



(pohled na) SDN

Radek Boch

Systems Engineer, CCIE #7095, rboch@cisco.com

5.10.2015

What is *SDN* ?

SDN

Many things to Many people

“An open solution for customized flow forwarding control in the Data-Center”

“A way to reduce the CAPEX of my network and leverage commodity switches”

“An open solution for VM mobility... the Data-Center”

“A means to scale my fixed/mobile gateways and optimize their placement”

“A solution to build virtual topologies with optimum multicast forwarding behavior”

“A way to distribute policy/intent, e.g. for DDoS prevention, in the network”

“A way to optimize link utilization in my network, through new multi-path algorithms”

“A platform for developing new control planes”

“A way to avoid lock-in to a single networking vendor”

“A way to define virtual networks with specific topologies for my multi-tenant Data-Center”

“A way to configure my entire network as a whole rather than individual devices”

“With SDN I can develop solutions to my problems far faster – “at software speeds”. I don’t have to work with my network vendor or go through length standardization”

“A means to do traffic engineering without MPLS”

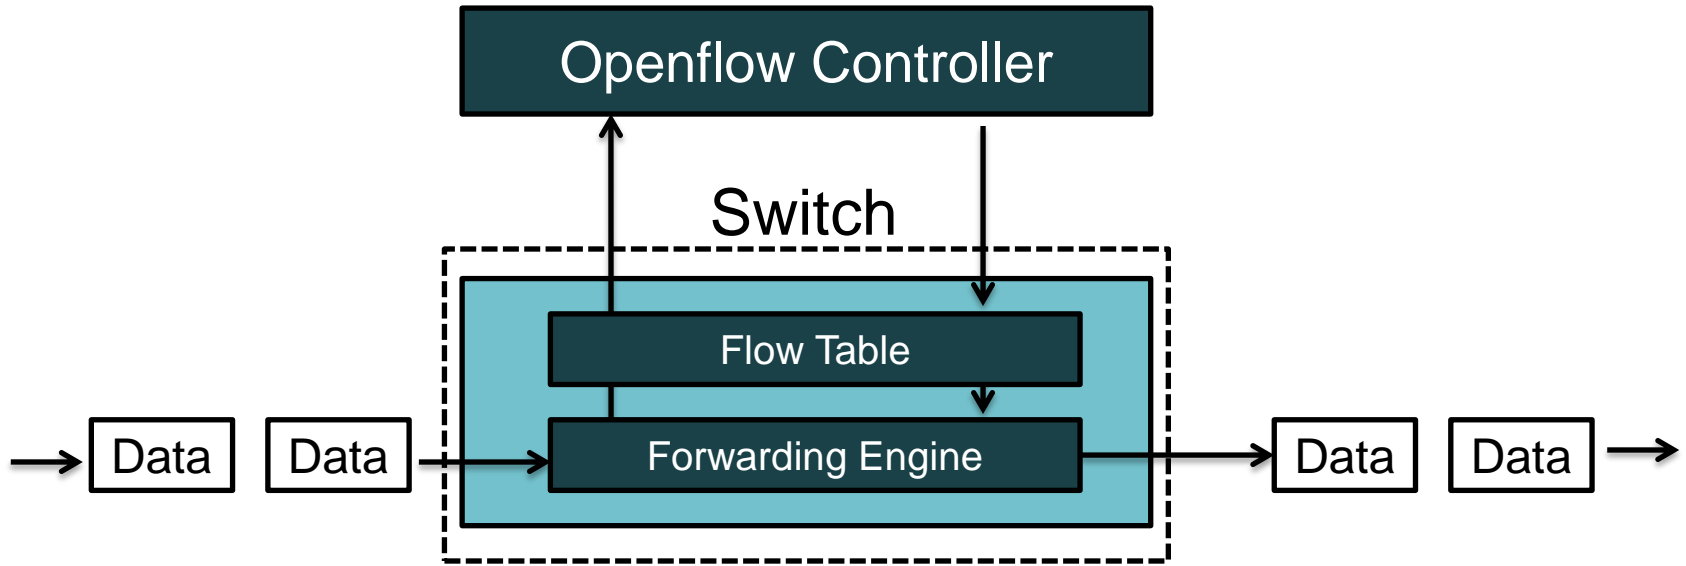
“A solution to build a very large scale layer-2 network”

“A way to build my own security/encryption solution, avoiding RSA”

“A way to scale my firewalls and loadbalancers”

“A solution to get a global view of the network – topology and state”

Software Defined Networking




Openflow Forwarding Model

(Does this look familiar?)

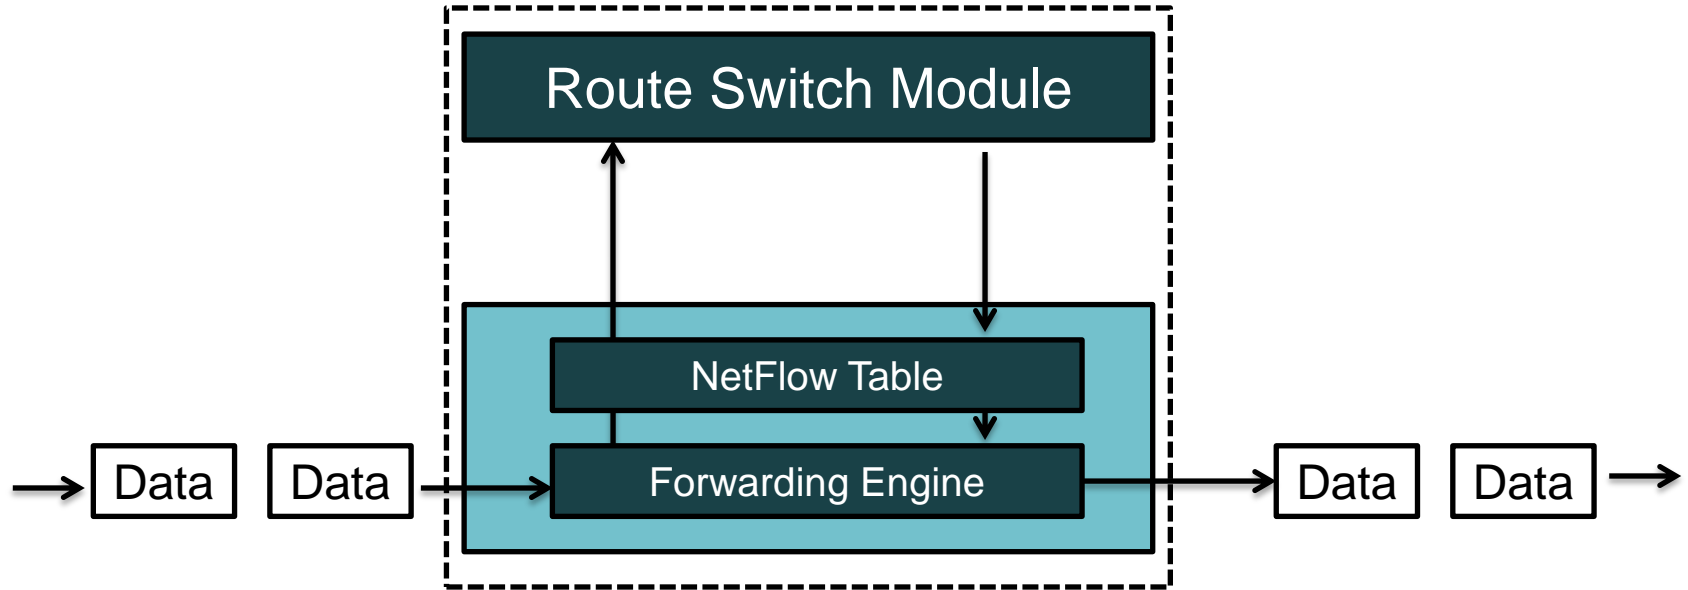
Is this a ***revolutionary*** idea?

Let us go back to the late 1990's

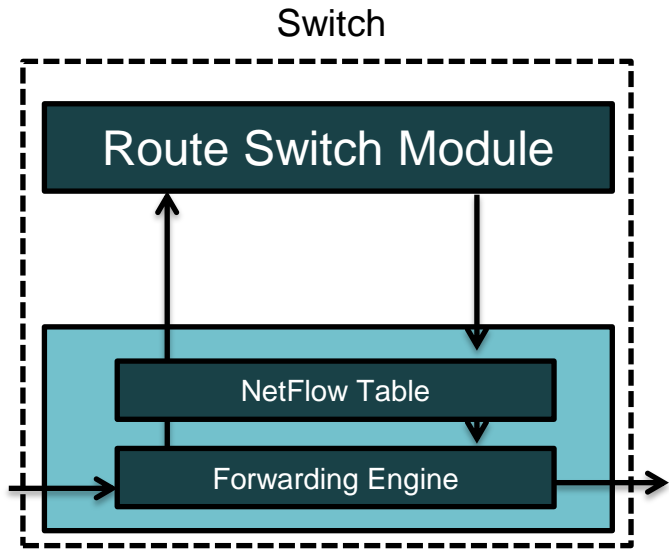


Remember the term
“Router on a stick”?

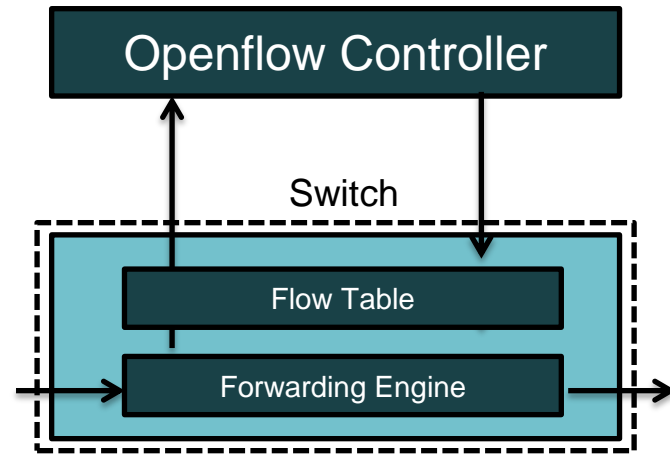
Switch



MLS – Multi Layer Switching



1997
Innovation
(MLS)




Today's
Innovation

A small problem with MLS


IT DID NOT SCALE

Cisco's *Forwarding* Evolution




Cisco moved to CEF to overcome
MLS *Scalability* issues

I will figure out where to send data and tell you



Controller




Switch



Imperative Model
(Openflow Forwarding Model)

You figure out where to send data and I will tell you how to handle it



Controller



Switch



Declarative Model
(ACI Group Policy Model)

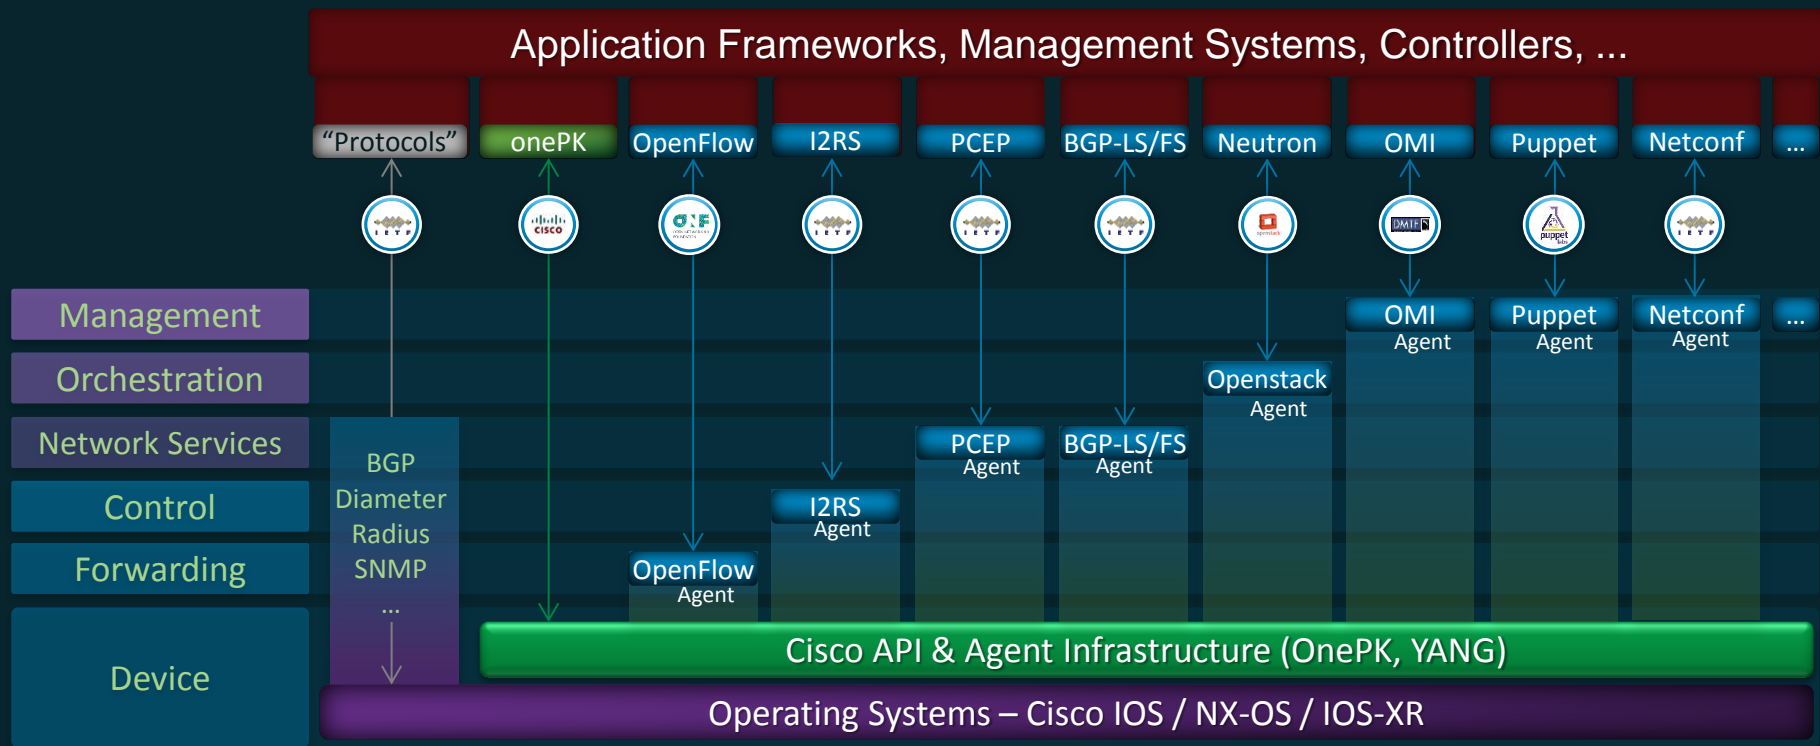


SDN

It's **Not About** Southbound Protocols

It's **About** Controlling Network Behavior
through Software in a Programmatic way

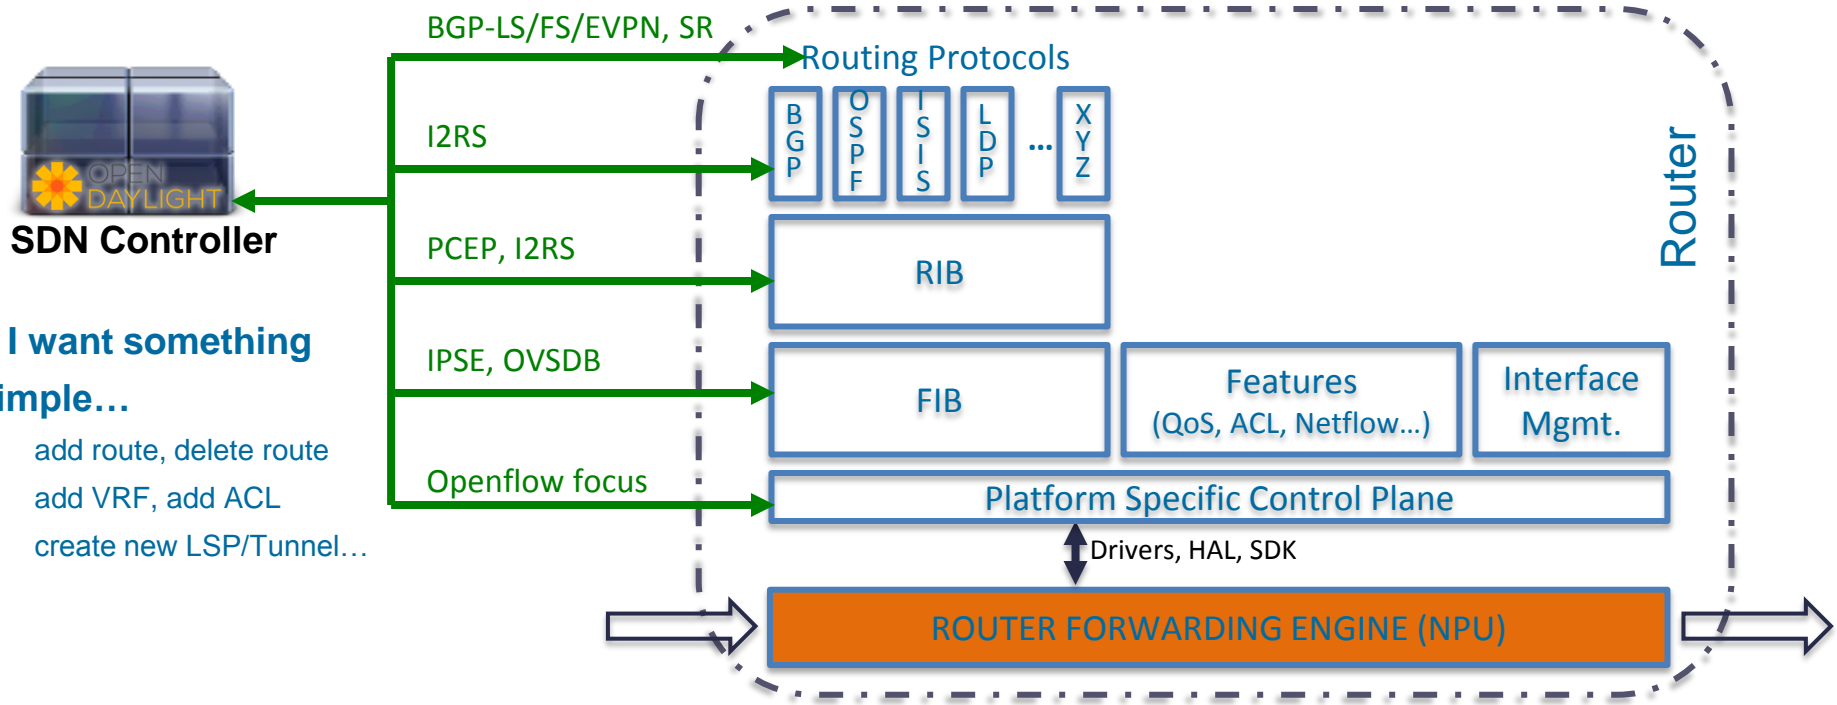
SDN Internet protocols



L3 device: Specific API's & Protocols

Router has different tables to control

- Higher-level abstraction

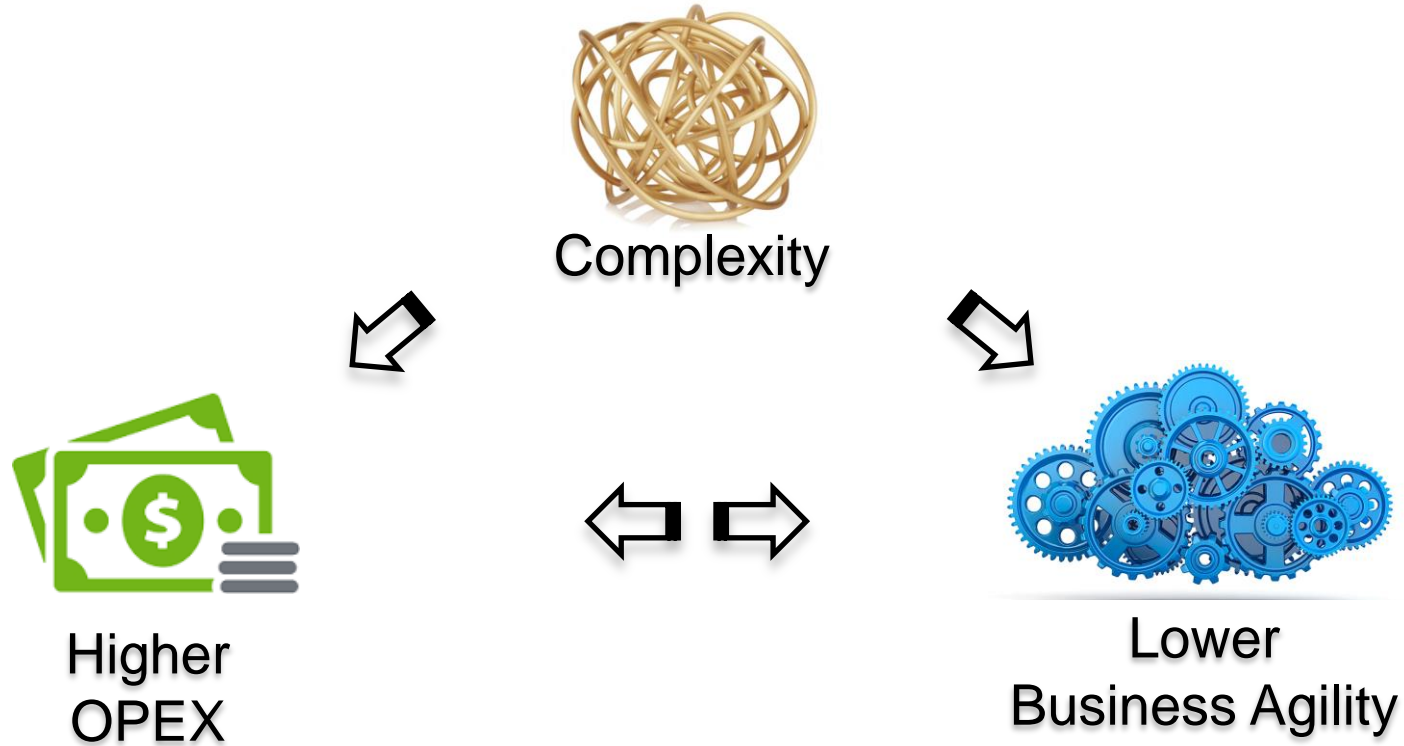


If I want something simple...

- add route, delete route
- add VRF, add ACL
- create new LSP/Tunnel...

Problem Statement

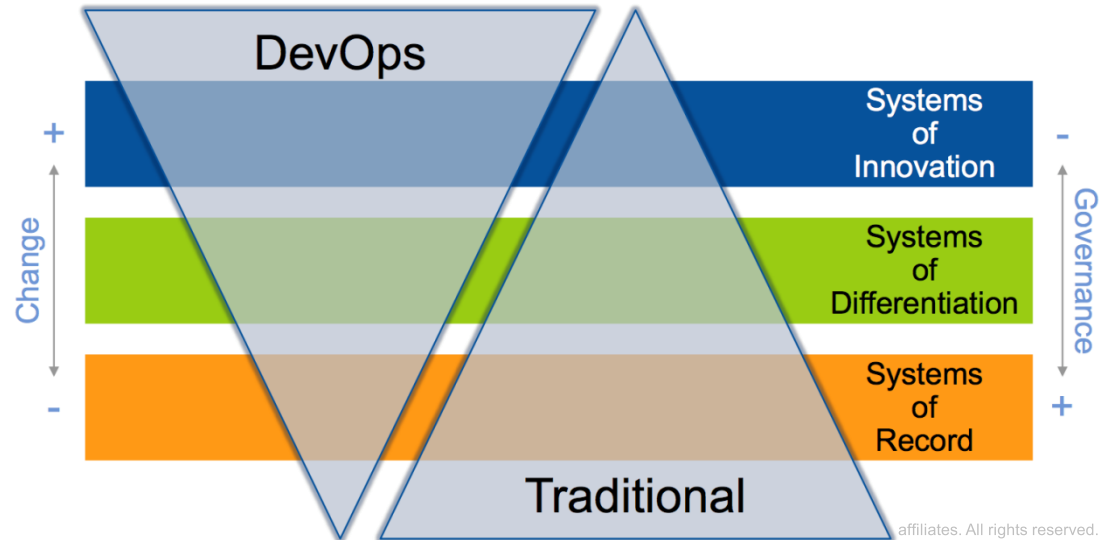
What Does It Really Mean?



The Pace of IT – Bimodal IT

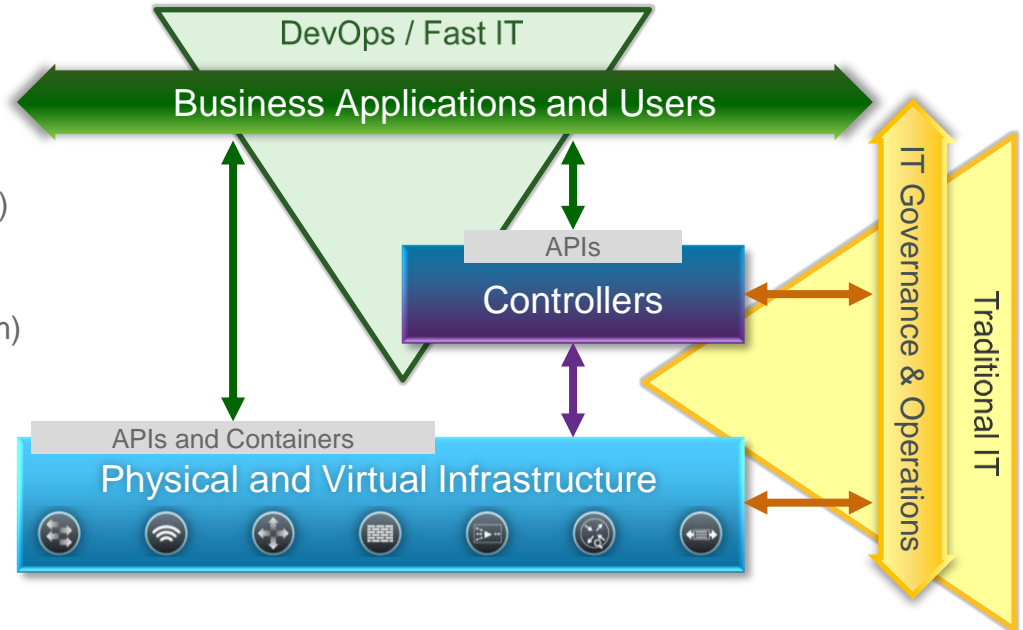
Problem: CIOs are challenged to keep running existing IT more efficiently and safely, while enabling business innovation and differentiation at a quickening pace.

Solution: Bimodal IT, enabling developers and enabling governance




Ingredients of Bimodal Architectures

- **Application Centricity**
- **Programmability** of
 - Infrastructure (RESTCONF + many)
 - Controllers (REST)
 - Services (apiconsole.cisco.com)
- **Virtualization** of
 - vAF: Application Functions
 - vMF: Management Functions
 - vNF: Network Functions

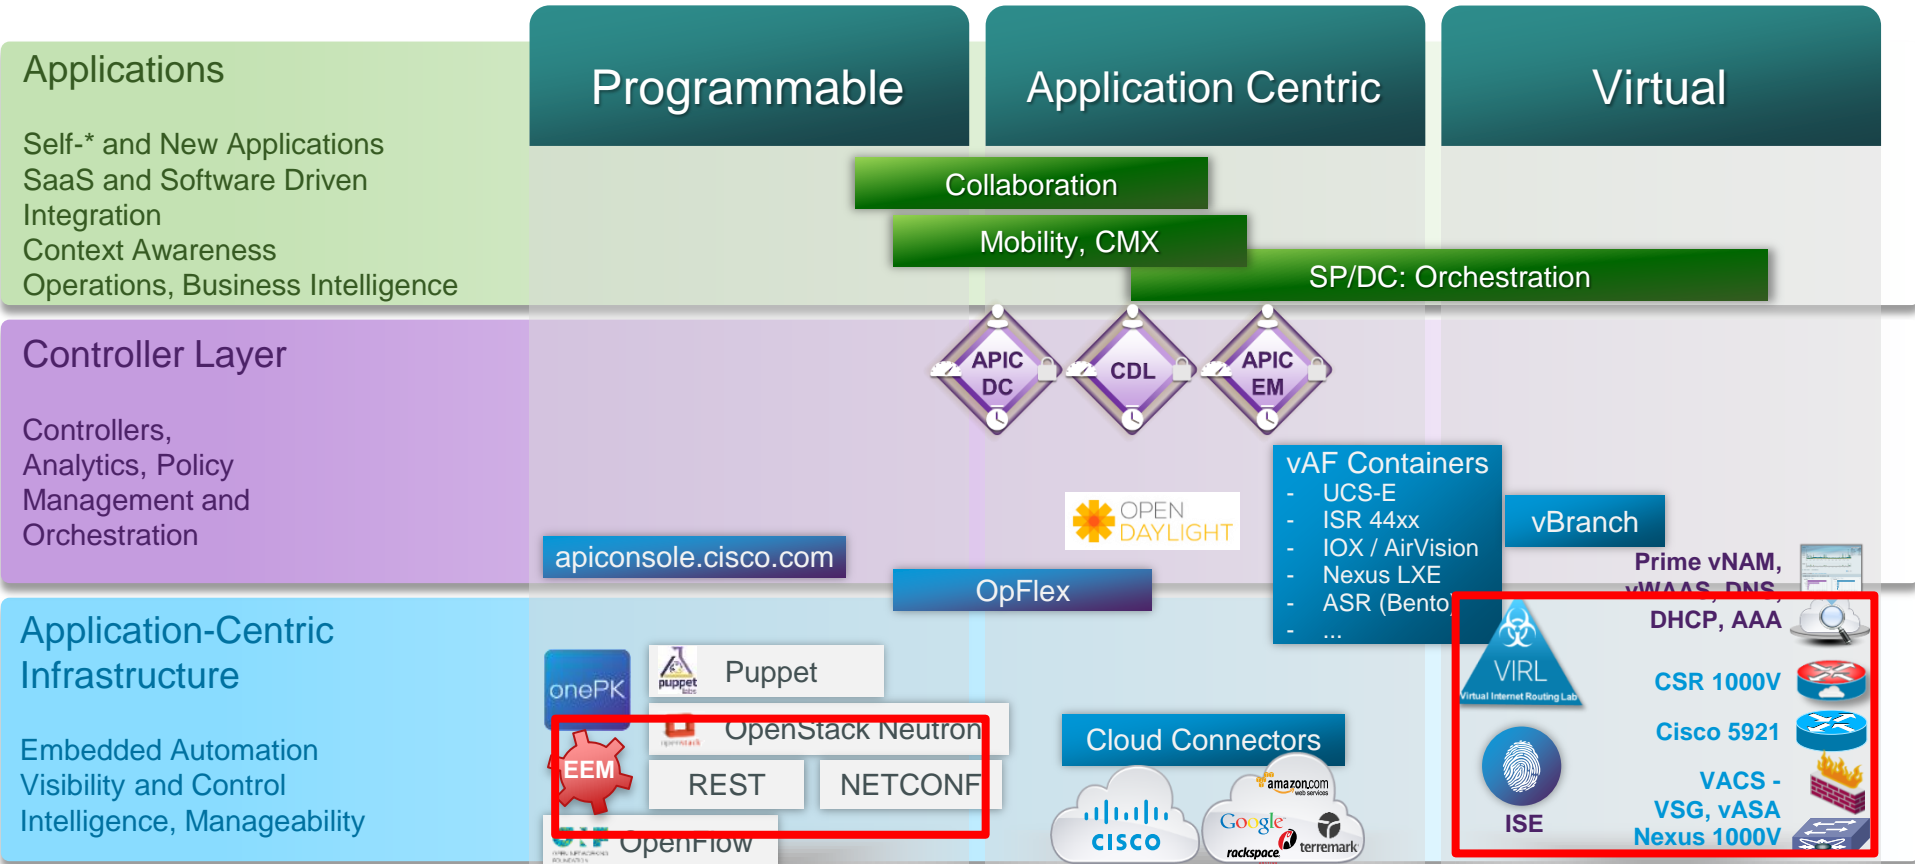


Bimodal IT Architectures to support Fast IT Business Needs

Cisco Enterprise ACI – 3x3 Portfolio (Subset)

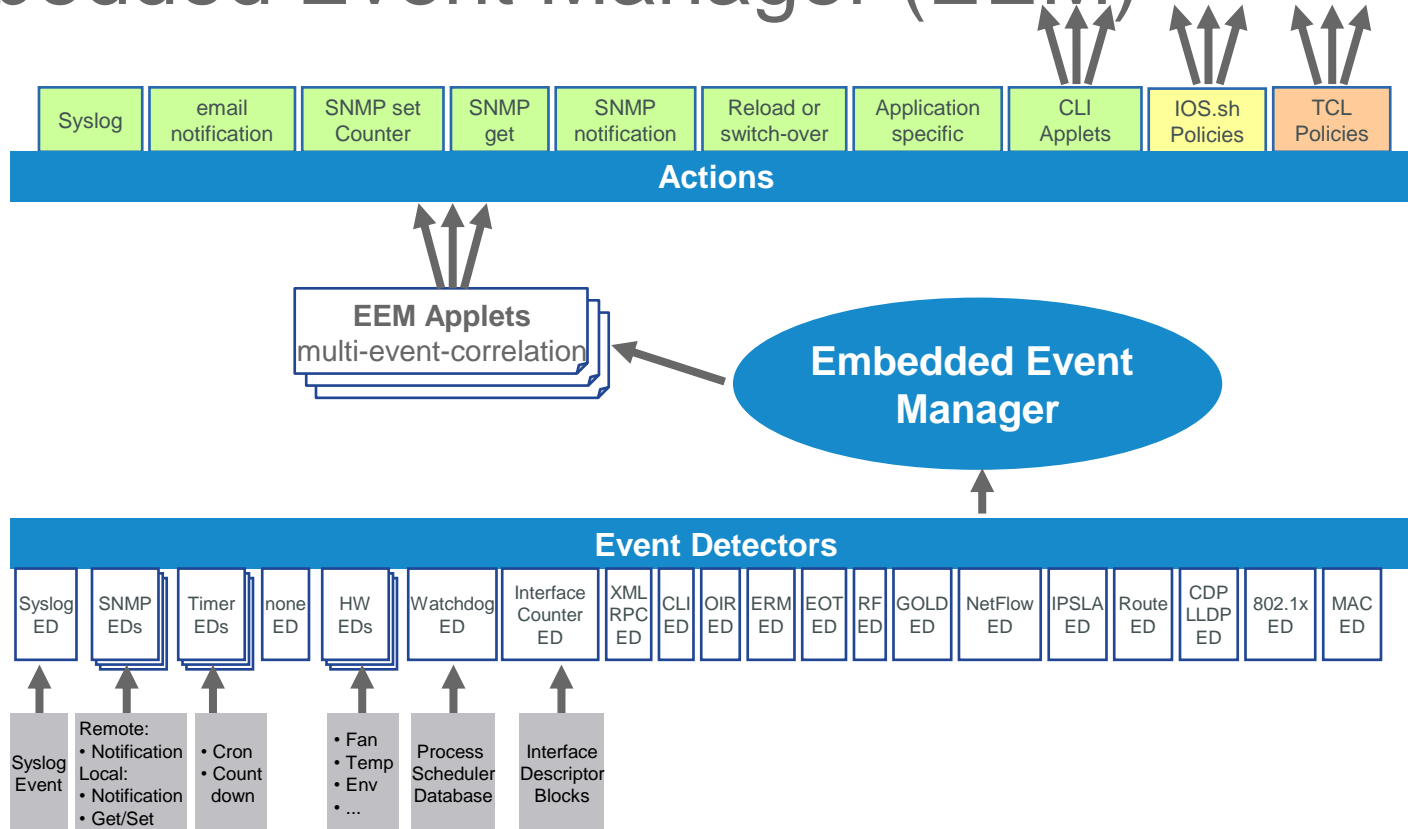


Cisco Enterprise ACI – 3x3 Portfolio (Subset)

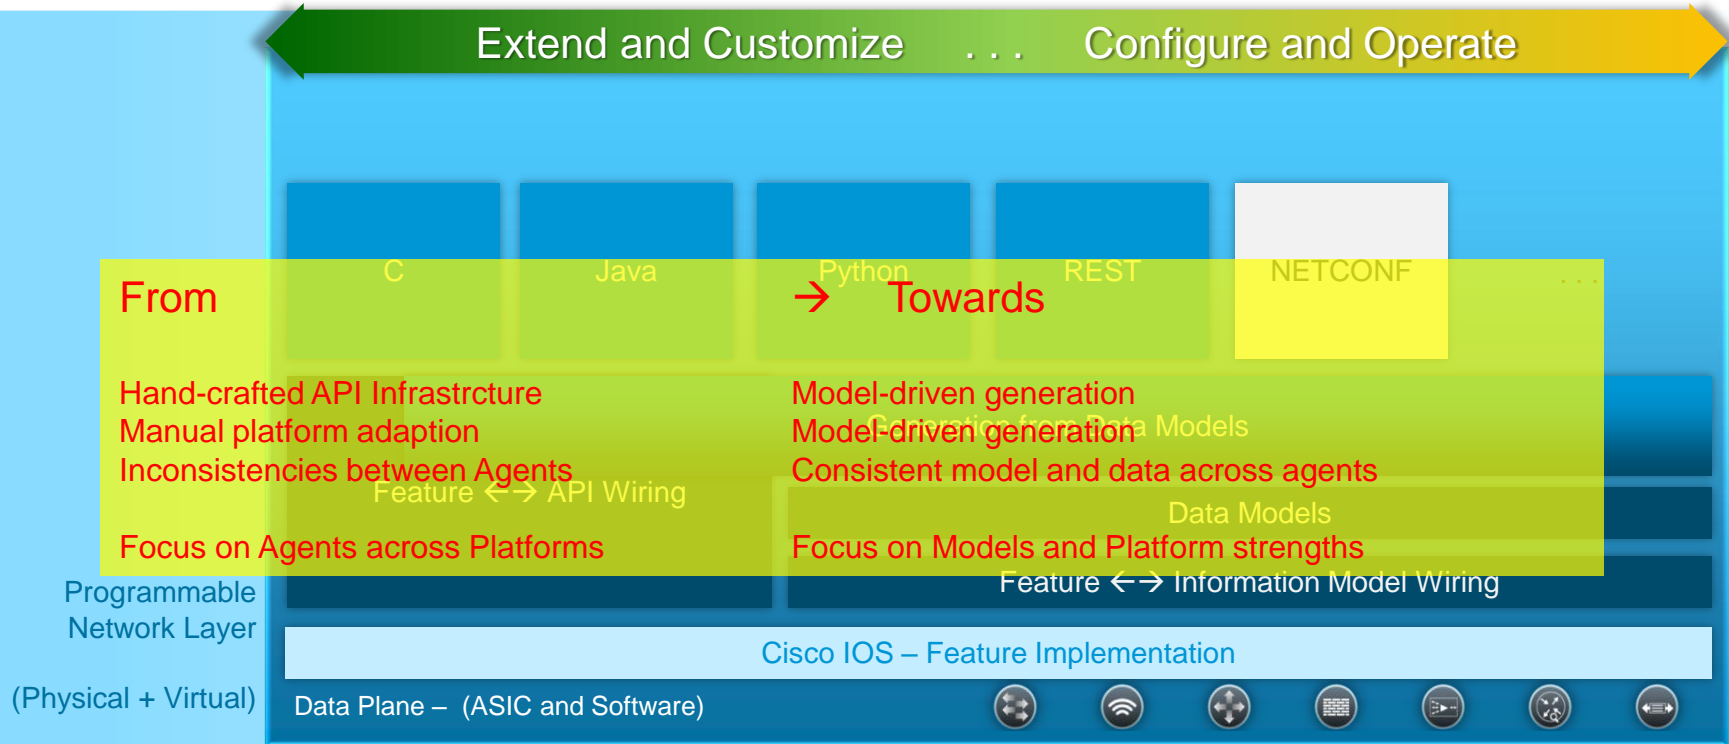


Focus: Programmable Network Layer

Embedded Event Manager (EEM)



Programmable Network Layer – Evolution



Programmable Network Layer – Hosting Models

Choice of 3
Hosting Models

“Process”

On the Node


- Shared memory/compute
- Very low latency and delay
- Available on select platforms



“Blade”

On A Hardware Blade

- Dedicated memory/compute
- Low latency and delay
- Requires modular hardware blade



“End-Node”

On An External Server

- Plentiful memory/compute
- Higher latency and delay
- Supported by all platforms



Programmable
Network Layer

(Physical + Virtual)

Device-Level Shell or API (such as guestshell, onePK, RESTCONF, etc)

Cisco IOS (Enterprise, Data Center, Service Provider)

Data Plane – (ASIC and Software)



vAF

Focus: Virtual


Virtual Containers – ISR 4400 Series

Service Containers

- Dedicated virtualized compute resources
- CPU, disk, memory for each service
- Easily repurpose resources
- Industry-standard hypervisor

Benefits

- Better performing network services
- Ease of deployment with zero footprint; no truck roll
- Greater security through fault isolation
- High reliability
- Flexibility to upgrade network services independent of router IOS® Software



vNF

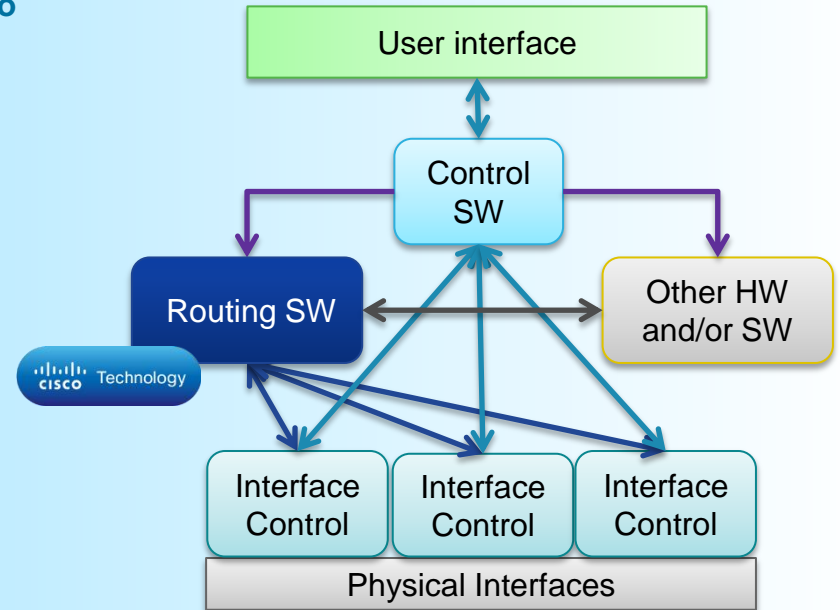
Focus: Virtual

vNF – Embedded Services Router – ESR 5921

The Cisco 5921 Embedded Services Router (ESR) is a Cisco IOS® software router application designed to operate on small, low power Linux-based platforms to extend the use of Cisco IOS into extremely mobile and portable communications systems.

- based on IOS 15.2(4)GC , synched to 15.2(4) M
- Includes special mobility features (Radio Aware Routing, ...)
- Up to 20 virtual Eth ports
- x86 compatible, 32 bit Linux application
- Can run on any x86 compatible hardware with sufficient resources:
 - x86 - e.g., Intel Atom and Intel Core i3/i5/i7
 - 512 MB RAM minimum
 - 300 MB Disk minimum
 - glibc compiled Linux
- Embedded by SI into final Product, using "Cisco Technology Inside" software
- Sensors, Portable Communications, Vehicular Communications, ...


Typical Use Case



Example: Secure VPN Gateway

Problem: How to securely connect to a virtual private cloud or virtual Data Centre where we can't deploy Hardware – across the public Internet?

Solution: Deploy VPN Gateway on Cloud Services Router 1000v



Challenges

- Inconsistent Security
- High Network Latency
- Limited Scalability

Solutions

- IPSec VPN, DMVPN, EZVPN, FlexVPN
- Routing and Addressing
- Firewall, ACLs, AAA

Benefits

- Direct, Secure Access
- Scalable, Reliable VPN
- Operational Simplicity

vNF – Network Simulations

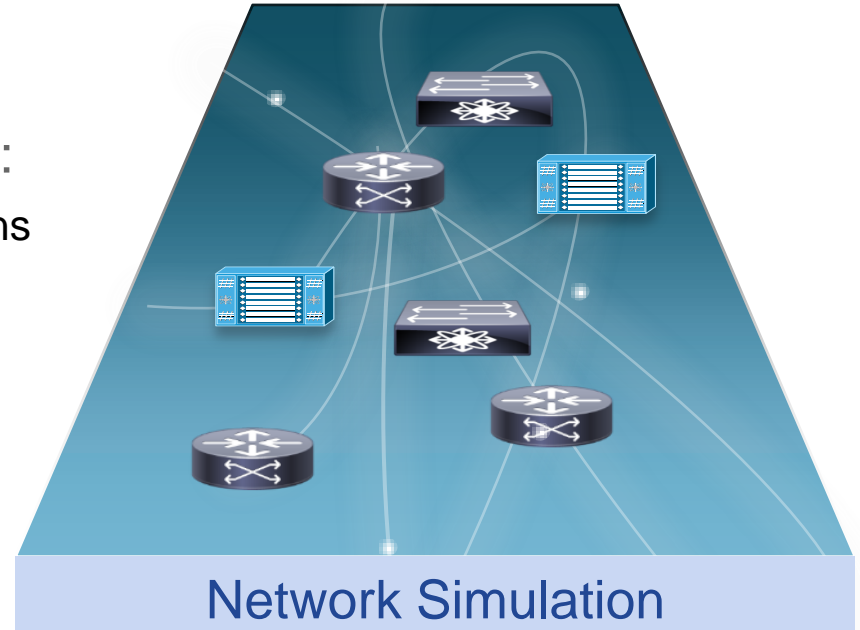
The Challenge

Developers have a compelling need to:

- Create new network applications and solutions
- Learn and test new features and facilities
- Innovate to solve business problems

To do this they need a Lab that is:

- Easy to build
- Easy to access
- Easy to (re-)configure
- Portable
- Easy to scale
- Inexpensive



Such a Lab did not exist ... until VIRL.cisco.com

vMMF

Focus: Virtual

vMF – Driving Visibility from Network Control

Problem: How to dynamically provide application visibility into remote sites or per virtualized tenant?


Solution: Deploy vNAM into the virtual workload POD

Service Assurance Actions (Examples)

- Apply Service Policies (Police, Mark, Shape, Queue) for reprioritization
- Implement custom routing optimized for specific application topology
- Set ACLs to establish the access rules

PROGRAMMABLE
Traffic Steering
Path Setup
Traffic Engineering

Application Services




Application Services

Network Application

Cisco Daylight (XNC)

REST/XML API



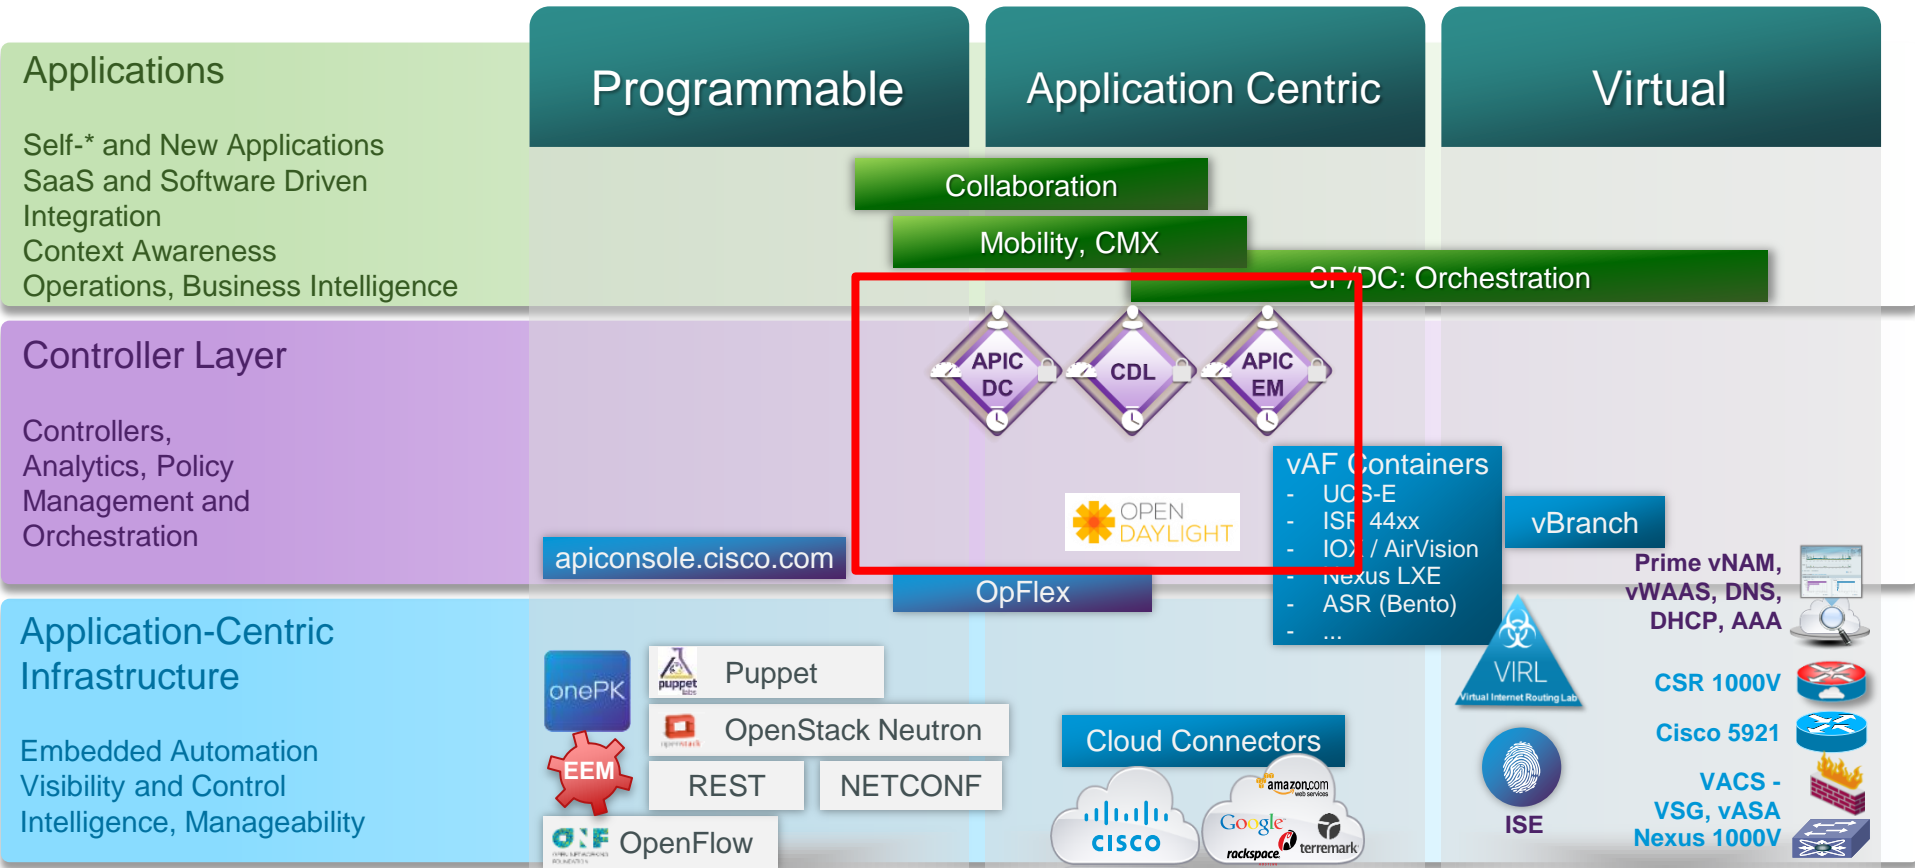
Hosted Workload for Tenant



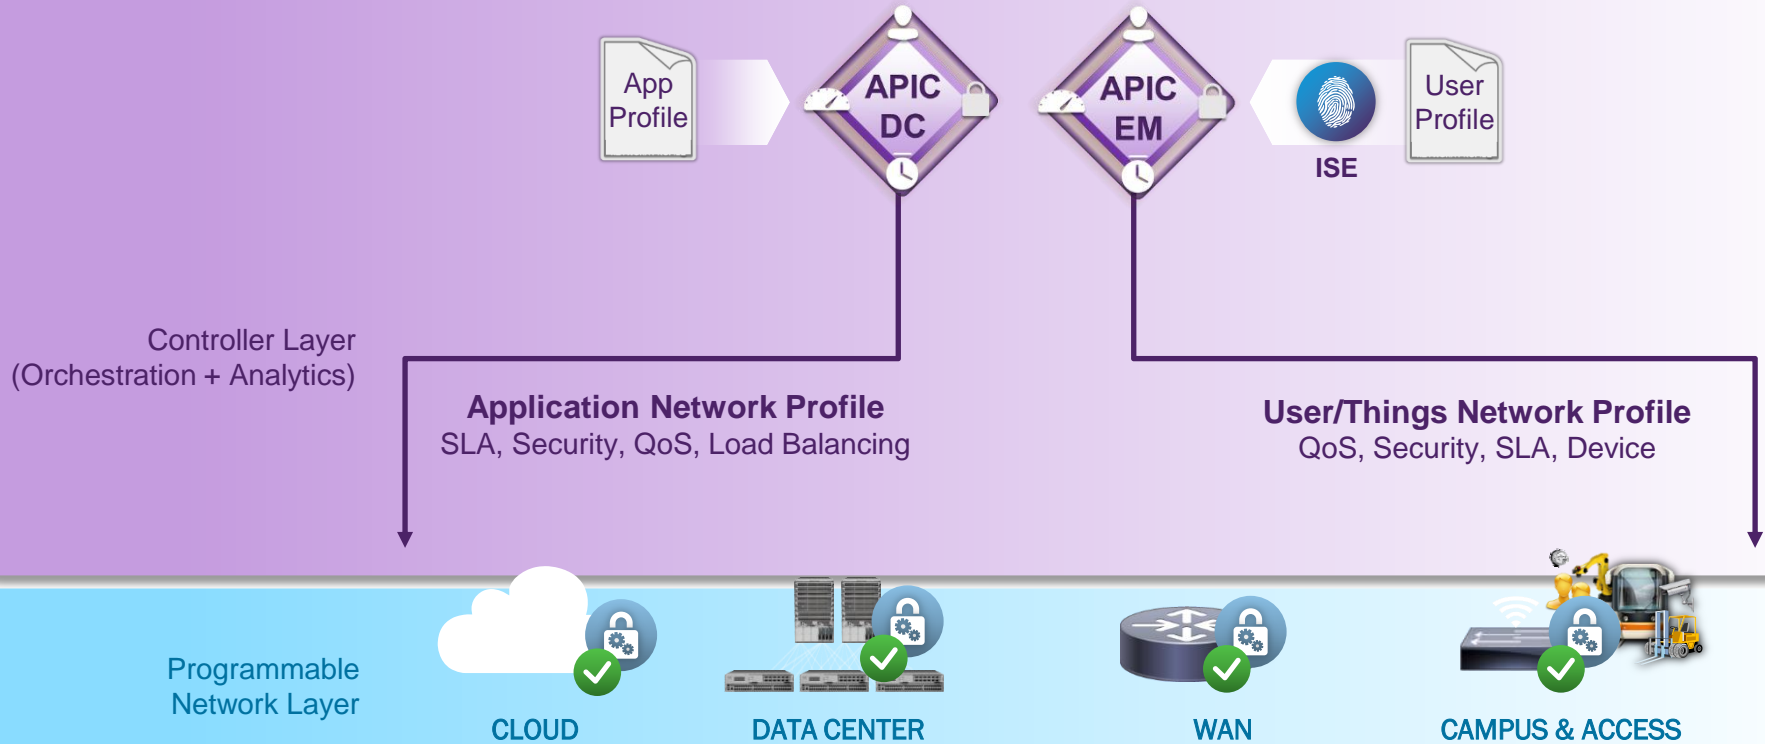
Extends Application Intelligence to the Virtual Infrastructure

Focus: Controller Layer

Cisco Enterprise ACI – 3x3 Portfolio (Subset)



Common Policy Across Domains




Controller Architecture

High Level



North Bound APIs

RESTful API
GET PUT POST DELETE




South Bound APIs

CLI, NetConf, REStConf, Openflow.....



Conventional

Admin
Driven



ACI Policy Model

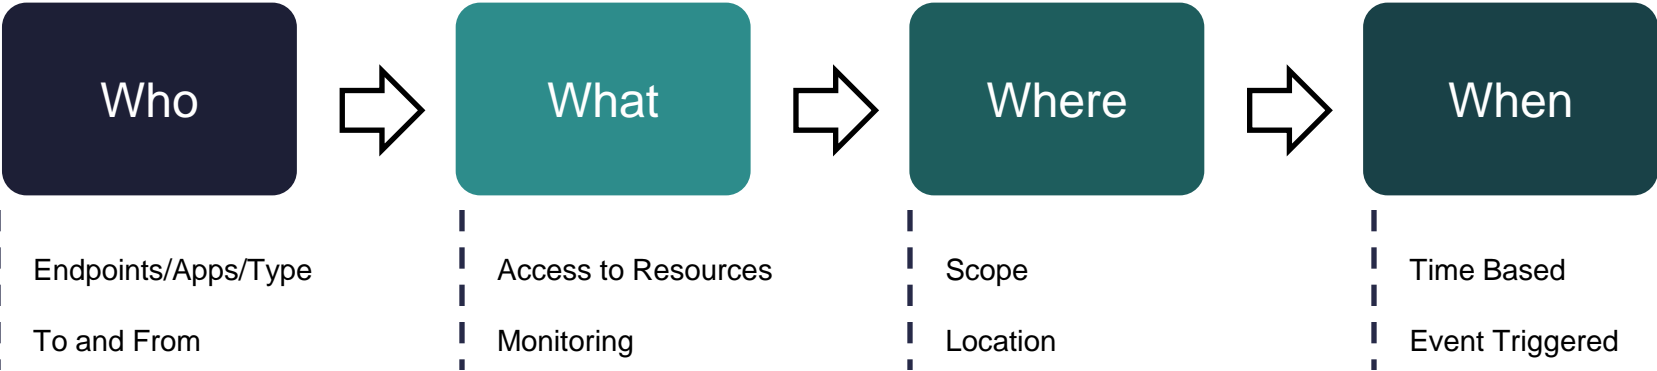
Admin Driven



APIC EM Driven

Lower OPEX and Better LOB Alignment Agility

What is a Policy...



Policy Examples....



Engineering Group (Who: From)



Engineering Applications (Who: To)



Laptop (Who: Device Type)



Permit (What: Action)



Properties: priority level - high, trust level – high (What: Action Properties)



Tom (Who: From)



Netflix (Who: To)



Permit (What: Action)



Properties: priority level – Low, trust level – low (What: Action Properties)



11AM-1PM (When: Time)



Under The Covers

Completely Hidden!

```
{"policyName":"tomweallow","policyOwner":"Admin","policyPriority":4095,"networkUser":{"userIdentifiers":["tom"]},"resource":{"ports":["80,80,tcp"]},"actions":["PERMIT"]}
```

```
CompositeNetworkPolicy [networkPolicy=NetworkPolicy [id=70be-adaf-4f41-bfb7-d1d9ee01e0f8, creatorUserId=Admin, policyName=bradweallow, policyPriority=4095, businessPolicyId=10d7e374-c1e0-4190-b35f-3c33aac312, flowId=7ba2034a-3cb0-4877-ae14-4a6c-84e5dc8d3cd, actionId=70fb3b4c-ccf8-4561-b490-684e5dc8d3cd, flow=Flow [flowId=7ba2034a-3cb0-4877-ae14-4a6c-84e5dc8d3cd, srcIp=10.10.30.2, srcIpMask=255.255.255.255, protocol=tcp, srcTptPortLower=0, srcTptPortUpper=0, dstIp=10.10.30.0, dstIpMask=255.255.255.0, dstTptPortLower=80, dstTptPortUpper=80], flowAction=FlowAction [actionId=70fb3b4c-ccf8-4561-b490-684e5dc8d3cd, action=permit, actionPropDscp=-1, ]]
```

```
CLI = config t, ip access-group User-Acl--8653840507576742282, 10 permit tcp host 10.10.30.2 any eq 80, interface GigabitEthernet1/0/4, ip access-group User-Acl--8653840507576742282 in, end 20:22:28.992 EST DEBUG c.c.c.qos.acl.AclPolicy - Acl Policy Created Successfully on the Device : d29d175f-aacc-4c9c-a290-2392fc80a0e3
```

Let's Look At Some Apps

Network Information Base

One Source of Truth



APIC - Enterprise Module



API

Sign Out

▲ 0

⊙ Filters

Layout: **Hardware** ▾

	Device Name	MAC Address	IP Address	IOS/Firmware	Platform	Serial Number	Config	Device Role
<input type="checkbox"/>	SDN-BRANCH-2960S	5C:50:15:BF:6D:C0	40.0.5.4	15.2(1)E2	WS-C2960S-24TS-L	FOC1612W32Z	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-2960S-STACK	00:26:52:7D:2C:C0	40.0.7.5	15.2(1)E2	WS-C2960S-48FPD-L	FOC1412Z2KG	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-3560CG	1C:AA:07:63:8B:40	40.0.5.7	12.2(55)EX	WS-C3560CG-8PC-S	FOC1516W4XR	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-3560X	60:73:5C:EF:13:40	40.0.5.9	15.2(1)E2	WS-C3560X-48U	FDO1634Z049	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-3650	F8:72:EA:0D:67:47	40.0.7.3	03.03.00SE	WS-C3650-24PD	FDO1733Q02X	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-3750X	2C:54:2D:93:1E:40	40.0.7.4	15.2(1)E1	WS-C3750X-48P	FDO1612P1XA	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-AP1252-1	00:26:CB:7E:D2:DC	40.0.5.39	15.2(20130113:221158)\$	AIR-LAP1252AG-A-K9	FTX133590NE	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-ASR1002	78:DA:6E:13:5E:00	40.0.3.6	15.2(4)S3	ASR1002	FOX1737GJVL	View	Border Router ▾
<input type="checkbox"/>	SDN-BRANCH-C2960S-L	70:10:5C:5A:47:C0	40.0.5.6	15.2(1)E2	WS-C2960S-48TS-S	FOC1709Z006	View	Access ▾
<input type="checkbox"/>	SDN-BRANCH-C4K	A8:0C:0D:98:CA:7F	40.0.5.2	03.03.00.XO	WS-C4510R+E	FXS1749Q1LC	View	Distribution ▾

10 ▾

25 Devices

[First](#) [Previous](#) 1 ▾ [Next](#) [Last](#)

Path Trace

Enhanced Application Flow Visibility

APIC - Enterprise Module

Path Trace Hosts: 212.1.10.20 → 207.1.10.20

Trace Results

View Small Show Reverse Scroll Lock Show Duplicate Devices

The diagram illustrates the path of traffic from Host A (212.1.10.20) to Host B (207.1.10.20). The path is as follows:

- Host A (212.1.10.20) connects to **CAMPUS-Access1** (212.1.10.1) via **Wired**.
- CAMPUS-Access1** connects to **CAMPUS-Dist1** (55.1.1.100) via **InterVlan Routing**.
- CAMPUS-Dist1** connects to **CAMPUS-Core2** (211.2.2.1) via **ECMP**.
- CAMPUS-Core2** connects to **CAMPUS-Router1** (210.1.2.1) via **ECMP**.
- CAMPUS-Router1** connects to a cloud icon via **OSPF**.
- The cloud icon connects to **Branch-Router2** (207.3.1.2) via **NetFlow**.
- Branch-Router2** connects to **Branch-Access1** (207.1.10.1) via **Connected**.
- Branch-Access1** connects to Host B (207.1.10.20) via **Switched**.

Reversed path (Host B to Host A):

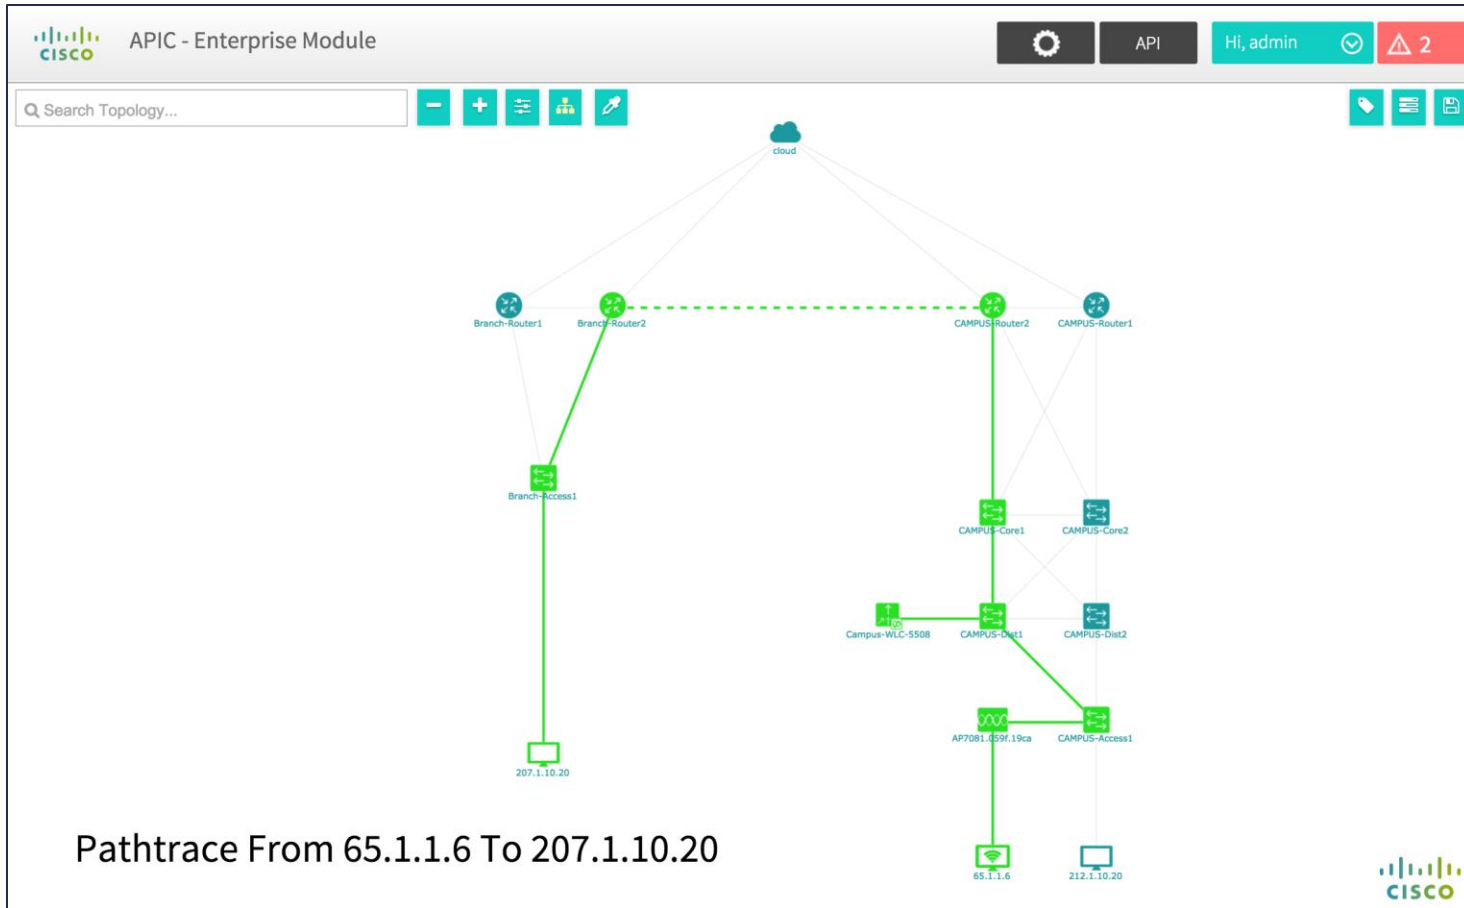
- Host B (207.1.10.20) connects to **Branch-Access1** (207.1.10.1) via **Switched**.
- Branch-Access1** connects to **Branch-Router2** (207.3.1.2) via **Switched**.
- Branch-Router2** connects to a cloud icon via **OSPF**.
- The cloud icon connects to **CAMPUS-Router2** (210.2.1.1) via **OSPF**.
- CAMPUS-Router2** connects to **CAMPUS-Core1** (211.1.1.1) via **ECMP**.
- CAMPUS-Core1** connects to **CAMPUS-Dist1** (55.1.1.100) via **OSPF**.
- CAMPUS-Dist1** connects to **CAMPUS-Access1** (212.1.10.1) via **InterVlan Routing**.
- CAMPUS-Access1** connects to Host A (212.1.10.20) via **Switched**.

Host Details:

- Host A (212.1.10.20):** IP: 212.1.10.20, Type: WIRED, Link Source: Wired
- Host B (207.1.10.20):** IP: 212.1.10.20, Type: WIRED, Link Source: Wired
- CAMPUS-Access1 (Left):** IP: 212.1.10.1, Type: SWITCH, Link Source: InterVlan Routing, Ingress Interface: GigabitEthernet1/0/47 [Vlan200]
- CAMPUS-Access1 (Right):** IP: 212.1.10.1, Type: SWITCH, Link Source: Switched, Ingress Interface: GigabitEthernet1/0/1

Path Trace

Topology View



Intent Based Policies - Controller View

Easily Express Business Requirements

APIC - Enterprise Module

Hi, admin 14

Create New Policy

Policy Name: End_of_Quarter

Source: 172.28.97.51 (1 apps selected)

Priority Level: 42

Scope: Select Tag

Policy Action: Permit Deny Copy

Destination: IP Address/Host User, Application

Copy Destination: Destination IP Address

Create Policy

Name	Scope	Source : Users	Source : Application	Destination : Users	Destination : Application	Status	Actions	Priority Level	Destination	Actions
Lync:video:172.28.97.51:2015-06-07 20:23:53.906065		172.28.97.51	29438;UDP			Active	Permit	31		
allow_VNC		172.28.97.51	adam-vnc1			Active	Permit	8		
Lync:audio:172.28.97.54:2015-06-07 20:07:38.382051		172.28.97.54	32486;UDP	172.28.97.51	3448;UDP	Active	Permit	46		
Lync:video:172.28.97.54:2015-06-07 20:07:38.403811		172.28.97.54	29120;UDP	172.28.97.51	29438;UDP	Active	Permit	31		
Lync:audio:172.28.97.51:2015-06-07 20:07:38.394602		172.28.97.51	3448;UDP	172.28.97.54	32486;UDP	Active	Permit	46		
allow_VNC-return		172.28.97.54	adam-vnc1			Active	Permit	8		

15 Policies

6 Policies

First Previous 1 Next Last

- Based on Users, Resources, Actions and Priorities
 - Integrates with IS/AAA/LDAP for Host user
 - Supports Tagging - e.g. can apply an ACL to a given site/branch

Plug and Play

Auto Device Provisioning

[Status](#) | [Sites](#) | [Image Management](#) | [Unclaimed Devices](#)

Site:

[Load](#) [Create](#) [Clone](#) [Delete](#)

Deploy devices that Do not Support Cisco PnP Protocol (Unsecure)

Specify additional site information

* * [Add Rule](#) [Refresh](#)

Building1 Devices

 Serial Number	 Device Name	 Product ID	 Config	 Bootstrap	Image	Details	Status	Delete
FAC1539W110	BLDG1_Floor2_Room1	WS-C3560CG-8PC-S	C3560CX_test1.cfg 		<input type="text" value=""/> 	Details	PENDING	

Displaying 1 of 1 Device

APIC-EM IWAN App

Network Wide Settings

Network wide settings

- System
- IP Address Pools
- Service Providers
- Hub Site WANs**
- Certified IOS releases

Select Hub Configuration

The are are attributes for the DMVPN hubs. Also, Performance Routing QoS, and a Trust certificate will be enabled in these hubs as well.

WAN Cloud: undefined			WAN Cloud: undefined
Service Provider: AT&T			Service Provider: AT&T
IP Address: 10.77.153.110			IP Address: 10.77.153.120
Bandwidth: 200 Mbps			Bandwidth: 50 Mbps
Router Type: ASR1K			Router Type: ASR1K
Management IP: 10.77.153.130			Management IP: 10.77.153.132

[Previous](#) [Next](#)

API example: APIC-EM Path Visualization and ACL Analysis

The screenshot shows the APIC-EM interface. On the left is a navigation menu with items: Home, Discovery, Device Inventory, Host Inventory, Topology, and Path Trace. The main content area is titled "Policy Analysis" and describes the "APIC-EM Service API based on the Swagger™ 1.2 specification". It lists available APIs: Discovery, Host Inventory, Policy Analysis, Role Based Access Control, Task, and Topology. The "Policy Analysis" section is expanded to show the "path : Path Computation API". The API method is "POST flow-path". It includes "Implementation Notes" (Method to post a 5-tuple and receive a task ID for the calculation of the path), "Response Class" (TaskIdResult, TaskIdResponse, TaskId), and "Parameters" (pathRequest). A text box shows the JSON for pathRequest: {"destIP": "207.1.10.20", "sourceIP": "65.1.1.6"}. The parameter content type is set to "application/json".

```
"response": {
  "request": {
    "sourceIP": "212.1.10.20",
    "destIP": "65.1.1.6"
  },
  "lastUpdate": "Thu Apr 23 01:23:21 UTC 2015",
  "properties": [ ],
  "networkElementsInfo": [
    {
      "id": "424621be-d2b4-4d42-ad16-92d4d5c19fa4",
      "type": "WIRED",
      "ip": "212.1.10.20",
      "linkInformationSource": "Wired"
    },
    {
      "id": "8beada2e-cd2c-421d-941f-3ba42696c489",
      "name": "CAMPUS-Access1",
      "type": "SWITCH",
      "ip": "212.1.10.1",
      "ingressInterface": {
        "physicalInterface": {
```

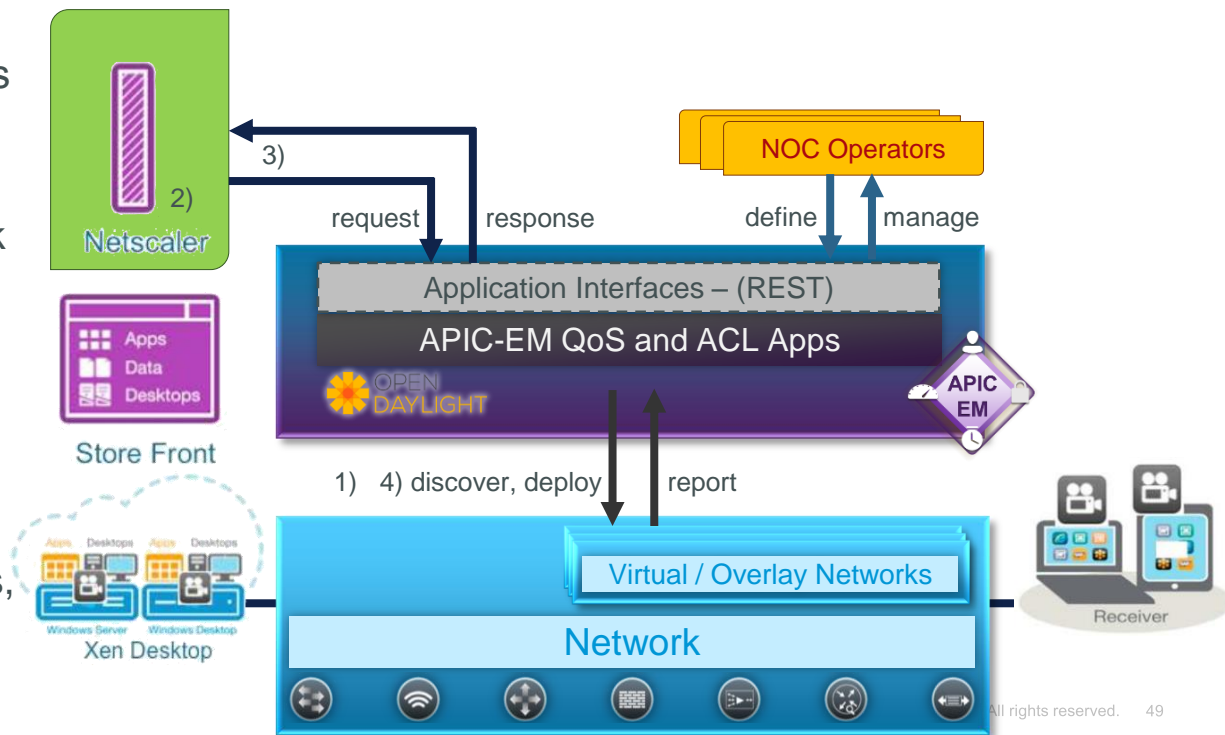
```
PathRequest {
  destPort (string, optional): Destination
  Port.
  destIP (string): Destination IP address.
  sourceIP (string): Source IP address.
  sourcePort (string, optional): Source
  Port.
  protocol (string, optional): Protocol
```

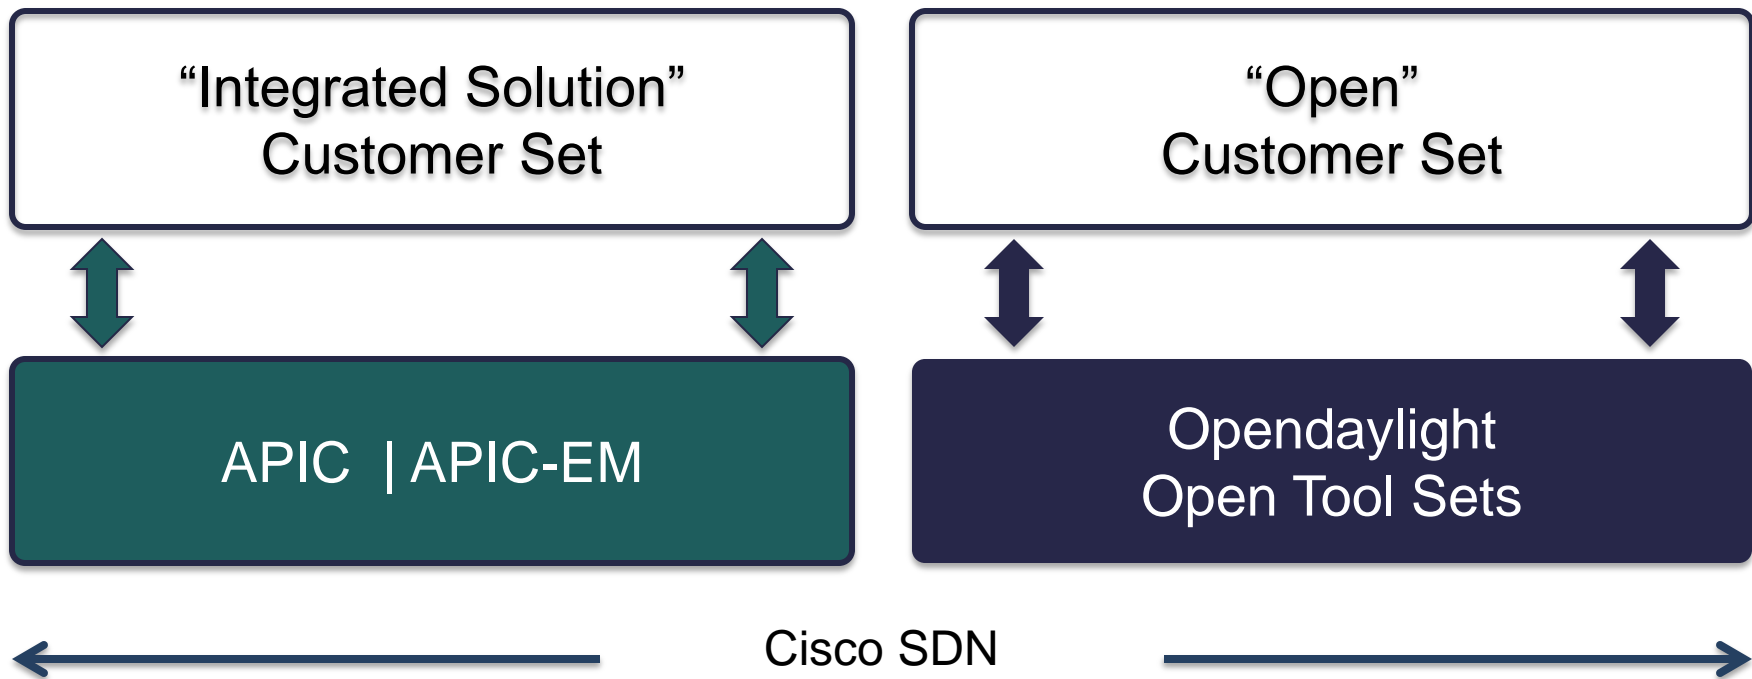
Example: Dynamic Policy for Citrix Clients

Problem: How to provide dynamic application-specific Policy to Citrix XenDesktop users ?

Solution: Use Citrix NetScaler's integration with APIC-EM:

- 1) APIC-EM discovers network and endpoints
- 2) NetScaler detects start of (video) data transfer
- 3) NetScaler requests QoS Policy via APIC-EM's API
- 4) APIC-EM validates, deploys, and reports the change





Controller Layer – Major Milestones

Major Milestones of Controller Development

Controller Layer
(Orchestration + Analytics)



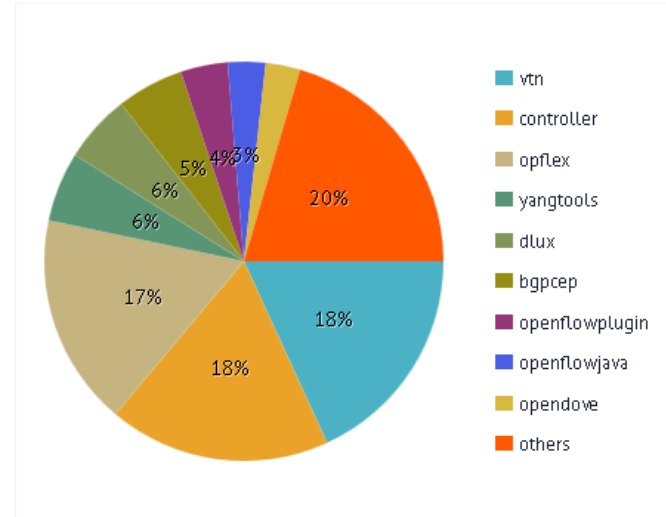

Programmable
Network Layer



OpenDaylight – Who is Contributing?

Question: Who are today's top contributors to Open Daylight?

Answer: Check OpenDaylight's Spectrometer (based on OpenStack Stackalytics)



Source: http://spectrometer.opendaylight.org/?metric=loc&project_type=opendaylight

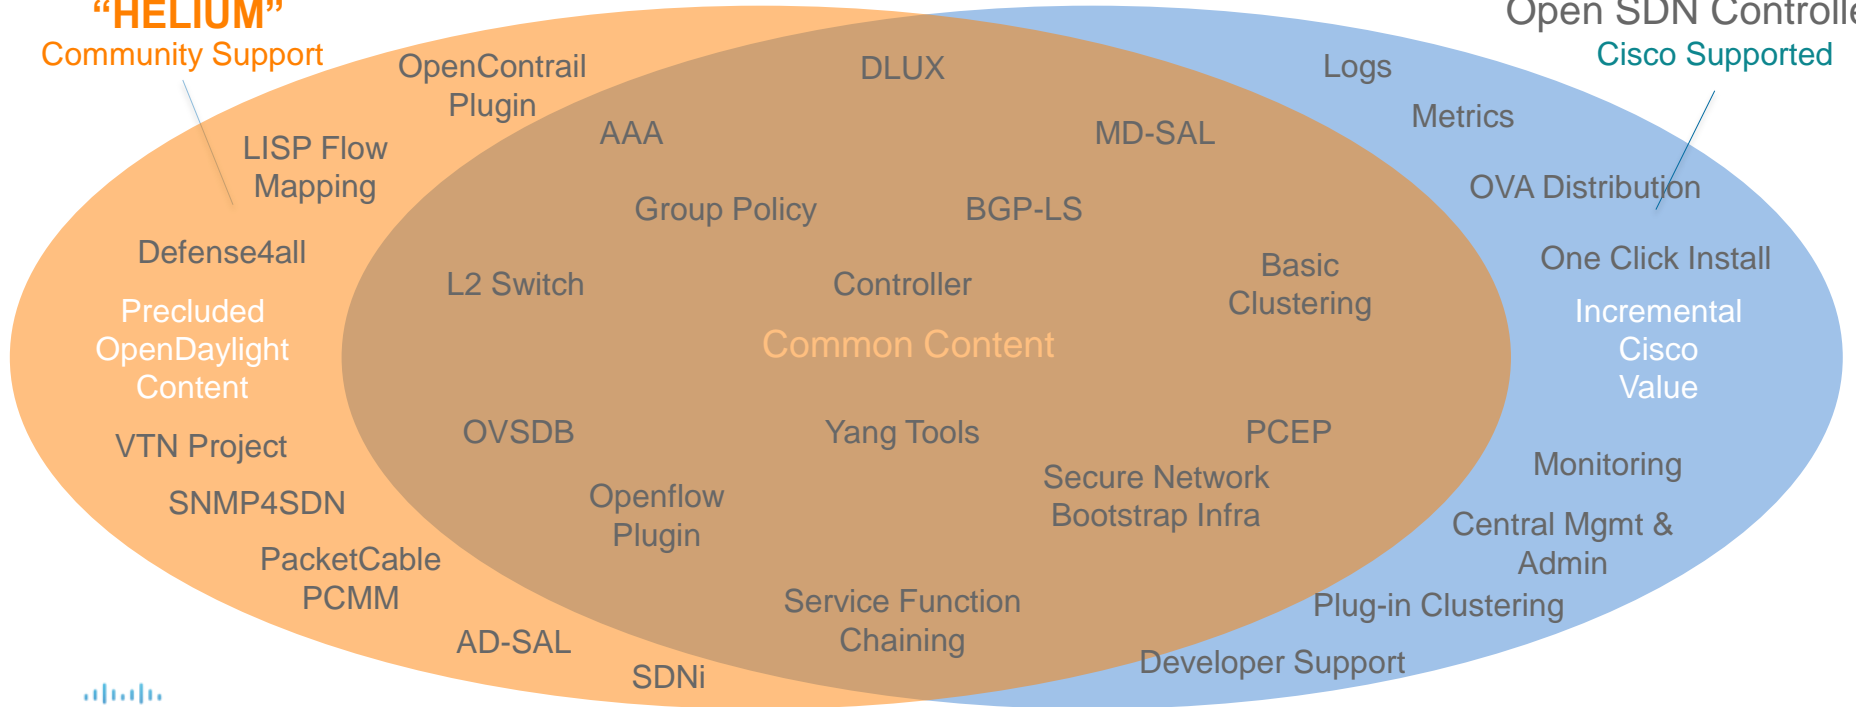
Cisco Open SDN Controller



Community Support

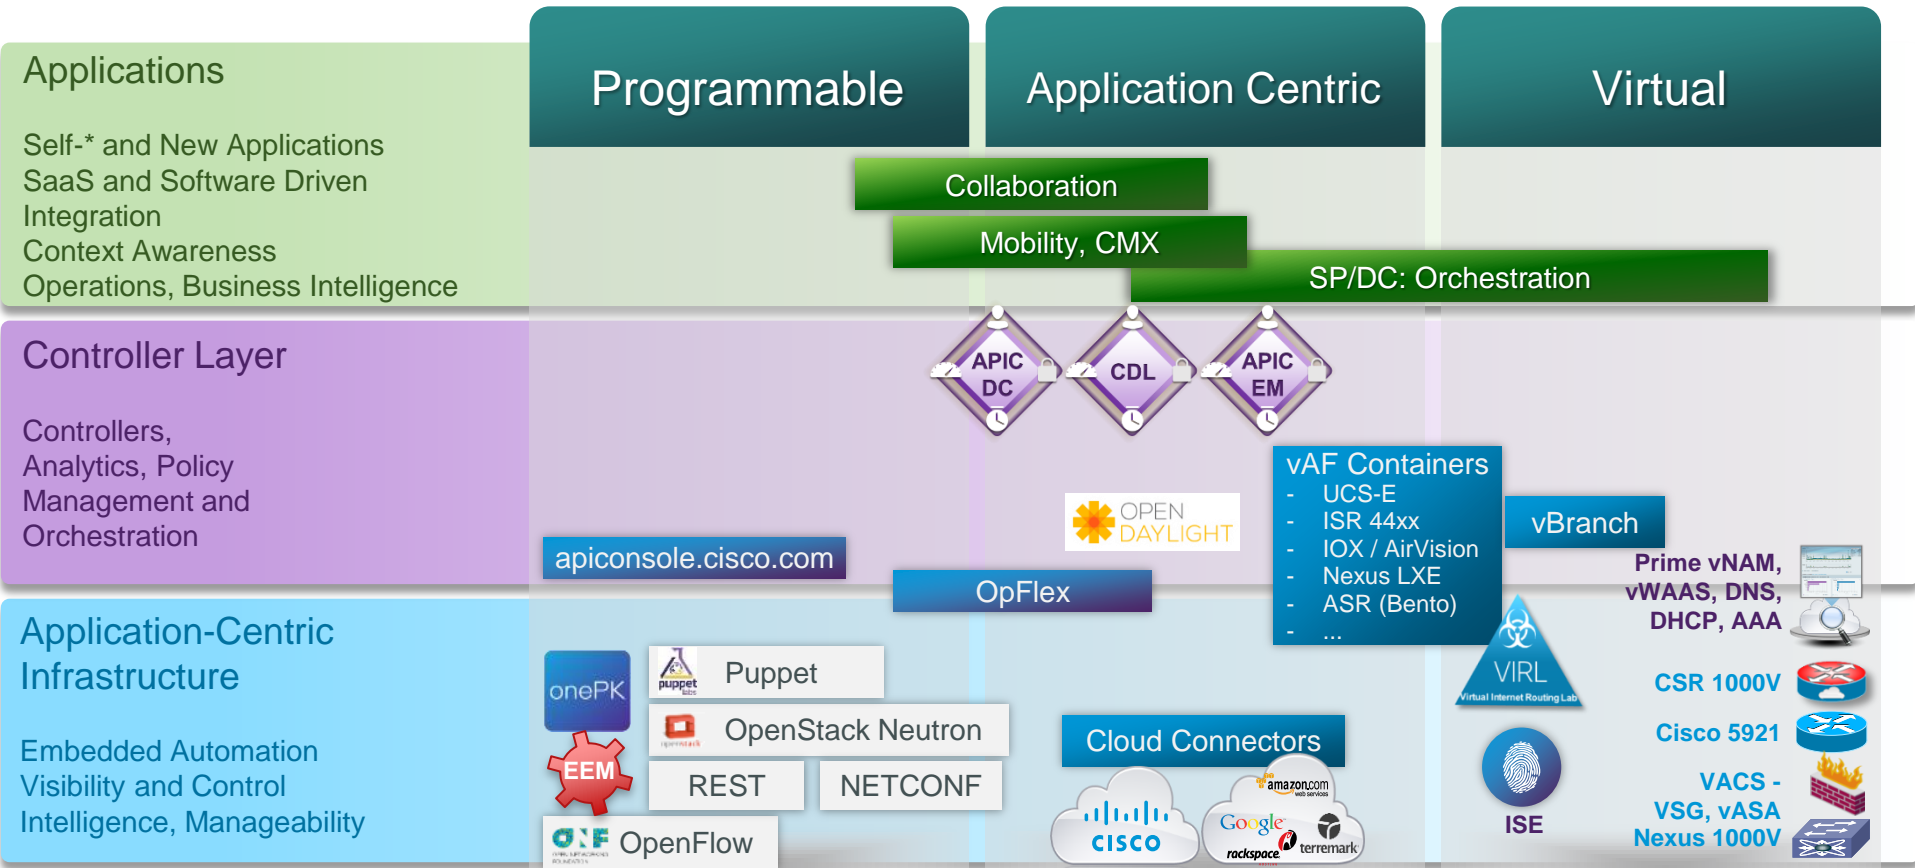


Open SDN Controller
Cisco Supported

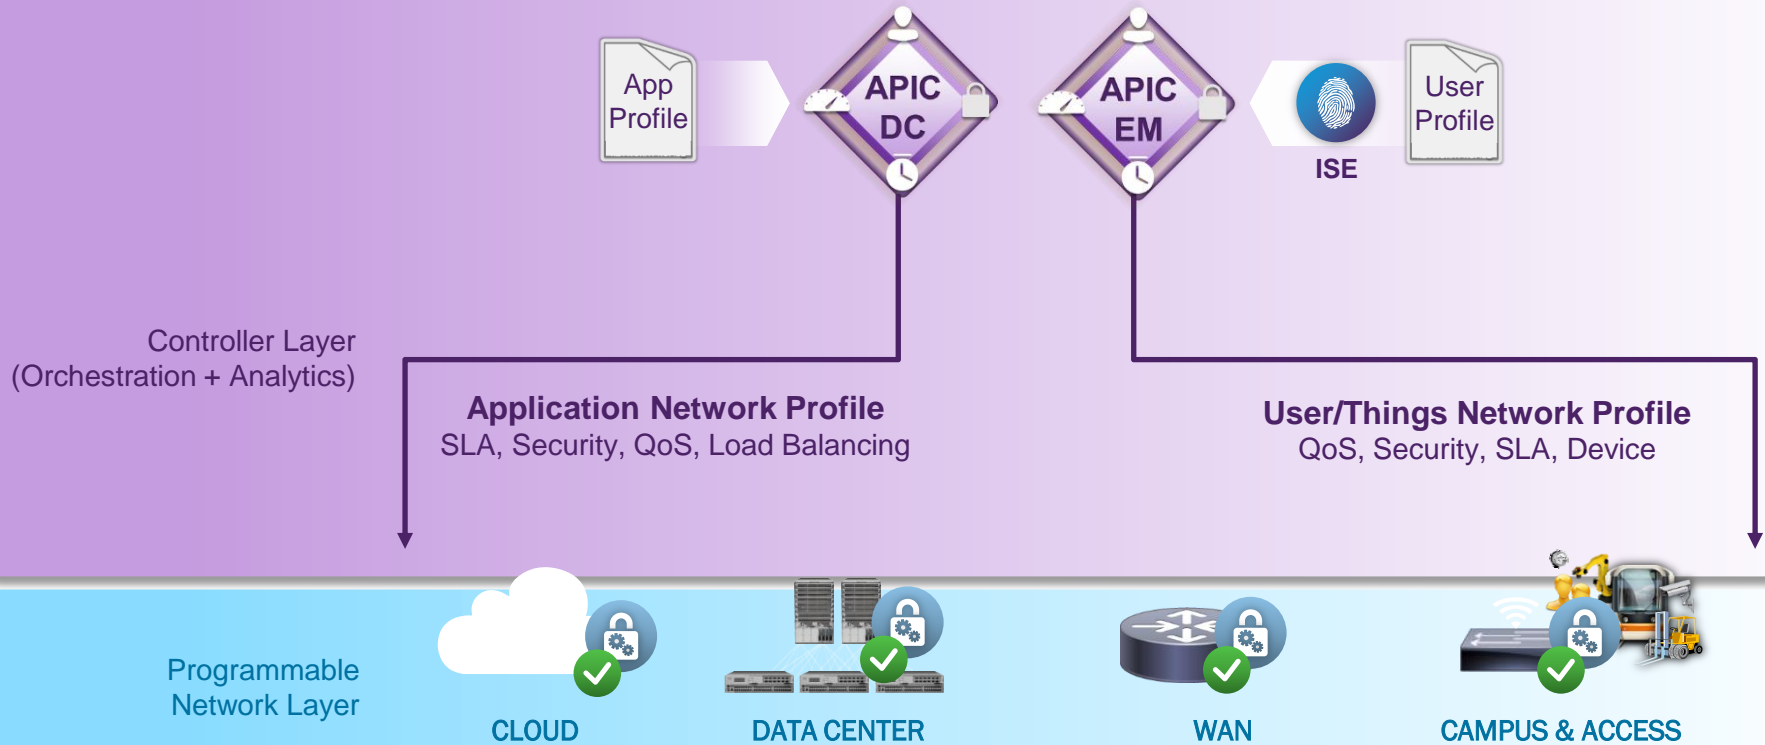


Závěrečná slova

Cisco Enterprise ACI – 3x3 Portfolio (Subset)





Common Policy End To End



Openness - OpenDaylight Contributions

OpenDaylight's Spectrometer (based on OpenStack Stackalytics)



Source: http://spectrometer.opendaylight.org/?metric=loc&project_type=opendaylight

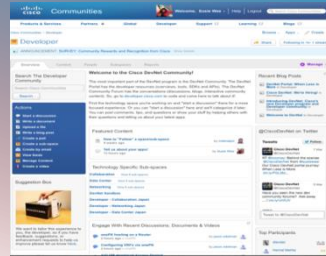
APIC-EM, OpenSDN in DevNet

DevNet Portal



Developer APIs,
SDKs

Community Forum



Integrated with
Cisco
Communities

DevNet Sandbox



Developer Lab in
the Cloud & SW
download.

DevNet Zone!










Cisco's
Developer
Conference



<https://developer.cisco.com/site/apic-em/>
<https://developer.cisco.com/site/openSDN>

Open SDN Controller Demos in dCloud

	Cisco Open SDN Controller Sandbox v1 Get early access to the new Cisco Open SDN Controller in this sandbox environment. Start/Schedule More Information <i>Added : 11/02/2015</i>
	OpenDaylight 2.0 Sample Apps with 8-Nodes v2 This OpenDaylight 2.0 demo showcases the adoption of BGP-LS, ACLs, and PathMan applications for software-defined networking (SDN) that utilize the OpenDaylight Open Source Platform. Start/Schedule More Information <i>Added : 30/09/2014</i>
	Cisco OpenDaylight v1.1 The OpenDaylight v1.1 Demo showcases the adoption of BGP-LS and ACLs, which are applications for software-defined networking (SDN) that utilize the OpenDaylight Open Source Platform. Start/Schedule More Information <i>Added : 16/06/2014</i>
	Cisco OpenDaylight 1.0 Sandbox v1 Control your own OpenDaylight environment and build your own LSPs Start/Schedule More Information <i>Added : 9/07/2014</i>
	Cisco WAN Automation Engine 6.0 with 8-Nodes v1.2 Show how WAE, a network modeling technology, allows for real-time analysis of traffic needs and placement in complex WAN topologies. Start/Schedule More Information <i>Added : 8/04/2015</i>
	Cisco Autonomic Networking Sandbox v1 Control your own Autonomic Networking environment with this Cisco Autonomic Networking Sandbox Start/Schedule More Information <i>Added : 24/11/2014</i>
	OpenDaylight Helium Sandbox v1.1 Control your own OpenDaylight Helium environment using the Cisco OpenDaylight (ODL) Helium Sandbox Start/Schedule More Information <i>Added : 30/03/2015</i>

- Cloud-based Demos (and Learning)
- Scheduled or on-demand
- Customize and Save your own
- Login to:
<http://dcloud.cisco.com>



CISCO

TOMORROW starts here.