


Presenter:

Vinit Jain, STSM, System Networking Development,
IBM System & Technology Group


A Case for Overlays in DCN Virtualization

**Katherine Barabash, Rami Cohen, David Hadas, Vinit Jain,
Renato Recio and Benny Rochwerger IBM**



10 years ago – networking was simpler

Physical Network



▪ Static workloads

- Applications ran on Operating Systems (OS)
- OS resided on physical server (vs moving from one server to another)

▪ Each workload has network state associated with it.

- Examples: VLAN, Access Control Lists - ACLs, Traffic rate limiting.

▪ The workload's network state:

- Resided in the physical network
- Associated with server and static

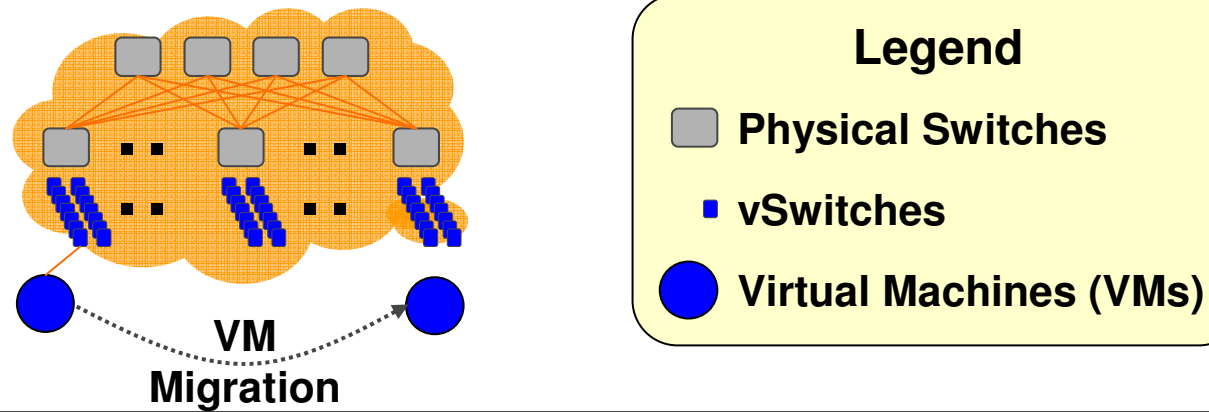
▪ Physical network is static & simple (per server addresses, per port state)

- Before (or after) physical server is installed, the network administrator configured the physical network once

Virtualization increased network complexity



Physical Network with Virtual Switches



▪ Server virtualization made workloads dynamic

- Virtual Machines (workloads) can be dynamically added to servers.
- VMs can move from an over-utilized server to less utilized server.

▪ The VM's network state:

- Now resides in both the server (vSwitch) & physical network
- Is no longer associated with the server, it's now associated with the VM & is dynamic (can move around).


▪ Physical network is dynamic and more complex

- Single “server port” now has multiple VMs, each with different network state.
- Requires network state coordination between Hypervisor and Physical network

Virtualization → Can it be simpler?

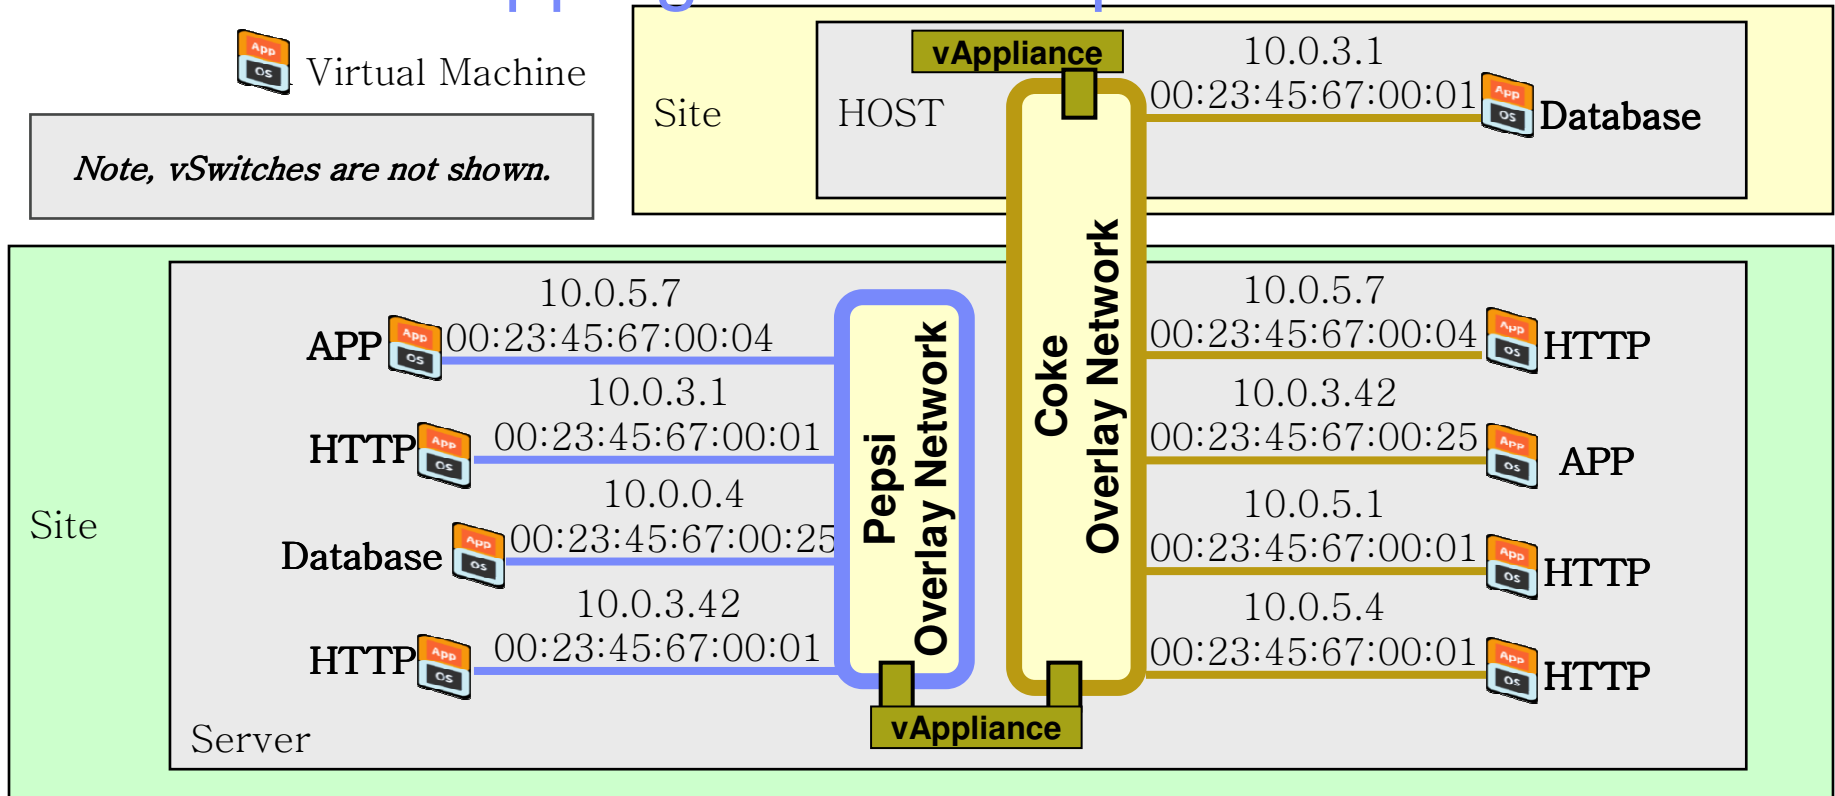


Overlay Networks



- **Concept: create an overlay network above the physical network.**
- **Workloads remain dynamic, but:**
 - Physical server state resides in physical network.
 - All VM network state resides in server's Distributed Overlay Virtual Ethernet (DOVE) network.
- **Physical network is static & simple**
(per server addresses, per port state)
 - Network administrator configures physical network once.
- **Virtual network is more efficient and supports multi-tenancy** (see next slide).

Multi-Tenant with Overlapping Address Spaces




- Multi-tenant, Cloud environments require multiple IP address spaces within the same server, within a Data Center and across Data Centers (see above).
 - Distributed Overlay Virtual Ethernet (DOVE) switches to enable multi-tenancy all the way into the Server/Hypervisor, with overlapping IP Address spaces for the Virtual Machines.

Network as a Service



- Logical description of the network
- Connectivity:
 - A Load Balancer is connected to the internet
 - A Load Balancer is connected to a set of Application servers
 - The set of Application Servers are connected to a database
- Security
 - All the incoming traffic from the Internet to the Load Balancer must pass through Firewall and an IDS
- Performance
 - All the traffic between the Application Servers and the Database must pass through a compression middle box
 - All the SSL traffic between the Load Balancer and the web servers must pass through SSL accelerator



Achieve same level of virtualization for networks as we have today for servers


Host virtualization should enable virtual machines

- To remain independent of physical location
- To remain independent of the **host** physical characteristics such as CPU, Memory, I/O, etc.
- To form isolated **compute** environments on top of the shared physical **host** environment

Network virtualization should enable virtual machines

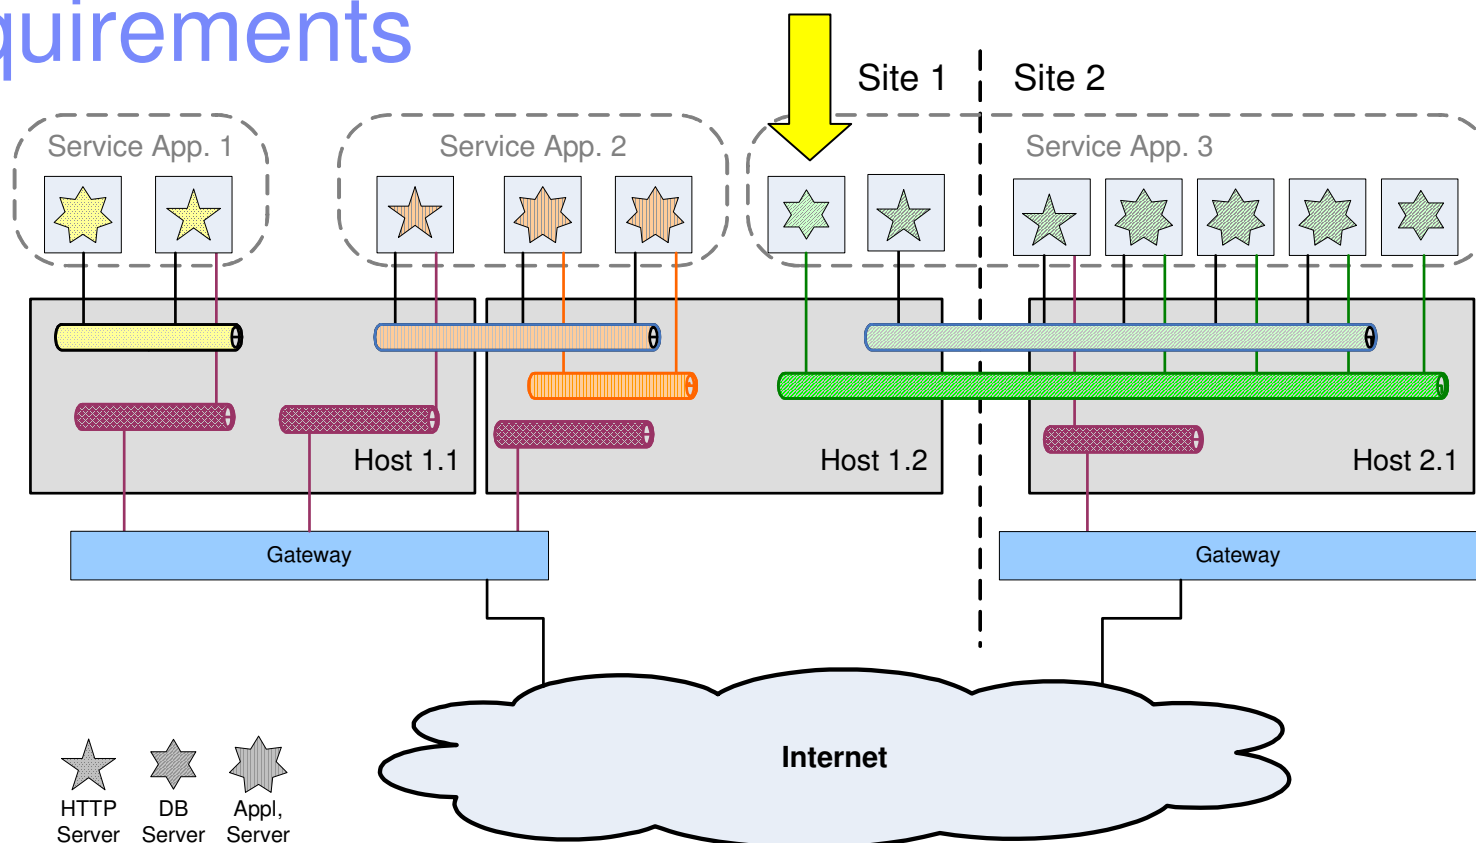
- To remain independent of physical location
- To remain independent of the **physical network infrastructure** characteristics such as network layer (2, 3), protocols, addresses, topology, etc.
- To form isolated **network** environments on top of the shared physical **network** environment serving the hosts

Network Virtualization for the Cloud - Requirements



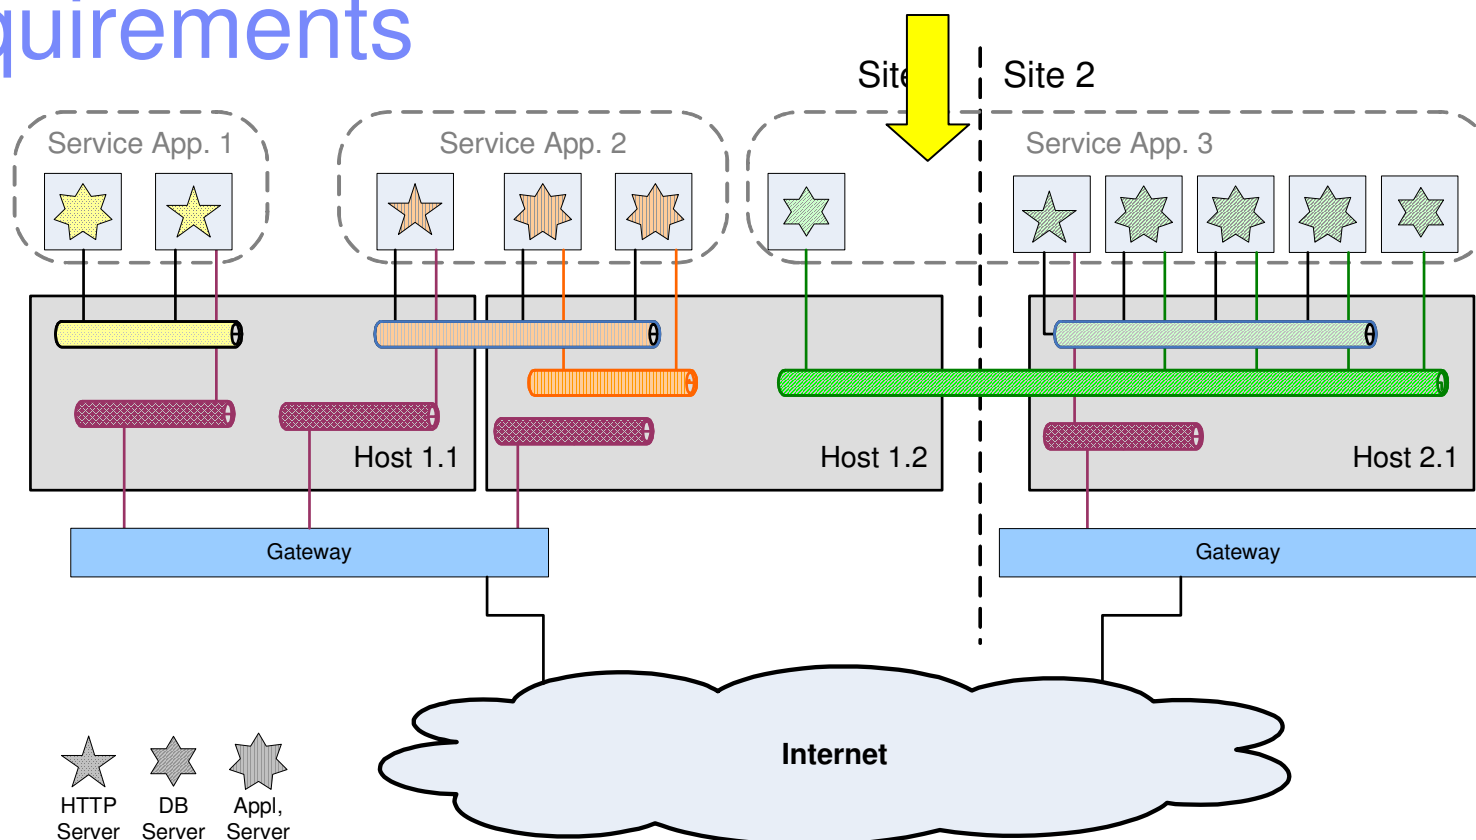
- Location and Topology Independence
- Isolation

Network Virtualization for the Cloud - Requirements



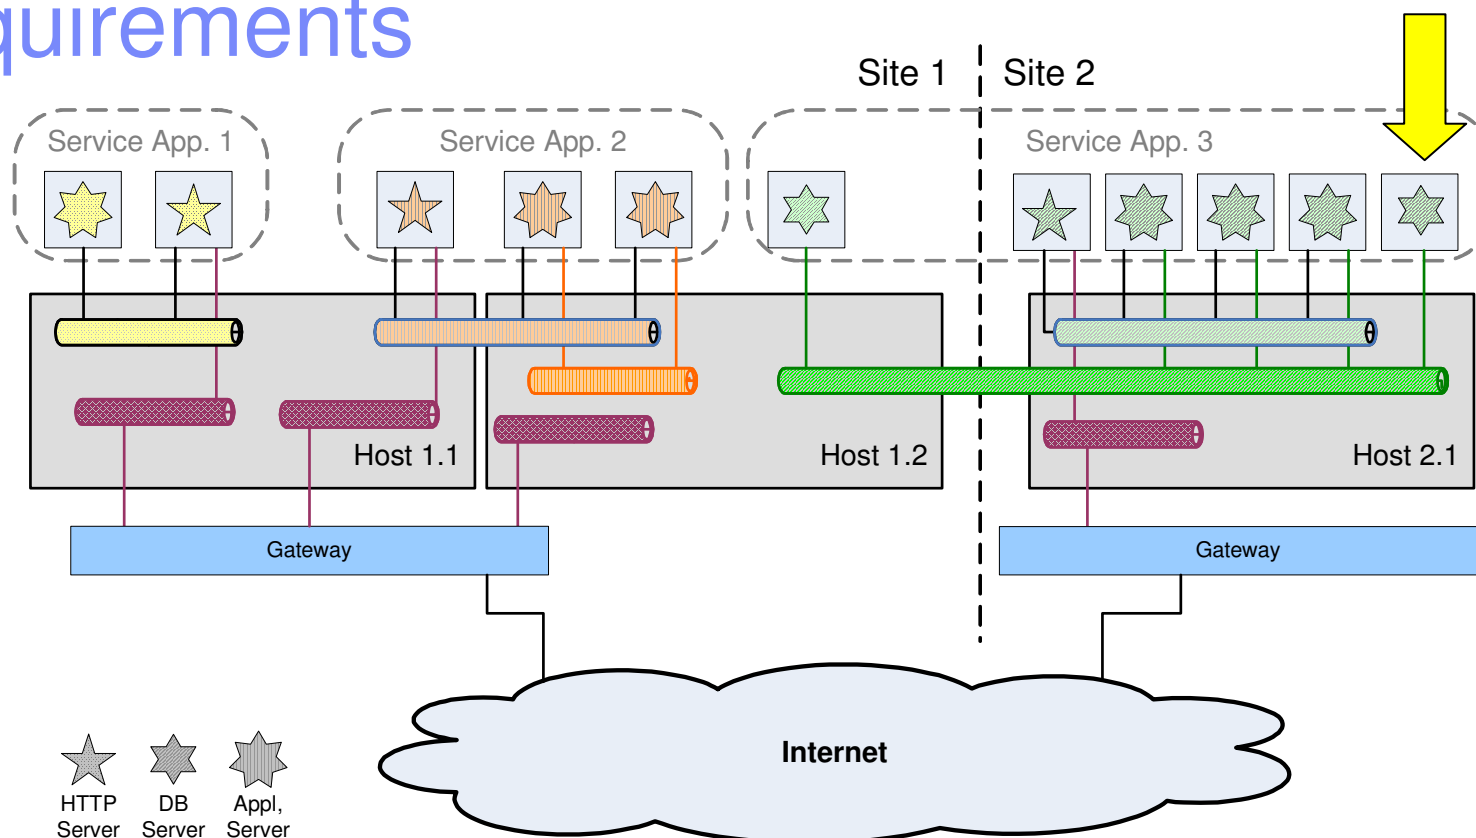
➤ Dynamically grow ...

Network Virtualization for the Cloud - Requirements



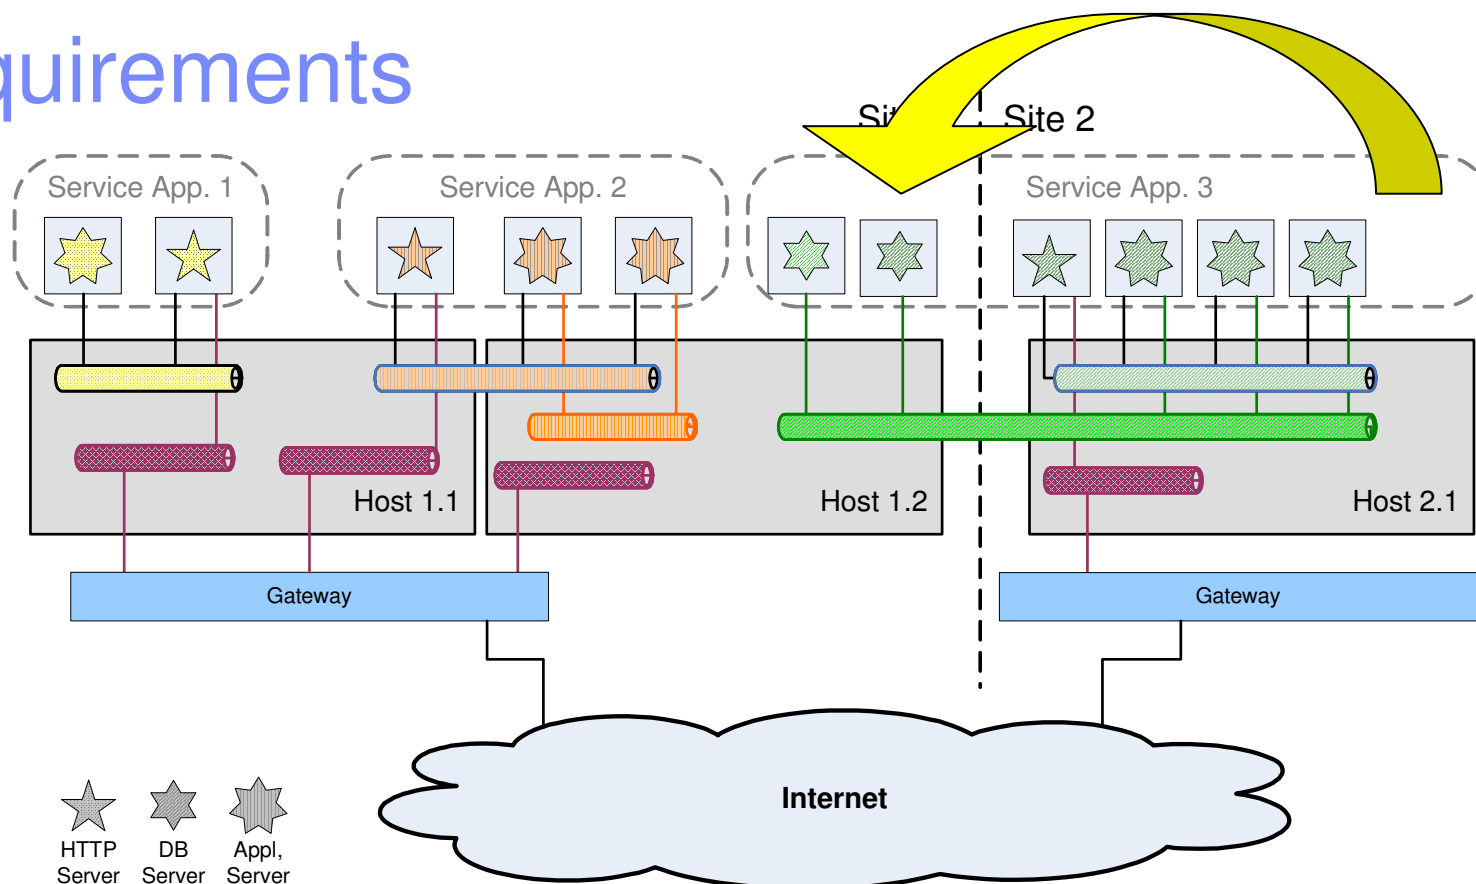
➤ Dynamically grow ... and shrink

Network Virtualization for the Cloud - Requirements



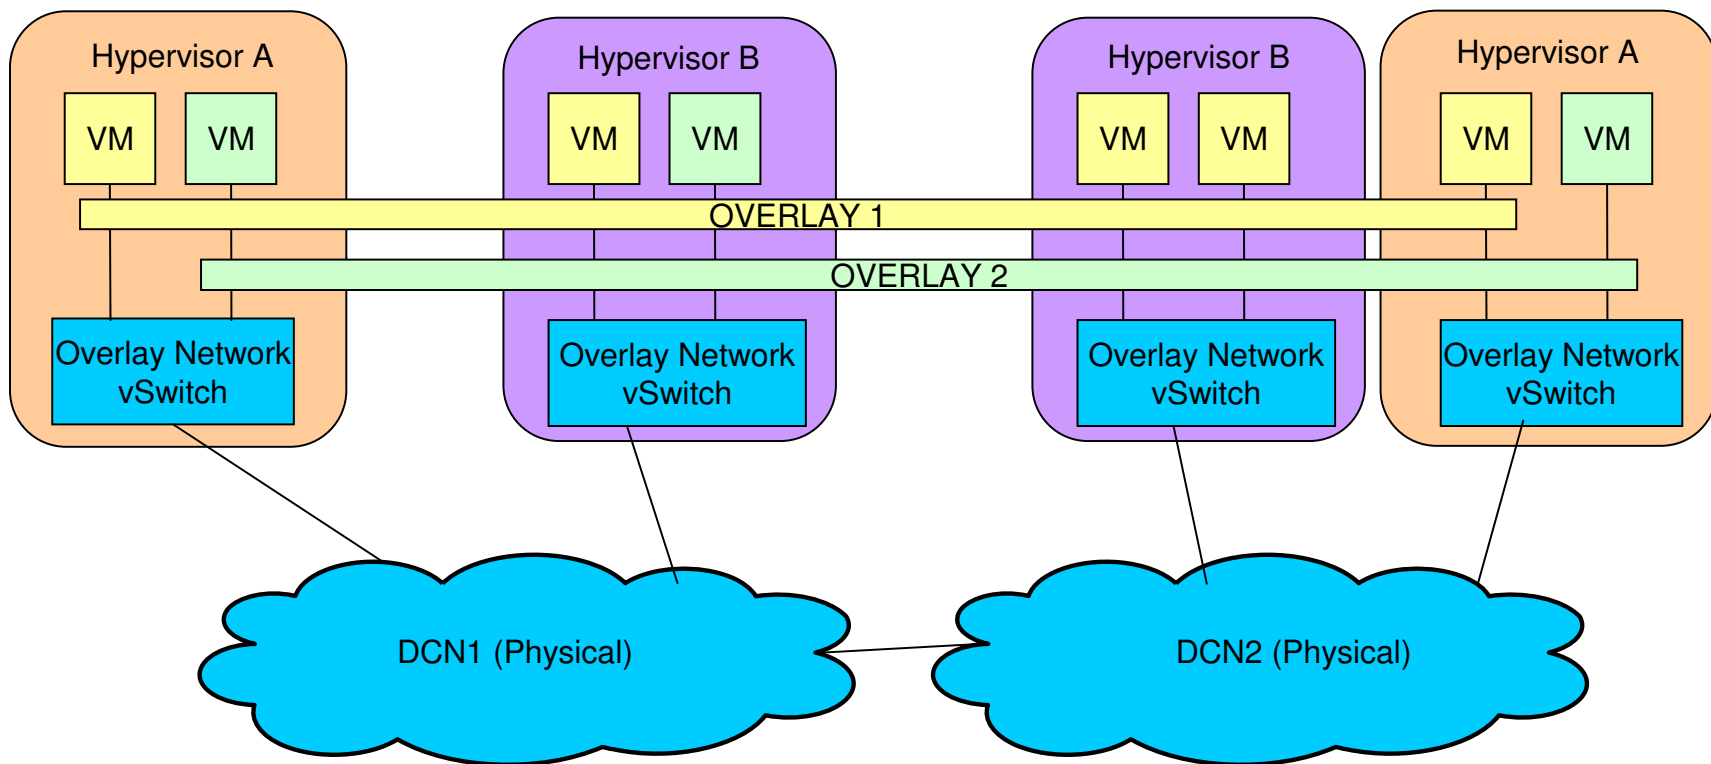
➤ Live migration “without borders”

Network Virtualization for the Cloud - Requirements

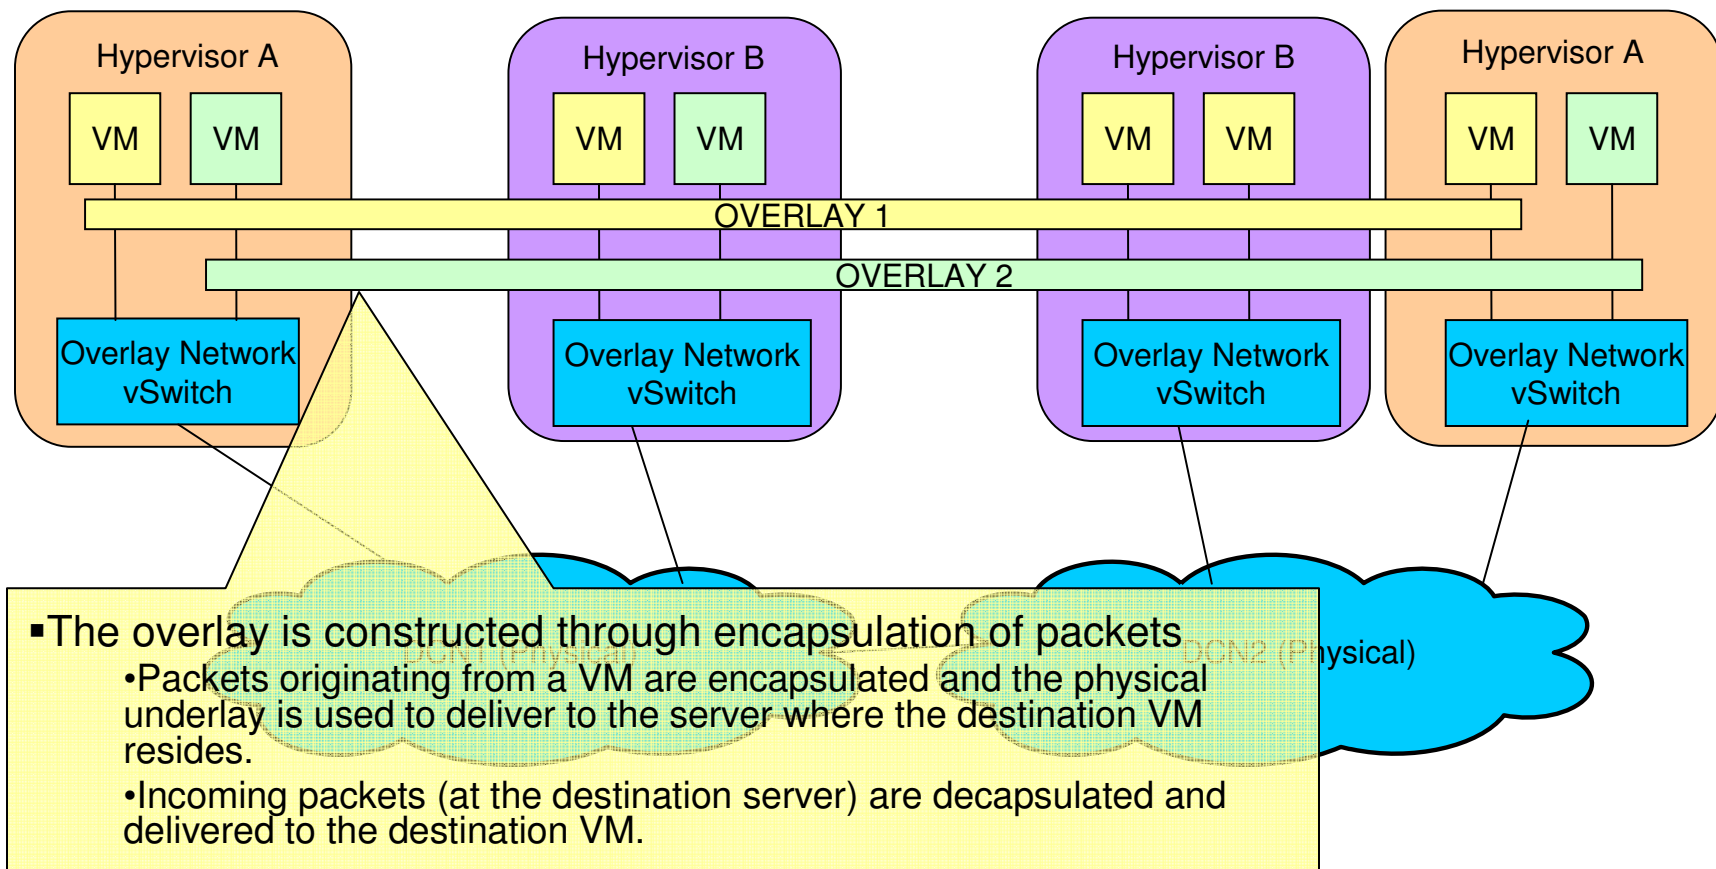


➤ Live migration “without borders”

The Distributed Overlay Virtual Ethernet (DOVE) approach:
 build the virtual network by creating an overlay networks between
 hypervisors, which can be connected to each other over an arbitrary
 physical topology



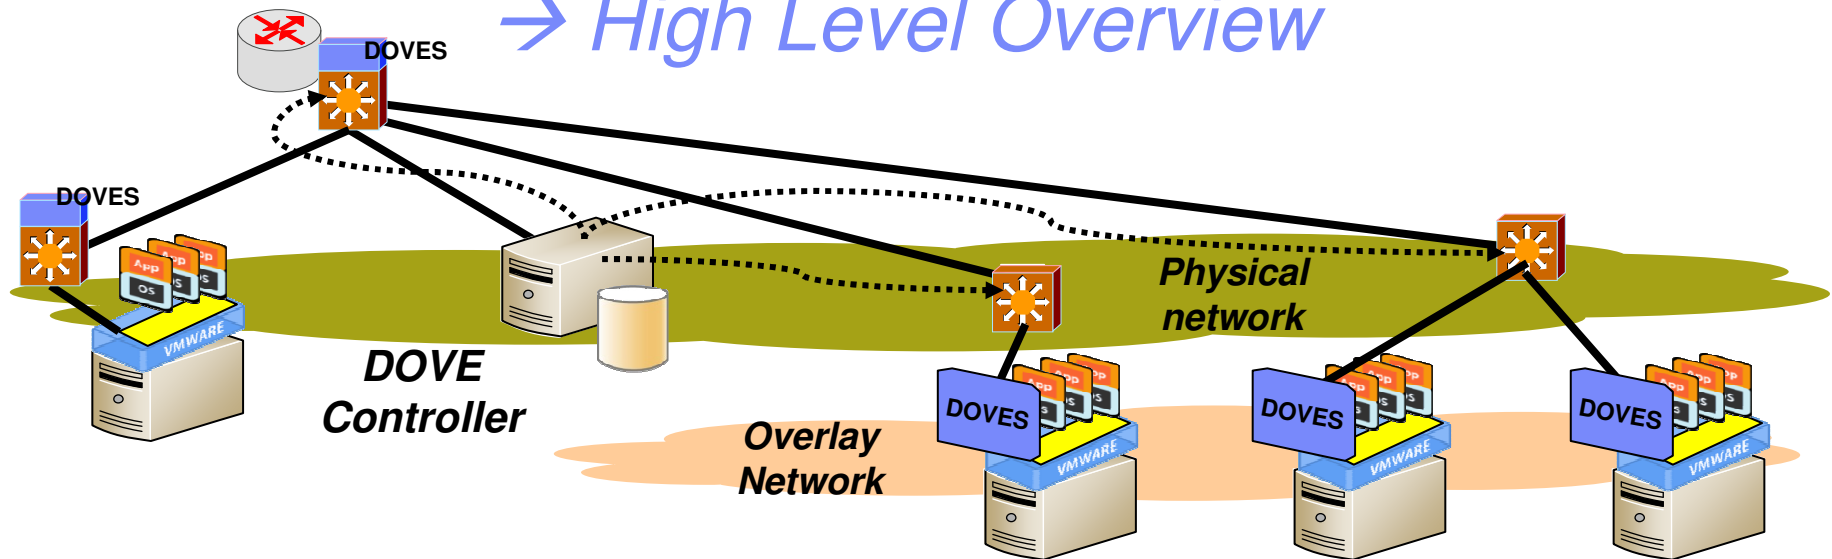
The Distributed Overlay Virtual Ethernet (DOVE) approach:
 build the virtual network by creating an overlay networks between
 hypervisors, which can be connected to each other over an arbitrary
 physical topology



DOVE Solution Elements




→ High Level Overview




- DOVE Controller
 - Performs management & a portion of control plane functions across DOVE Switches
- DOVE Switches (DOVES)
 - Provides layer-2 over UDP overlay (based on OTV)
 - Performs data and some control plane functions
 - Run in Hypervisor vSwitch or gateways
 - Provides interfaces for Virtual Appliances to plug into (Analogous to appliance line-cards on a modular switch)

DOVE Encapsulation (OTV + Extension)

Original Packet

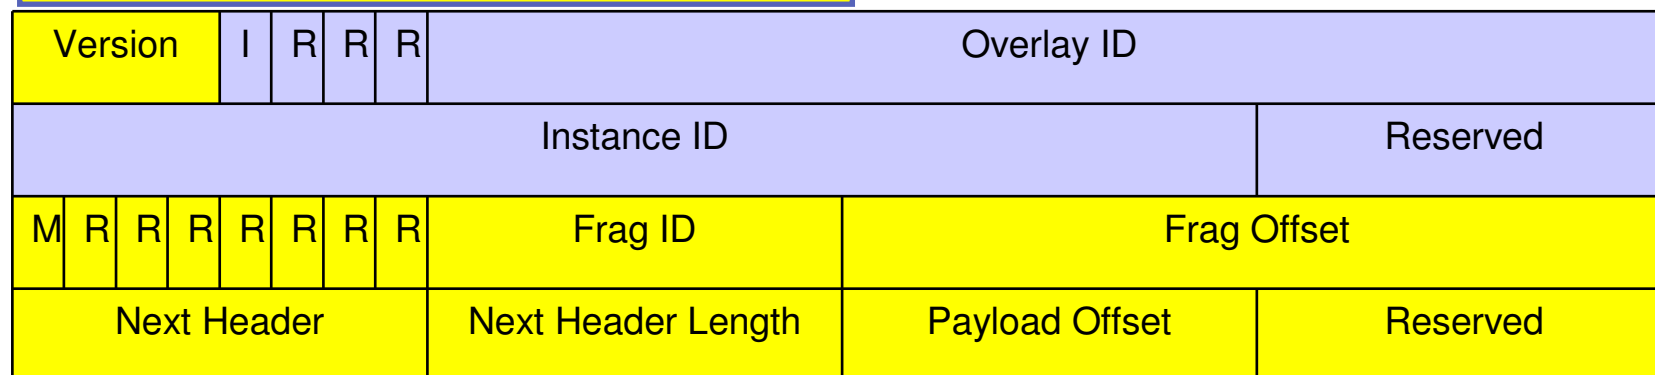


Encapsulation Options



Encapsulation Protocol (EP) Header

(Yellow is possible extensions to OTV)



DOVE's advantages



- Independence & Transparency
 - Using DOVE a virtual network can be deployed on any physical infrastructure
 - e.g. Ethernet, InfiniBand, IPv4, IPv6
 - Each infrastructure may utilize a different implementation (e.g. using Openflow in IP/Ethernet based network)
 - Using DOVE the network topology is flexible
 - VM can move from anywhere to anywhere
 - Each virtual network can be configured independently
- Scalability
 - Using overlay, DOVE reduces the forwarding table size both on switches and routers
 - Addressing only physical server
 - Reduces cost and improves performance
 - DOVE does not require forwarding entities configuration upon migration
 - not based on VLAN
 - Number of virtual network is not limited
 - Not based on VLAN

