Towards Building Secure and Reconfigurable Virtual Networks on Multi-Tenant Data Centers

Jinwoo Kim[1] and Jaehyun Nam[2]

[1]Kwangwoon University

[2]Dankook University

Network Virtualization (NV)

- Key to today's multi-tenant data centers[1]
  - To implement Infrastructure-as-a-Service (IaaS)
- Enables tenants to build ***virtual networks*** over VMs
  - For a custom network topology and policies
  - E.g., an SDN network, BGP VPN, and NFV service chain

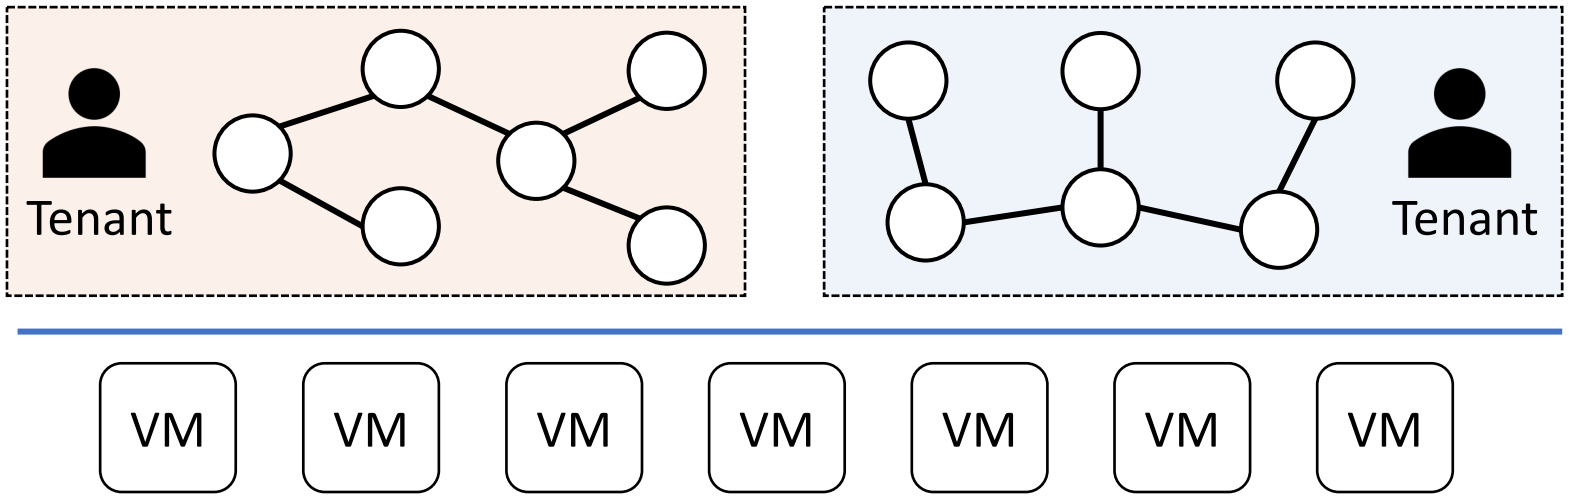

1 Network virtualization in multi-tenant datacenters, NSDI 2014

Existing NV Solutions

- Rely on network segmentation
  - Splitting a physical network into network segments
- Utilize tunneling protocols
  - Encapsulating a packet with a different header
  - E.g., NVGRE and VxLAN

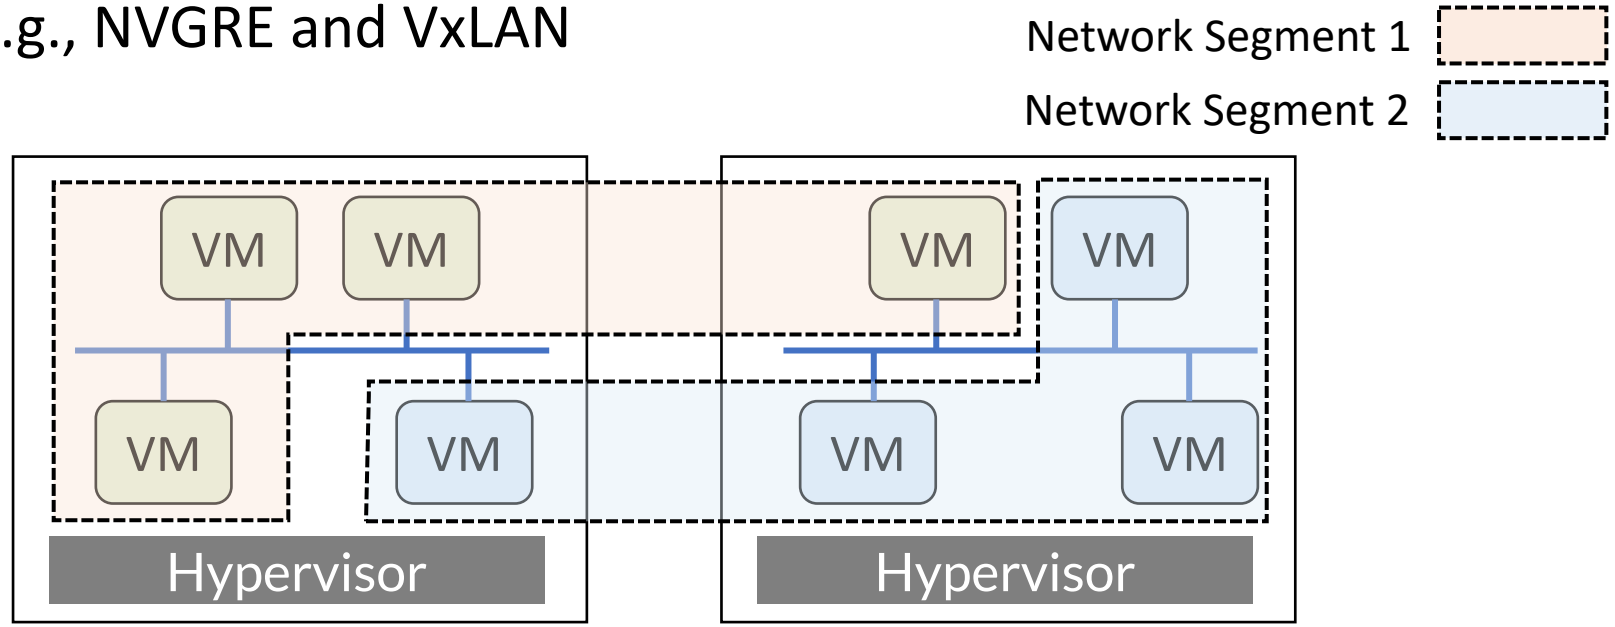

Challenge 1: Configuration

- Need for setting tunnel endpoints (TE)
  - E.g. IP addresses and interfaces
- Mostly rely on tenant's ***manual labor***
  - Time-consuming and error-prone

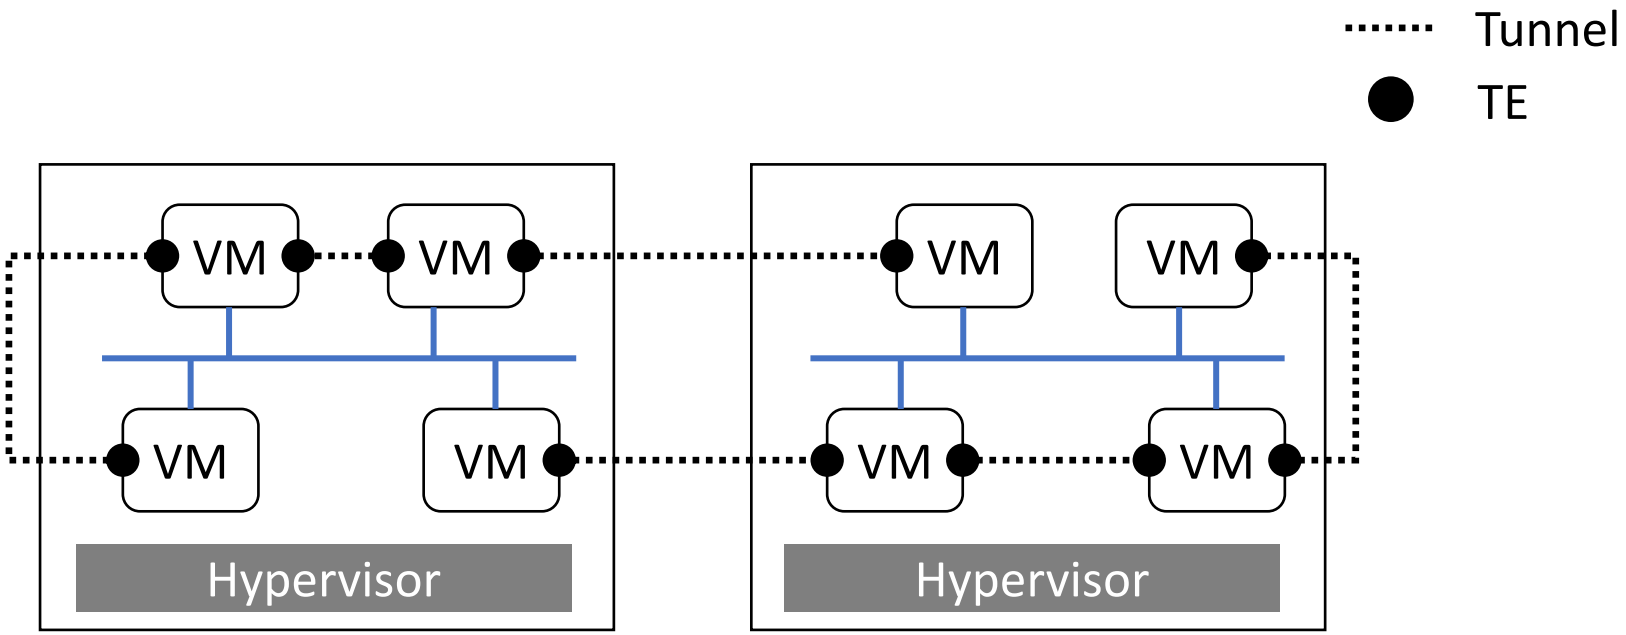

Challenge 2: Management

- Need for updating rules dynamically
  - If a topology or policy changes
- Requires to manage rules of distributed TEs
  - Can cause rule conflict by accident

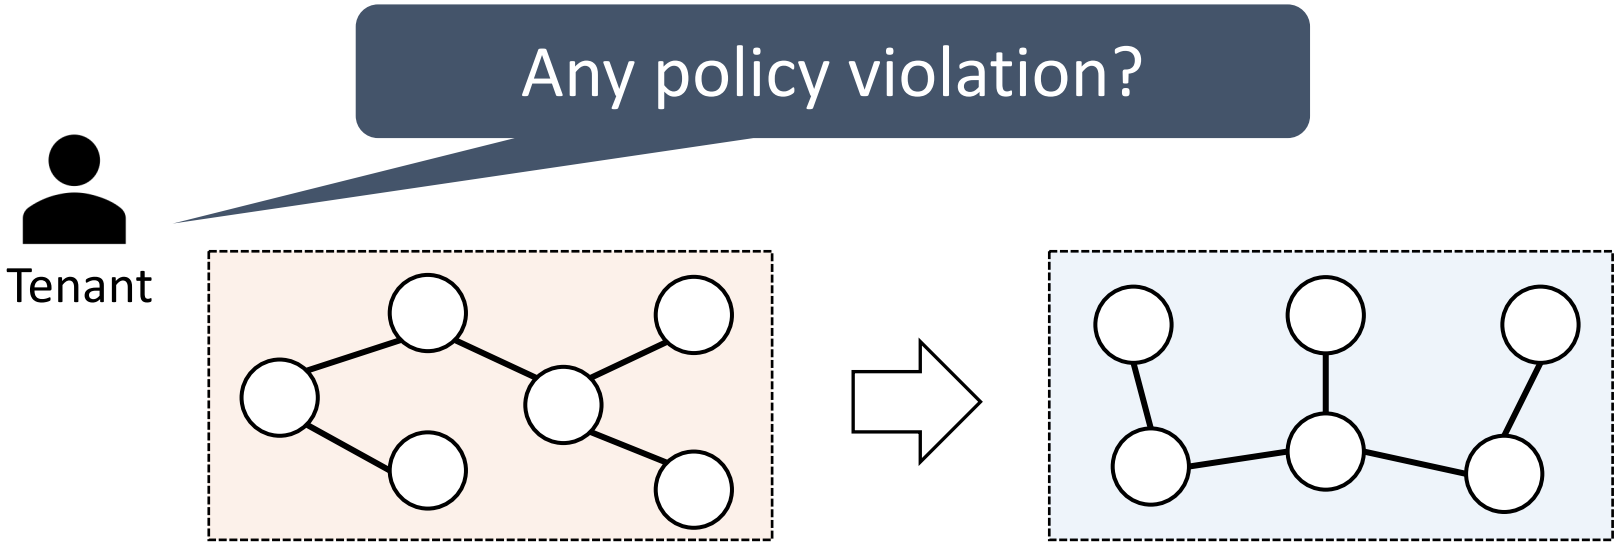

Challenge 3: Security

- No traffic isolation between VMs
  - Only performed for between segments
- Vulnerable to traffic ***eavesdropping*** and ***tampering***
  - Due to the bridge-based architecture of hypervisors

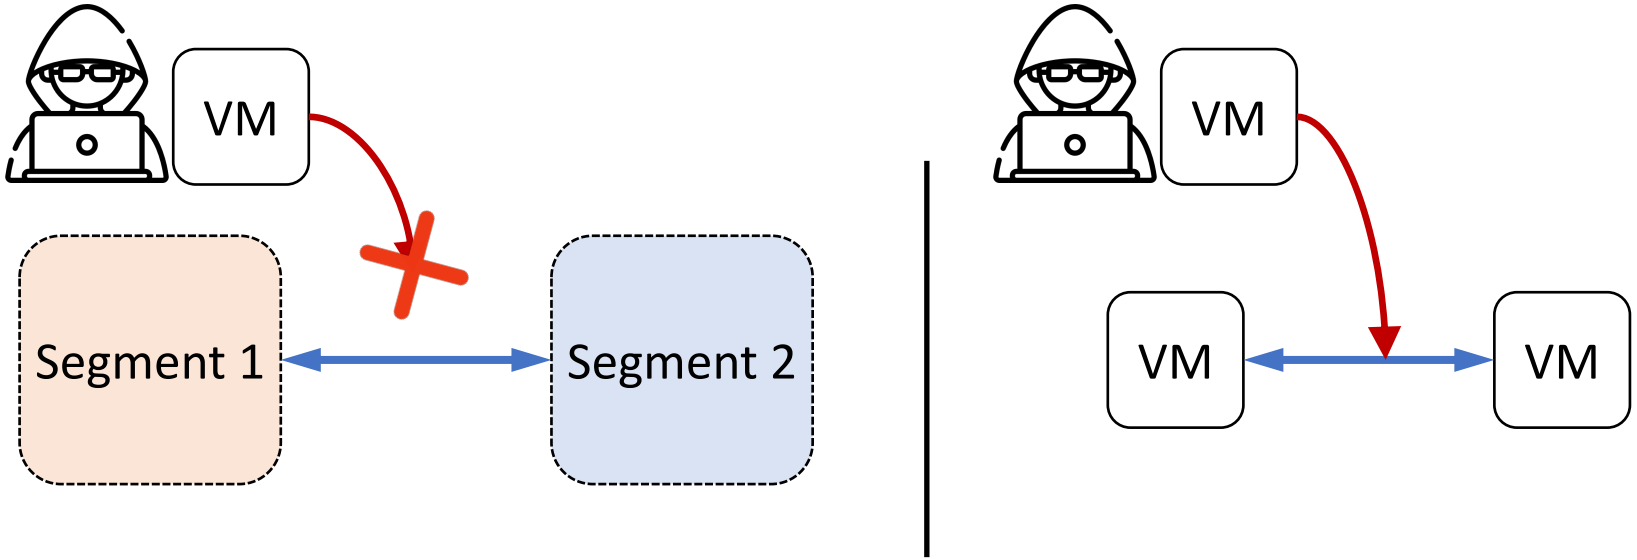

Related Work

- ***Tunneling-based*** network virtualization
  - VxLAN, NVGRE
  - Require manual configurations
- ***SDN-based*** network virtualization
  - FlowVisor [OSDI 2010], Koponen et al. [NSDI 2014]
  - Not applicable to VMs

None of them addresses all the challenges

LinkWire

- A system for building ***secure*** and ***reconfigurable*** virtual networks on multi-tenant data centers
  - Generates configurations automatically
  - Manages virtual networks with centralized architecture
  - Updates virtual networks on-the-fly

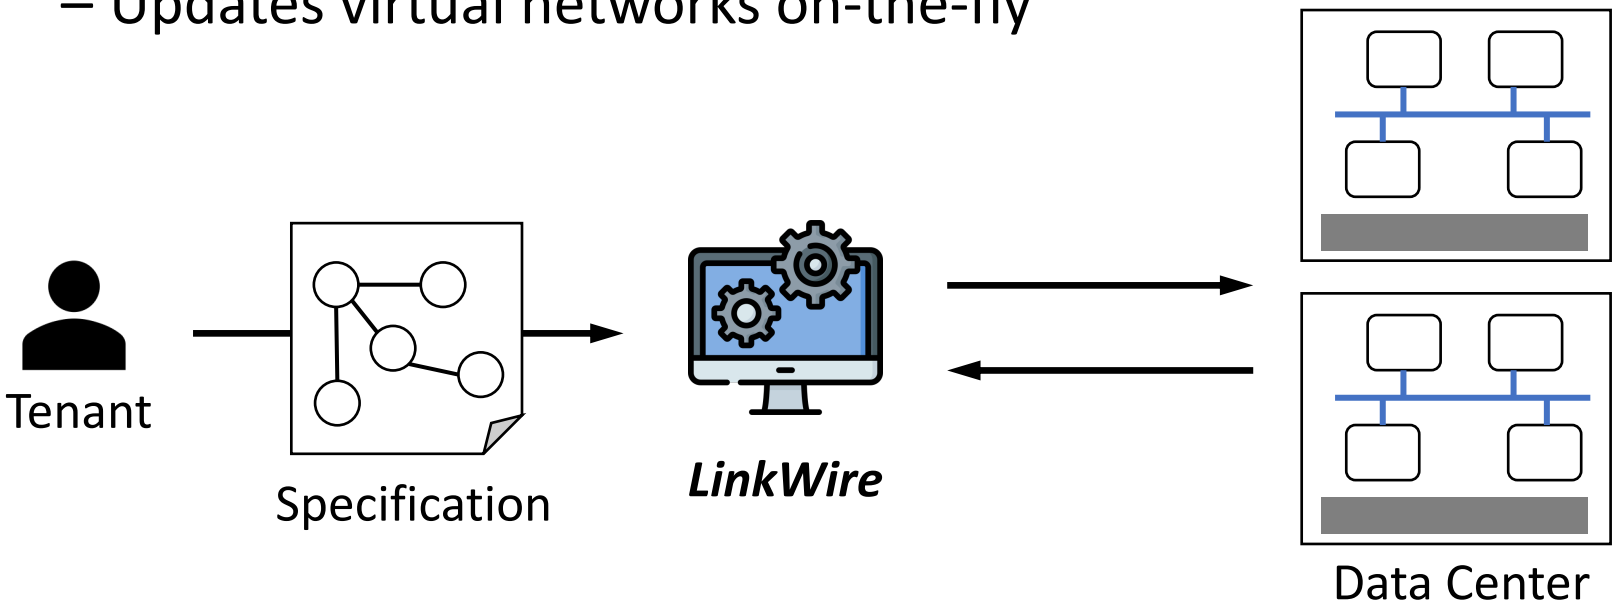

LinkWire Architecture

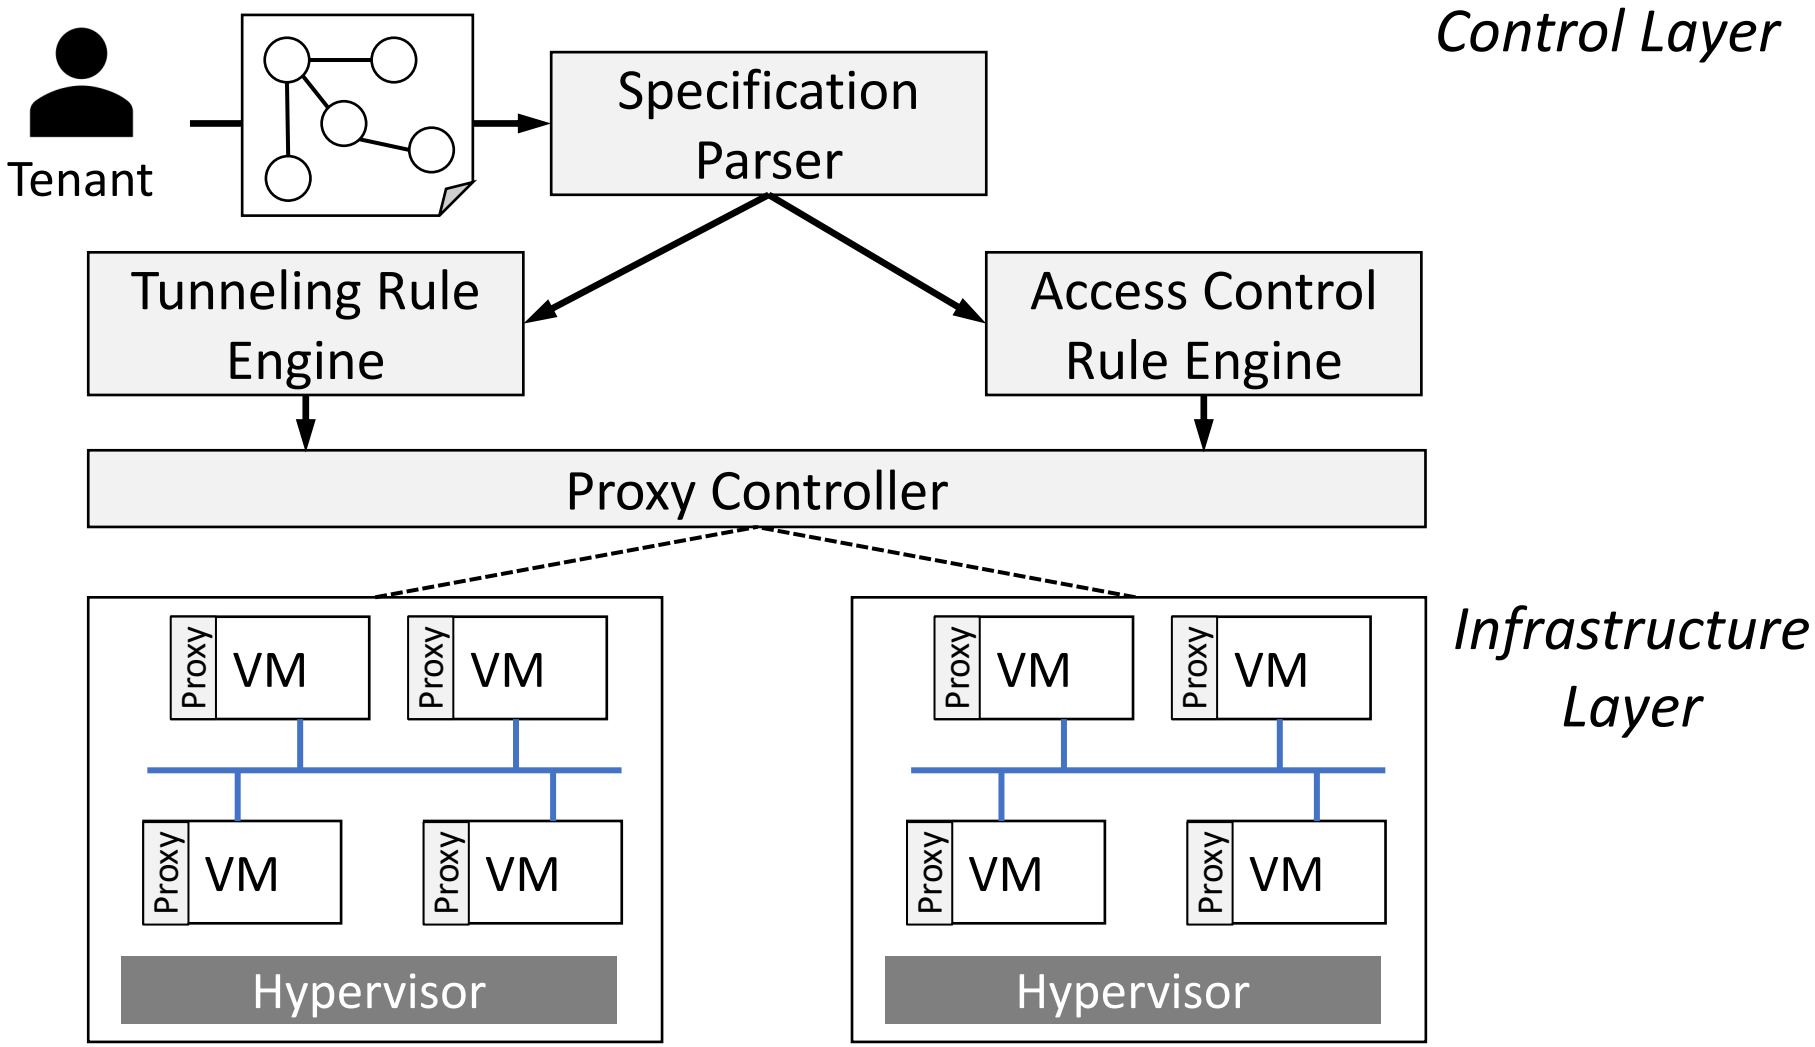

Control Layer

- Provides a tenant with ***graph abstraction***
  - Vertex: virtual node (i.e., VM)
  - Edge: virtual link (tunnel)
  - Edge label: a pair of tunnel endpoints
- Translates the graph into rules
  - I.e., tunneling and access control

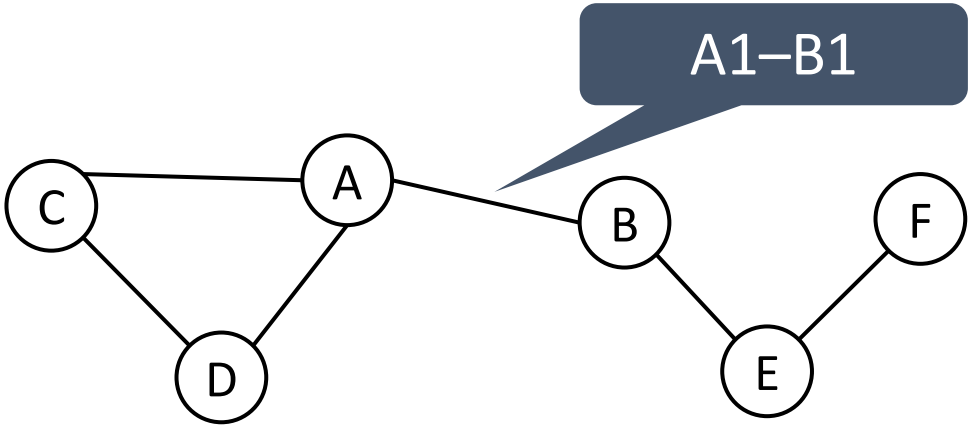

Infrastructure Layer

- Utilizes XDP/TC hooks for ingress and egress interfaces
  - To implement access control and tunneling

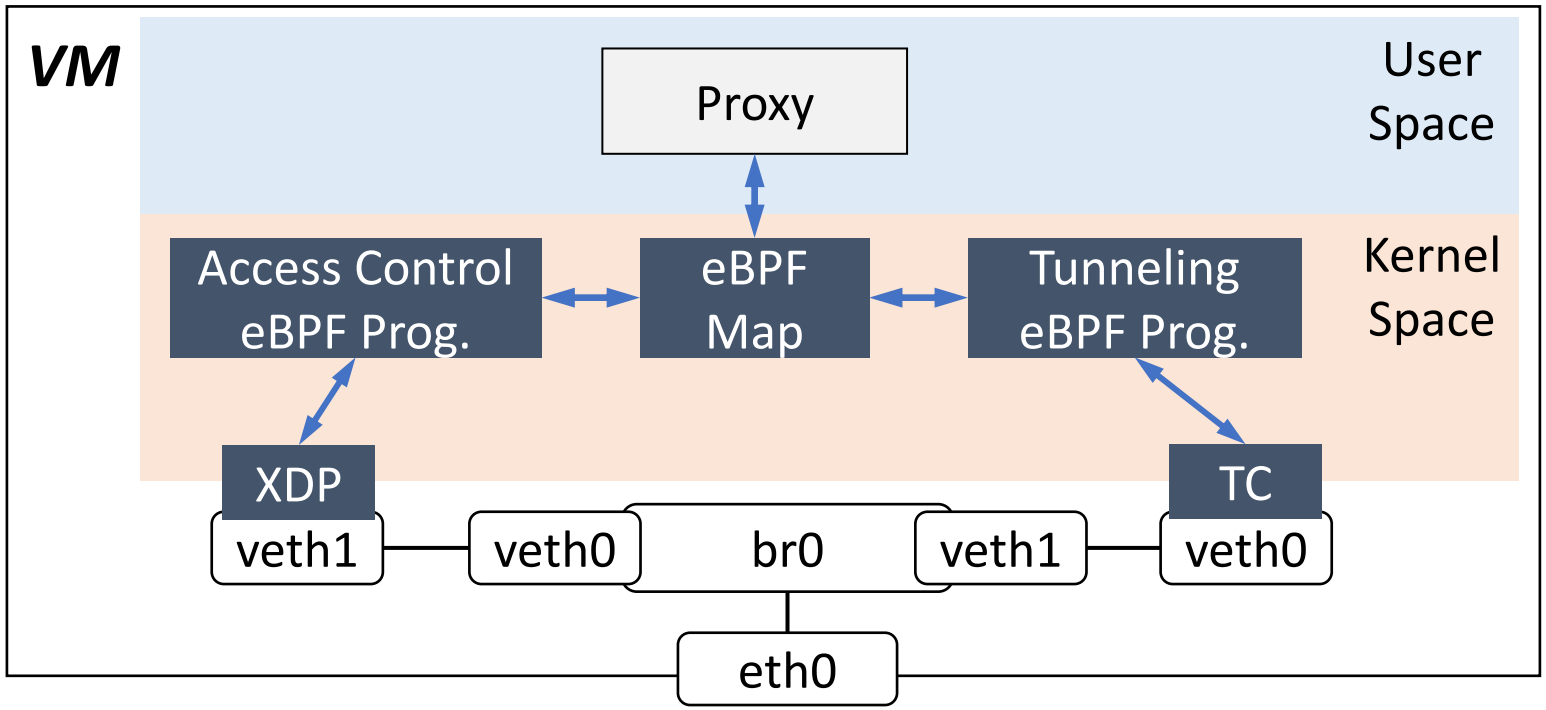

Use Case 1: SDN Network

- Deploying SDN controller and switches over VMs
  - E.g., ONOS[1] and OVS[2]
- Difficult to build a control channel due to complexity
  - Can be constructed by ***LinkWire*** easily

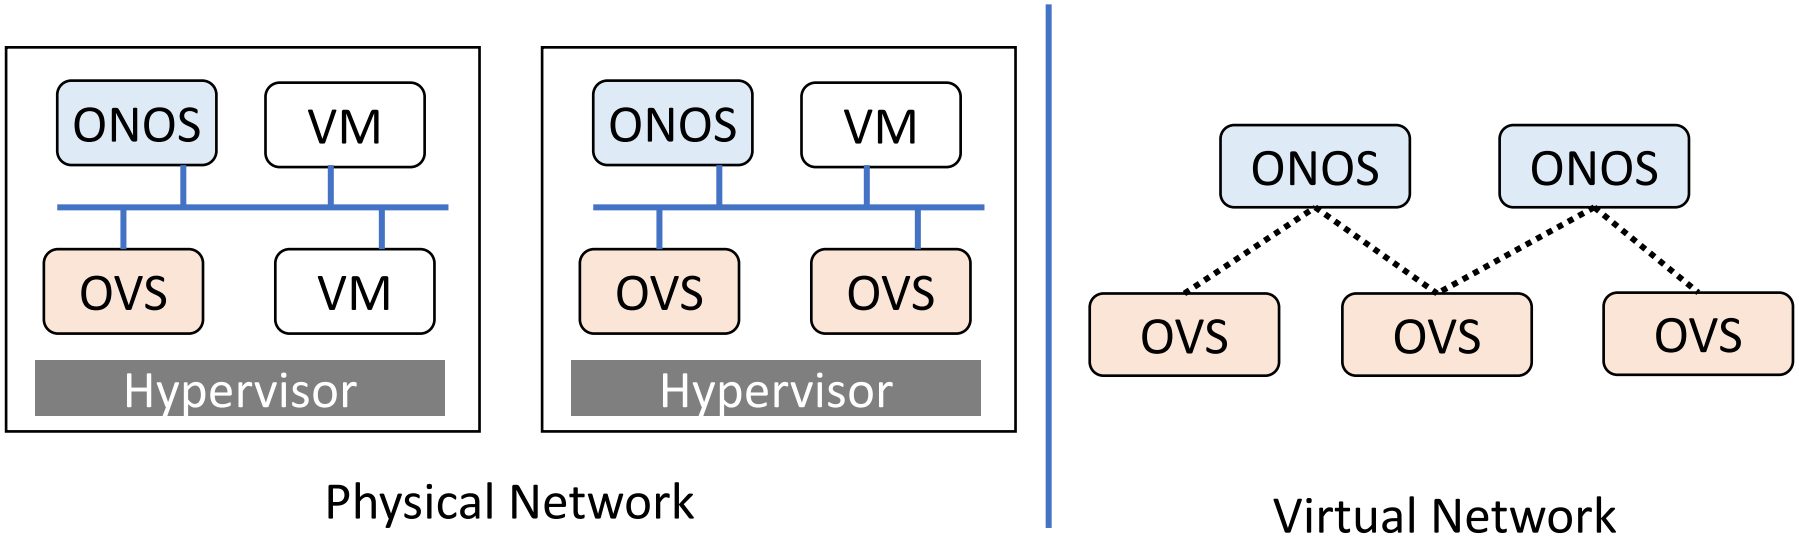

1 https://opennetworking.org/onos/
2 https://www.openvswitch.org/

Use Case 2: BGP VPN

- Connecting a remote site to tenant VMs via BGP[1]
- Weak to BGP poisoning attacks
  - Can be protected by ***LinkWire*** access control

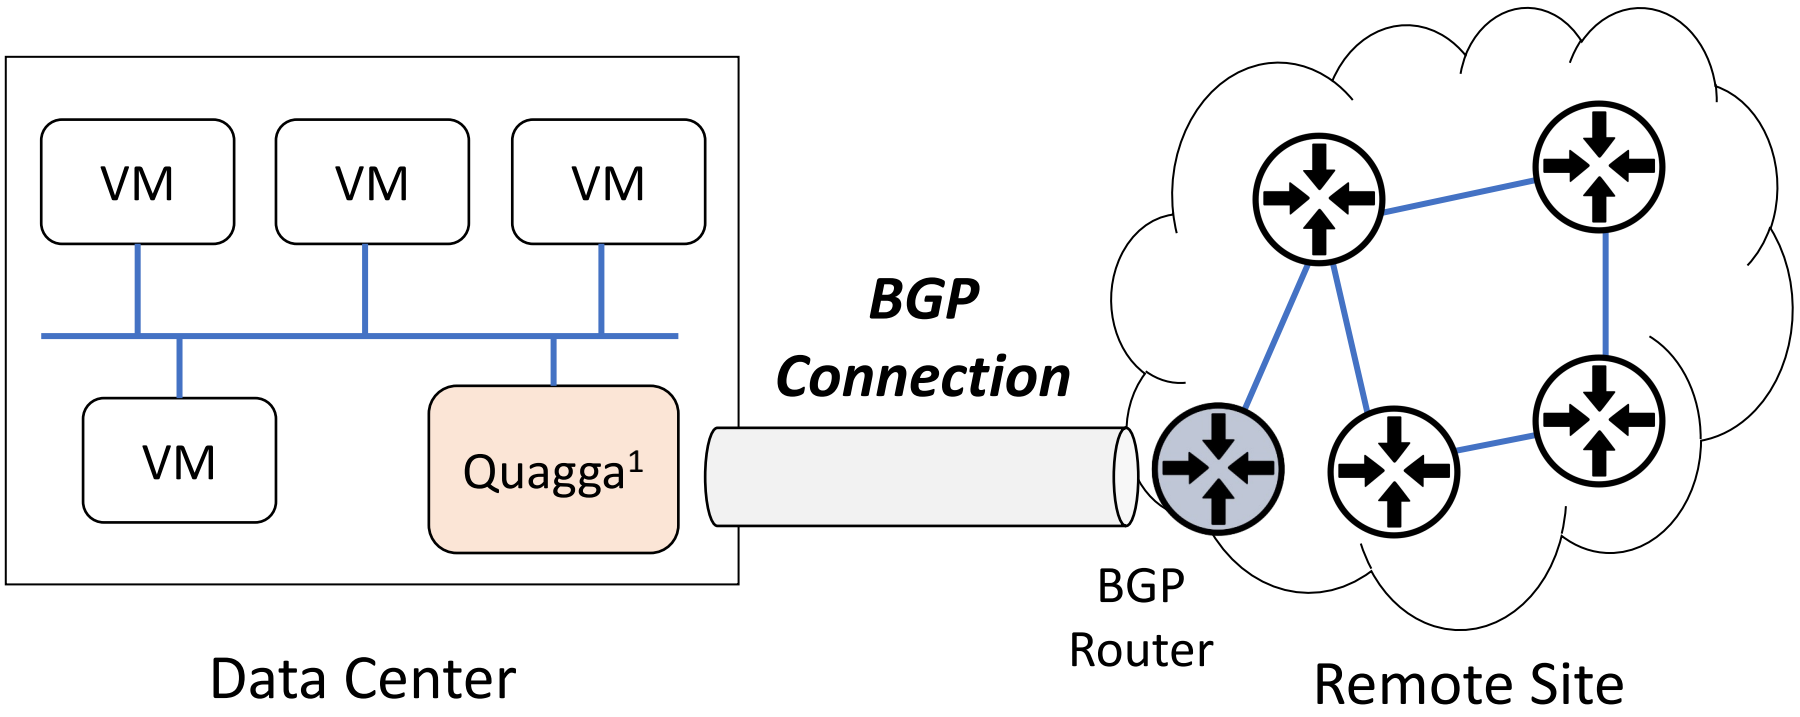

1 https://docs.openstack.org/networking-bgpvpn/ocata/overview.html

Conclusion

- Existing network virtualization solutions
  - Mostly rely on tunneling protocols
  - Require manual configurations and distributed management
  - Vulnerable to traffic eavesdropping and tampering

- ***LinkWire***: a secure and reconfigurable system for virtual networks on multi-tenant data centers

- Future work
  - eBPF-based implementation
  - Evaluation in real cloud environments

Thank you for listening

(jinwookim@kw.ac.kr)