Vijoy Pandey, Ph.D.
CTO, Network OS
IBM Distinguished Engineer
vijoy.pandey@us.ibm.com

Panel : Future Data Center Networks

Networking folks were poor

- Custom silicon or poor functionality
- Low bandwidth ASICs
- Poor topologies
- Immature protocols
- Non-robust control plane software

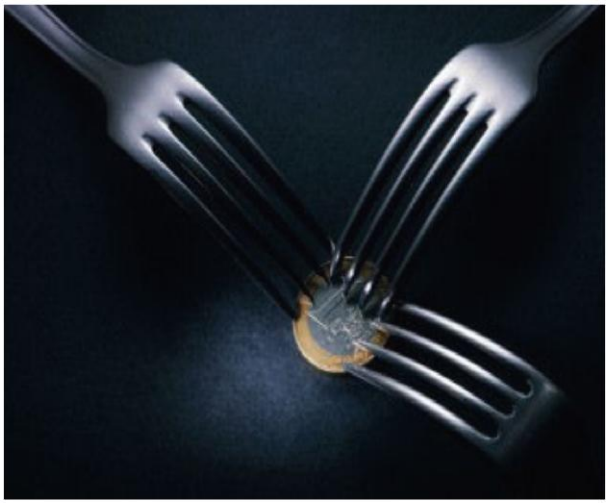

Data Center network over the past decade

First wave of Data Center (network) deployments were all about **Deployment Velocity (Time To Value)**

- How quickly can you deploy infrastructure?
- How scalable is the infrastructure?
- How easily can you manage this scale-out infrastructure?

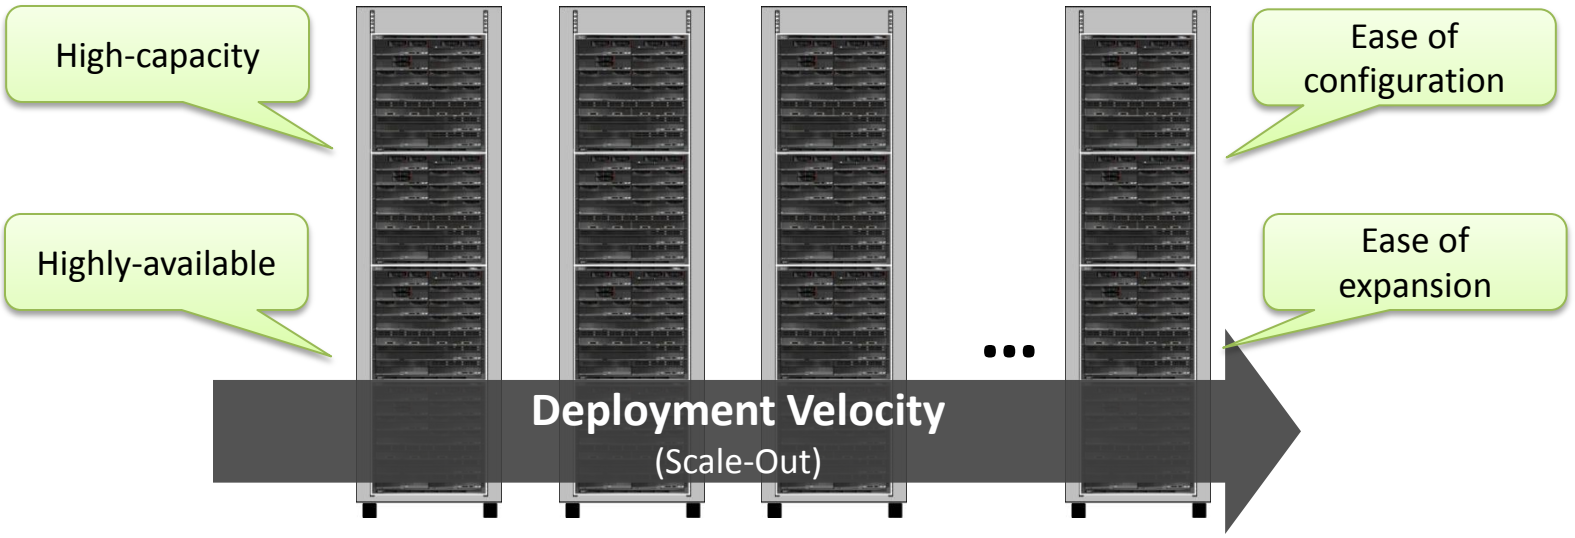

More Bandwidth

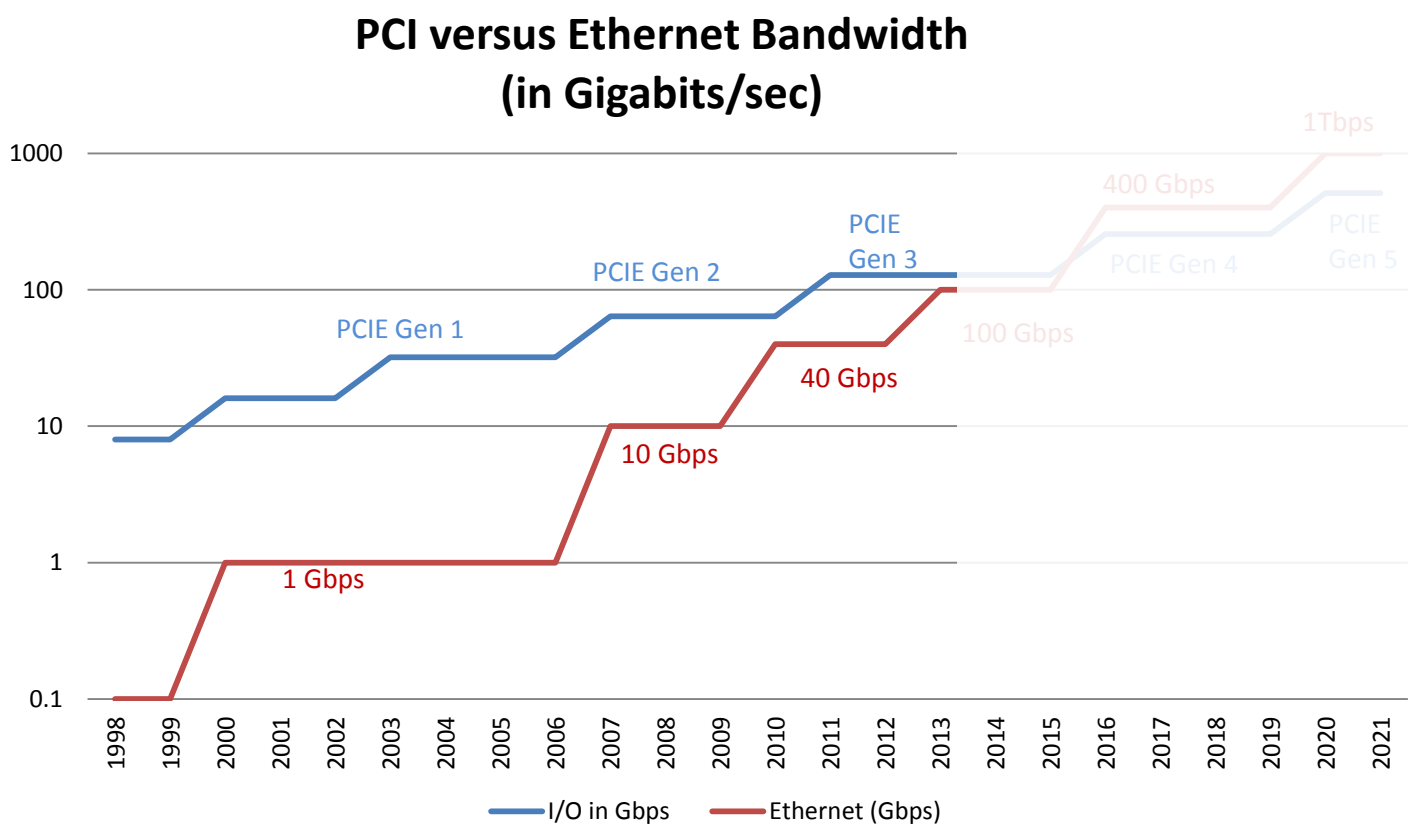

Better Topologies

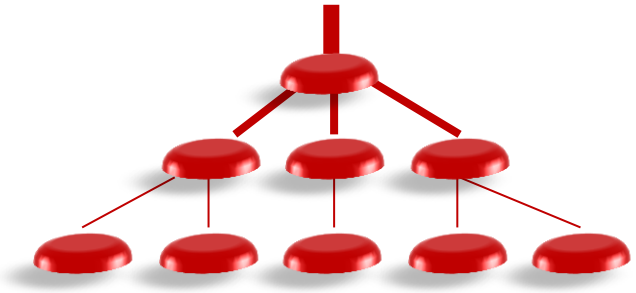

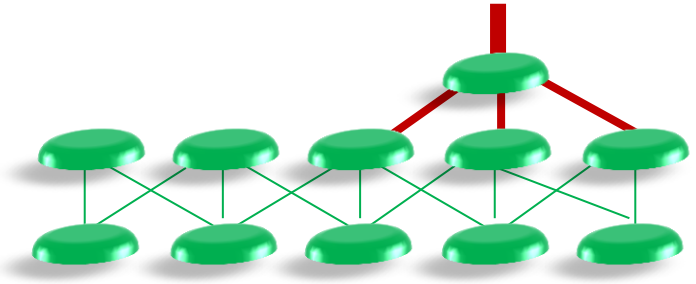

- Multi-tiered tree topologies
- High oversubscription
- Expensive, high bandwidth uplinks
- Robustness of higher tier product has been a concern

- 2-tiered mesh or Clos topologies
- Oversubscription only to WAN/core
- Large cross sectional bandwidth (TOR bandwidth is cheap)
- Mature Layer 2/3 software

We are entering an era of Network Affluence

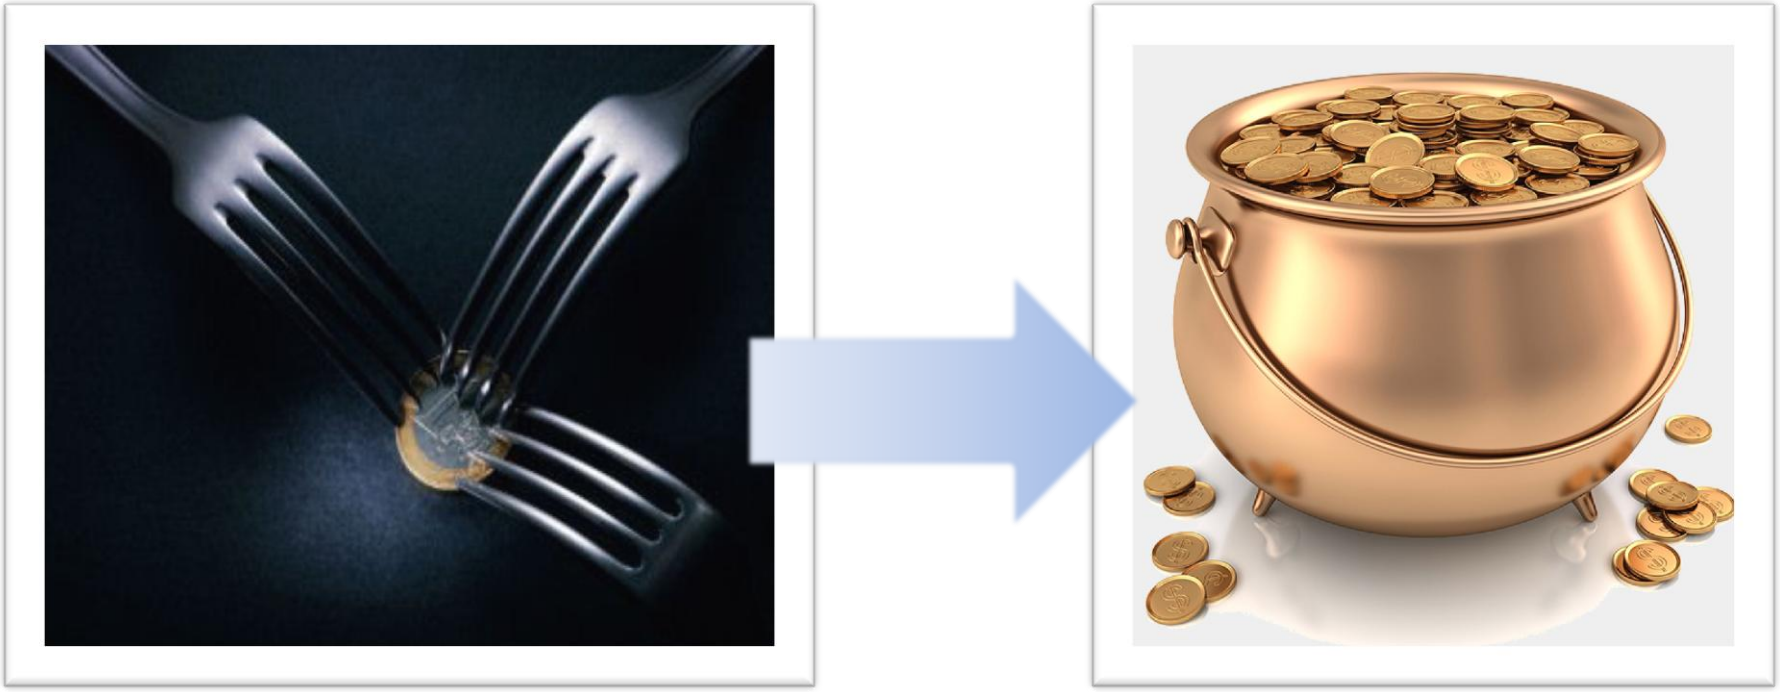

With affluence comes a demand for Quality of Life:

- Can you ease my provisioning headache?
- Can you hide all complexity of the physical infrastructure?
- Can my applications talk to my network?
- Can you simplify how I monitor my network?

- Can “this particular communication” be of “Platinum” service
  - Can you guarantee certain latency characteristics? End to End?
  - Can you guarantee certain bandwidth? End to End?

Application Velocity

Next-Gen Data Center Network deployments will demand both:

- **Application Velocity**
  - Can you provision virtualized network resources along with compute/storage
  - Can the network be smarter due to application awareness
  - Can you quickly and effectively enable newer network services

- **Deployment Velocity**
  - How quickly can you deploy infrastructure
  - How scalable is the infrastructure
  - How easily can you manage this scale-out infrastructure

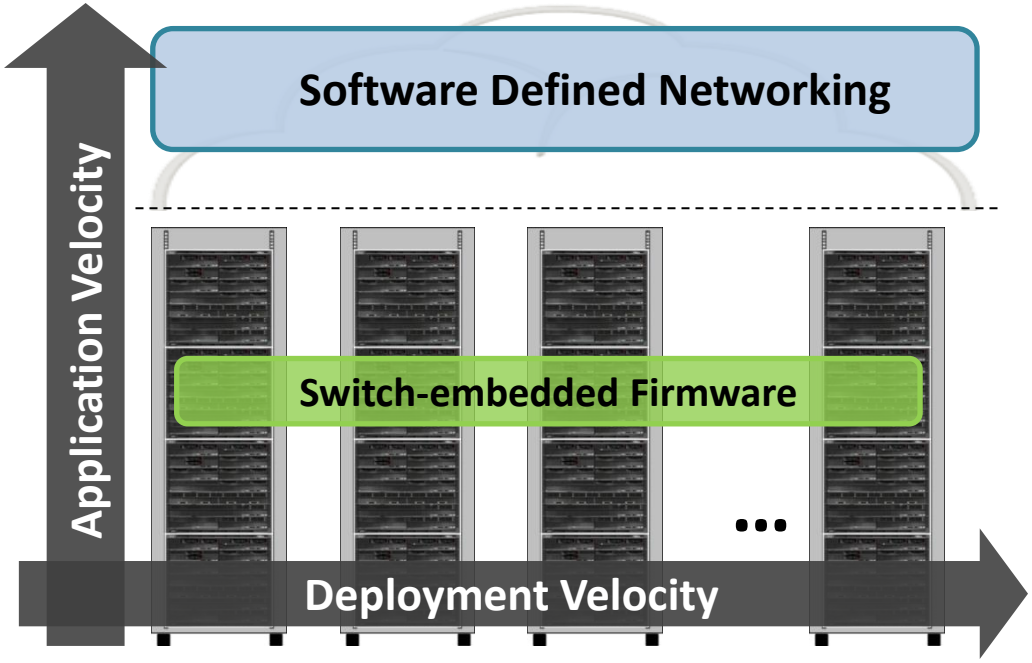

Virtual Provisioning of an Integrated System

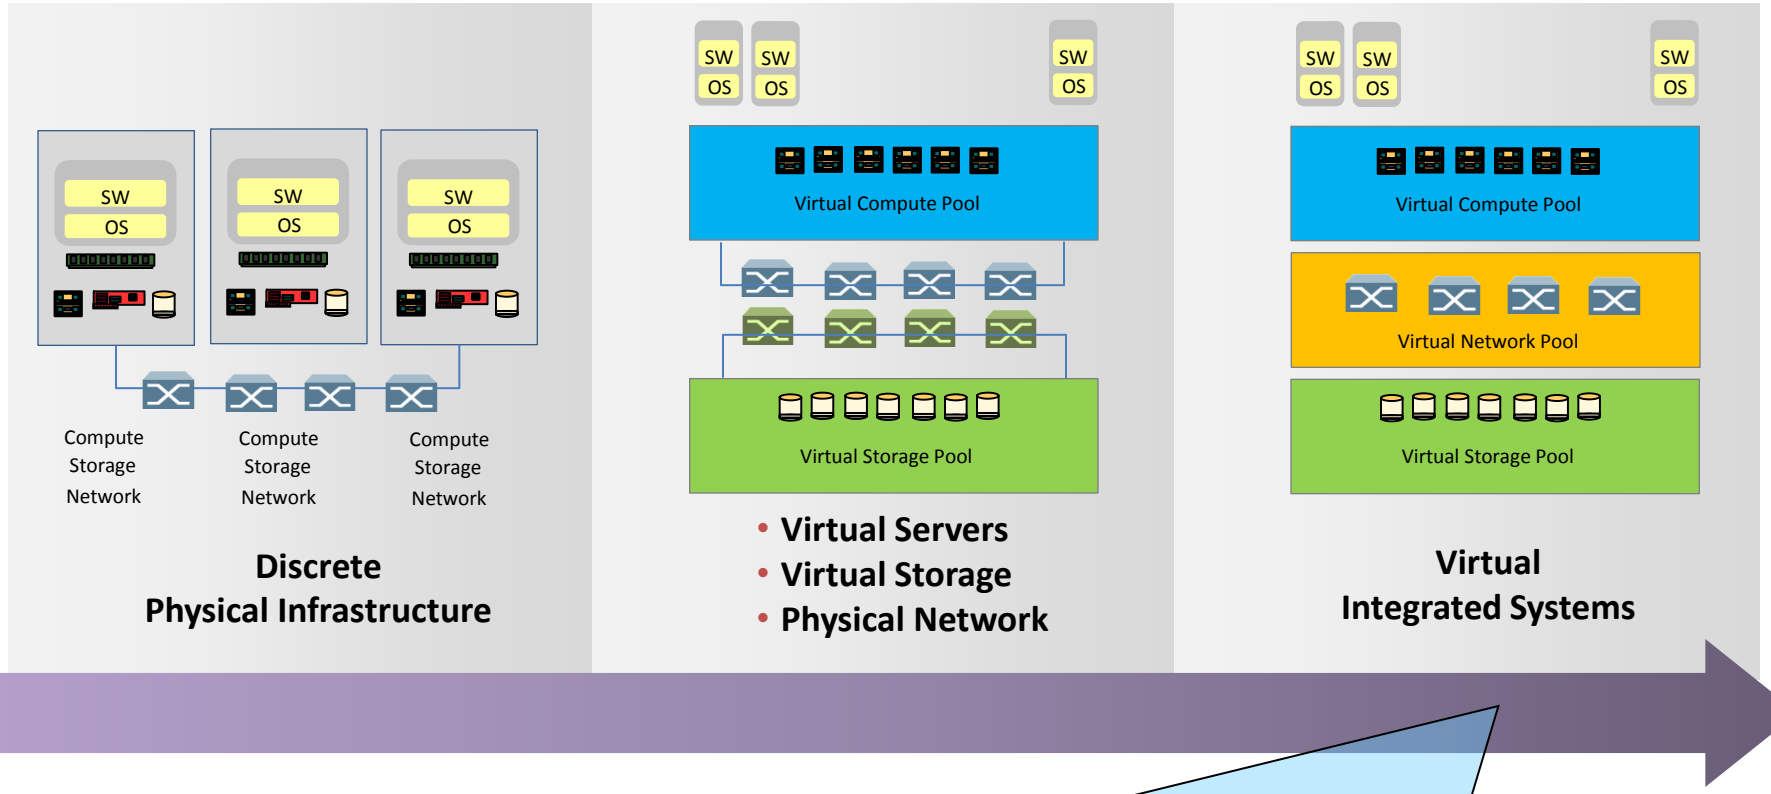

Managing a group of systems – servers, storage, network with the simplicity of a single system

- Virtualized System comprises servers, storage and networking
- End to end experience
  - Initial set up
  - Provisioning of new workloads, including image management
  - Continuous optimization through mobility, etc.

Virtualized Provisioning : Example

IBM

Application connectivity services		
Enterprise or Cloud	• allow users to declaratively specify logical application topologies • path attributes, security rules, and service traversal • instantiate paths, rules, etc. using SDN (virtual or physical)	• seamless integration between application deployment and required network configuration • removes need for separate network admin handoff • services can be constrained / specified by networking team

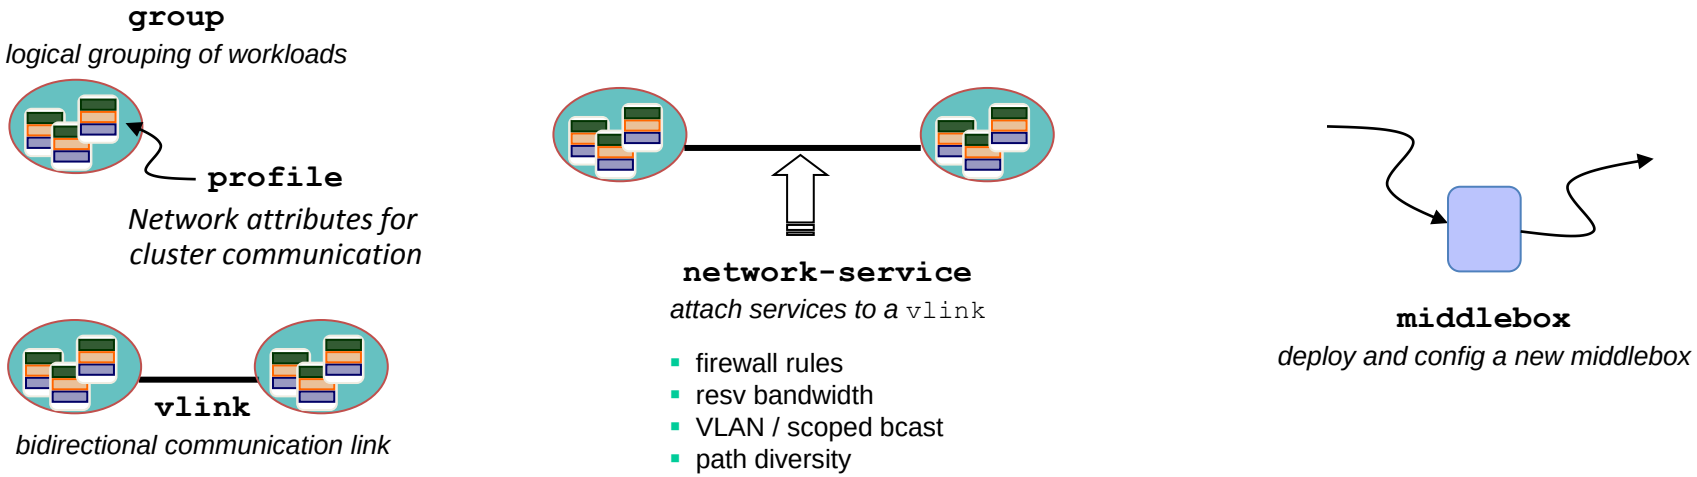

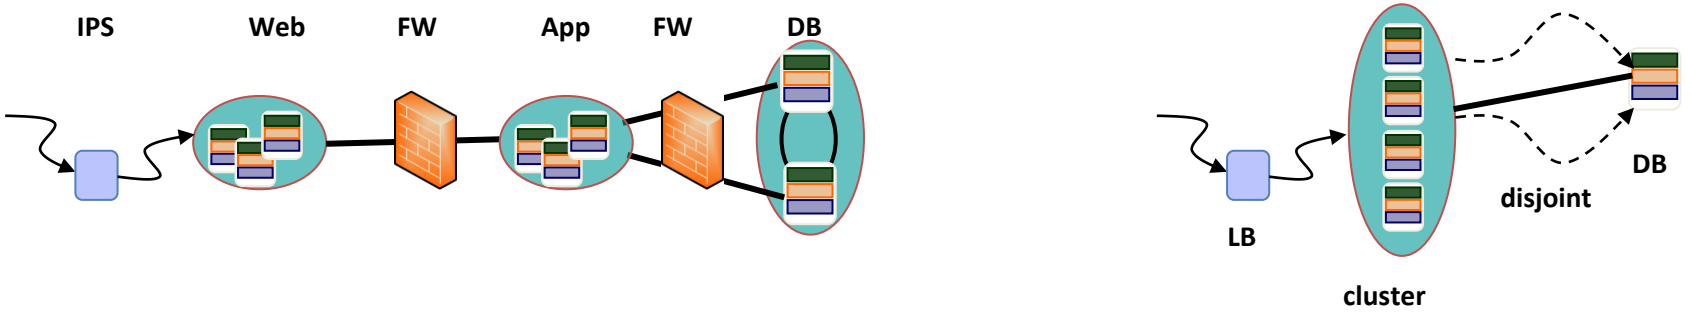

Applications influence the network

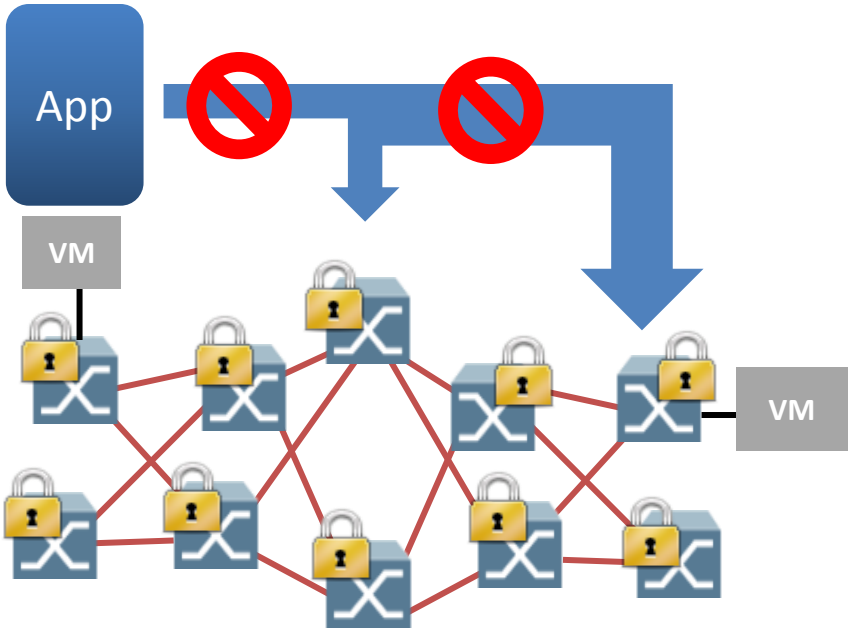

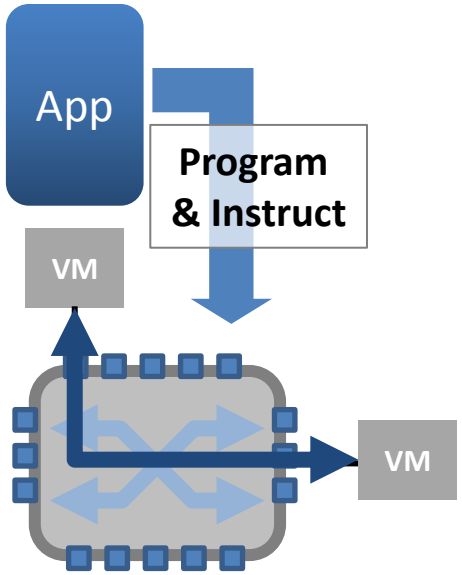

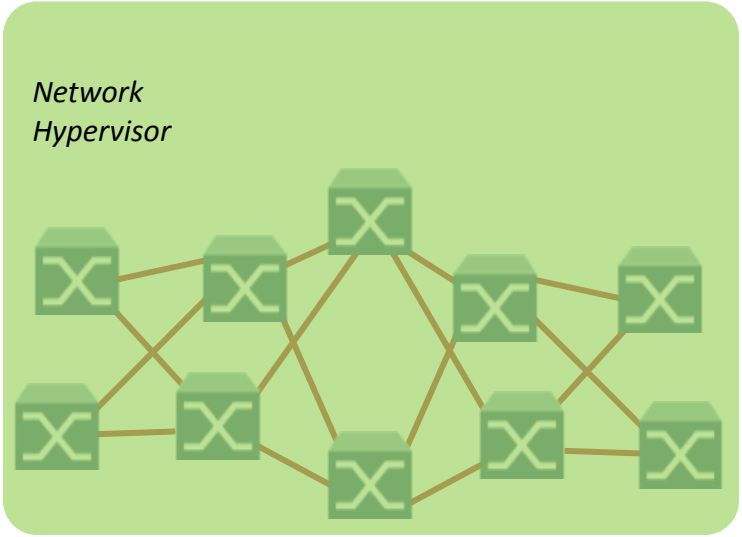

What is Software Defined Networking?

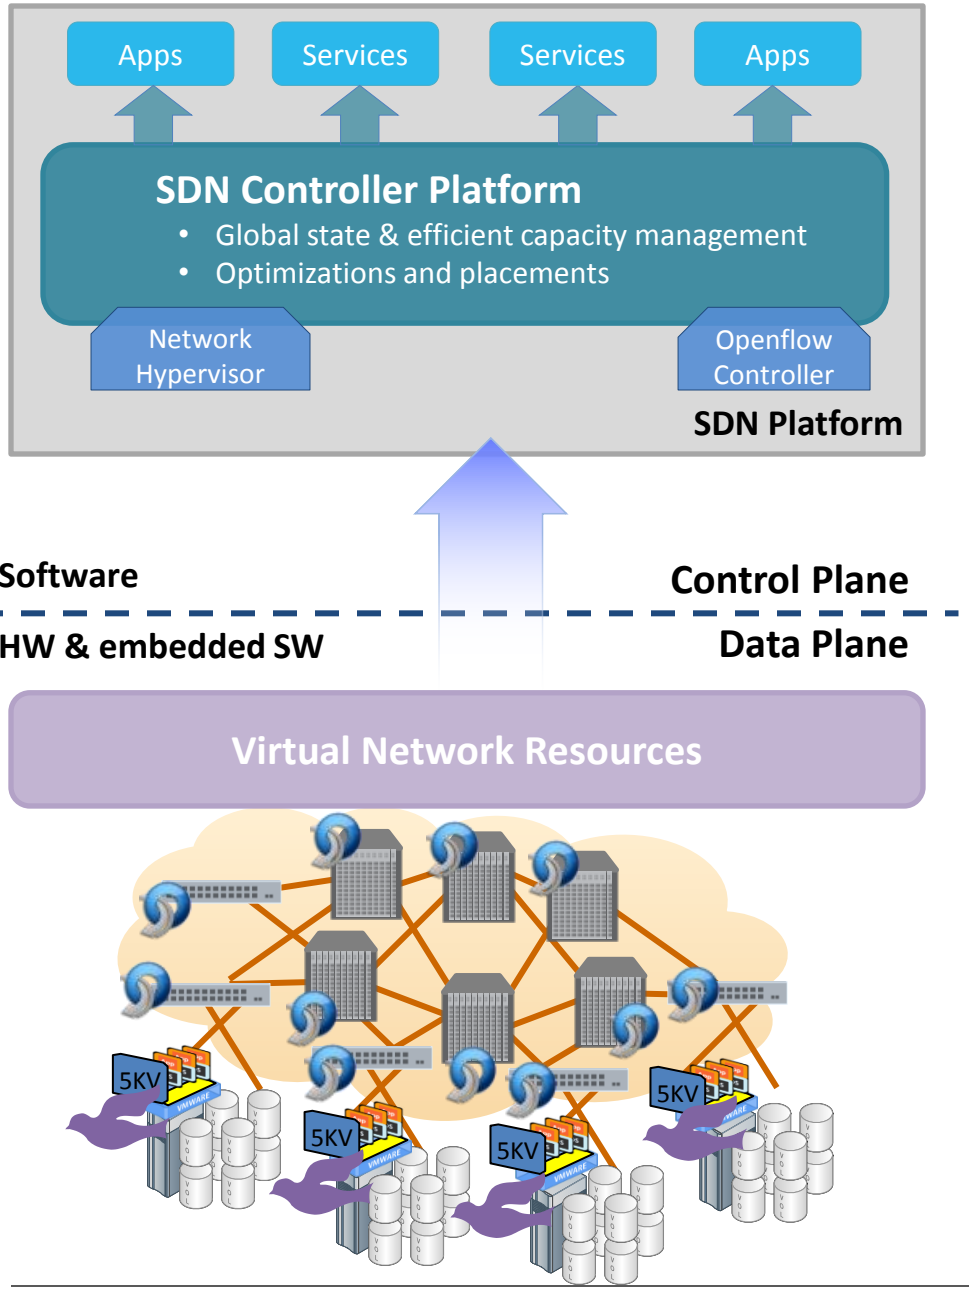

- **Applications-awareness benefits:**
  - Business Applications and Services can program and influence the network
  - Create and deploy new applications and protocols quickly
- **Network Hypervisor benefits:**
  - Virtualized network resource provisioning
  - De-couples virtual network from physical network
  - Simple "configure once" network
  - Cloud scale (e.g. multi-tenant)
- **Control-Data separation benefits:**
  - End-2-end Semantics and Guarantees
  - Simpler to deploy, debug and monitor
  - Fine grained control for each client-server pair(s)
  - Openflow (protocol and various controllers) are a standard way of achieving this

Openflow as defined today

Mishmash of concepts within the SAME data center

Distributed or Centralized

- Ethernet topologies were built distributed

  Scalable but hard to monitor

- Openflow topologies (today) are centralized

  Control-data separation forces this model

Packet or Flow Switched

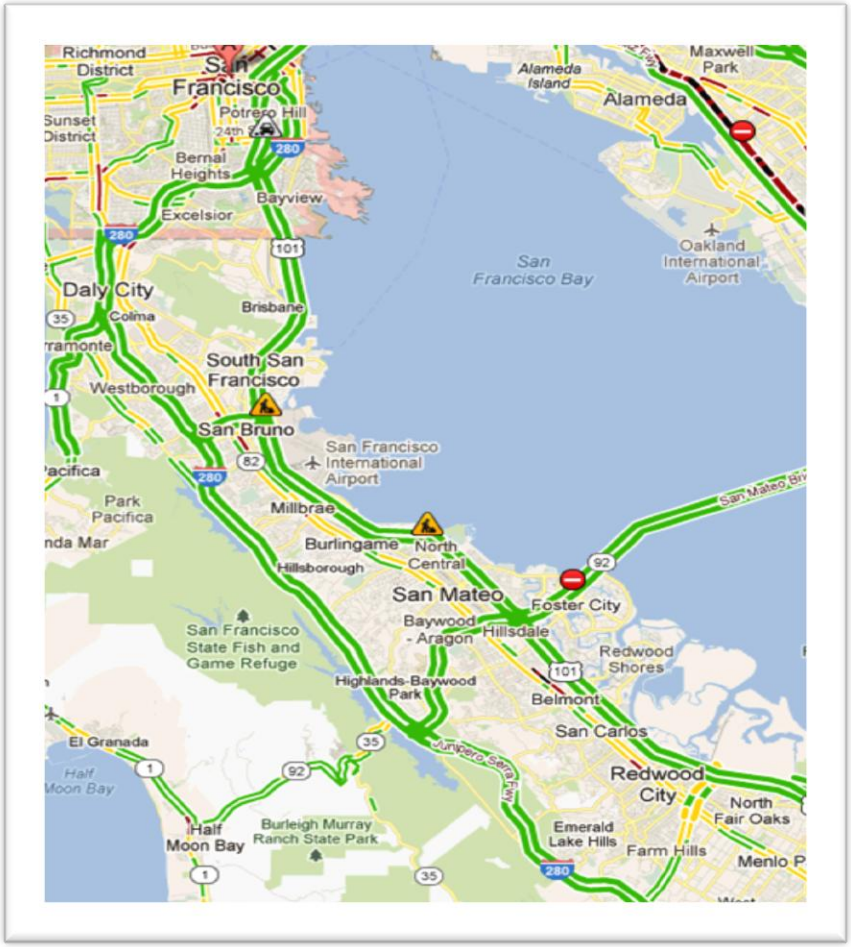

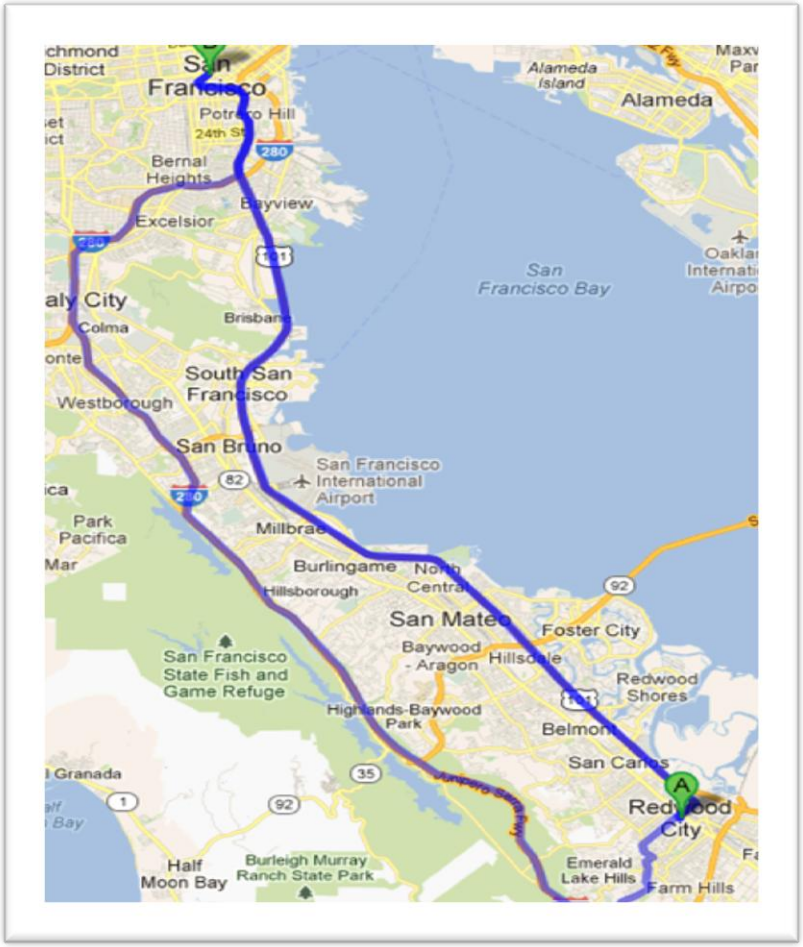

- Ethernet topologies are packet switched

  Statistical link utilization

- Openflow topologies (today) are flow switched

  Application level network control

Open questions for the Research & Openflow Communities

- No customer pays for the re-invention of the wheel
- Customers do pay for a smoother ride

Open Questions for the Research and Openflow Communities

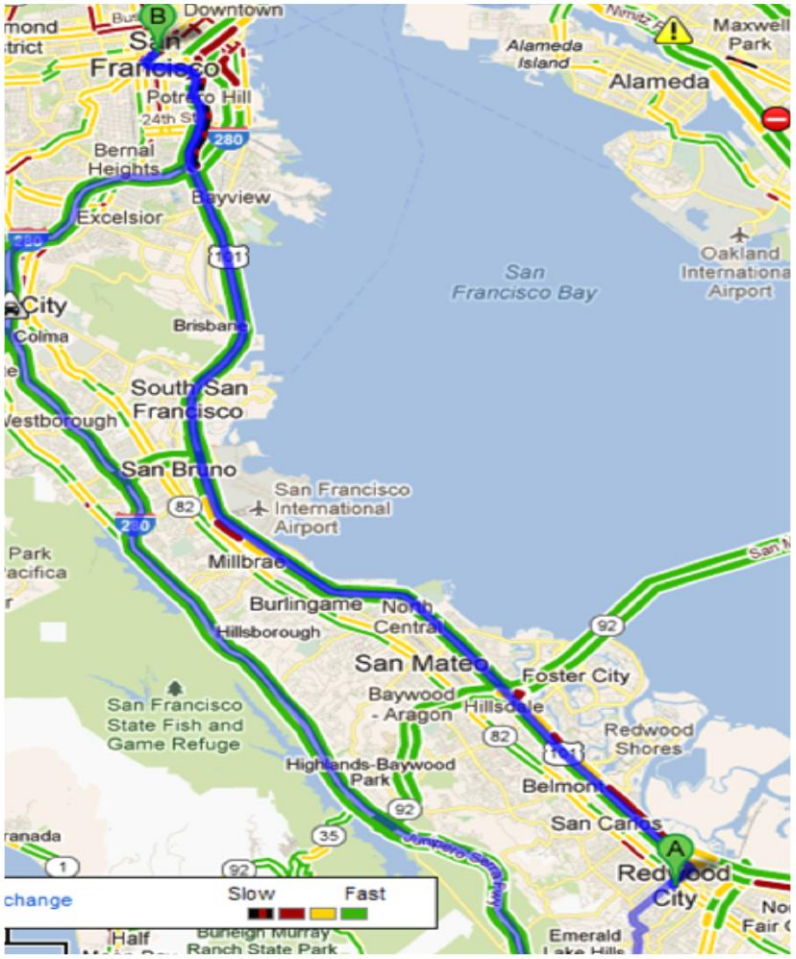

Isolate a few (long) flows for preferential treatment by applications

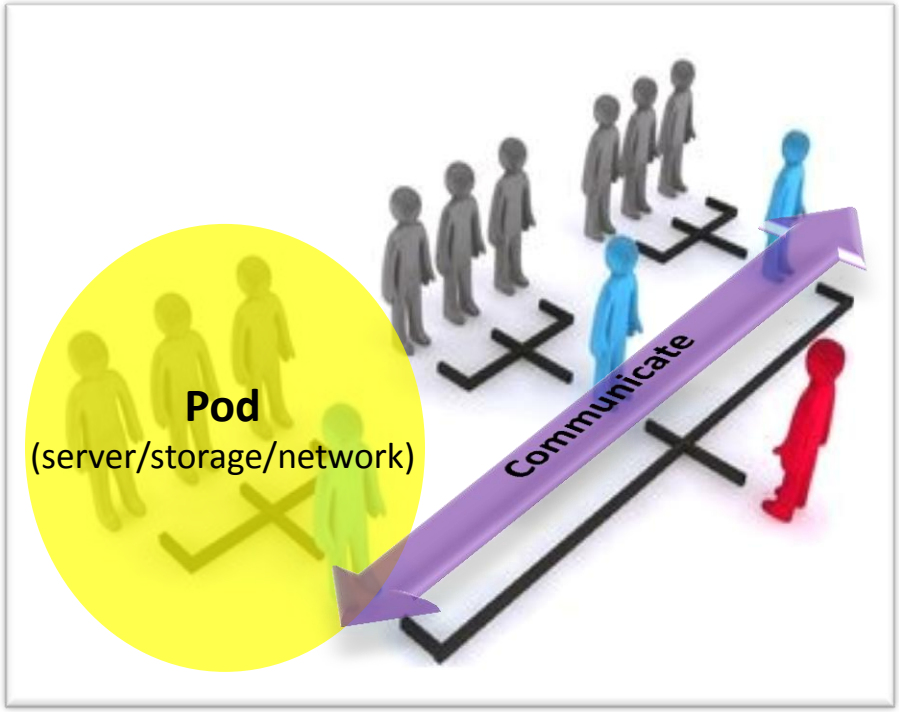

Federation of controllers with each controller handling (smaller) integrated system (pod)