OpenFlow/SDN for IaaS Providers

Open Networking Summit 2011
Stanford University

Paul Lappas & Ivan Batanov

The Public Cloud

Our Definition

- Shared infrastructure operated by a service provider where no single client uses a significant percentage of the available capacity.
- Provides many infrastructure services: compute, storage, etc
- Metered billing and elastic consumption

Used by

- SMBs for test/dev and production workloads
- Few enterprise users today
- Very little integration with on-premise networks
  - Some VPN stuff happening – ie Azure VPN and Amazon VPN

Public Cloud Networks Need

Infrastructure

- Large-scale Datacenters ~10k hosts, ~1M VMs
- API provisioning + high rate of change
- Add new capacity with minimal operation overhead
- Mix of “soft” and “hard” devices

Connectivity

- Any-to-any connectivity between arbitrary VMs
- L2 isolation between very large number of domains
- Identity-based routing instead of IP-based routing
- “Fair” or guaranteed resource distribution
- High cross-sectional bandwidth

Cost-effective

Public Cloud Networks Need

Business

- 80% utilization of resources
- Grow on a ‘just in time’ basis
- Minimal customizations with broad broad market reach

Product

- Unlimited ‘burst’ capacity
- Unlimited number and size of L2 domains
- VM mobility
- QoS

Price competitive

Limitations of traditional networks

- Distributed model is broken
  - Complexity of distributed configuration increases exponentially
  - Current/Legacy protocols have inherent protocol limitations
  - There are limitations in # of routers in OSPF, number of cascaded switches, etc
  - As # of nodes increase, so do convergence times
- L2 “pods” are not the solution – need some form of L3 aggregation, but
  - Hardware does not scale well
  - Lack of programmatic access
  - Multi-vendor/platform integration is difficult/impossible
  - Vendor solutions are expensive/proprietary and incomplete

Business needs are not being met

- You end up segmenting your capacity into “pods”
  - And maintaining spare capacity per pod is expensive
- Maintaining one L2 domain per customer is difficult
- Migration of VMs is expensive, sometimes impossible
- Adding more capacity around the DC requires significant planning and we usually get it wrong

- Customers run out of capacity their ‘pod’
- There is really no way to enforce QoS over the entire network

The Paradigm Shift

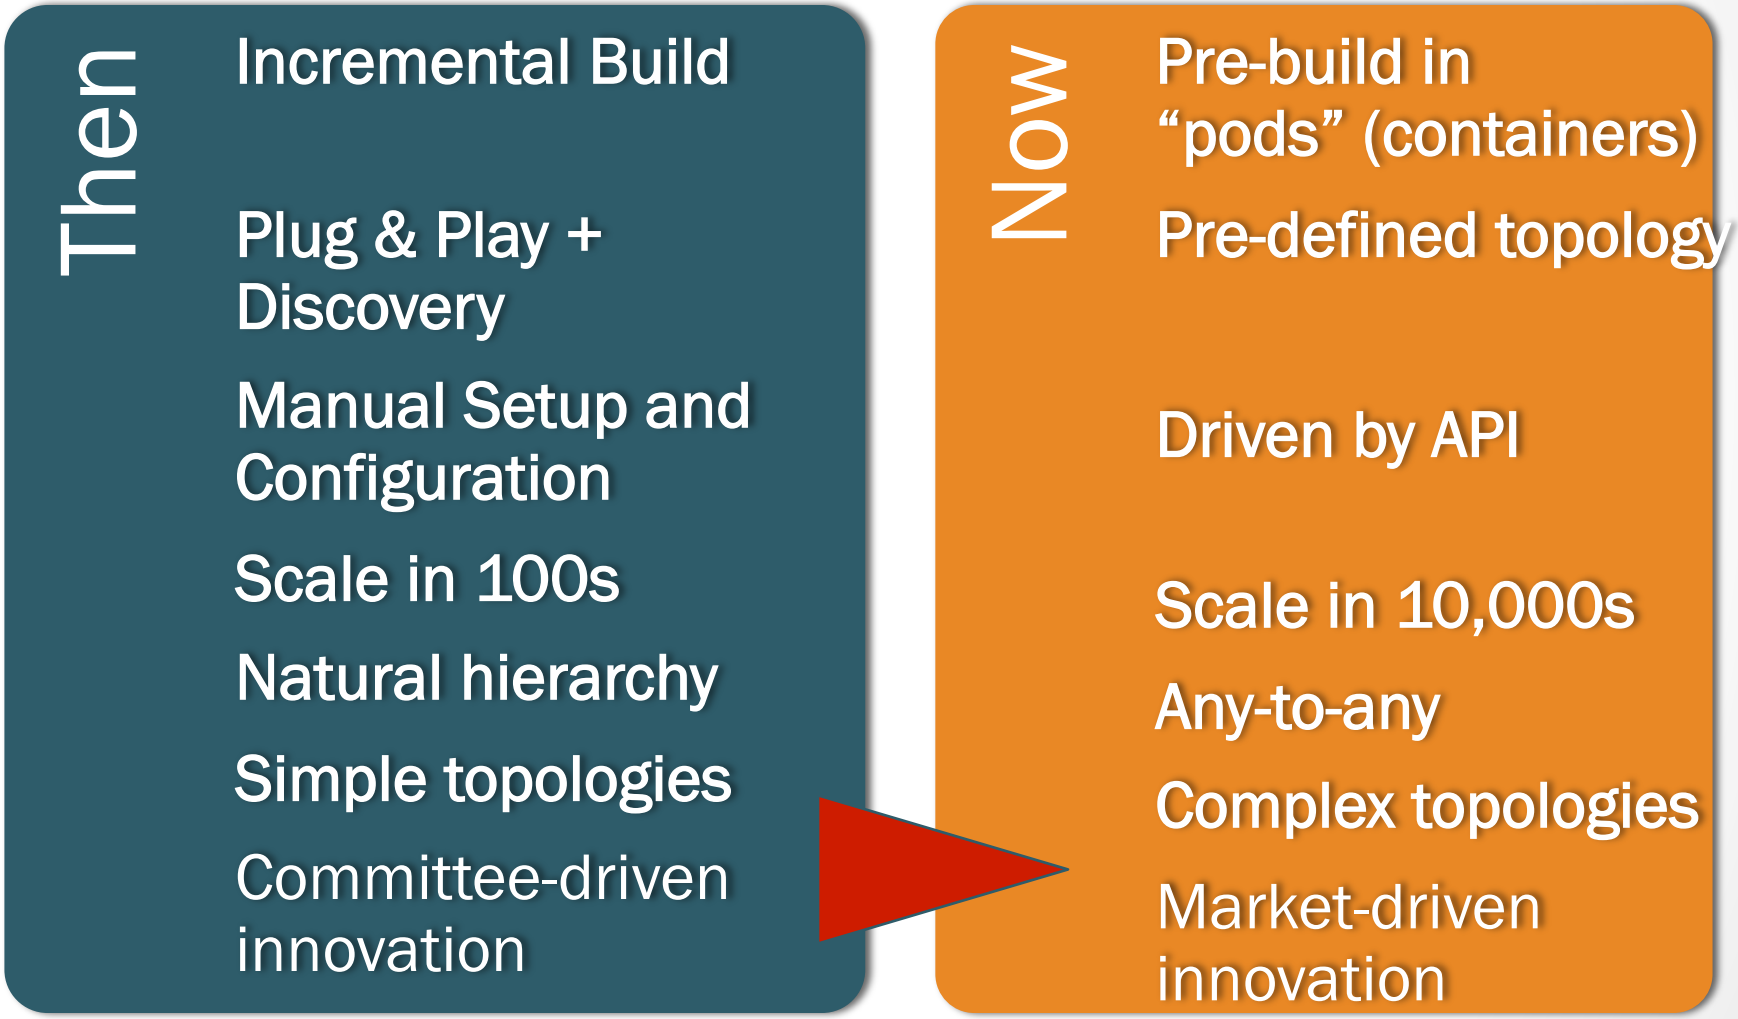

With OpenFlow / SDN

Management and orchestration	Protocol Limitations	Vendor-independent abstraction	Scalability
Centralized API to OF controller replaces the CLI of multiple devices	Allows innovative protocols like L2-over-L3 tunneling at scale	Hardware vendors	Scaling OF controller is solvable problem
Implementation of transactional semantics is trivial	Provides centralized control plane	Software implementations like OVS	Reduces complexity of network devices

Key capabilities enabled by SDN

- “Network infrastructure as code”
  - Configuration version control and automatic provisioning
  - Networks portable across public/private clouds
- Disaster Recovery
  - Quickly provision networks without manual configuration
  - VM migration across WANs (private->public, public->public)
- Customers can use arbitrary IP space
- Let customer’s manage their own VLANs, subnets
- End-to-end QoS

Other Technologies to Watch

- Microsoft VL2
- ConteXtream
- Juniper Qfabric

Contact

Paul Lappas

http://www.linkedin.com/in/paullappas
paullappas@gmail.com

Ivan Batanov

http://www.linkedin.com/in/ivanb
ivan@ivanb.net