# Cloud Computing and the Internet

Conferenza GARR 2010

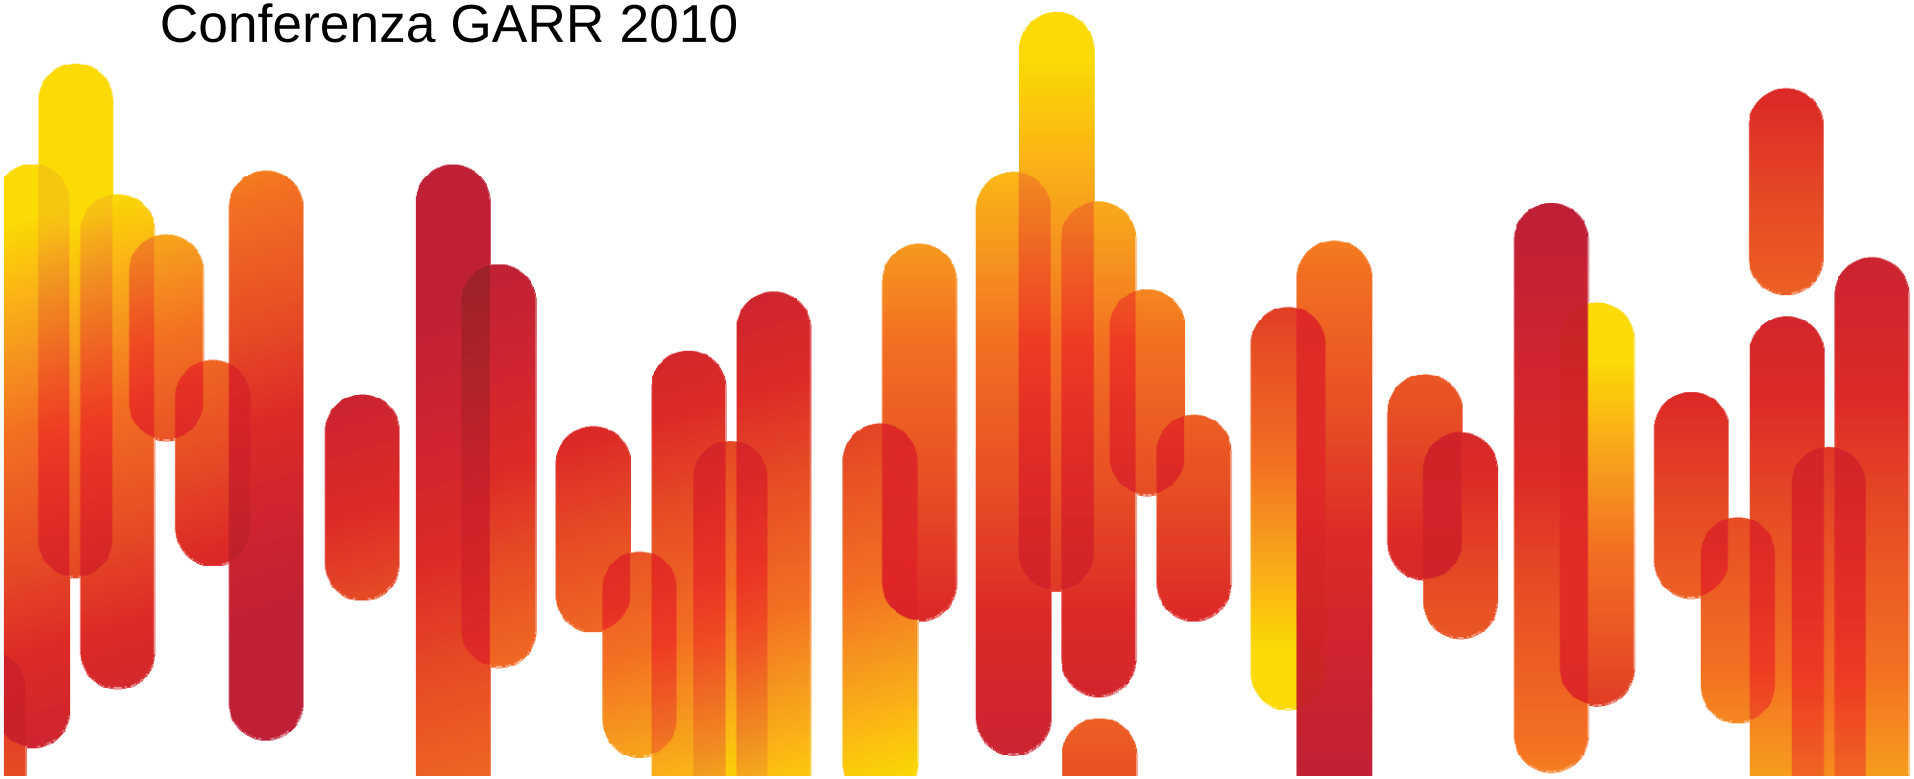

# Cloud Computing

- The current “buzzword” ;-)
- Your computing is in the cloud!
- Provide computing as a utility
  - Similar to Electricity, Water, Phone service, etc.
- The Internet is the delivery infrastructure
- Limited or no capital expenditure
  - Consumption-based billing, or
  - Usage-based billing
- The model is applicable both inside a company and as a public service


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Key Building Blocks

- The Internet
- Virtualization
- Provisioning portal
- Provisioning/Monitoring/Management software


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Differences/Similarities with previous efforts

- Client–server model – Client–server computing refers broadly to any distributed application that distinguishes between service providers (servers) and service requesters (clients)
- Grid computing — a form of distributed computing and parallel computing, whereby a 'super and virtual computer' is composed of a cluster of networked, loosely coupled computers acting in concert to perform very large tasks
- Utility computing — the packaging of computing resources, such as computation and storage, as a metered service similar to a traditional public utility, such as electricity


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# New Data Centers

- Data Centers are evolving rapidly
  - From virtualization to cloud computing
- Physical border of a data center disappears
  - Public and hybrid clouds
  - The Internet is a key enabler
- More bandwidth and new protocols are required
  - Address duplication
  - Overlay networks
  - Independence of address from locator


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Virtualization

- Unbundling of software from hardware
  - The end of mainframes, minicomputers, and proprietary Unix architecture
- The X86 platform as the only hardware platform for computing
  - Superior performance and memory
  - Constantly increasing reliability
  - *Angelo Raffaele Meo was right!*
- Virtualization software (hypervisors) on bare metal
  - Vmware, KVM, Hyper-V, XEN
- Virtual Machines (VMs) instead of physical servers


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Virtualization is … well, not exactly new

- Nothing new! Concept known to mainframes back in the ’70s

  Virtualization is not a new concept

  Mainframe of the ‘70s were underutilized and over-engineered


http://www-07.ibm.com/systems/my/z/about/timeline/1970/


© 2010 Cisco and/or its affiliates. All rights reserved.

# Virtualization Characteristics

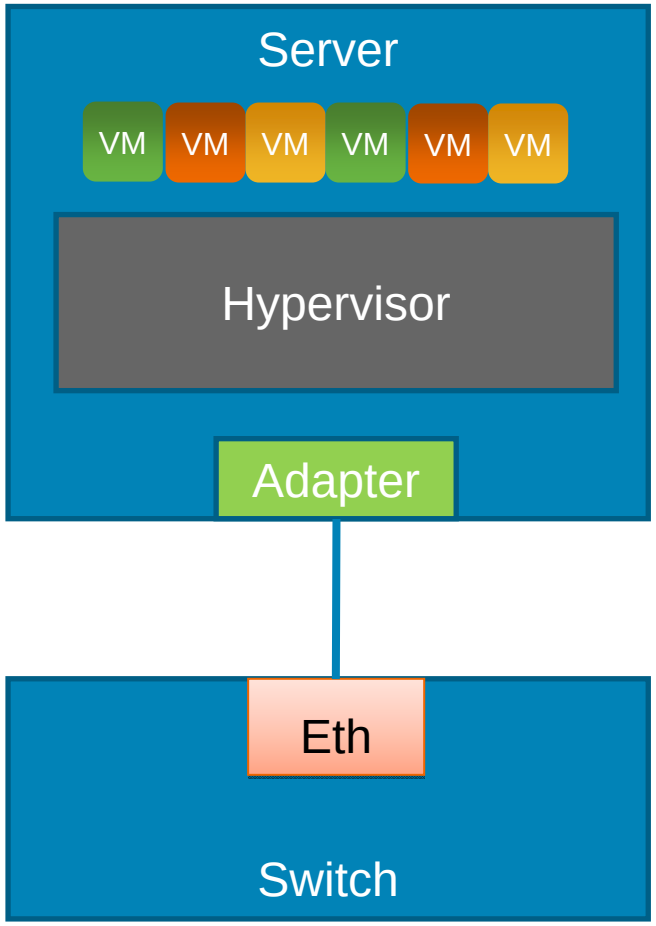


**Server CPU (x86)**
- Today: 2 – 4 cores/CPU common
- Trend: 6-12+ cores, crypto kernels

**Server**
- Today: mostly dual sockets
- Trend: 4 sockets (high-end servers)

**Virtual Machines**
- Today: ~10 VMs per server
- Trend: 10s to several-100 VMs

**Adapter**
- Today: 1G for rack server, 10G for blade chassis
- Trend: 10G

**Networking Options:**
1. Bring VM traffic out to external switch
2. Bring network functionality into Hypervisor


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Ethernet Is the Only Surviving Network

At least in the data center!


- Servers connected to 10GE
- Backbone at 40GE and 100 GE
- Requirement for high bisectional bandwidth

Image Credit:
Flickr user **Catherine Rankovic**
Creative Commons License


© 2010 Cisco and/or its affiliates. All rights reserved.

# Wire Once Architecture

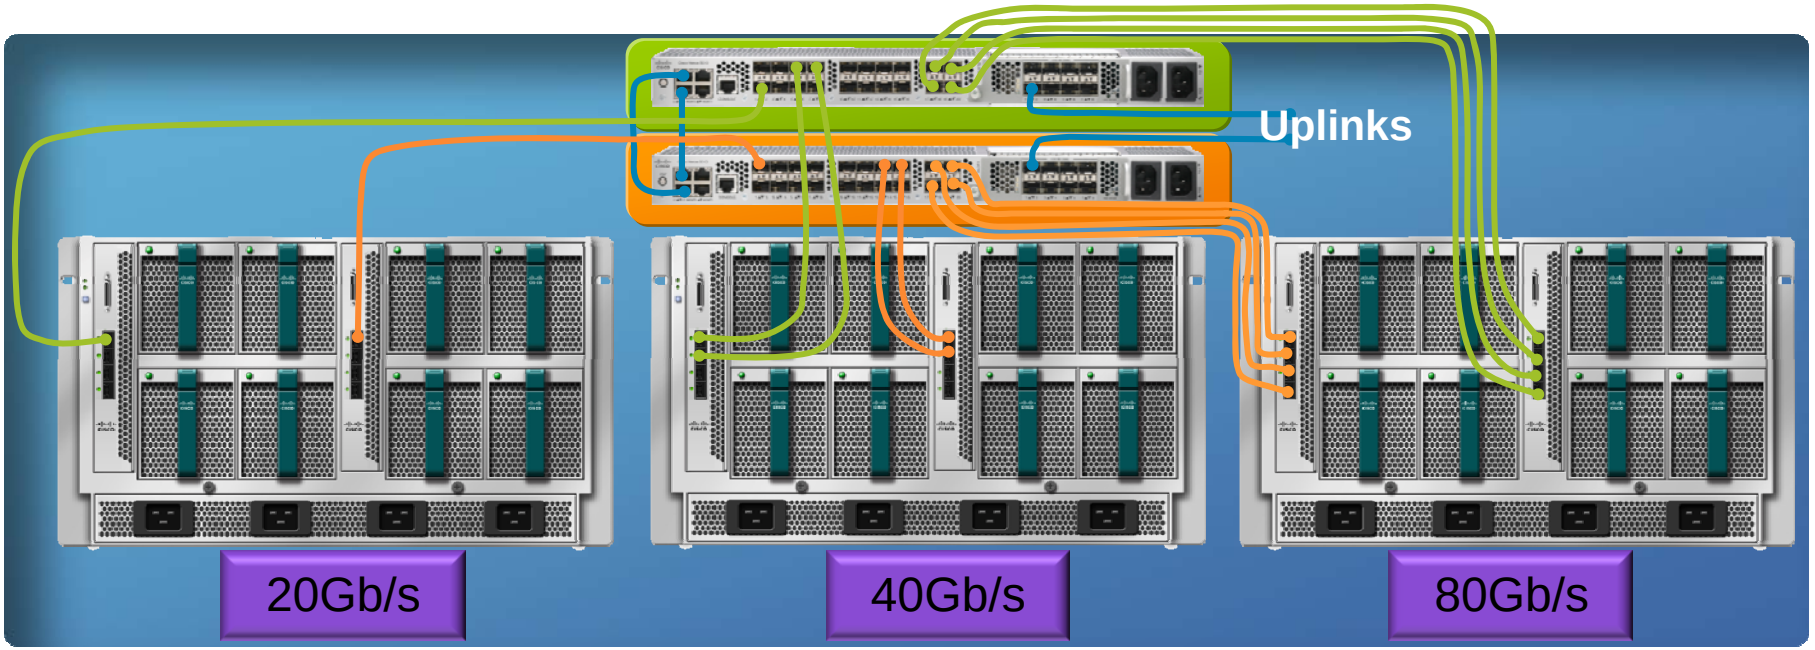


- Wire once for bandwidth, not connectivity
- All links can be active all the time


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Ethernet Switches and X86 Servers

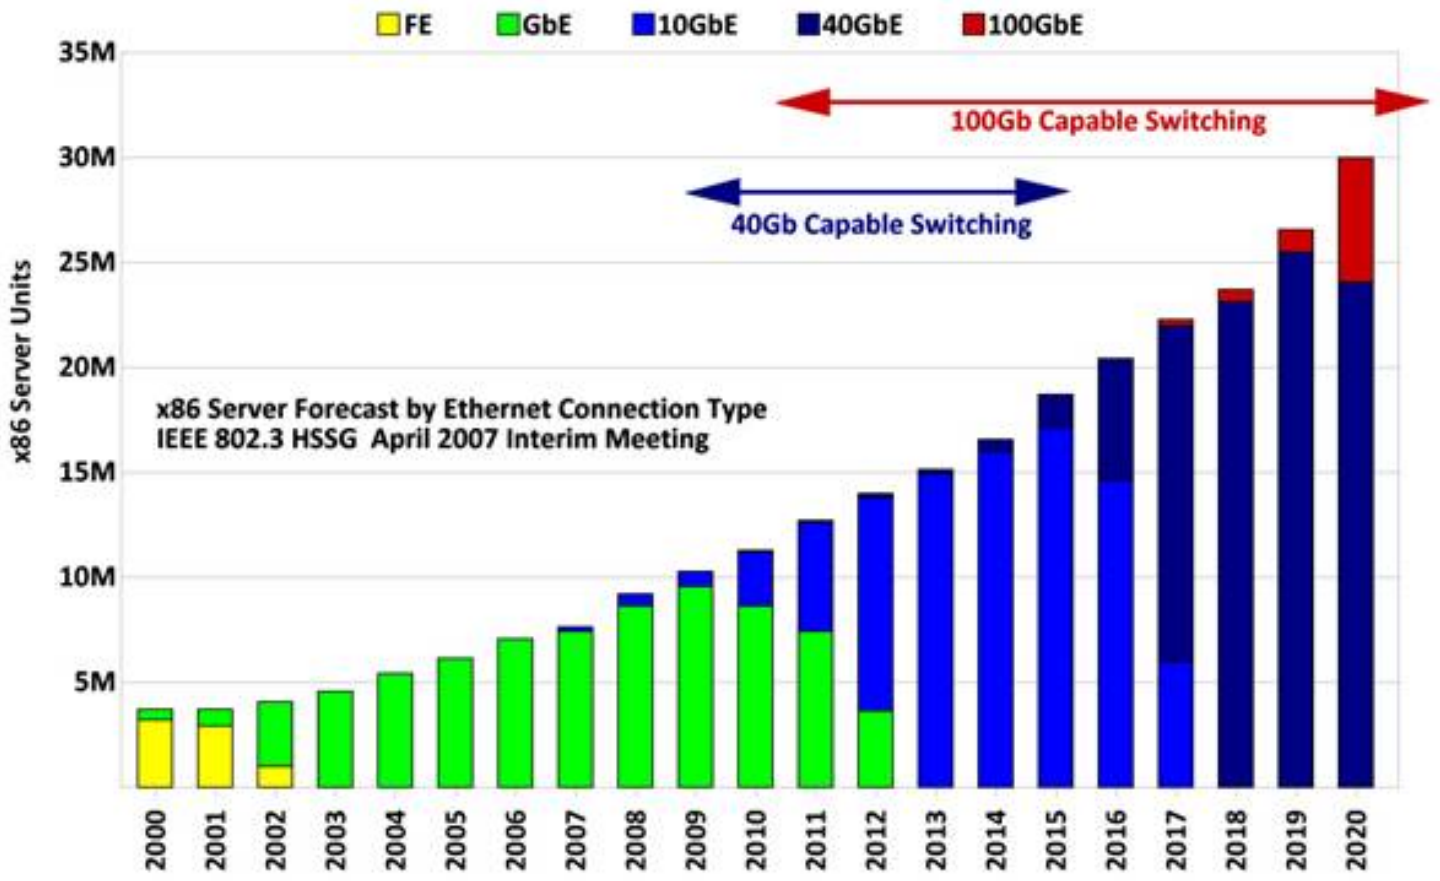



© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# FC Over Ethernet (FCoE)

## FCoE

- Mapping of FC Frames over Ethernet
- Enables FC to Run on a Lossless Ethernet Network

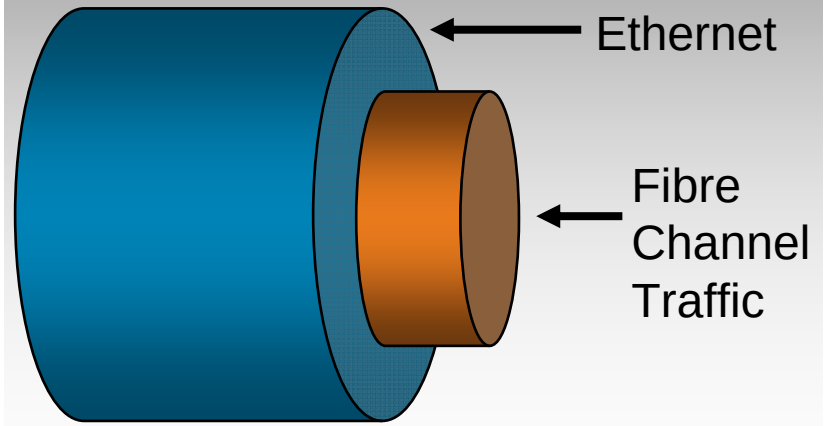


## Benefits

- Fewer Cables
  - Both block I/O & Ethernet traffic co-exist on same cable
- Fewer adapters needed
- Overall less power
- Interoperates with existing SAN's
  - Management SAN's remains constant
- No Gateway

# RoCE (RDMA Over Converged Ethernet)

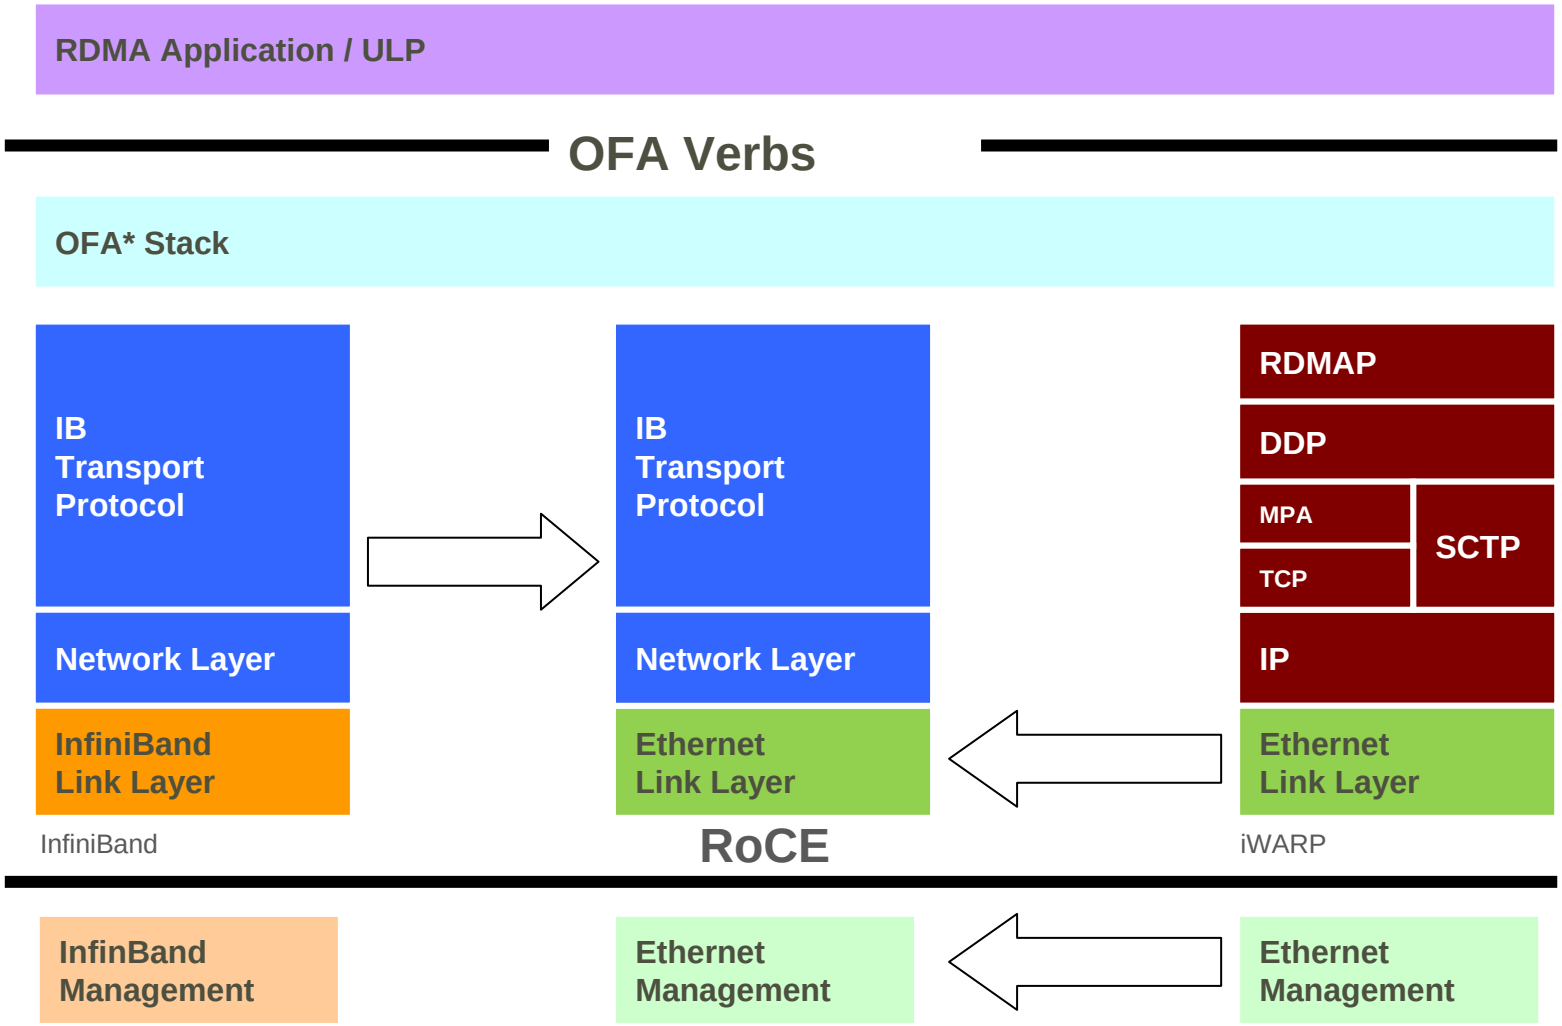


Source: Mellanox Corp.


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Virtual Machine Must Be Able to Move

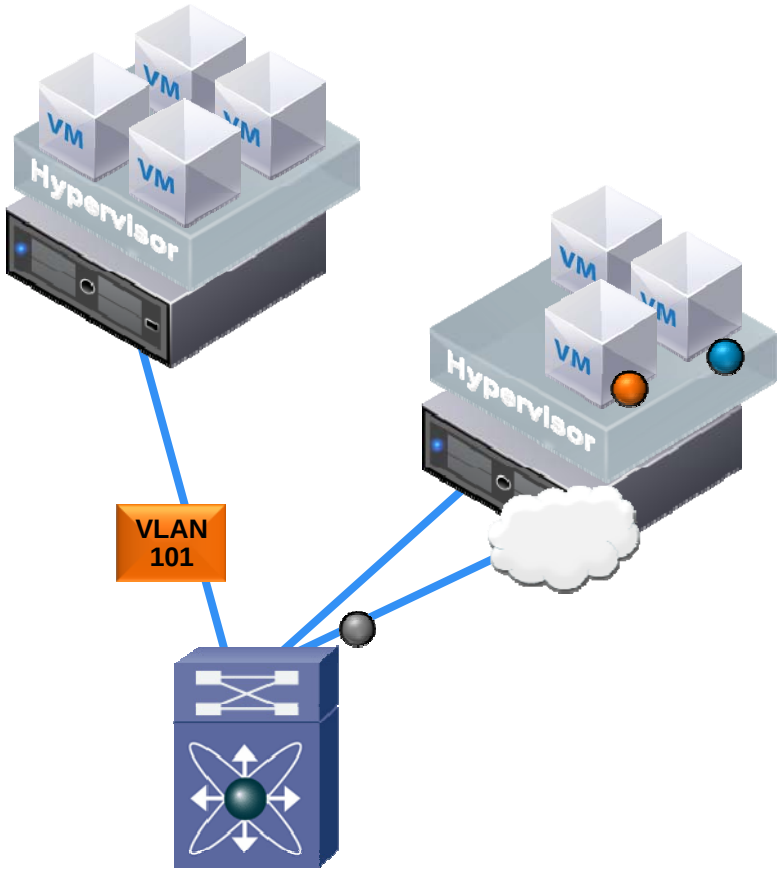


**Problems:**

- VMotion may move VMs across physical ports—policy must follow
- Cannot correlate traffic on physical links—from multiple VMs

**Solution:**

- An network architecture VM-aware
- Extends network to the VM
- Consistent services
- Coordinated, coherent management


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# VM movement corollaries

- VMs move and keep their MAC and IP addresses
  - Large Layer 2 network in the datacenter
- VM cloning (move and keep the original)
  - Duplicated MAC and IP addresses
    - Different VLANs?
- VM move across datacenter
  - Huge bandwidth requirement (need to move storage also)
  - Layer 2 networks must span multiple data centers
    - OTV, VPLS


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Cloud types

- Private (a flexible computing infrastructure inside a company)
- Public (a public provider of VMs)
- Hybrid (a part of a private cloud that is off-sourced to a public cloud)


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Example of public clouds

- Amazon Elastic Compute Cloud (EC2)
- IBM's Blue Cloud
- Oracle Cloud Computing
- Google AppEngine
- Windows Azure Services Platform.


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Hybrid Clouds

IP address duplication
MAC address duplication
VLAN-ID duplications

Extension of a
private cloud
into a public cloud

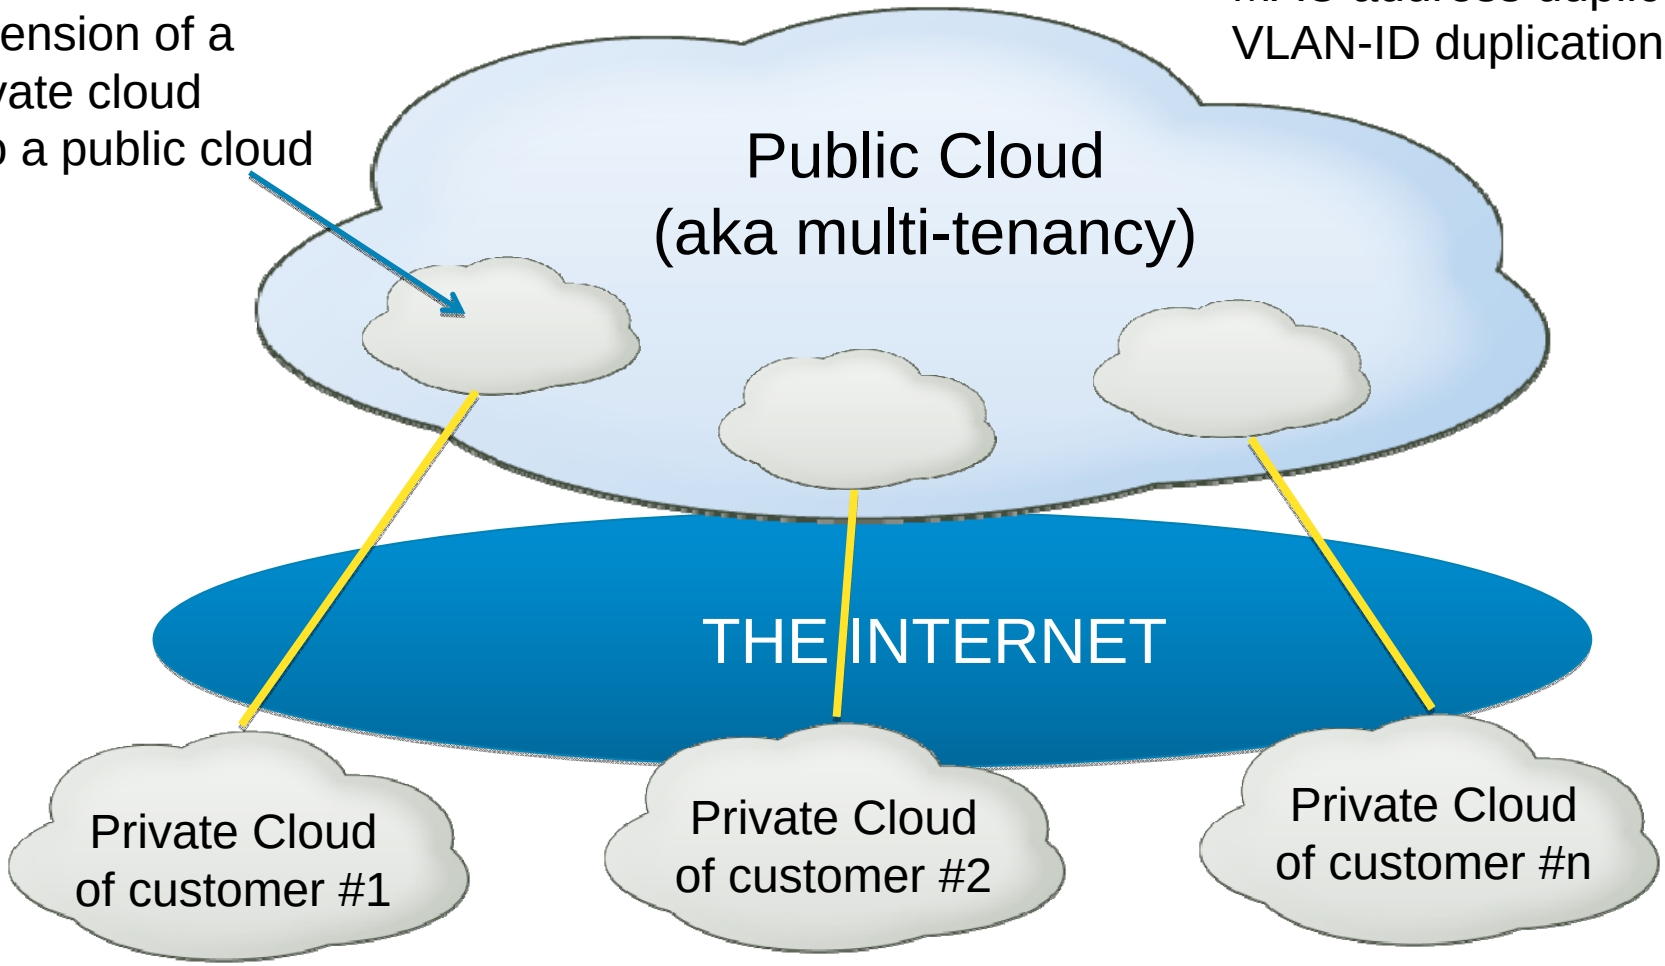



© 2010 Cisco and/or its affiliates. All rights reserved.

# Possible Solutions

| | | Protocol Encapsulated | |
|---|---|---|---|
| | | Ethernet | IP |
| Encapsulation | Ethernet | TRILL | N/A |
| | IP | OTV/VPLS | LISP |


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# TRILL (Transparent Interconnection of Lots of Links)

- IETF standard
  - Computes topology and forwarding via IS-IS
  - Provides optimal pair-wise unicast forwarding
  - Provides multi-pathing for unicast and multicast frames
  - Provides Seamless interoperability with existing devices
- Approach alternative to spanning tree
  - Eliminates Spanning Tree from the backbone
- Important when there is limited or no differentiation in speed between access links and backbone links
- Reduces latency


Image Credit:
Flickr user Clint M Chilcott Creative Commons License


© 2010 Cisco and/or its affiliates. All rights reserved.

# Overlay Transport Virtualization (OTV)

IP-based (MACinIP) multi-point ethernet extension

- Ethernet LAN Extension over any Network
  - Works over dark fiber, MPLS, or IP
  - Multi-data center scalability
- Simplified Configuration & Operation
  - Seamless overlay - No network re-design
  - Single touch site configuration
- High Resiliency
  - Failure domain isolation
  - Seamless Multi-homing
- Maximizes available bandwidth
  - Automated multi-pathing
  - Optimal multicast replication

Many physical sites –
One logical Data Center

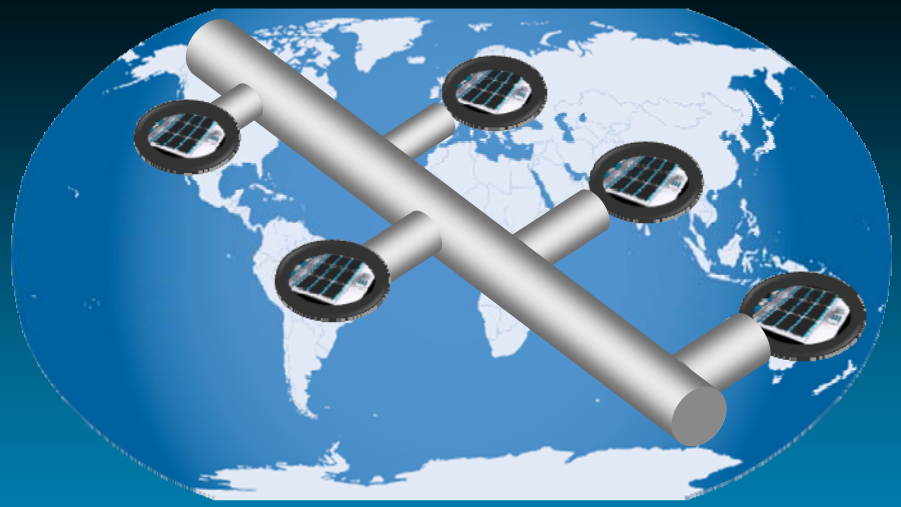

**Any Workload, Anytime, Anywhere**
Unleashing the full potential of compute virtualization

# LISP – A Level of Indirection for IP Addressing

- Creates a "Level of indirection" by using two namespaces – EID and RLOC

- **EID (Endpoint Identifier)** is the host IP address
- **RLOC (Routing Locator)** is the infrastructure IP address of the LISP router
- **Mapping Database (M-DB)** is the distributed database and policy repository

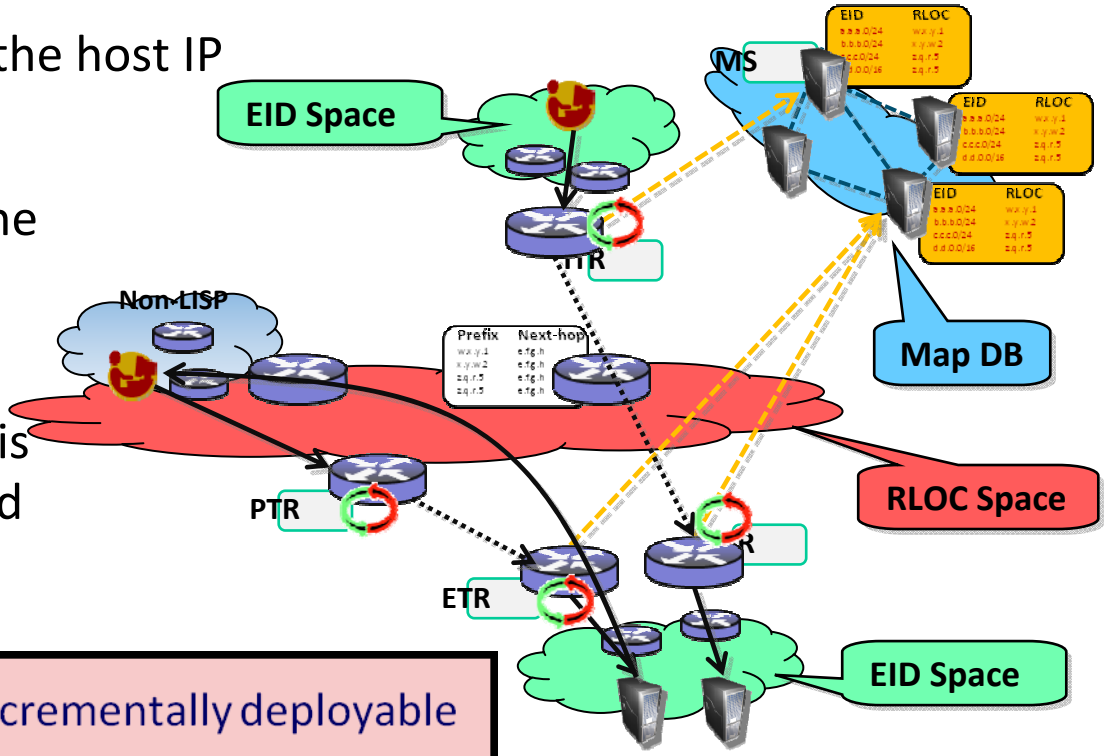


- Network-based solution
- No host changes
- Minimal configuration
- Incrementally deployable
- Support for LISP-MN


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Do I use clouds?

- Sure, my last tax return was done on a cloud application
- I am thinking to run a version server on Amazon to write a new book with two other authors located around the world
  - I don’t have enough uplink bandwidth at my house to run a server myself
- I use a cloud application to share files over the Internet


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

# Conclusions

- Cloud Computing is a new application of The Internet
  - It promises to reduce computing costs, in particular capital expenditure
- It requires to standardize on the X86 platform (either Windows or Linux)
- It requires significant bandwidth
- It requires new network technologies especially for hybrid clouds and multi-tenancy


© 2010 Cisco and/or its affiliates. All rights reserved. Cisco Confidential


CISCO