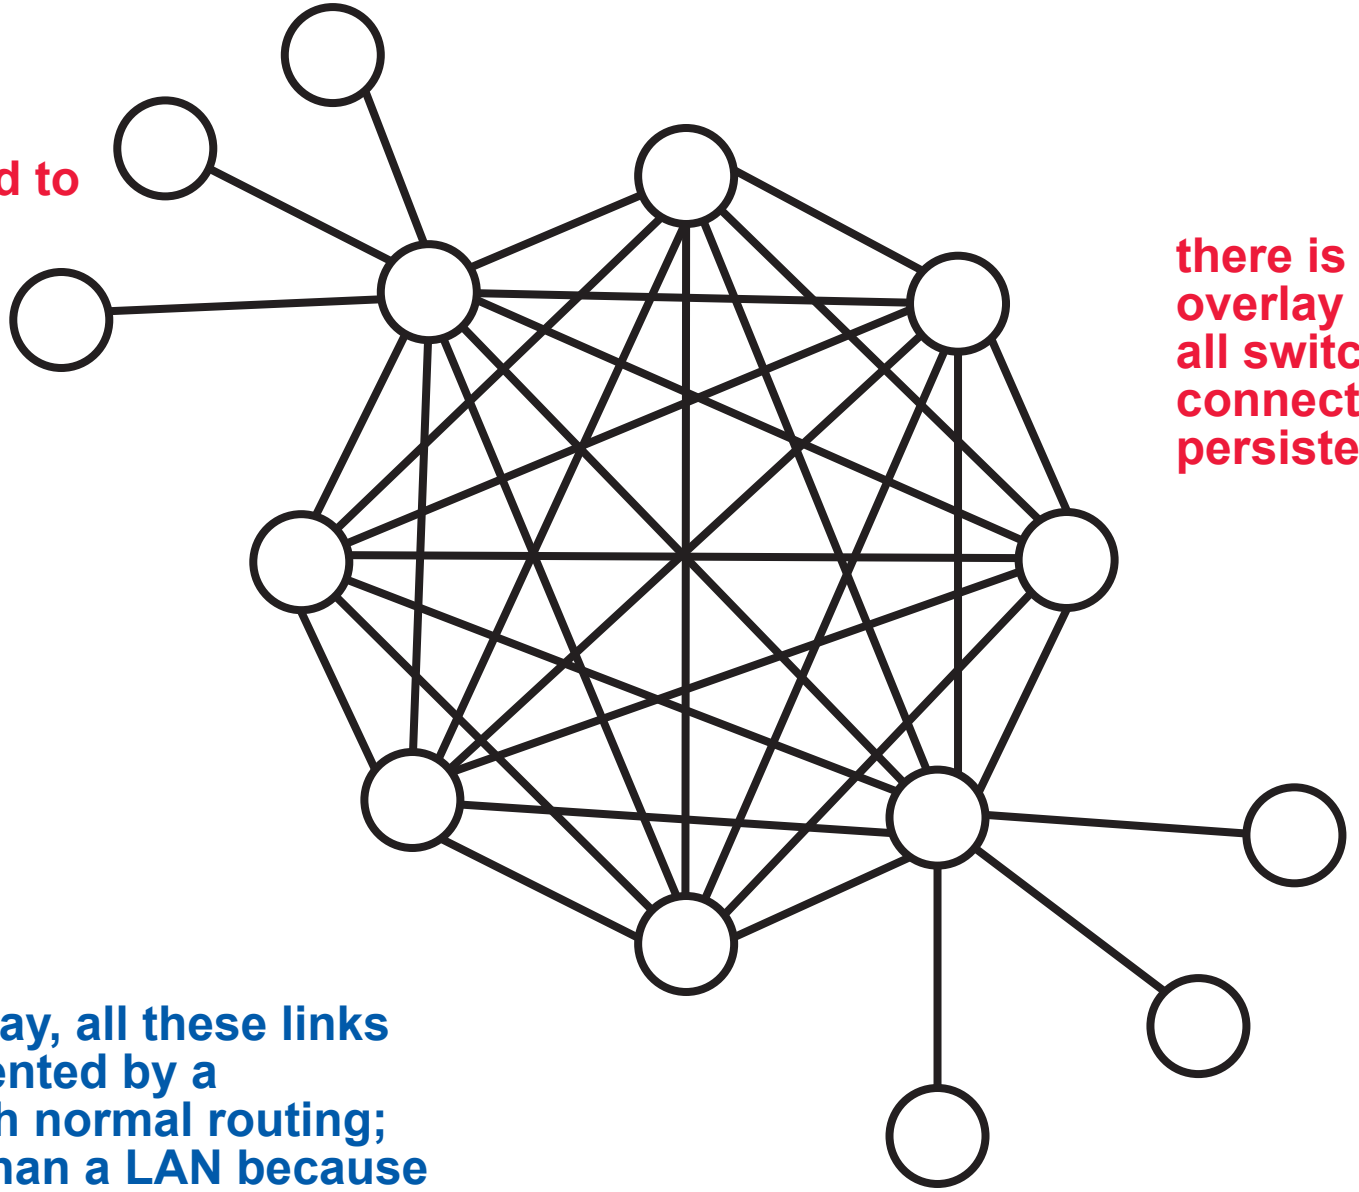

there is an overlay in which all switch pairs are connected by direct, persistent links

in an underlay, all these links are implemented by a network with normal routing; it is better than a LAN because links need not be confined to a spanning tree

SEATTLE ARCHITECTURE

as in a normal Ethernet each switch has a sparse routing table, containing only entries for destinations it is currently communicating with

as in a normal Ethernet, a switch gets a route to a new destination when it needs one

unlike a normal Ethernet, a switch gets a route to a new destination by looking it up in a directory

(it cannot flood, and this is more efficient)

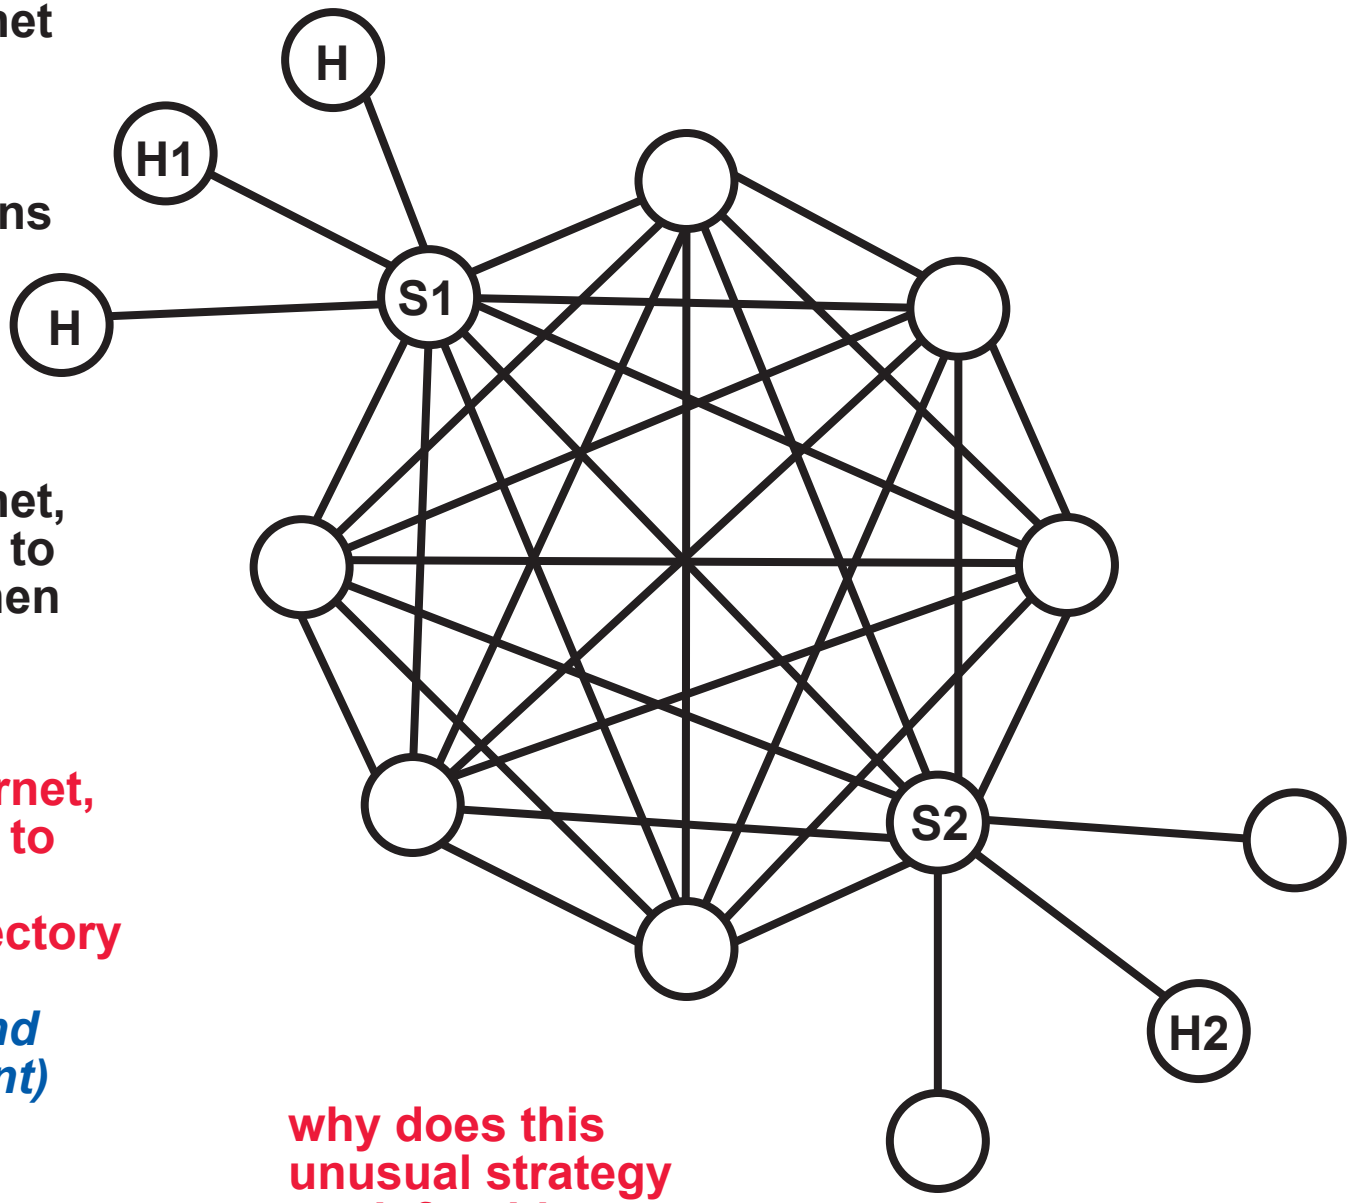

why does this unusual strategy work for this architecture?

SEATTLE ARCHITECTURE

unlike a normal Ethernet,
a switch gets a route to
a new destination by
looking it up in a directory

why does this
unusual strategy
work for this
architecture?

1 normally, a network has
many different routes to
a destination, used by
different sources

2 in this case each outlink from
a switch is identified with the
MAC address of the switch
at the other end

with this scheme, each switch has
the same route to a particular switch,
and also the same route to a host on it

for all switches, the route to
H2 is S2

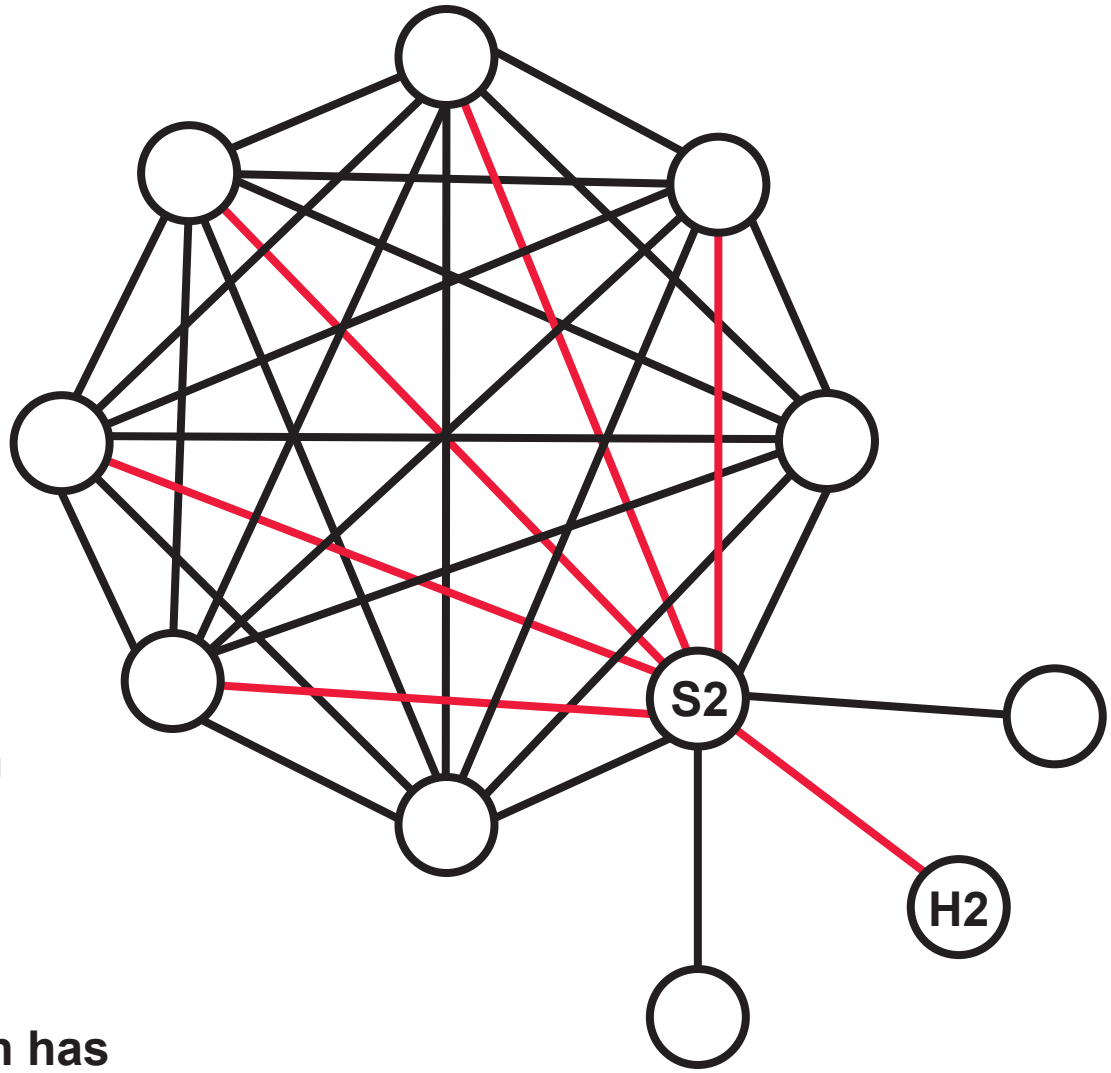

MIDDLEBOXES IN CLOUD COMPUTING

LOGICAL PROBLEMS OF SERVICE CHAINING

- routing loops
- large number of switch-level forwarding rules
- session affinity
- middleboxes that modify the 5-tuple used to identify packets
- middleboxes that classify packets

PROBLEMS OF DEPLOYMENT AND DYNAMIC RESOURCE ALLOCATION

- how is service chaining deployed in a cloud data center?
- what happens when load must be redistributed?
- what happens when a virtual machine migrates?

A CLOUD DESIGN

DESIGN GOALS

- accommodate clouds of the largest size

  10 data centers
  100 K hosts per data center,
  100 M virtual machines

- put in *all* the capabilities desirable in large-scale, multi-tenant clouds

NEW SOURCES AND COMPARISONS

- SIMPLE
- Stratos

SOME SOURCES

- CloudNaaS
  [Benson, Akella, Shaikh 11]

  tenant-specific address spaces, policy links

- VL2 *[Greenberg et al. 09]*

  identifier/locator split, IP routing in cloud layer

- WL2 *[Chen, Liu, Liu, Loo, Ding 14]*

  multiple data centers, VM migration

- OpenStack

  tenant-specific links

NETWORKS CONTRIBUTING TO THE CLOUD

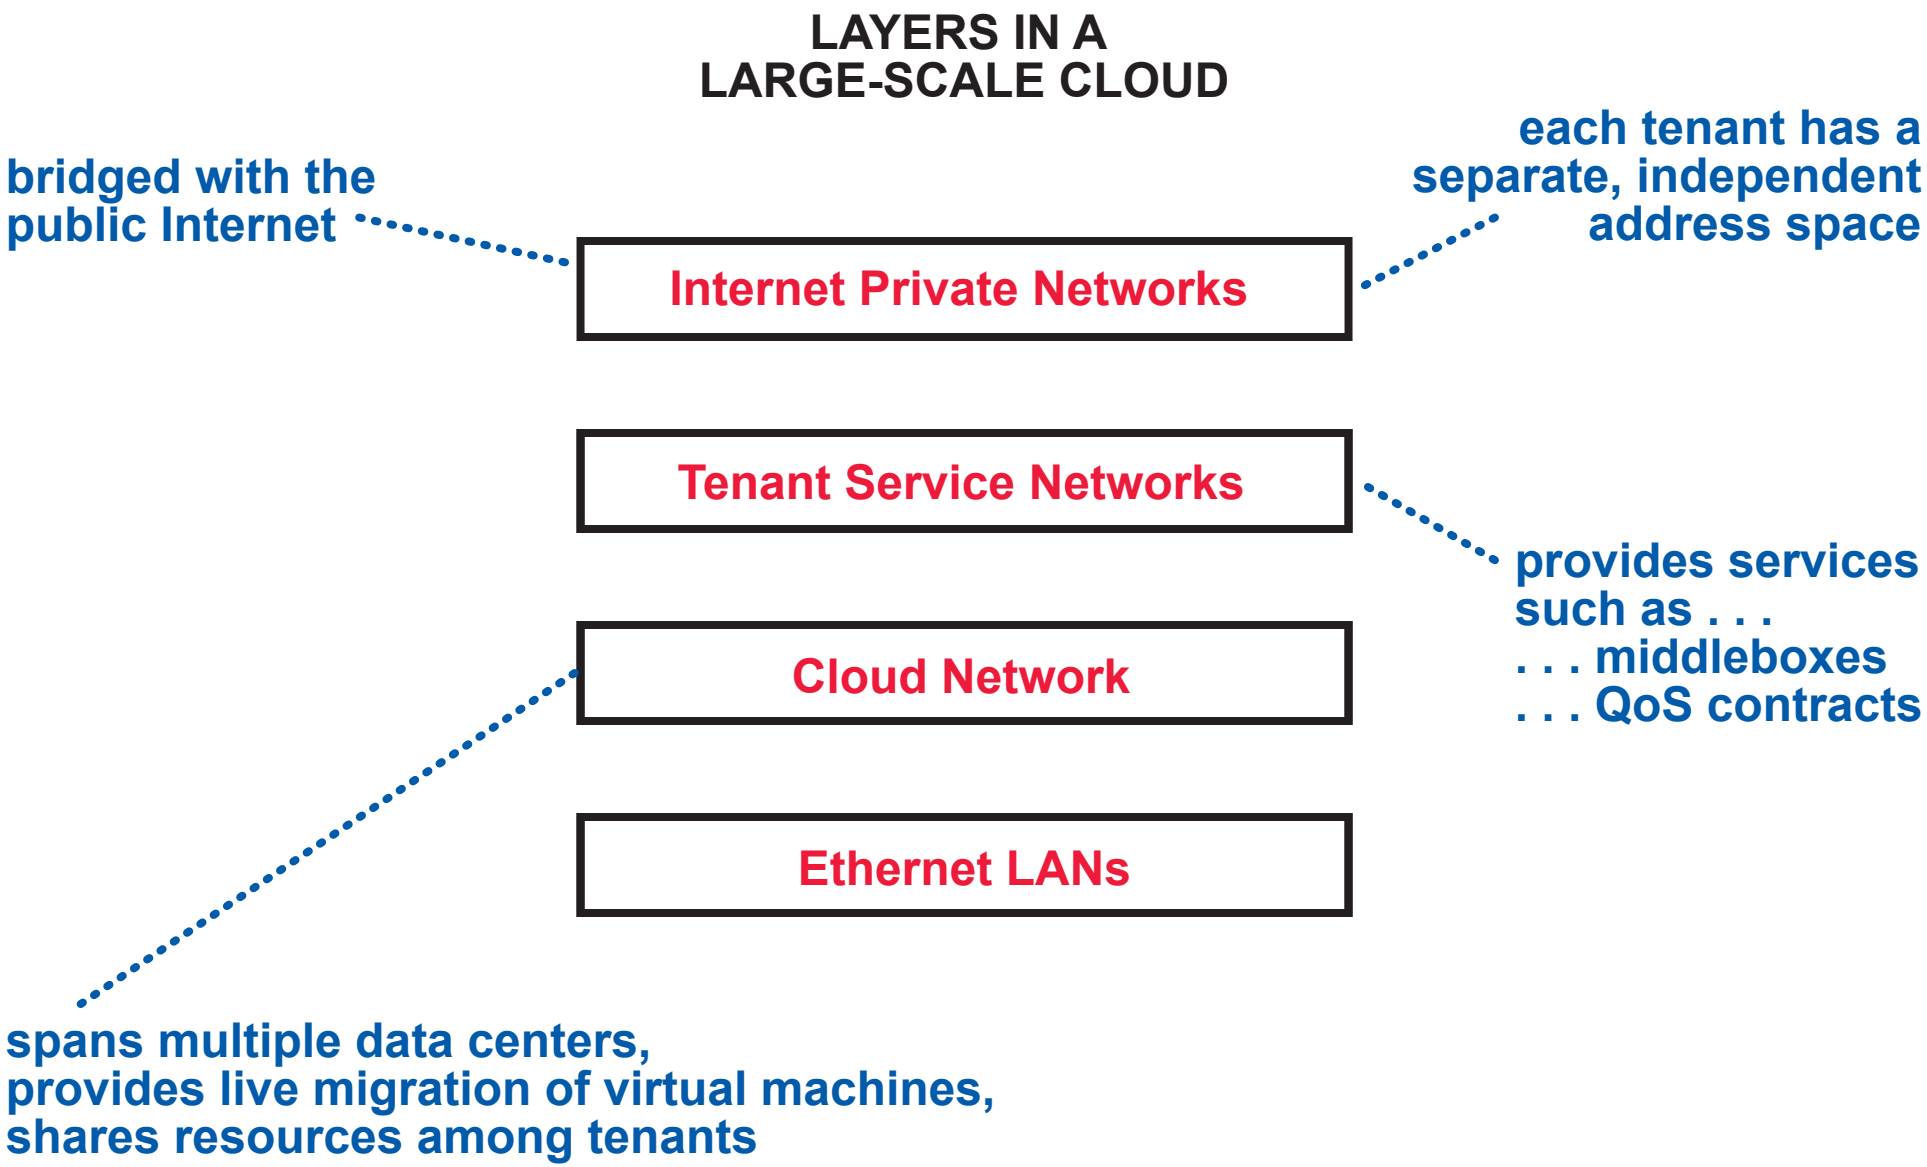

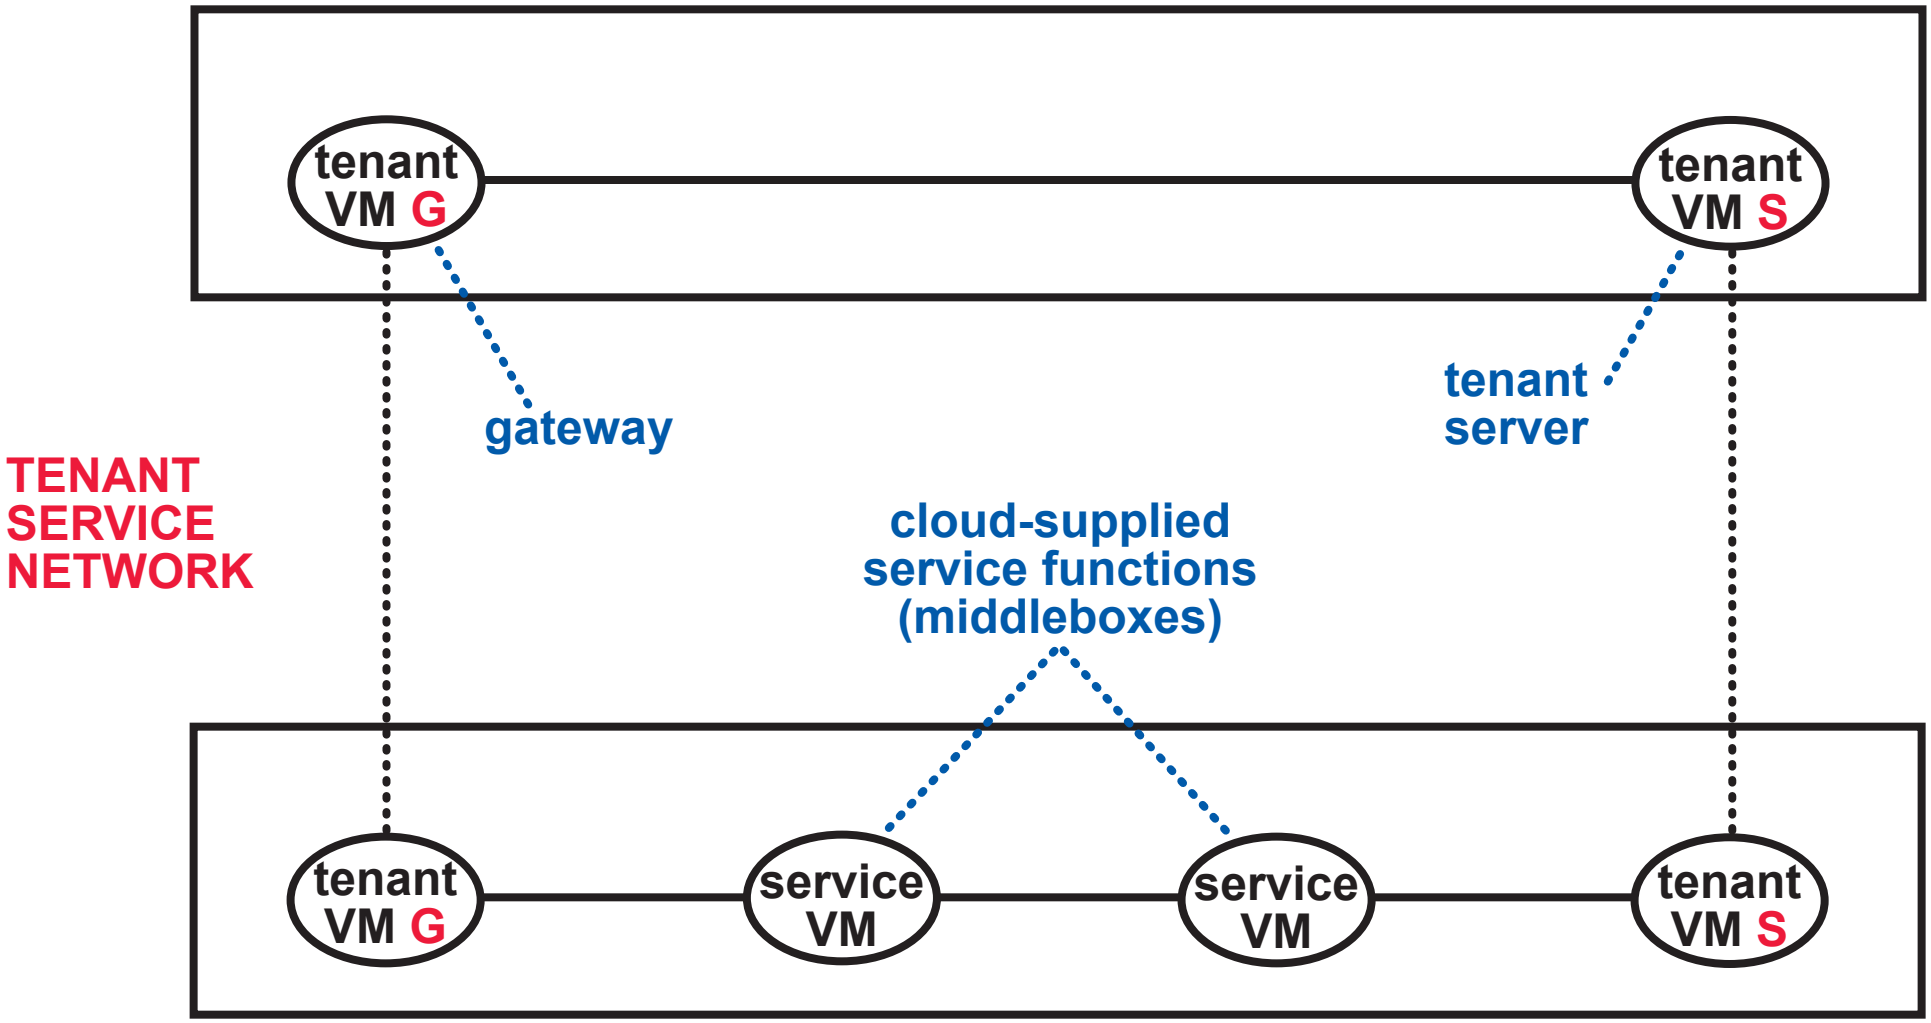
TENANT PRIVATE
INTERNET NETWORK
tenant
VM G
tenant
VM S
gateway
tenant
server
TENANT
SERVICE
NETWORK
cloud-supplied
service functions
(middleboxes)
tenant
VM G
service
VM
service
VM
tenant
VM S

TENANT PRIVATE
INTERNET NETWORK

for each tenant, VL2 lumps the two networks together

VL2 paper does not say how tenant is provided with expected IP addresses

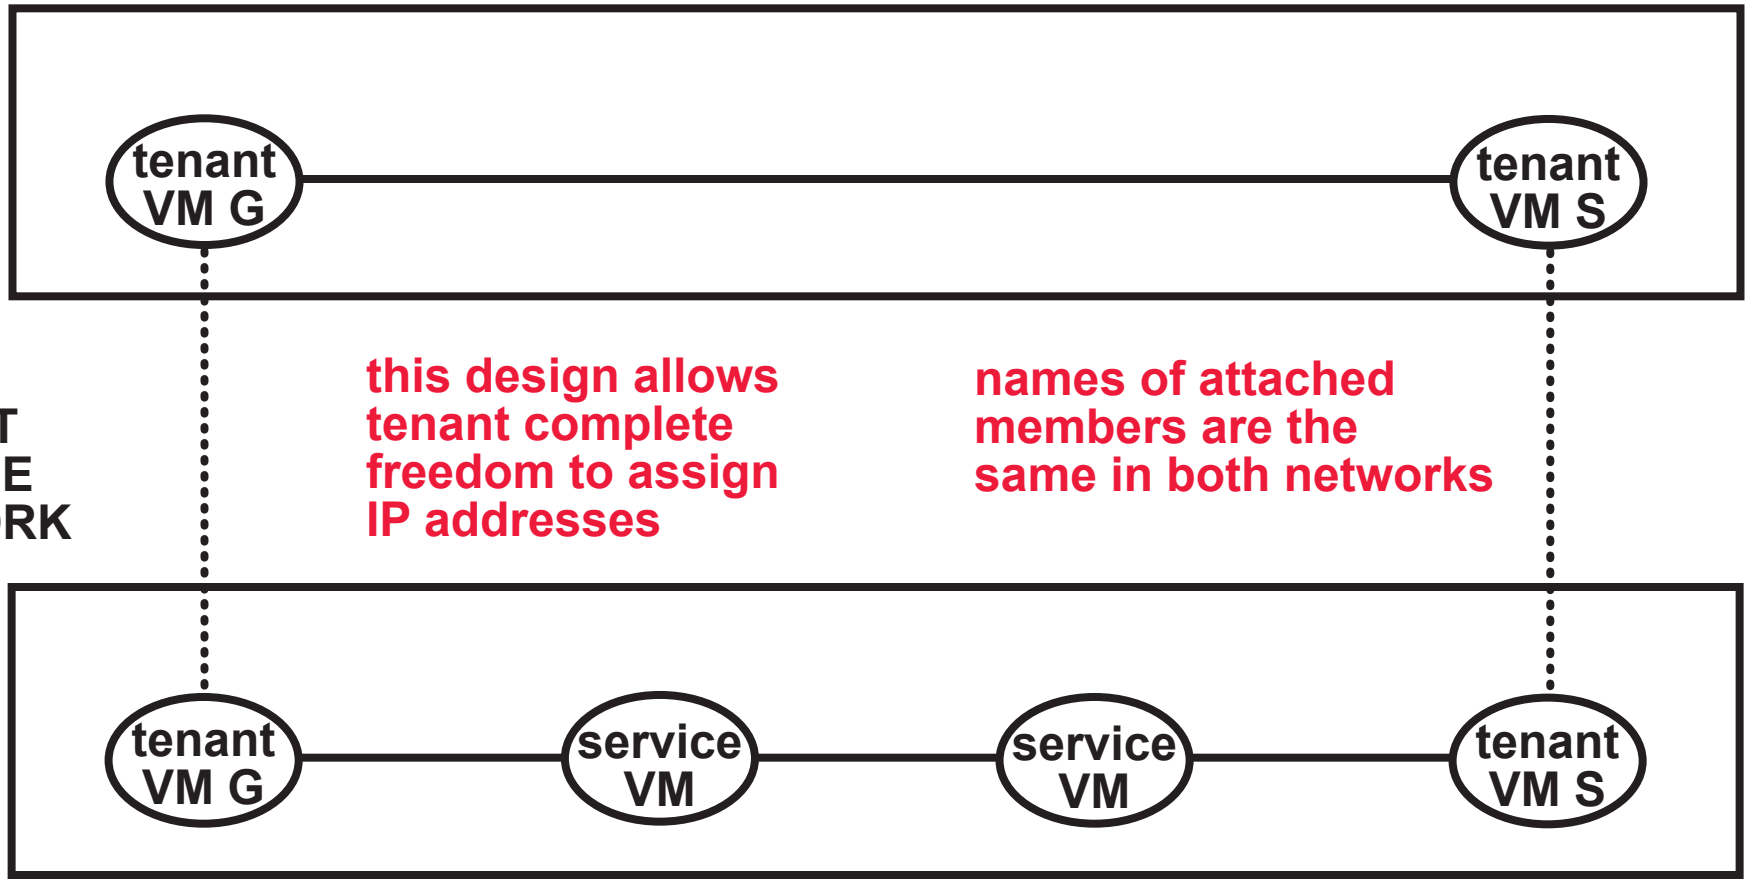

Stratos lumps both networks, for all tenants, together

Stratos paper does not say how IP addresses are shared by tenants

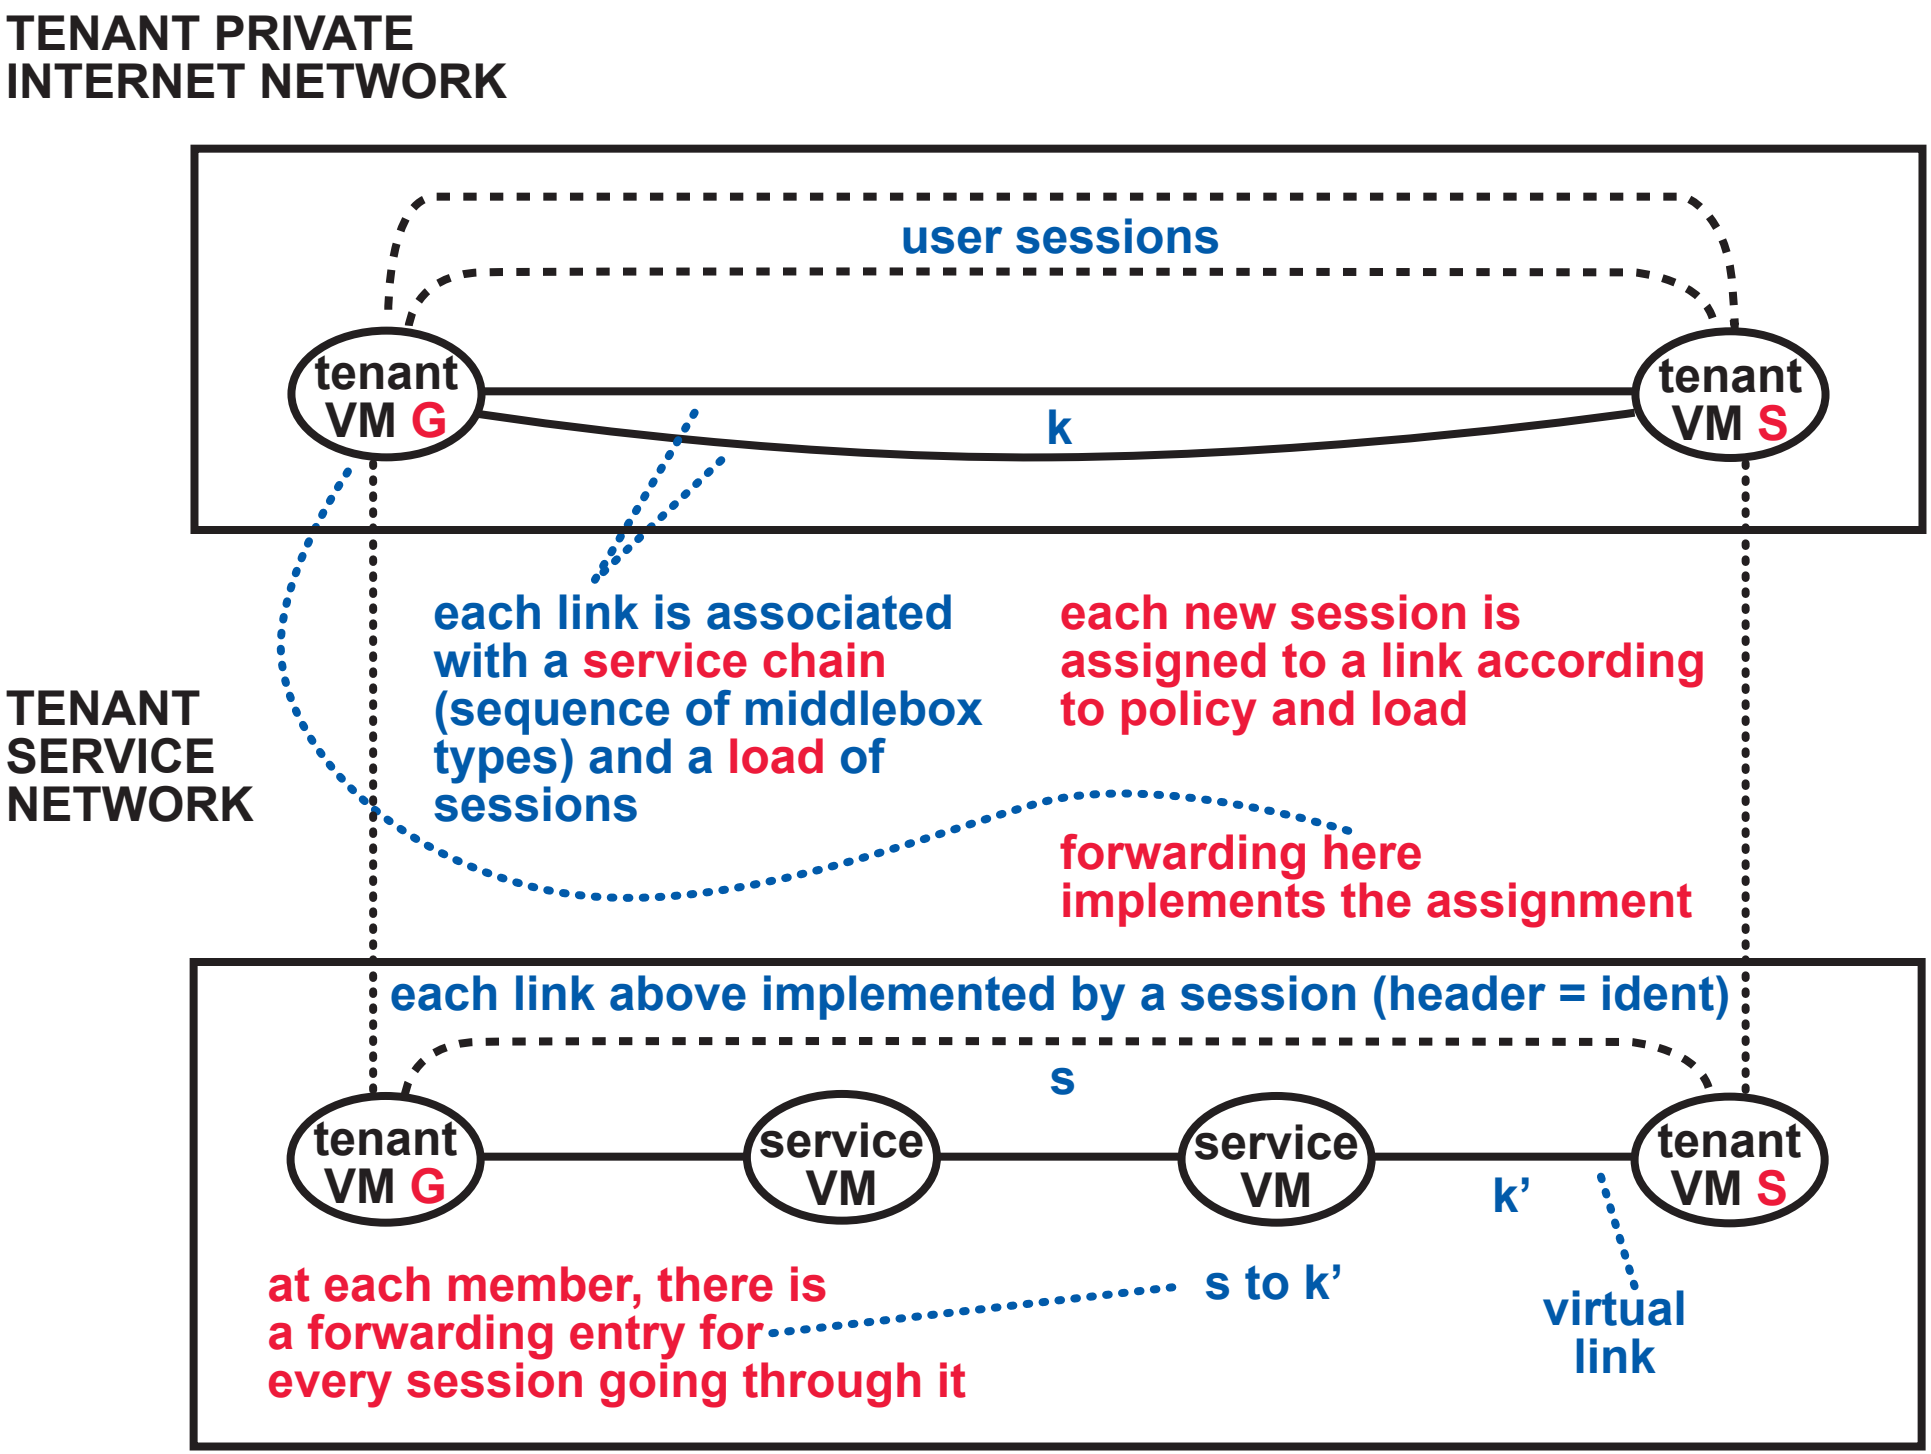
TENANT PRIVATE
INTERNET NETWORK
user sessions
tenant
VM G
k
tenant
VM S
each link is associated
with a service chain
(sequence of middlebox
types) and a load of
sessions
each new session is
assigned to a link according
to policy and load
TENANT
SERVICE
NETWORK
forwarding here
implements the assignment
each link above implemented by a session (header = ident)
s
tenant
VM G
service
VM
service
VM
tenant
VM S
k'
at each member, there is
a forwarding entry for
every session going through it
s to k'
virtual
link

TENANT PRIVATE
INTERNET NETWORK

VL2 paper does not have
service chaining

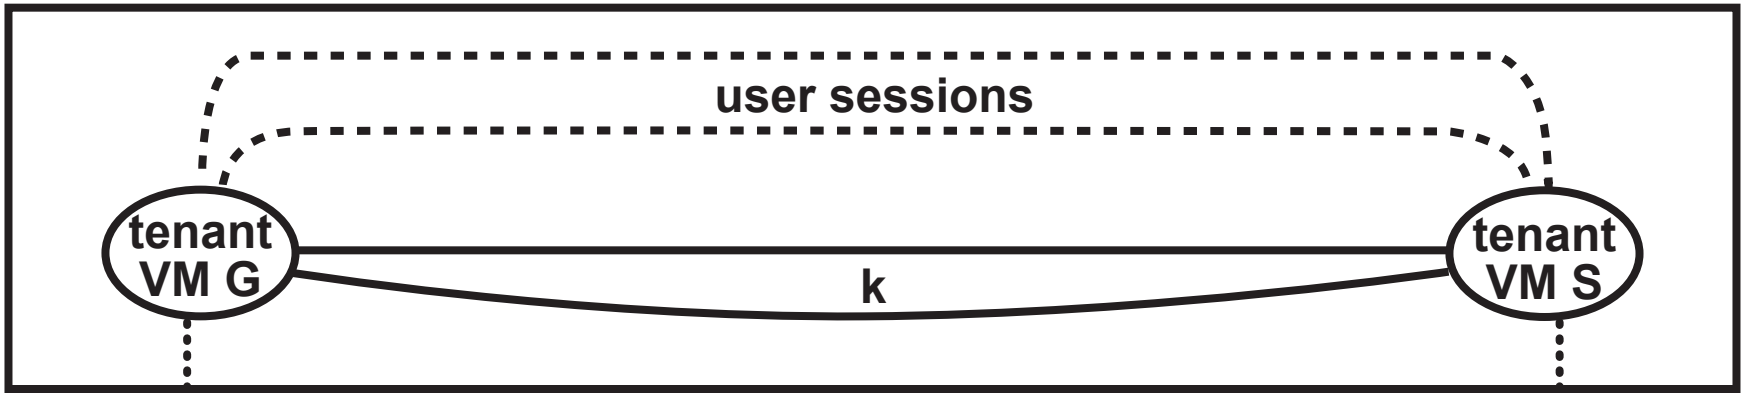

TENANT
SERVICE
NETWORK

assignment of individual
user sessions to the
“flow” that is session s
provides redistribution
of load with session
affinity

middleboxes cannot
do packet classification

only Dysco allows this

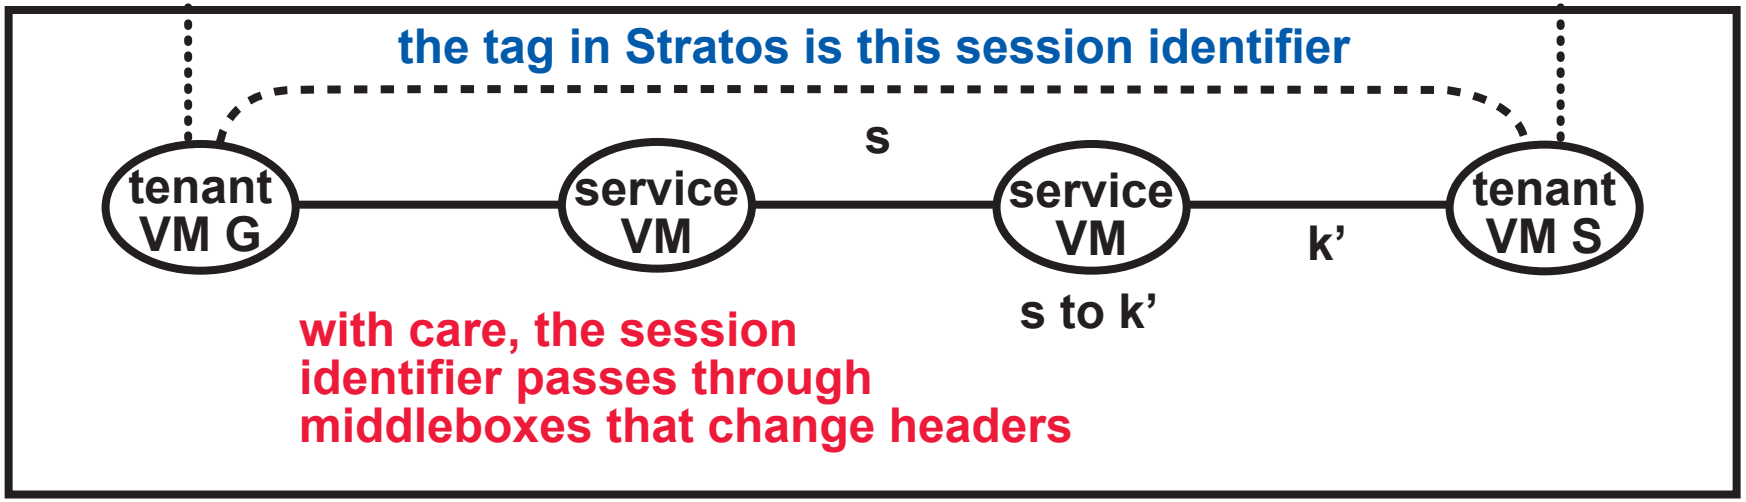

TENANT SERVICE NETWORK

in the cloud design, virtual links go between middleboxes, so there are no routing loops

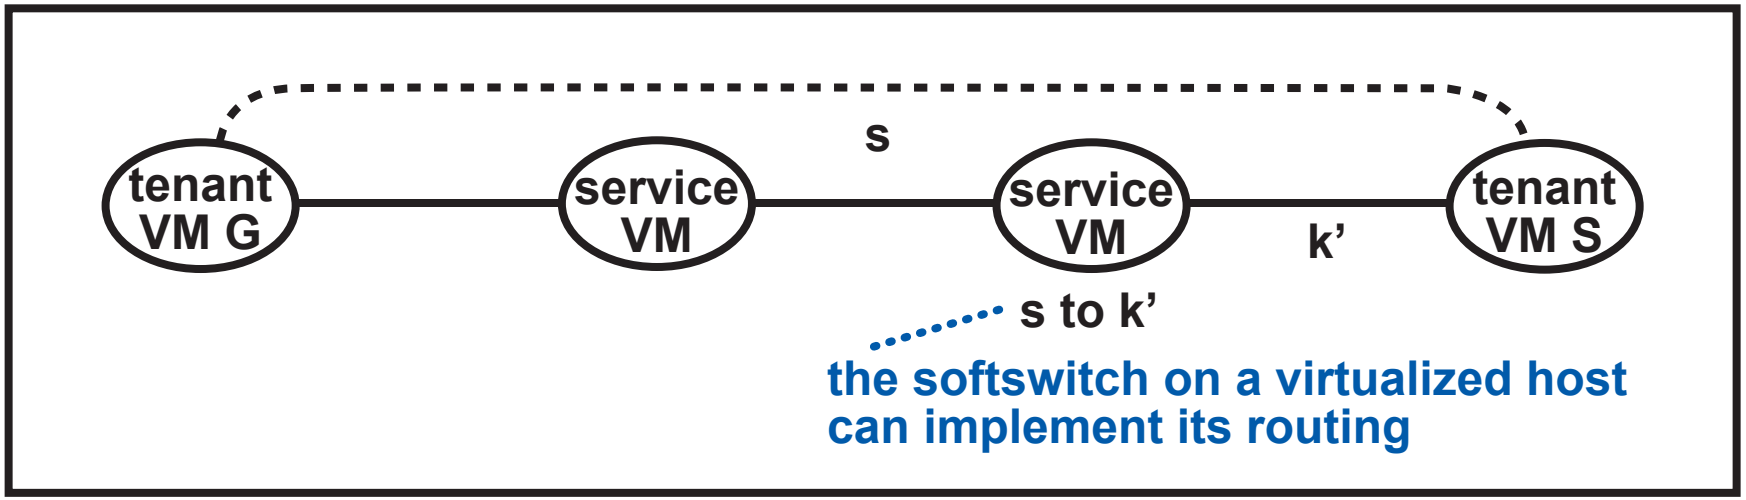

Stratos paper does not say what the switches are (soft-switches? TOR switches?)

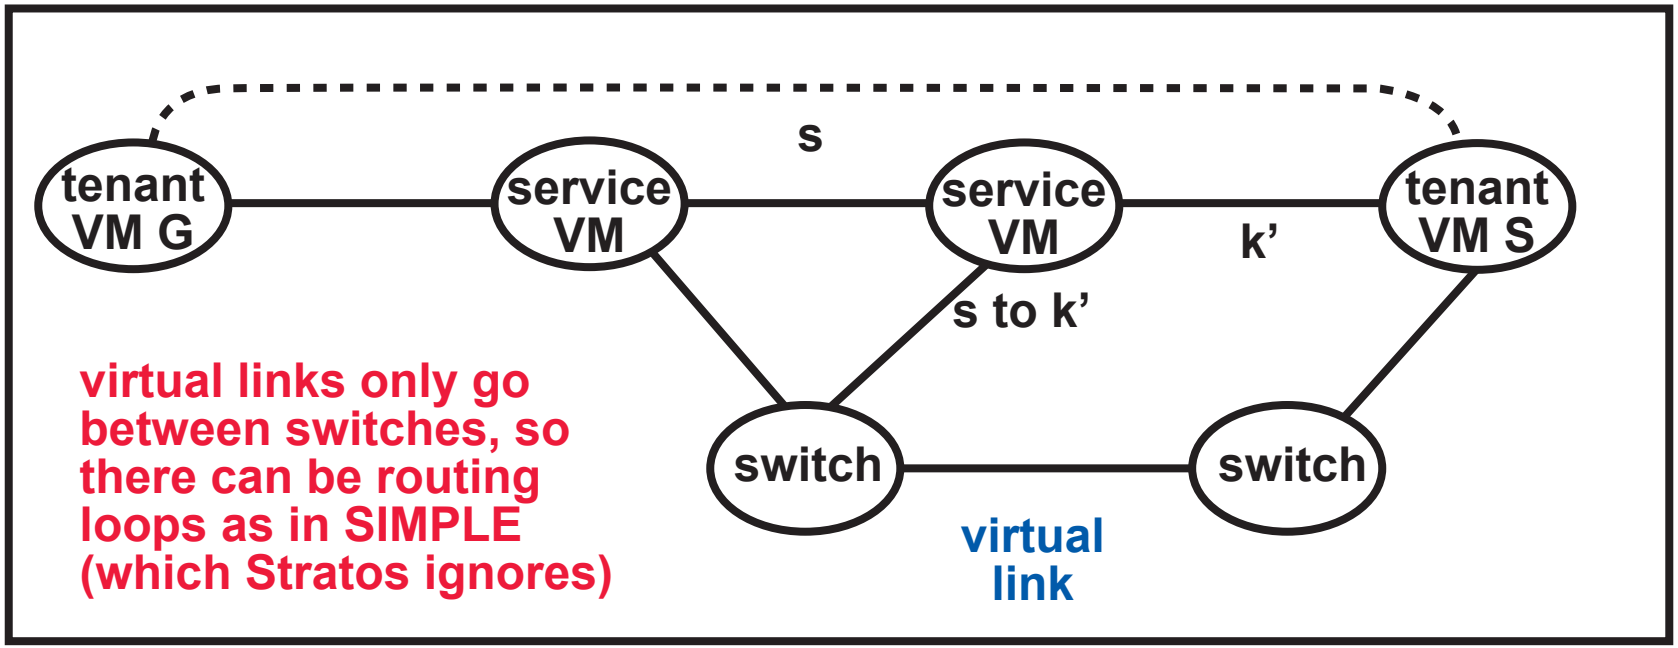

TENANT SERVICE NETWORK

Stratos has an underlay implementing virtual links between switches but it does not extend to middleboxes and does not provide for migration of VMs

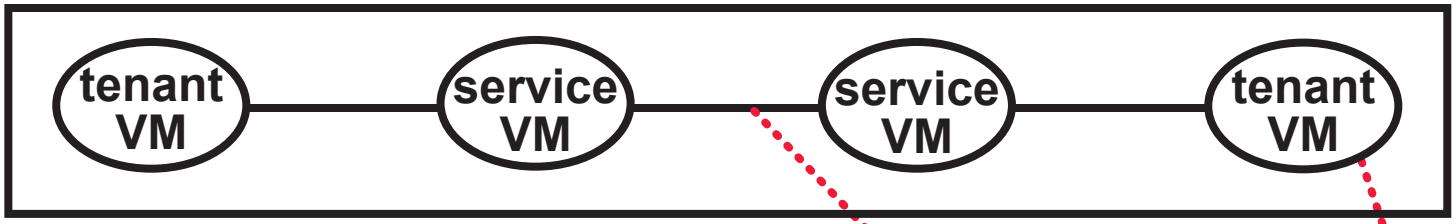

this is like VL2, except . . .

. . . location lookup is by (tenant, name)

implementation

location

. . . VL2 has much more detail about efficient communication within a data center

SHARED CLOUD NETWORK

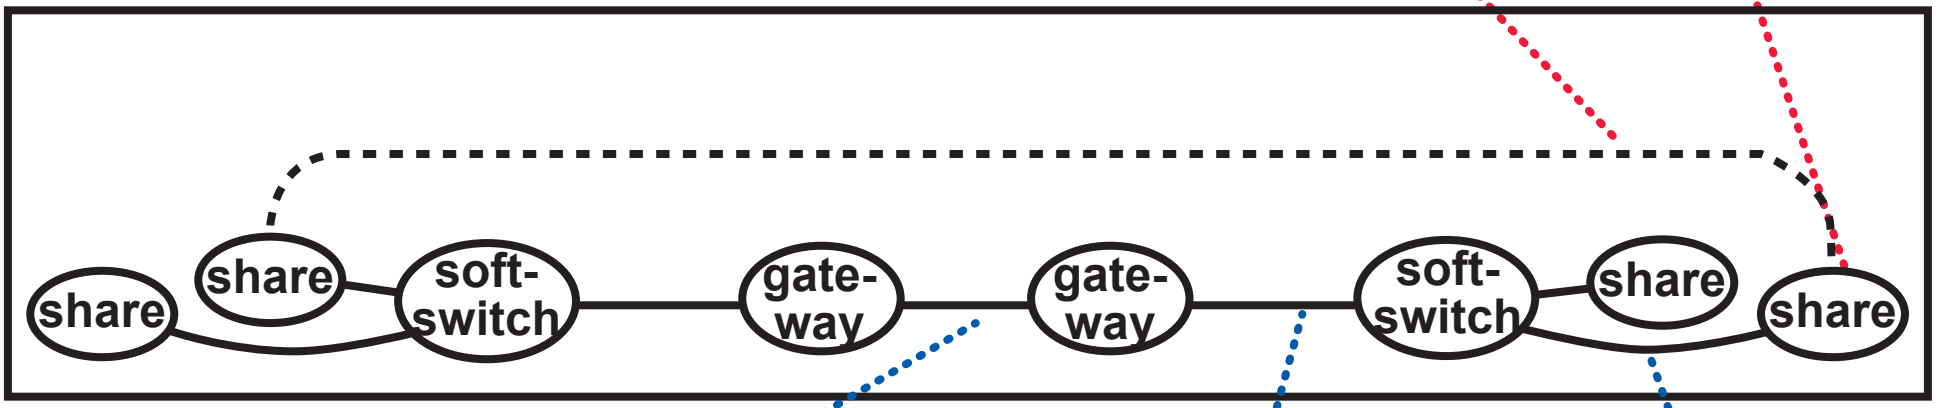

trunk between data centers

link inside a data center (TORs and other switches can be here, too)

implemented by hypervisor of shared machine

VM MIGRATION: CLOUD LAYER HIERARCHY

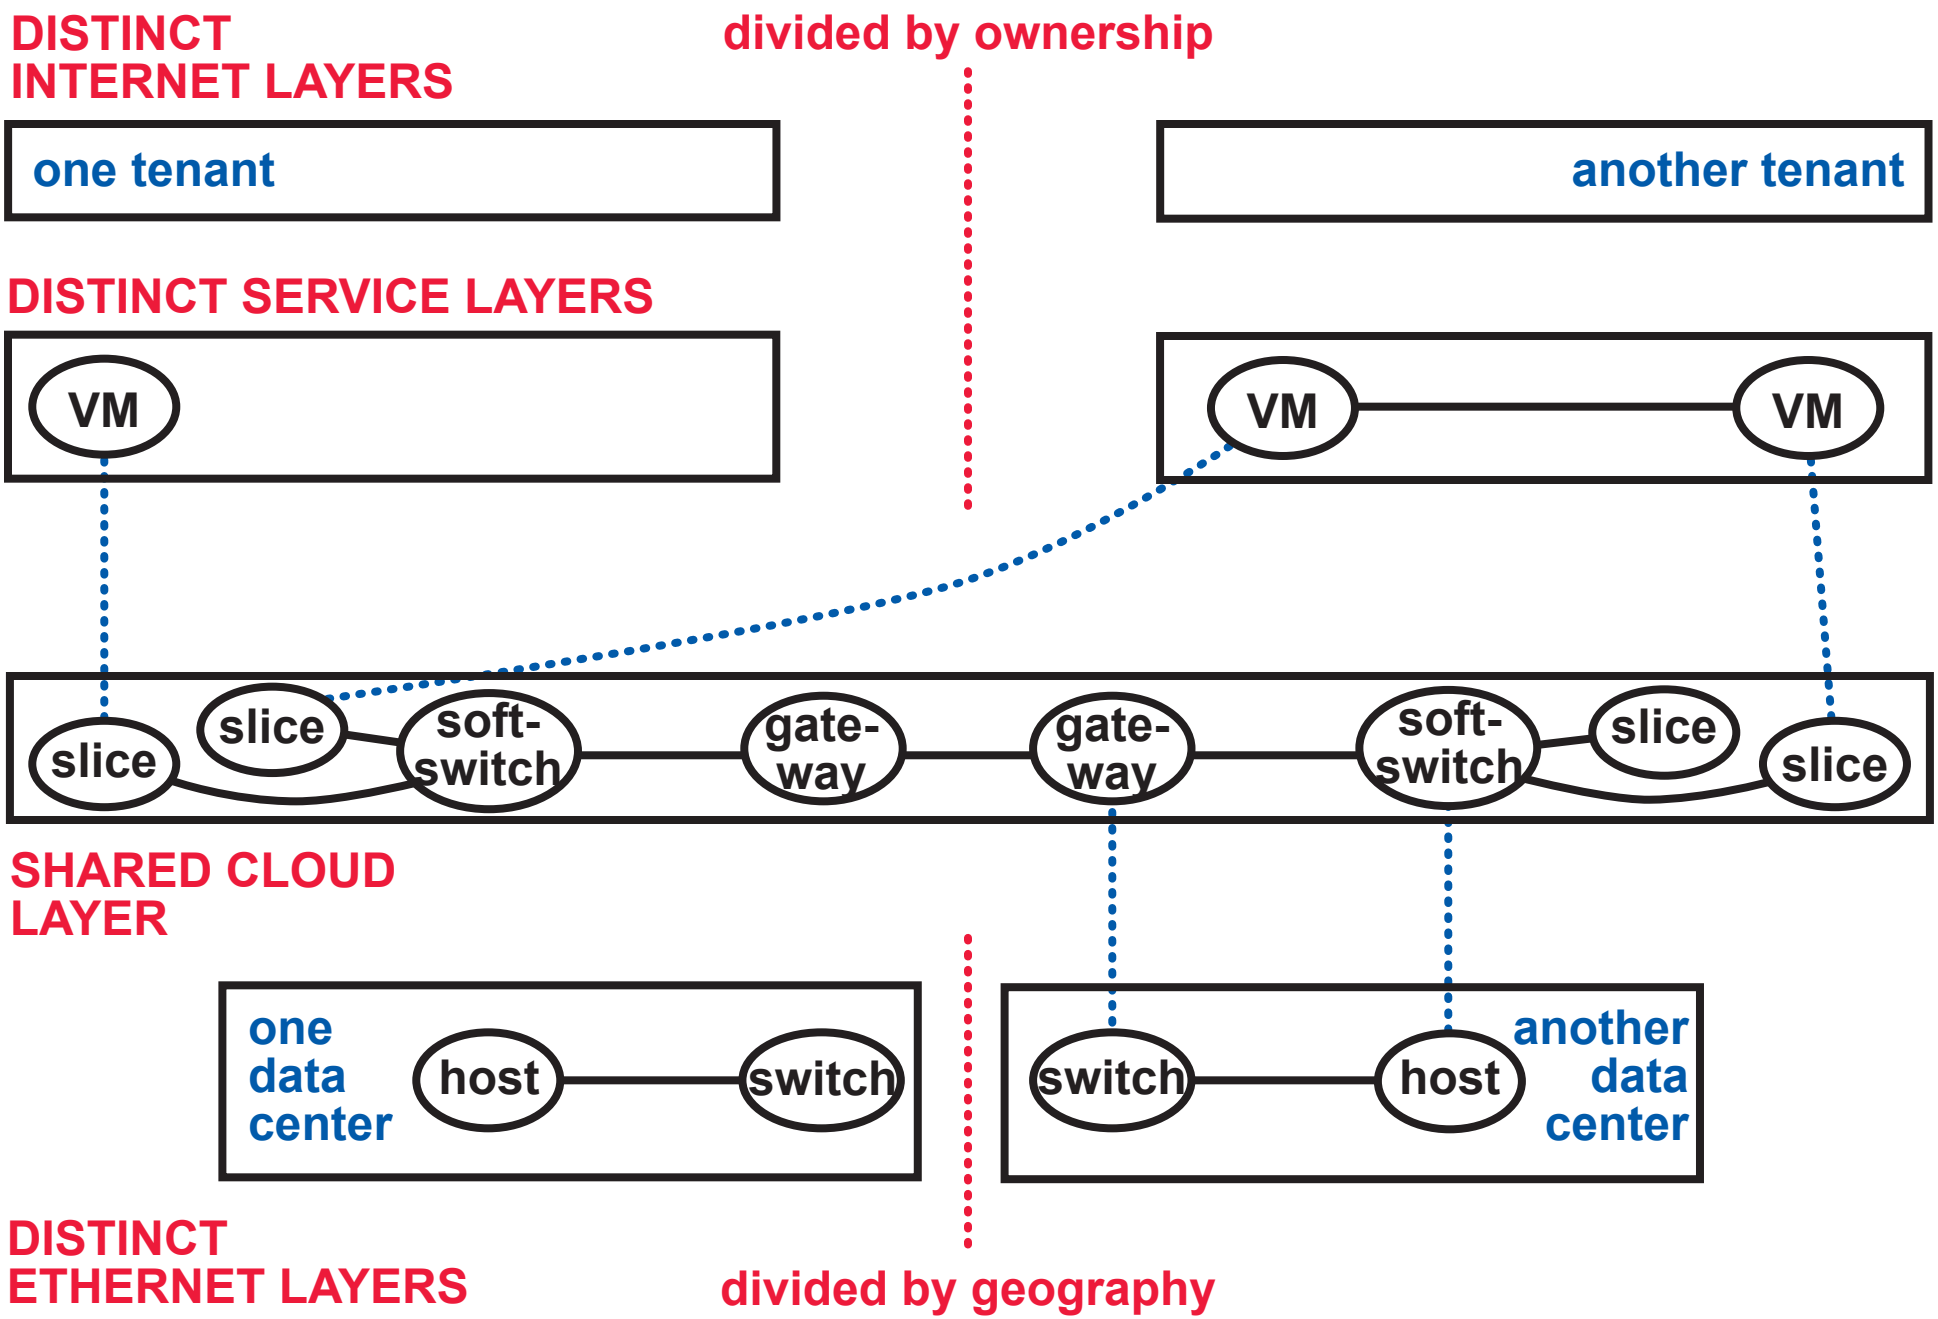

VM MIGRATION: MOBILITY

this shows a link in the service layer, and the session in the cloud layer that implements it

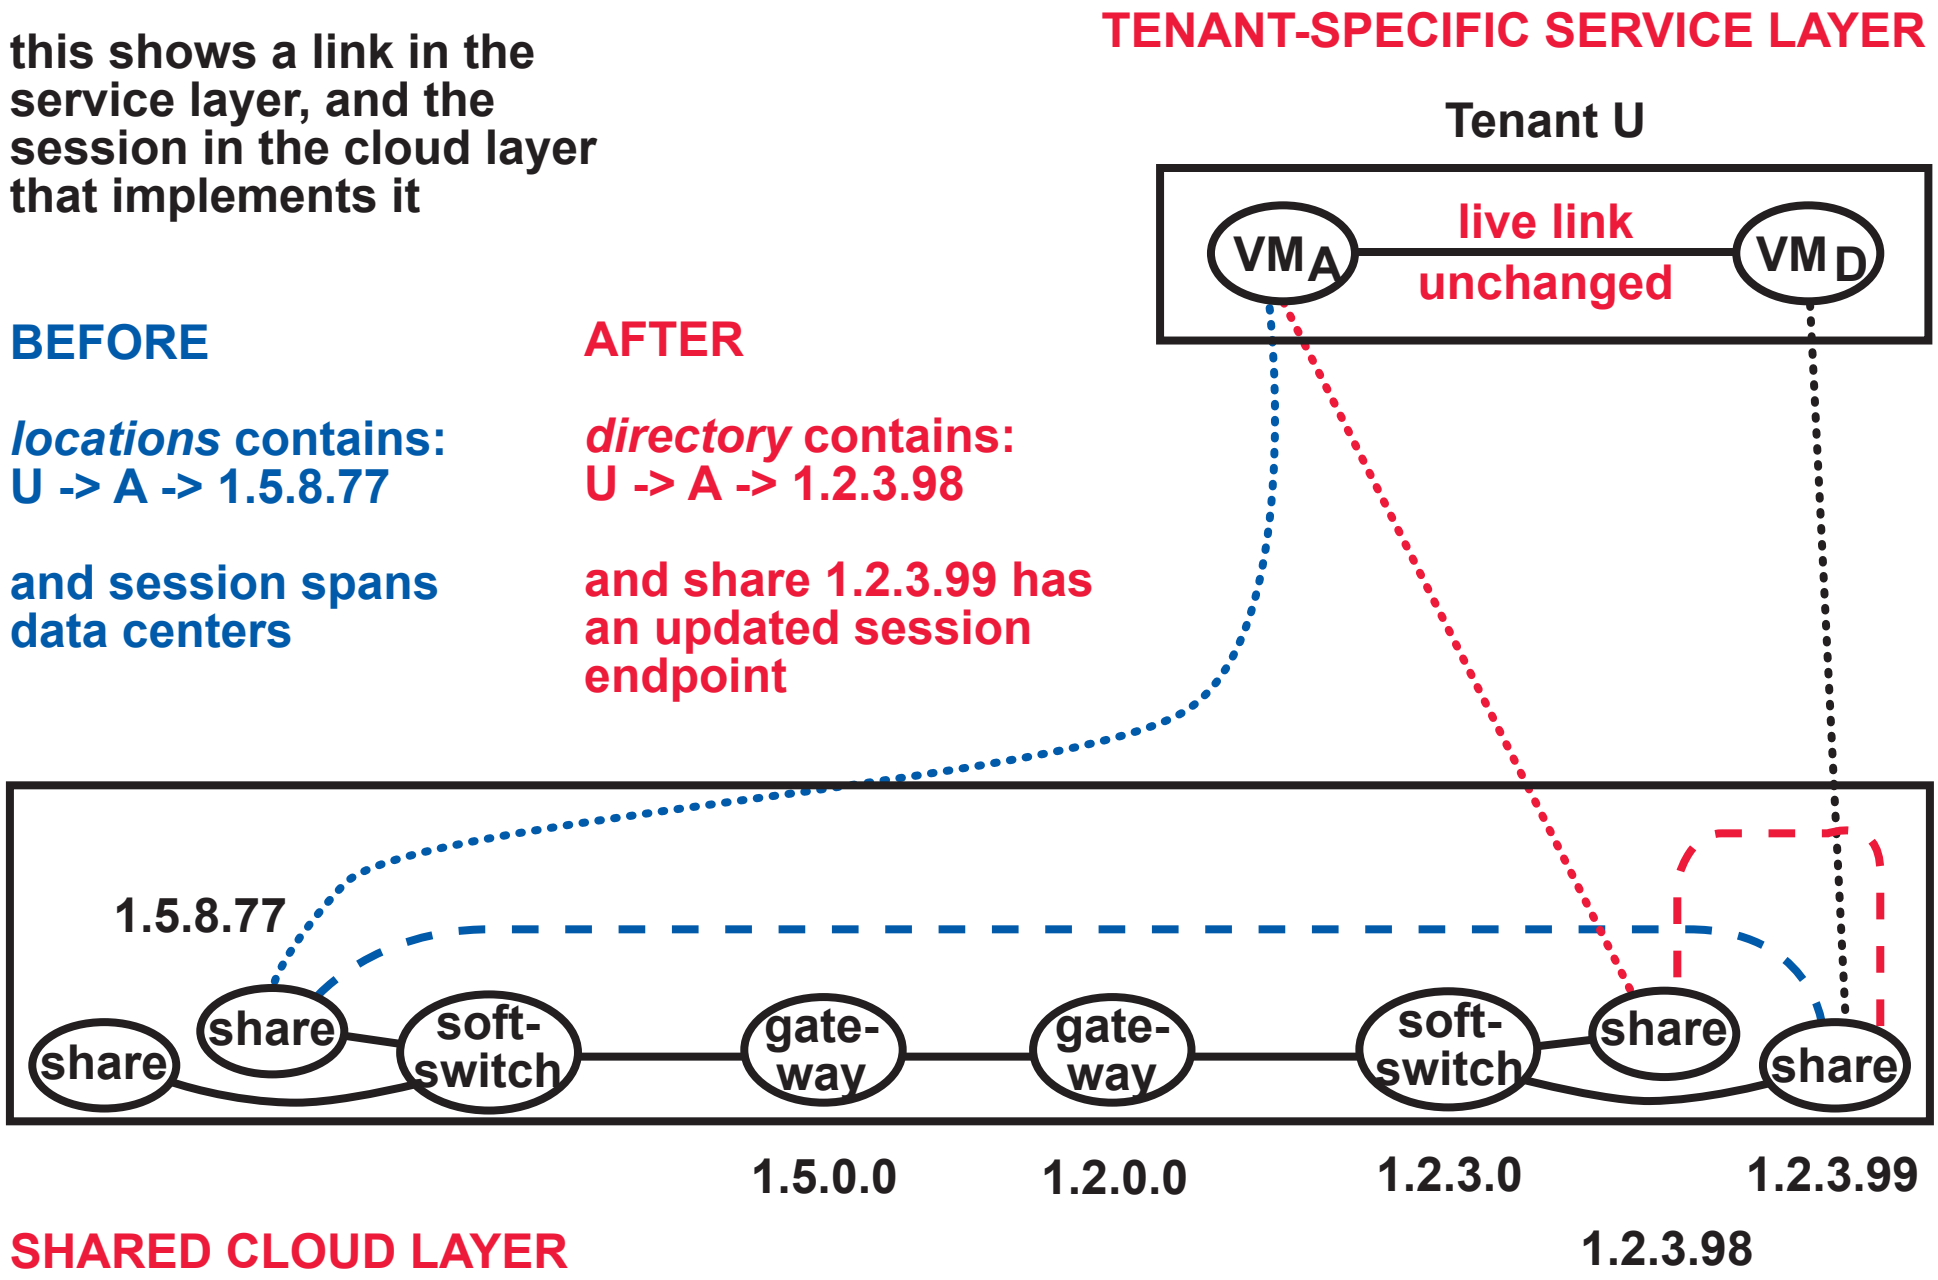

VM MIGRATION: A THREAT TO TENANT ISOLATION

we want to verify that a tenant's VM can never receive messages from another tenant's VM

enforcement is by means of tenant-specific links in the cloud layer (implemented on shared links in the Ethernet layers)

if it is proved that forwarding is limited to chains of links of the same tenant, tenant isolation should be guaranteed by this layer

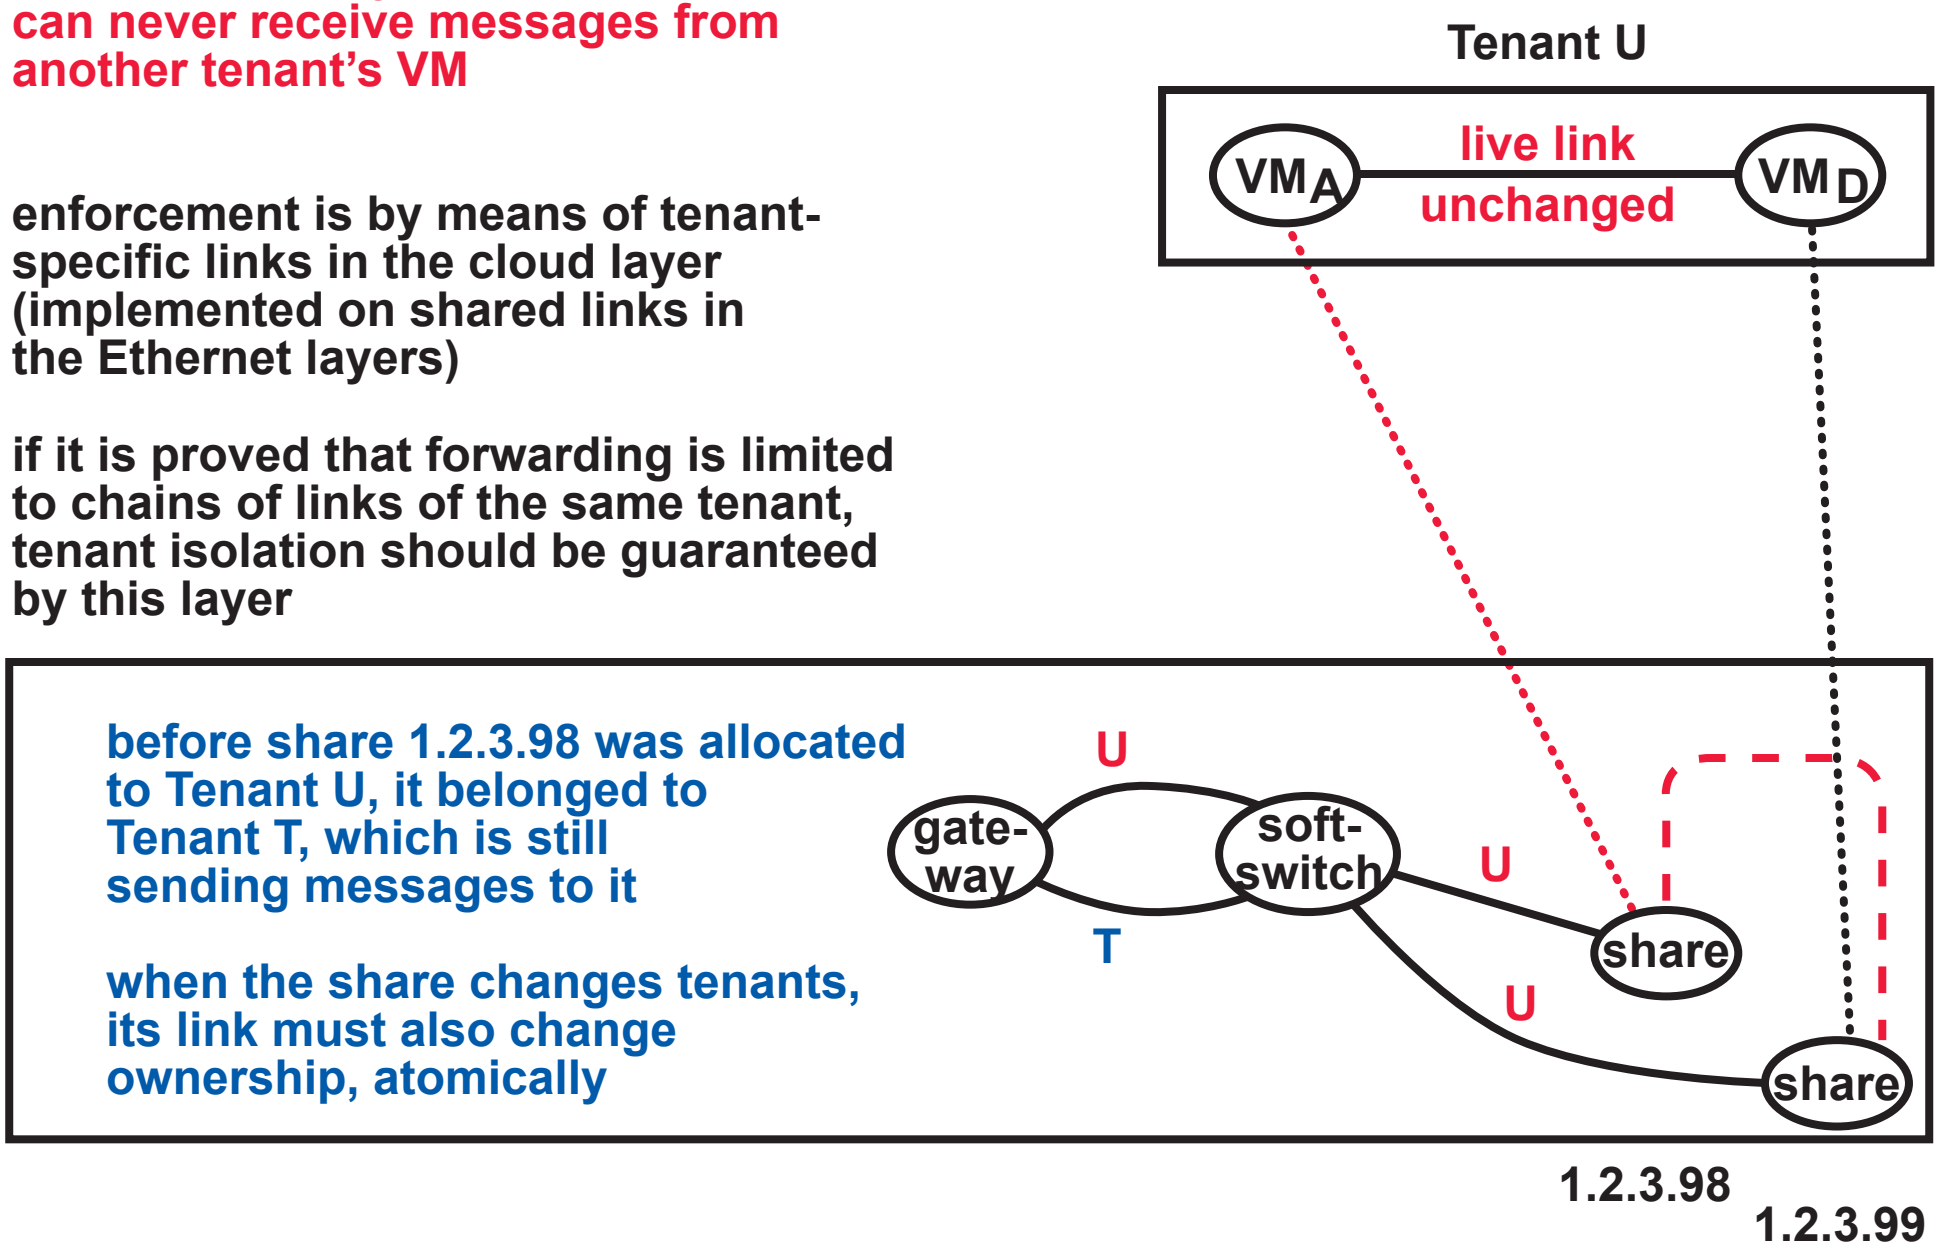

MIDDLEBOX POLICIES: UPDATES ARE CONSISTENT BY CONSTRUCTION

LAYERED CONTROL PROGRAMS

detect need for more capacity

create new policy link

create new service session

allocate middleboxes, create links and forwarding for session

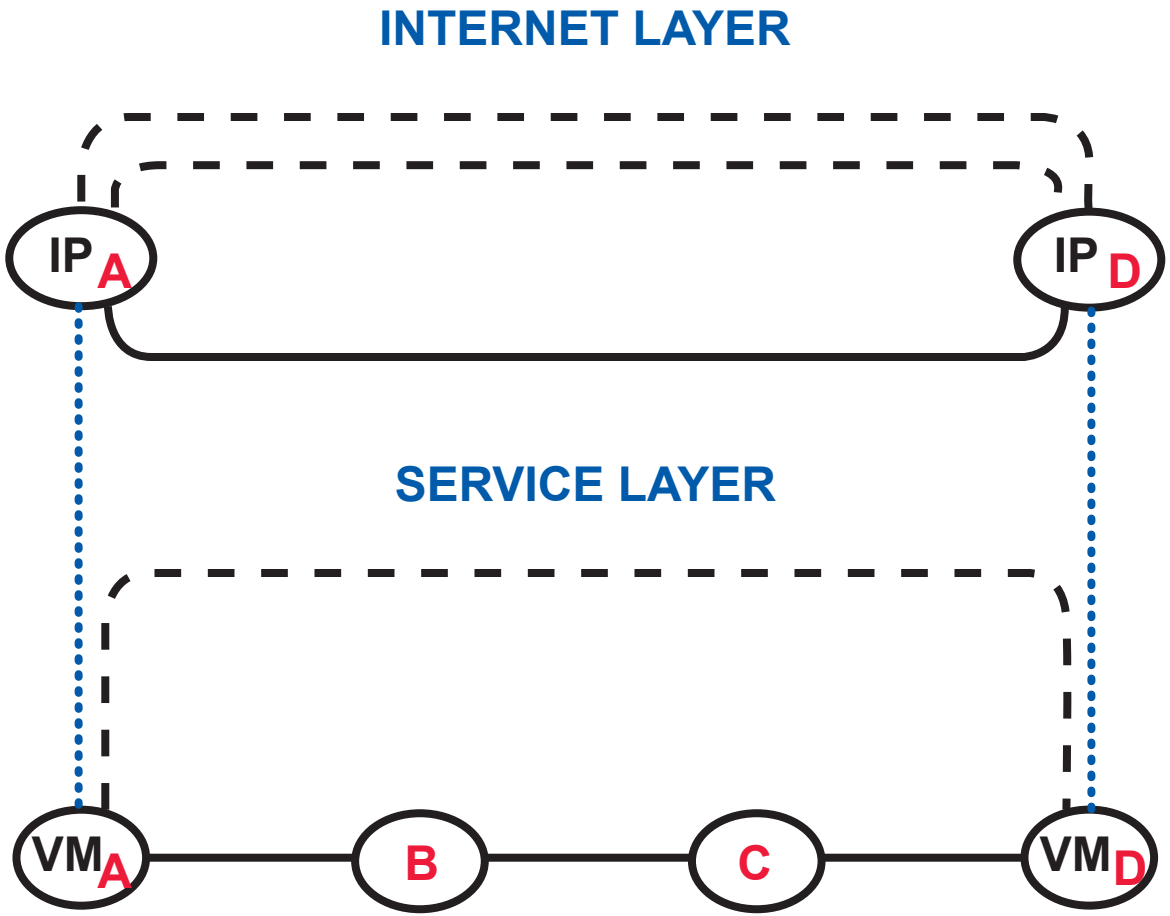

application sessions are not allocated to new policy link until this call returns

HEADER OPTIMIZATION

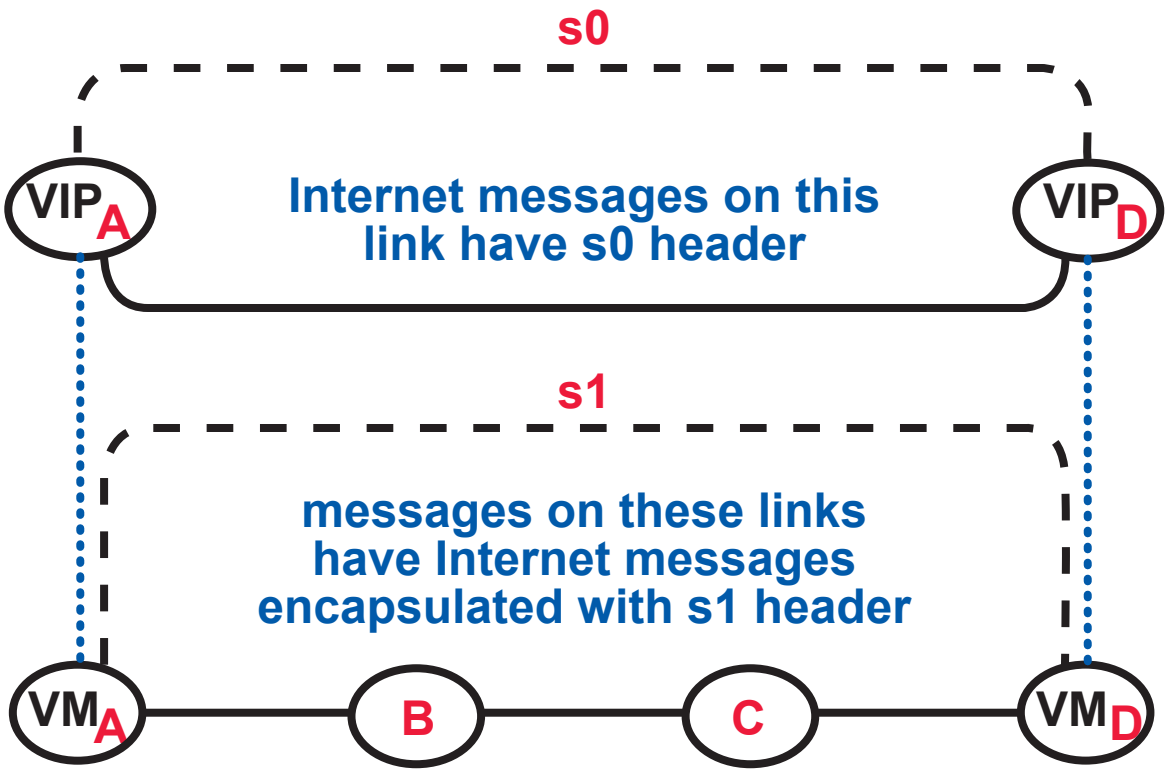

soundness of optimizations is easy to reason about in Alloy

if you optimize, you know what generality you are losing

HOWEVER, . . .

- if names or link/session identifiers coincide in two layered networks, they can be omitted from one of the headers

. . . and,
if sessions are set up by control plane (rather than by exchange of messages). . .

- if there is no more than one session between two endpoints, header can omit identifier

- if there is no more than one hop (link) in a session path, header can omit names

SUMMARY: REASONING WITH THE FORMAL MODEL

LOGICAL EFFECTIVENESS OR REACHABILITY

- legitimate destinations are reachable from legitimate sources

verified separately for each layer

- the mobility mechanism always succeeds in the cloud layer

even without central control, both endpoints moving simultaneously

SECURITY

- only allowed and authenticated messages are delivered

verified separately for each layer

- middlebox policies are enforced by the service layer
- one tenant's VM cannot receive messages from another tenant's VM in the cloud layer

UPDATE CONSISTENCY

- for propagation of top-down changes due to tenant configuration, policies, or load

consistency by construction, using informal hierarchical reasoning

- for propagation of bottom-up changes due to mobility, resource failure, or resource reconfiguration

verification and informal reasoning, both hierarchical

HEADER OPTIMIZATION (WHEN POSSIBLE TO OMIT FIELDS)

verified separately for each layer

BANDWIDTH TRACEABILITY (SUPPORT FOR QoS CONTRACTS)

load from each tenant is formally defined and traceable

NETWORK VIRTUALIZATION IN MULTI-TENANT DATACENTERS

by Teemu Koponen and 24 others, mostly from VMware

NSDI `14

This is believed to be the ultimate cloud design, but no one understands the paper. Good time to try again.