

Manage the Data Center Network as We Do the Servers

Albert Greenberg

With Deepak Bansal, Changhoon Kim
Windows Azure, Microsoft

Agenda

- Why Software Defined Networking
- Data Center server management
- Data Center network management
- Next steps

Why Software Defined Networking

- Get innovation to the network faster
 - Reducing the cost+time from H/W realization (ASICs, switches, chassis) to high confidence deployment
 - Goal: Scale deployment soon after unit tests pass
 - Enabling new features as needs arise, without a new switch OS image
 - Goal: APIs/protocols to simple agents running onboard hosts and devices, extending features as needs arise
- Run a more robust network
 - Zero human touch, predictable, highly reliable
 - Goal: automated deployment, triage, RMA
 - S/W at edge to cope with extreme churn in service and VM creation
 - Simpler, higher capacity physical network
 - Goal: reliably move bits really really fast
 - Detect error and route around them quickly

Cloud Data Centers

Key Components

- H/W: large #s of devices; small #s of SKUs
- S/W: systems control of all resources; single extensible end-host OS
- High pressure to innovate in the infrastructure
 - New H/W+S/W, driving functionality up and costs down
 - Phenomenally high rate of service introduction, increasingly impacting the network
- Principles
 - Software control
 - End-host networking stack / vswitch provides versatility and agility in S/W to meet high pressure from new scenarios
 - Scale, simplicity, reliability

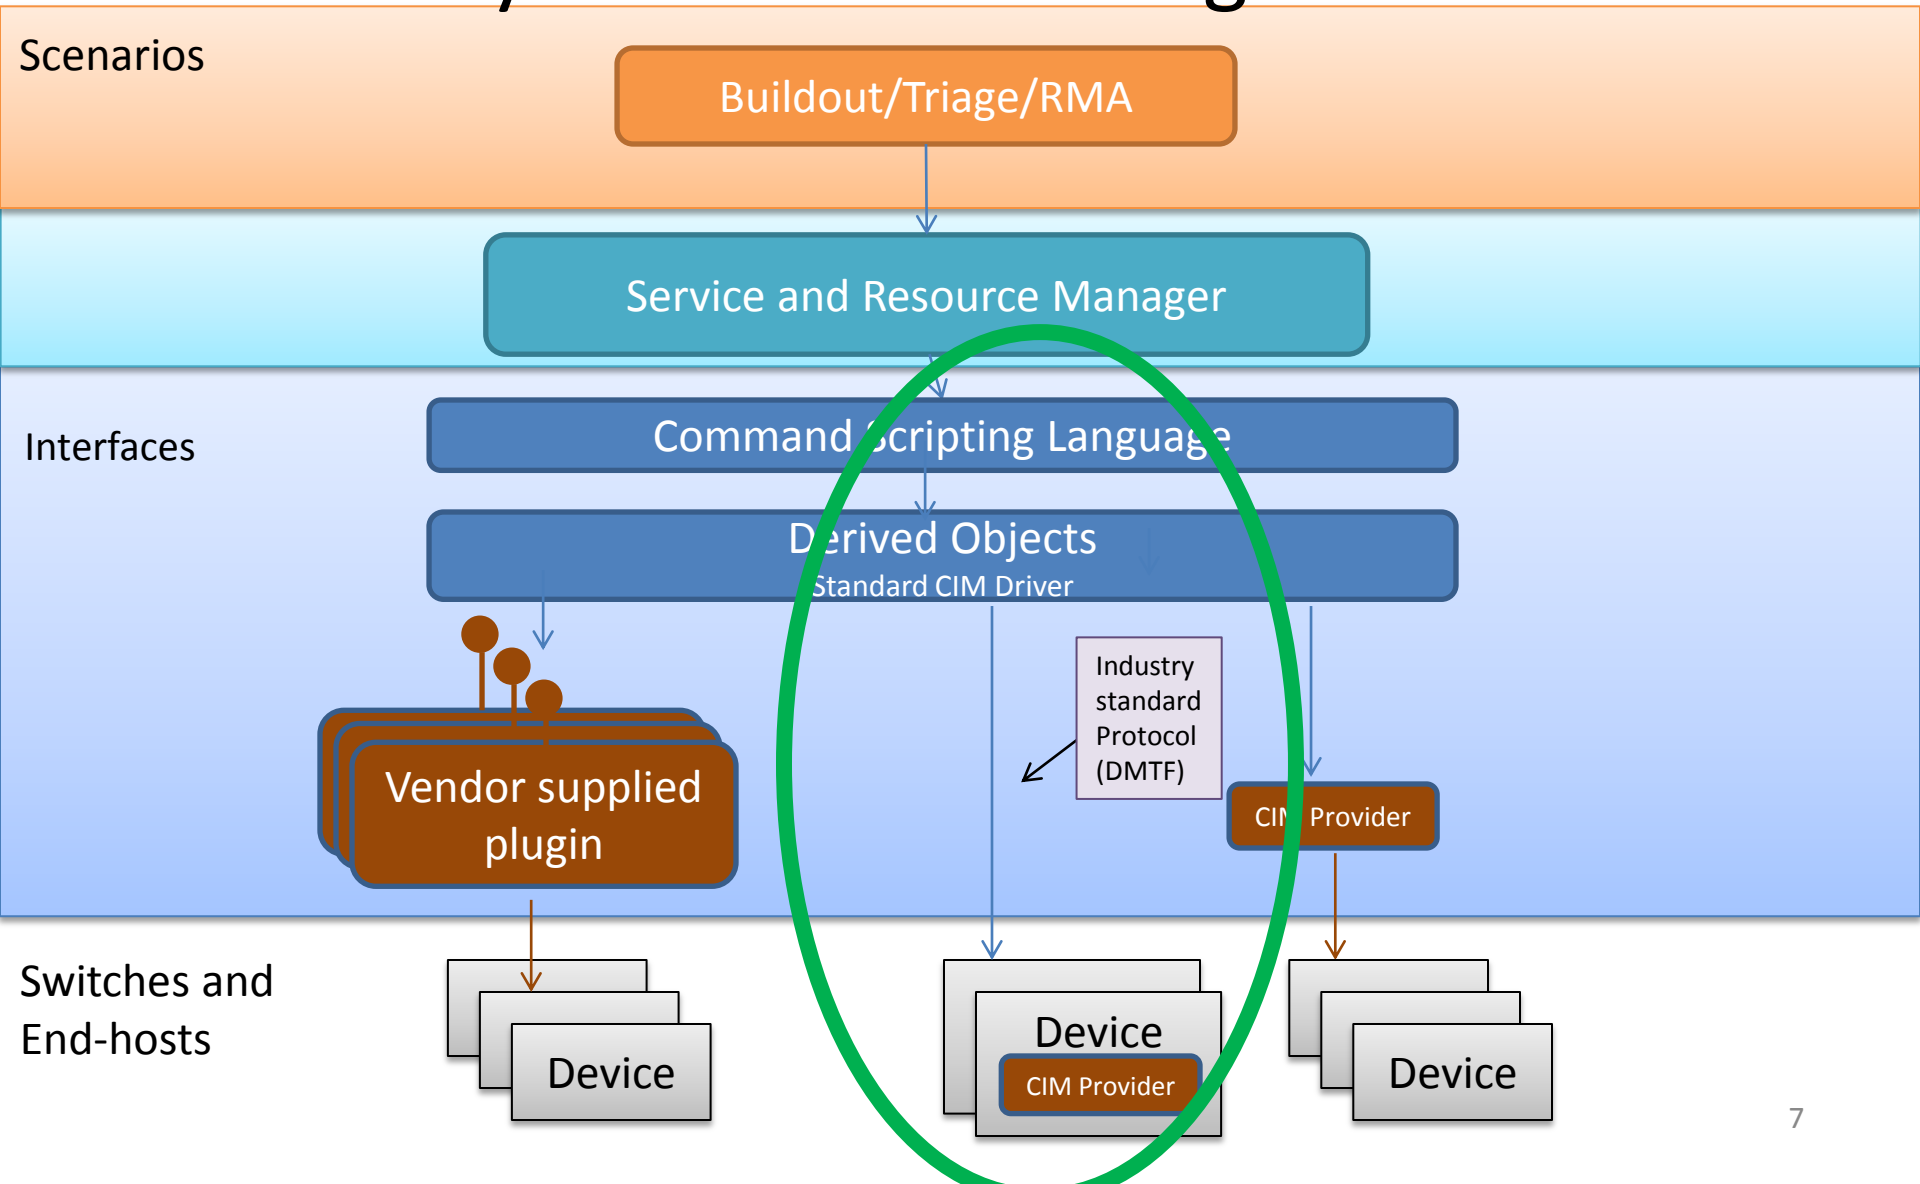
Data Center Server Management Architecture

- Service Manager: Goal State Generator
 - Generates **goal states (desired policies)** for resources, meeting tenant requests
 - **Integrated mgt of all compute/storage/network resources**
 - Logically centralized (modest scale-out)
- Resource manager: Goal State Driver
 - Drives current state for resources to goal state
 - Logically centralized (modest scale-out)
- End-host agents: Simple/local Goal State Drones
 - **The foot soldiers** who execute will of resource manager (high scale-out; monitoring the most intensive task)

Data Center Network Management Architecture (Adheres to Same Principles!)

- Service Manager generates goal states
- Network Resource Manager drives goal states
 - End-host networking stack/vswitch (fine grained)
 - Physical Network devices (coarse grained)
- Host/Device Agents drive to goal states
 - Agents on each end-host (again, high scale-out) executing will of the network manager

Concrete SDN Approach Physical Device Configuration



Next Steps

- Now
 - Apply principles proven for sound management of servers, storage, power to network
 - Select protocols and mechanisms already prevalent in the server world (Windows, Linux) and widely used in data centers: such as WS-MAN
 - Apply and reuse!
 - Continue to innovate in the end-hosts / hypervisor
 - Performance predictability, low latency, ...
- Next few years
 - Engage industry on APIs for direct SDN control of physical devices, for realization of new features
 - Opportunities for better monitoring or direct control