

lecture 17:

virtualization with OVS

5590: software defined networking

anduo wang, Temple University

TTLMAN 401B, R 17:30-20:00

Open vSwitch

network virtualization

network virtualization

virtual networks over the same physical network

- each with independent service models
- topologies
- addressing architectures

network virtualization

virtual networks over the same physical network

- each with independent service models
- topologies
- addressing architectures

the creation and management

- done through global abstractions, rather than pieced together through box-by-box configuration

virtualization

benefits

- quick provisioning
- cloud
- improved availability

virtualization

benefits

- quick provisioning
- cloud
- improved availability

imposing requirements far beyond physical deploy

- seamless handling of mobility: VM migration
- scaling: datacenter can host hundreds of thousands VMs
- strong isolation: multi-tenant

virtualization

benefits

- quick provisioning
- cloud
- improved availability

imposing requirements far beyond physical deploy

- seamless handling of mobility: VM migration
- scaling: datacenter can host hundreds of thousands VMs
- strong isolation: multi-tenant

virtualization

benefits

- quick provisioning
- cloud
- improved availability

imposing requirements far beyond physical deploy

- seamless handling of mobility: VM migration
- scaling: datacenter can host hundreds of thousands VMs
- strong isolation: multi-tenant

providing features that ease networking


- multicast membership
- topology more tractable

new challenges and opportunities

standard Ethernet switching

- L2 switch and IP router
- *neither satisfies requirements nor leverages the available advantages*

network in virtual environments



simple L2 switch within the hypervisor

- VM connected to virtual interface (VIFs)
- virtual switch connects VIFs and physical interfaces (PIF)
- *a new networking layer: ... 40 VMs on each server ...*

virtual switch, opportunities

VM introspection

- tight integration with the virtualization software allows
 - characterize VIF: MAC addresses, IP allocation, multicast listening
 - infer host events

the network virtualization layer — leaf of the physical network topology in datacenter

- free from routing protocols

virtual switch, challenges

isolation

- multiple clients, VM migration

dynamic virtual overlays at L2

- broadcast domains to each client

distributed QoS and policing

- VMs from one client distributed across many physical hosts

Open vSwitch (OVS)

a switch platform

- take advantage of the virtual environment
- flexible enough to implement the various solutions

Open vSwitch (OVS)

purpose-built for virtualization, exposing

- interface for fine-grained control of the forwarding
 - QoS, tunneling, filtering rules
- a remote interface for migrating configuration state
 - attach policy to VMs

flexible, table-based forwarding engine

- logically partition forwarding plane

OVS interface

manipulating the forwarding state

- allow (remote process to) write to the forwarding table

managing configuration

- allow (remote process to) read, write
- set up triggers to receive events

OVS (local) interface

connectivity management


- allow virtualization layer to manipulate the topology
 - creating switches, managing VIF/PIF connectivity

flow-table forwarding model

similar to OpenFlow, rational

- near-arbitrary logical partitioning of forwarding


architecture and implementation



Open vSwitch implementation consists of two components

- kernel-resident “fast path”
- userspace “slow path”

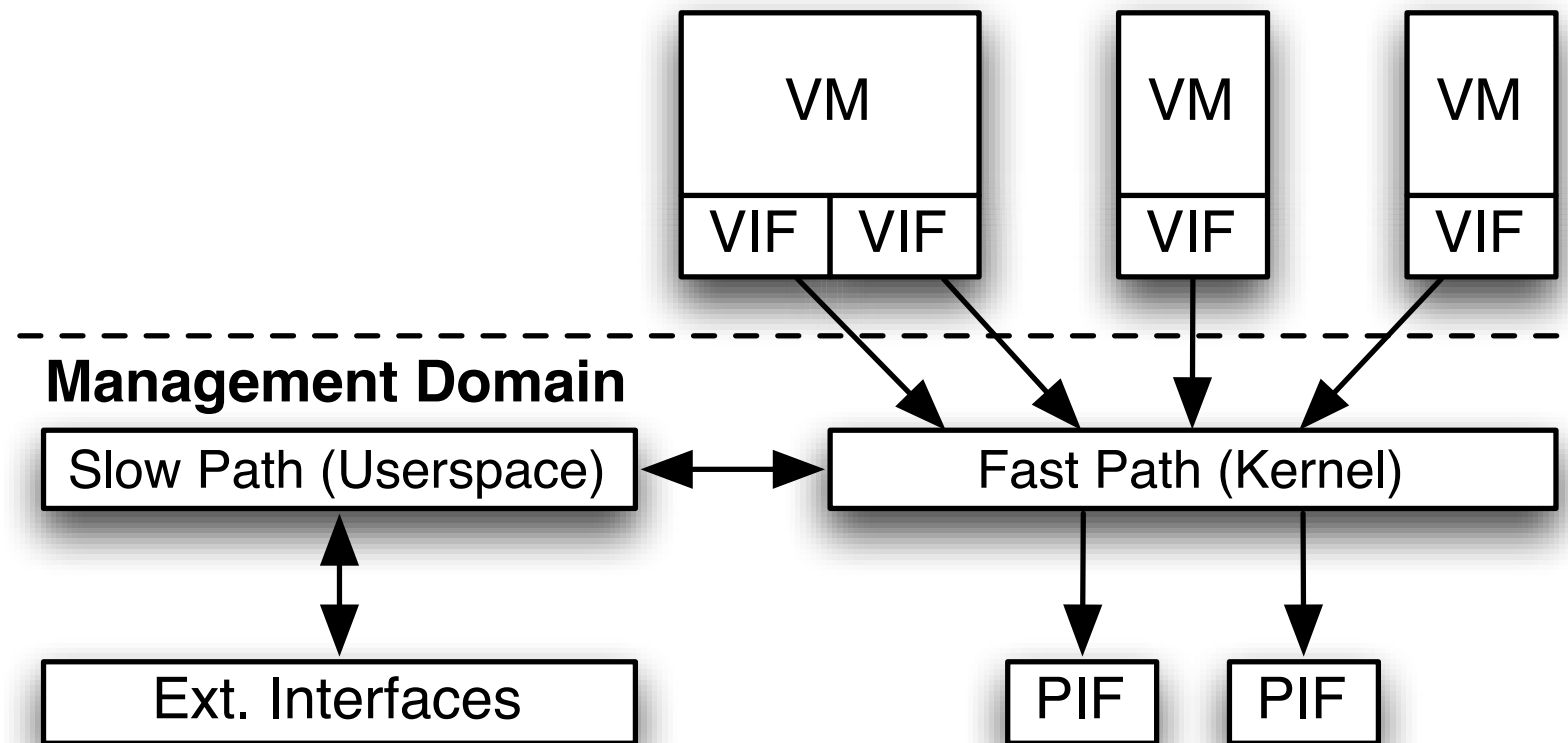
architecture and implementation



fast path

- implements forwarding: per-package look-up

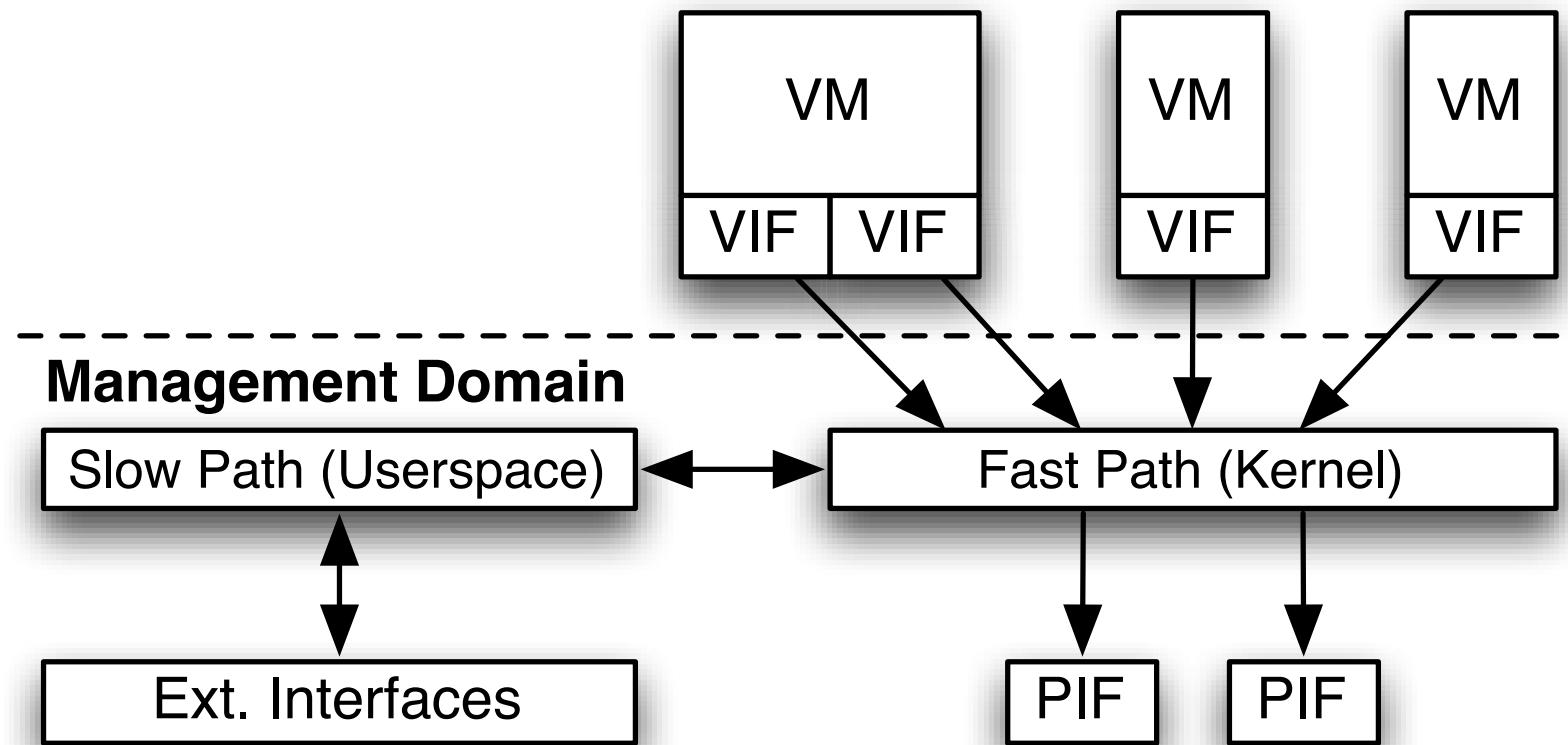
architecture and implementation



fast path

- implements forwarding: per-package look-up
- system-specific, keep small (3000 lines) for easy porting

architecture and implementation



slow path (30,000 lines)

- implements forwarding *logic*
 - MAC learning, load balancing
- remote visibility and configuration interfaces
 - NetFlow, OpenFlow, remote management protocols

usage — one big switch

query and configure a collection of virtual switches as a single switch

- create a single logical switch over multiple Open vSwitches on separate physical servers
- a global management process synthesizes a logical view from the switches' configuration
- lets admins operate on that view
- build CLI for configuring the network as a whole
- each virtual port corresponds to a unique VM

OVS recap

networking layer *at the end host*

- resembles physical switch in its simplest deployment
- exposing interfaces for globally managing configuration and forwarding state
- enables distribution of switch functions across multiple servers
 - decouples the logical network topology from the physical one

a larger perspective

a growing tension between simplicity and requirements of modern enterprises and datacenter

- simplicity has been losing out to functionality

virtualization with OVS — a resolution?

- virtualization extends to the host (edge)
- a clean distinction between the simple core and the edge
- OVS implements the requirements
 - security, visibility, mobility