
TOMORROW starts here.
cisco
Cisco live!

Cisco Modelling Labs - Lessons from a Virtual World

BRKRST-2646

Joel Obstfeld

Distinguished Engineer

Cisco live!

Network Virtualisation

The challenge -

- How do you deploy new services, make changes, or troubleshoot in the live network
  - Quickly, Consistently, Efficiently
  - And at scale

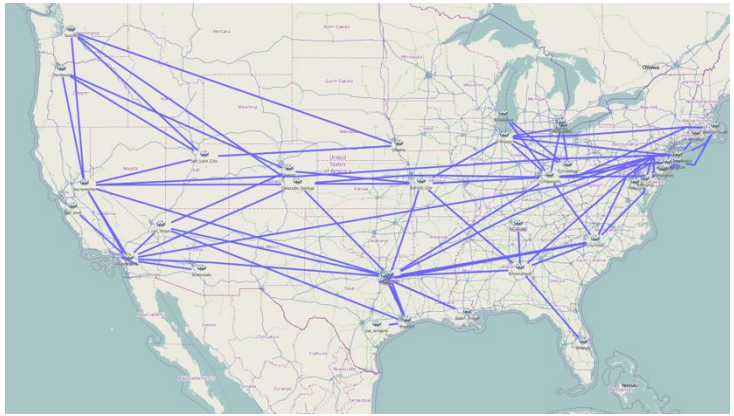

Network Virtualisation

Lab equipment comes at premium

- More equipment requires the more power, space and cooling
  – and there's never enough equipment for all of the people who want to use it!

Similar challenges have already been faced in the Data-centre world

- Improve the utilisation of compute hardware while reducing power, space and cooling

Can we apply a similar approach to network devices?

Why use Virtual Networks?

Technical Opportunities

- Build and deploy networks at scale – virtually
- Verify designs and validate configurations
- Prototyping of new capabilities
- Reduce risk and errors through improved training

Benefits

- Decrease time to deployment for new services
- Replay “events” for training or fix verification
- 100 router testbed - 1 physical and 99 virtual
- Test combinations of new services

Training & Education

- Teach and train the next generation
  - Network engineers
  - Operators
  - Designers
  - Architects
- Students need ‘hands-on’ experience but the challenge is access to hardware
  - Learn by doing!
  - 10 students to 1 router or 1 student to 10 routers?
  - Real-world operating systems or Open-source?

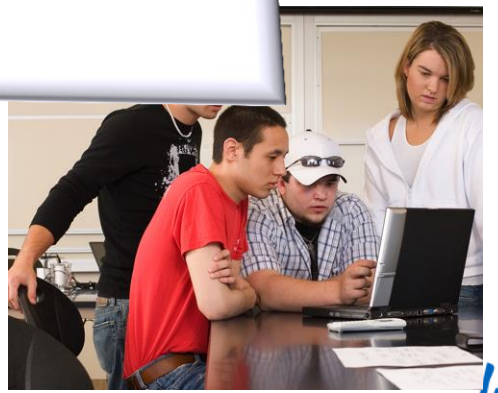

New Features, New Capabilities

- Build, test, experiment and learn SDN technologies – new paradigms create new challenges
- Cisco OnePK network programmability framework – developers must be able to test and validate applications before deploying to the real network
- How do you test and trial new control-plane capabilities such as ‘Segment Routing’ if you don’t have a lab?

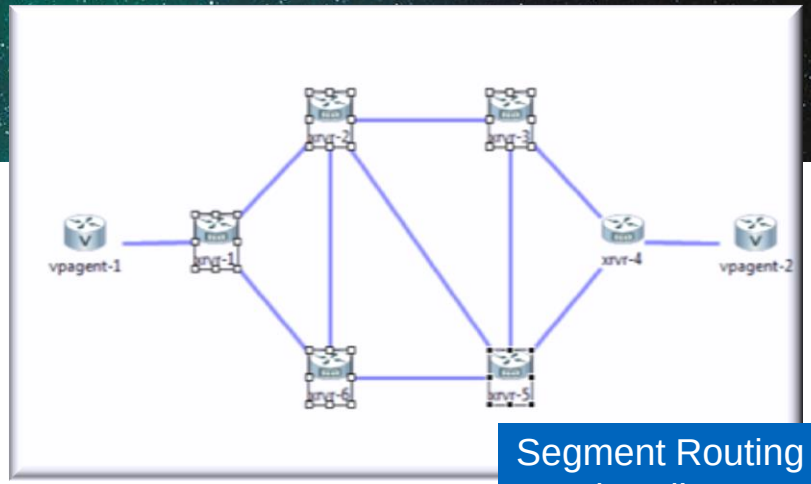

Segment Routing – hand’s on

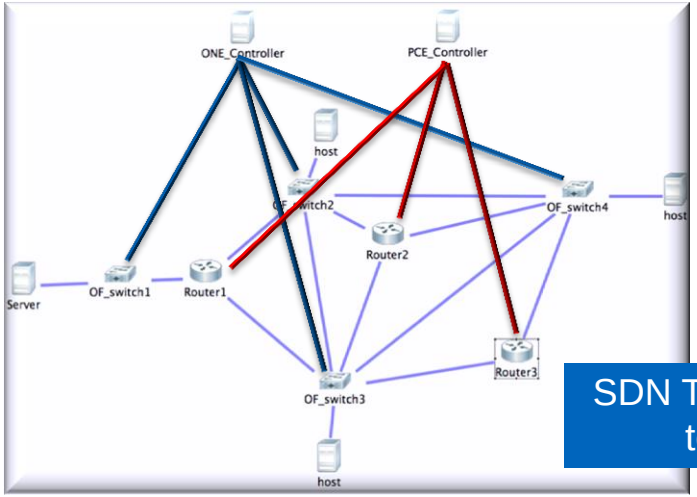

SDN Technologies testbed

Agenda

- Virtualisation - a brief tour
- Building the virtual world
- Student testbed – cutting your teeth
- Topology building
- Cisco Modelling Labs in action…Live demo
- Case studies
- Summary
- Q & A

Virtualisation – A Brief Tour

Virtualisation

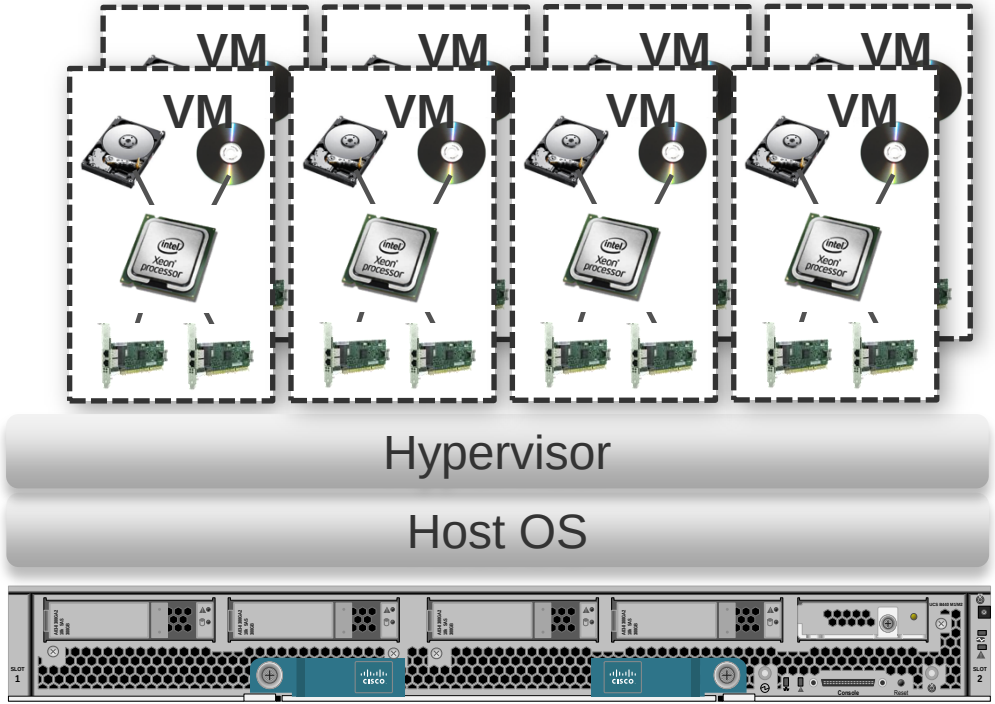

- Virtual Machines are presented with a set of virtual hardware resources
- Resources can be shared between the virtual machines or can be reserved

- What is the performance expectation?
- Is performance predictable and consistent?

Ubuntu / KVM Machine Virtualisation Suite

- Virtualised devices (CPU, I/O, memory) enable a single host to support many virtual machines
- KVM / QEMU provides a kernel-based Hypervisor / host-virtualisation facility
- Ubuntu 12.4 provides the basic host operating system
- Intel VT-x / AMD-V capable CPUs expose hardware-virtualisation functions to Host O/S and Hypervisor

Virtual Machines

Hypervisor

Host O/S

Physical Host

Virtualised Platform Operating Systems

IOS XR

Virtualised in IOS XRv

NX-OS

Virtualised in NX-OSv

IOS XE

Virtualised in CSR1000v

IOS

Virtualised in IOSv

Servers

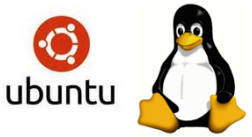

Ubuntu, Cirros, 3rd party appliances

Virtual Machines run the operating system but are NOT representations of a particular hardware platform – no fans, no switch fabric, no ASIC models

Building the Virtual World

Cisco Modelling Labs workbench

- The graphical topology design tool used to create XML topology files
- Powerful topology design functionality
- Drag and drop ease
- Attributes attached to network element define:
  - ✓ Addressing schema
  - ✓ IGPs
  - ✓ BGP
  - ✓ Additional control-plane elements
- Provides simulation management and console access to virtual routers

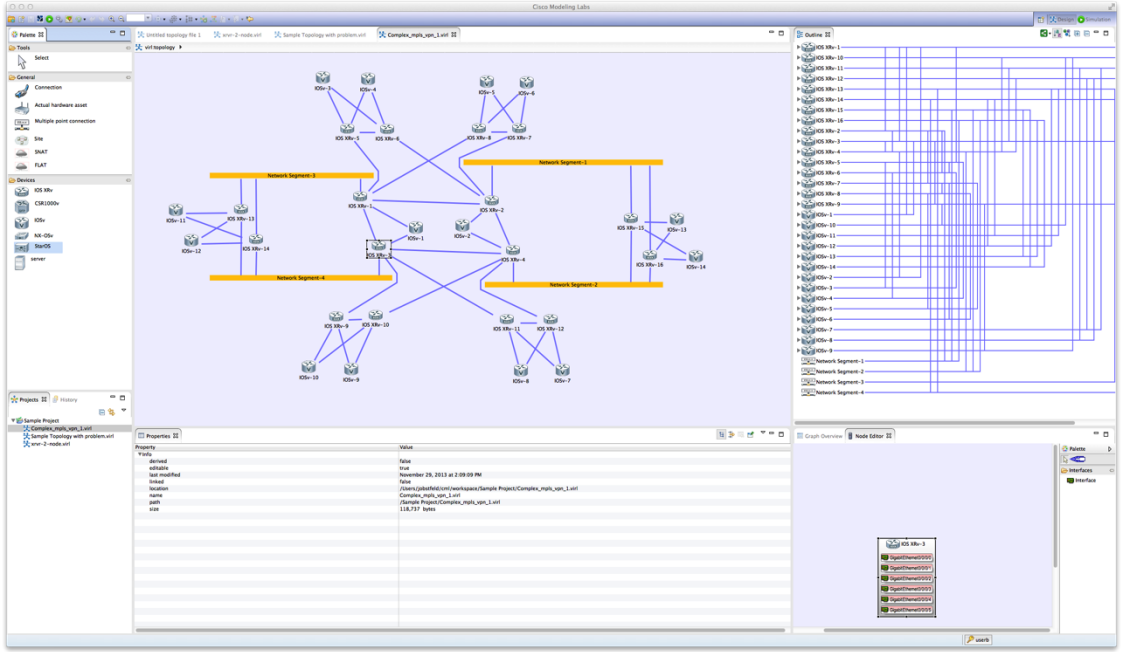

Automated Network Configuration / Rendering

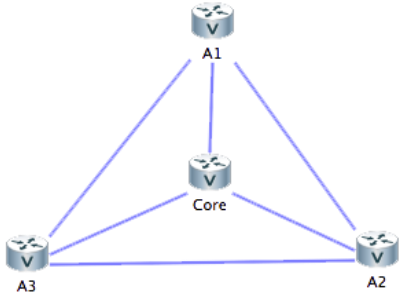

XML Topology Definition

AutoNetKit

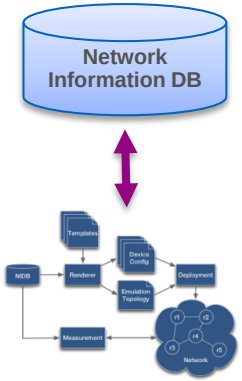

Configurations

```
line vty 0 4
 exec-timeout 720 0
 password cisco
 login
line con 0
 password cisco
!
!
interface Loopback0
  description Loopback
  ip address 192.168.0.4 255.255.255.255
!
interface GigabitEthernet0/0
  description OOB Management
  ! Configured on launch
  no ip address
  duplex auto
  speed auto
  no shutdown
!
interface GigabitEthernet0/1
  description to IOSv-1
  ip address 10.0.128.2 255.255.255.252
  ip ospf cost 1
  duplex auto
  speed auto
  no shutdown
!
```

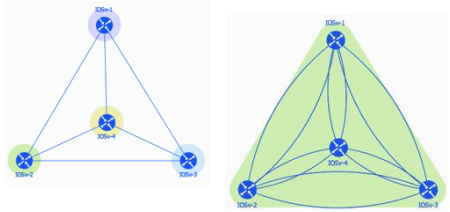

Topology Renderings

Automated configuration Engine

- Uses a DB containing device- and OS-specific information to create configurations for each virtual device
- Presents graphical representations of many topology-specific attributes – links, interfaces, adjacencies, areas
- Creates OS-type specific configurations

Openstack for VM Orchestration

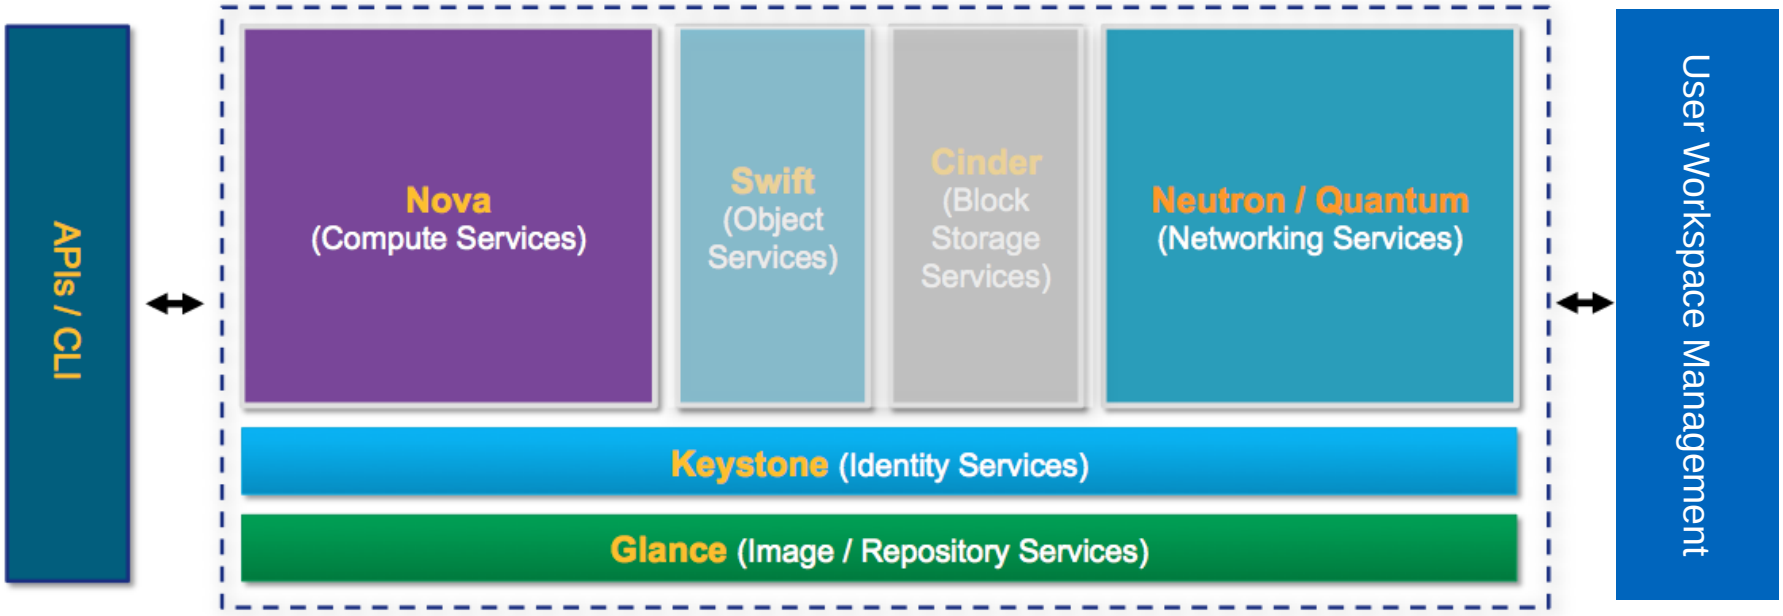

Openstack creates, links, and deletes virtual compute and network resources according to API- or CLI-based instructions

From XML to Virtual Machine

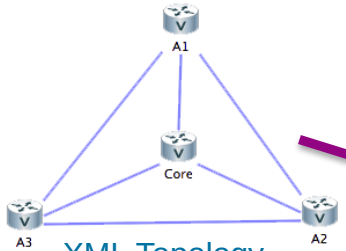

XML Topology Definition

Create Routers

- Identify Type / Flavor
- Associate Image (**Glance**)
- Identify / Assign Resources
- Associate Configuration file
- Launch the VM (**Nova**)

Create Networks / Links

- Identify Links and End-Points
- Assign End-Points to VMs
- Assign Network / Link Characteristics
- Launch the Switch (**Quantum**)

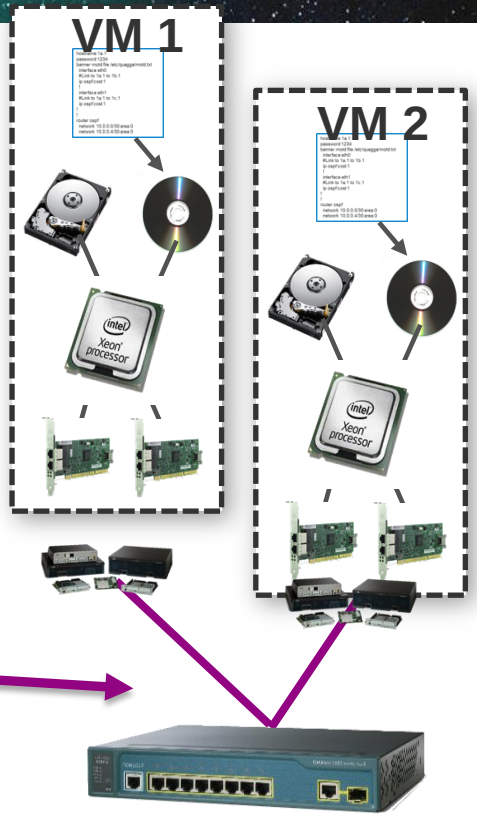

- The Services Topology Director generates Openstack calls for the creation of virtual routers and links based on the XML topology definition
- Injects configuration into the virtual routers

Workflow

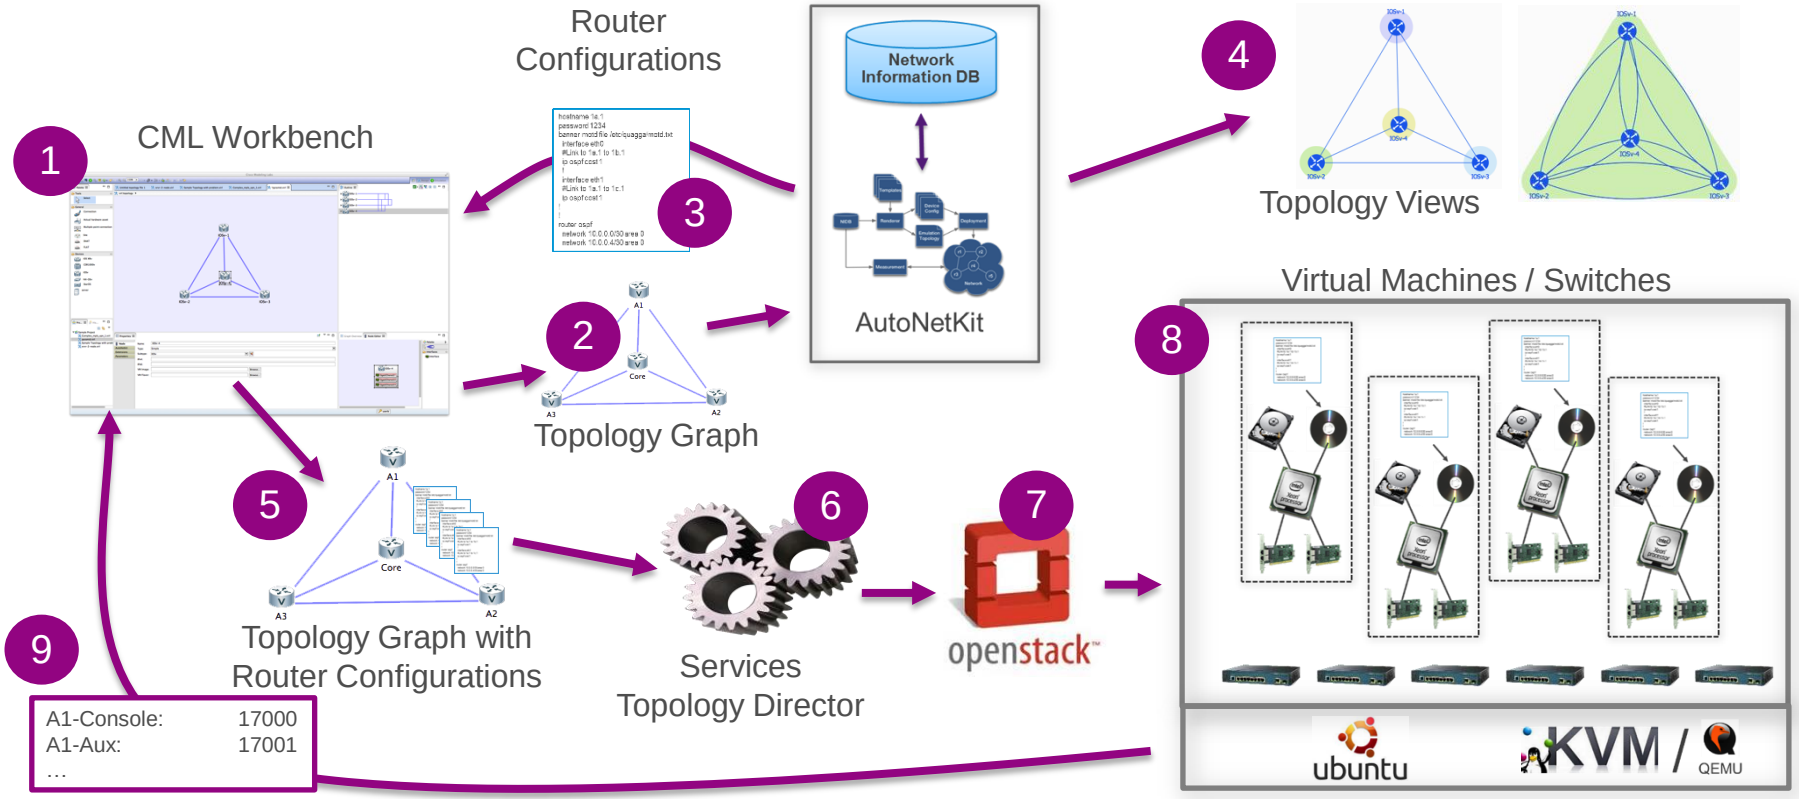

Topology Files

- Complete topology definition together with device configurations held in flat XML file
- Files are highly portable to enable easy sharing
- Integrated GIT repository support enables multi-user file sharing and version control

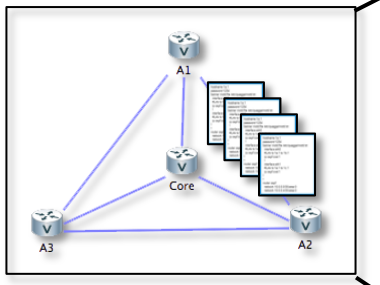


```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<topology xmlns="http://www.cisco.com/VIRL" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" schemaVersion="0.6" xsi:schemaLocation="http://www.cisco.com/VIRL http
    <extensions>
        <entry type="String" key="AutoNetkit.IGP">ospf</entry>
    </extensions>
    <node location="373,292" subtype="IOSv" type="SIMPLE" name="iosv-1">
        <interface name="GigabitEthernet0/1" id="0"/>
        <interface name="GigabitEthernet0/2" id="1"/>
        <interface name="GigabitEthernet0/3" id="2"/>
    </node>
    <node location="392,170" subtype="IOSv" type="SIMPLE" name="iosv-2">
        <interface name="GigabitEthernet0/1" id="0"/>
        <interface name="GigabitEthernet0/2" id="1"/>
        <interface name="GigabitEthernet0/3" id="2"/>
    </node>
    <node location="269,368" subtype="IOSv" type="SIMPLE" name="iosv-3">
        <interface name="GigabitEthernet0/1" id="0"/>
        <interface name="GigabitEthernet0/2" id="1"/>
        <interface name="GigabitEthernet0/3" id="2"/>
    </node>
    <node location="475,388" subtype="IOSv" type="SIMPLE" name="iosv-4">
        <interface name="GigabitEthernet0/1" id="0"/>
        <interface name="GigabitEthernet0/2" id="1"/>
        <interface name="GigabitEthernet0/3" id="2"/>
    </node>
    <connection src="/virl:topology/virl:node[3]/virl:interface[1]" dst="/virl:topology/virl:node[4]/virl:interface[1]"/>
    <connection src="/virl:topology/virl:node[1]/virl:interface[1]" dst="/virl:topology/virl:node[2]/virl:interface[1]"/>
    <connection src="/virl:topology/virl:node[3]/virl:interface[2]" dst="/virl:topology/virl:node[2]/virl:interface[2]"/>
    <connection src="/virl:topology/virl:node[2]/virl:interface[3]" dst="/virl:topology/virl:node[4]/virl:interface[2]"/>
    <connection src="/virl:topology/virl:node[3]/virl:interface[3]" dst="/virl:topology/virl:node[1]/virl:interface[2]"/>
    <connection src="/virl:topology/virl:node[1]/virl:interface[3]" dst="/virl:topology/virl:node[4]/virl:interface[3]"/>
</topology>
```

Topology Visualisation

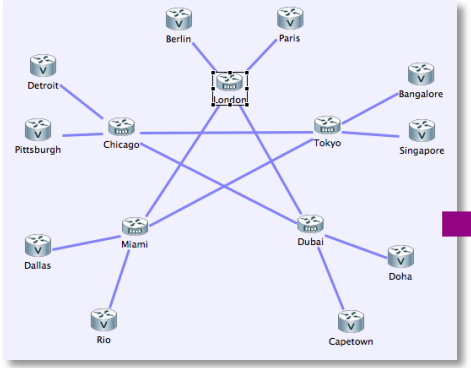

Network topology design as input

Physical layer view

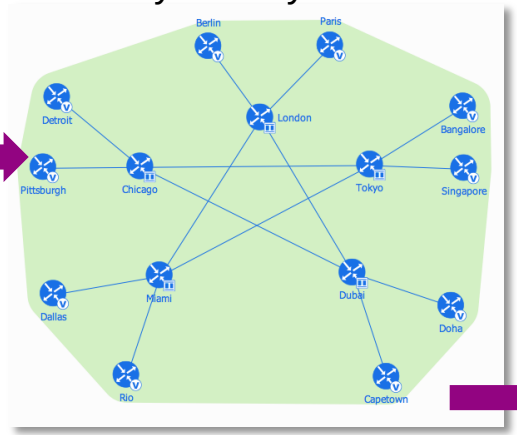

BGP view

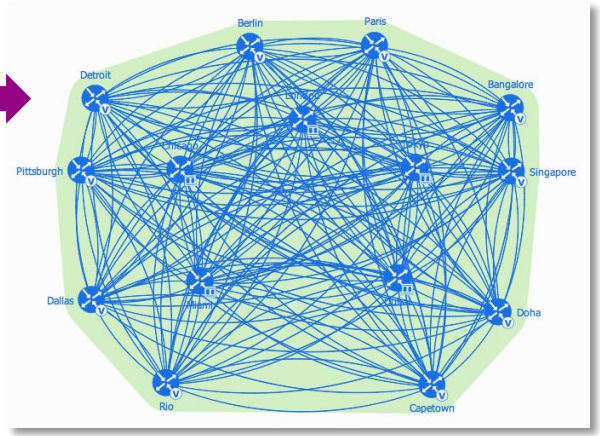

- Configuration Engine generates two outputs – per-node configuration files and network visualisation diagrams
- ‘Layer views’ show protocol-centric topologies

Configuring Attributes in AutoNetkit

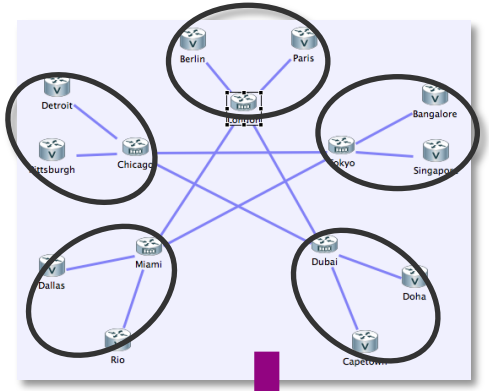

Setting ASN values on groups of nodes

BGP view – five AS, all nodes are BGP speakers

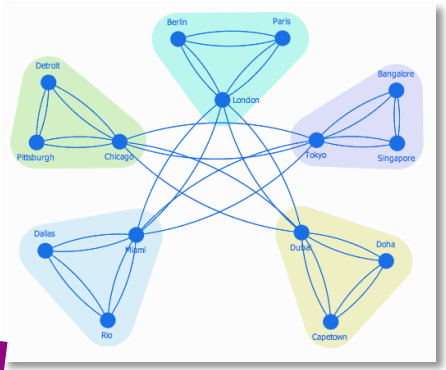

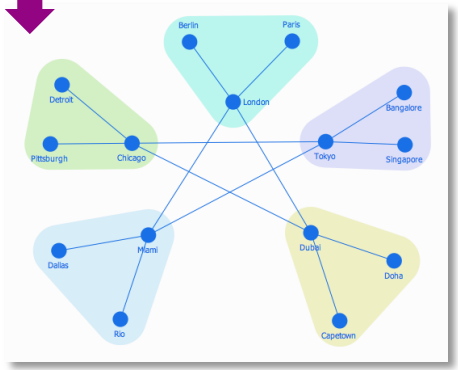

Physical layer view

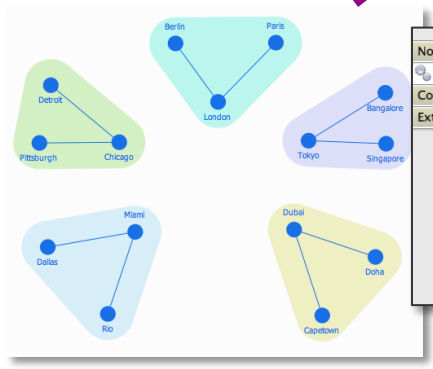

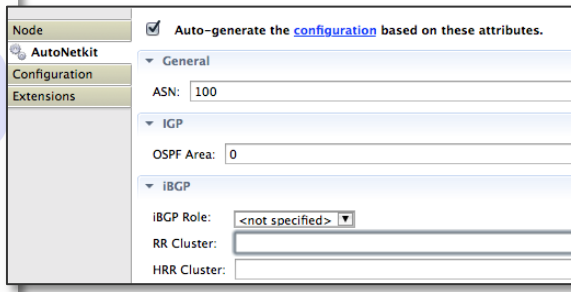

IGP view – five AS's, five area 0's

Configuring Attributes in AutoNetkit

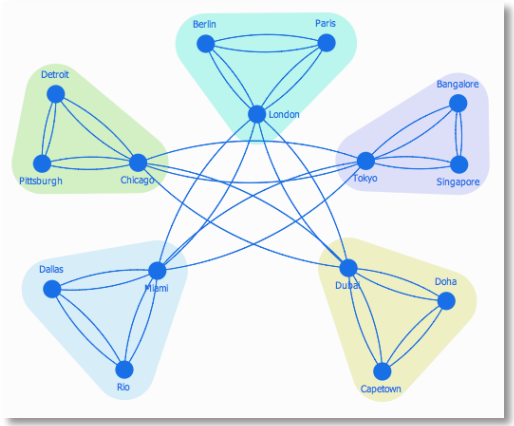

BGP full mesh within each AS

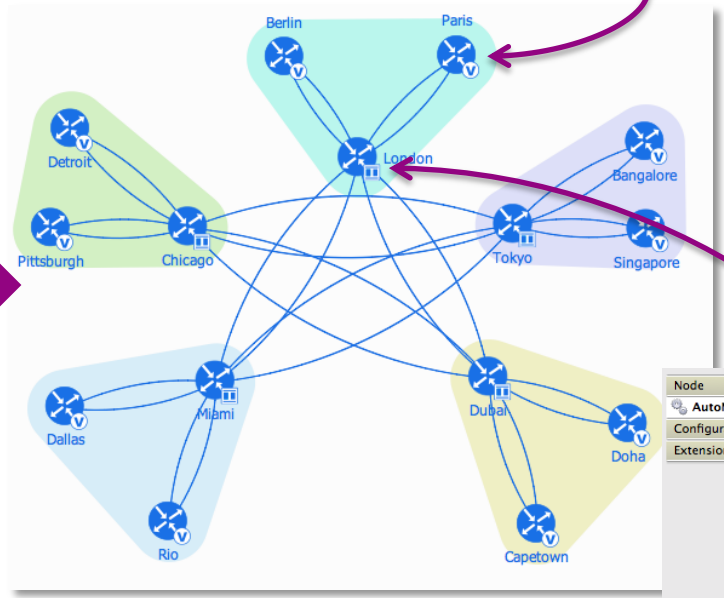

BGP route-reflector clusters and AS's configured

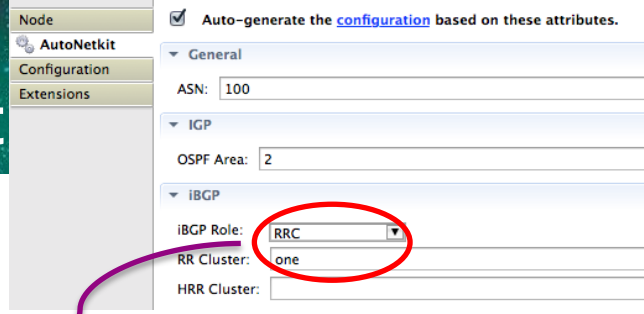

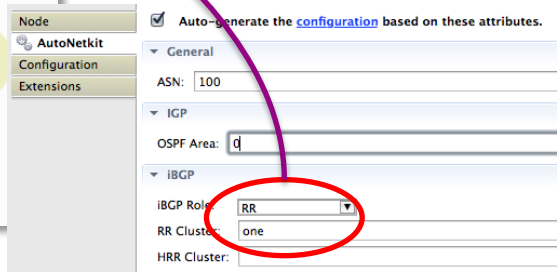




Student Testbed – Cutting your Teeth

Study Topology

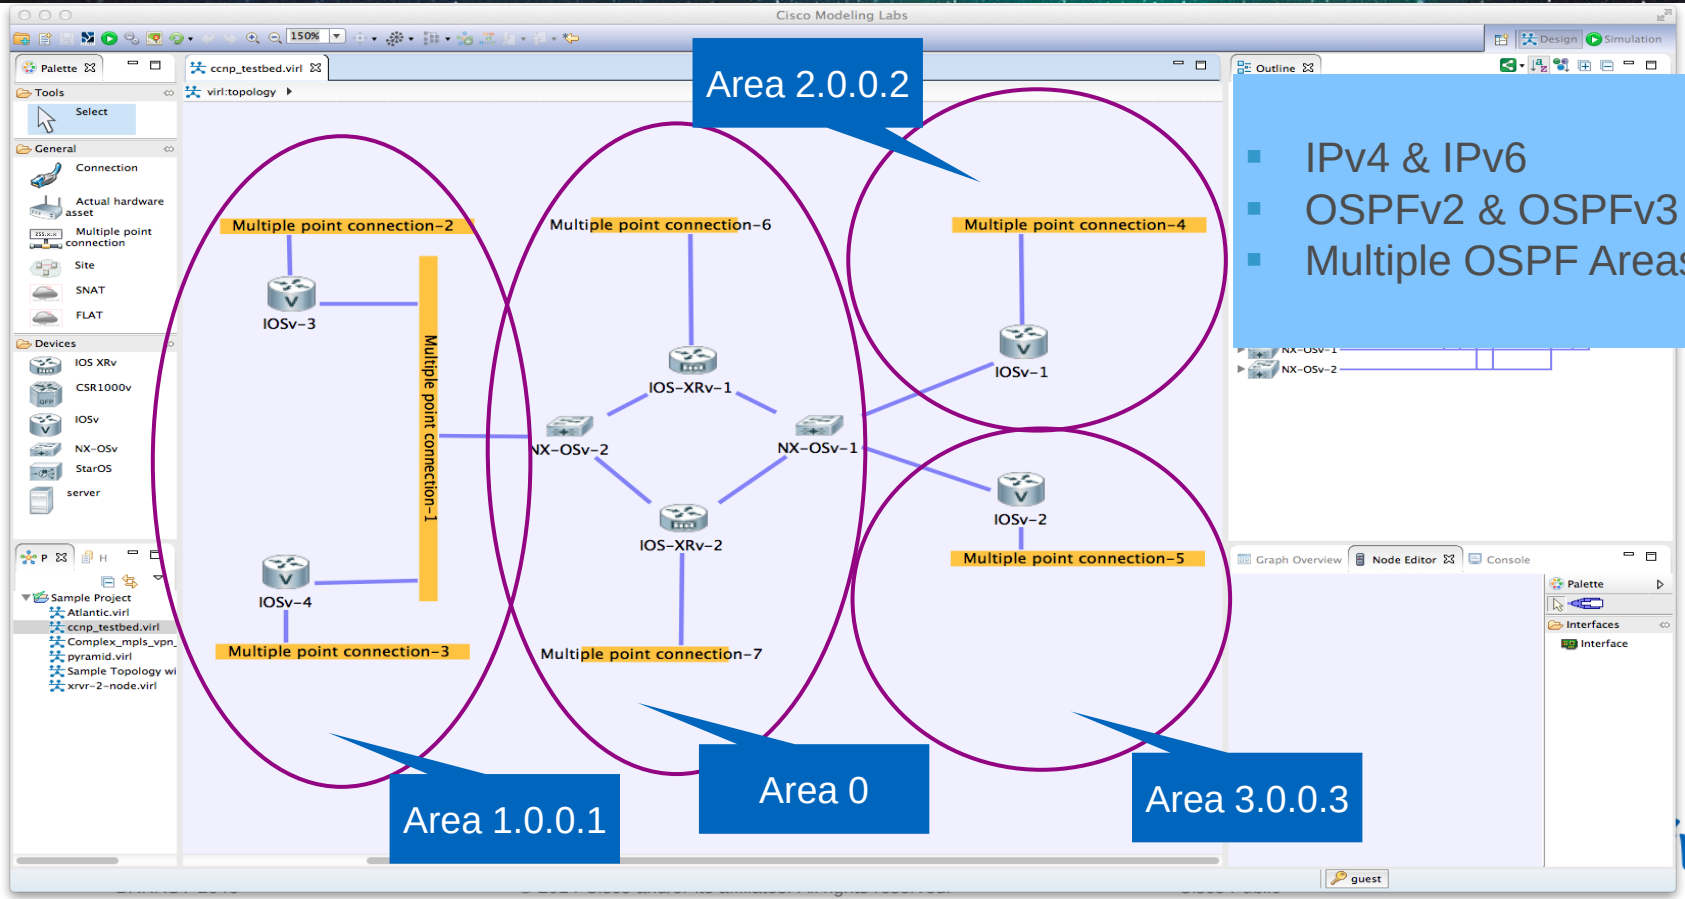

Study Topology – Setting Address Families and IGP

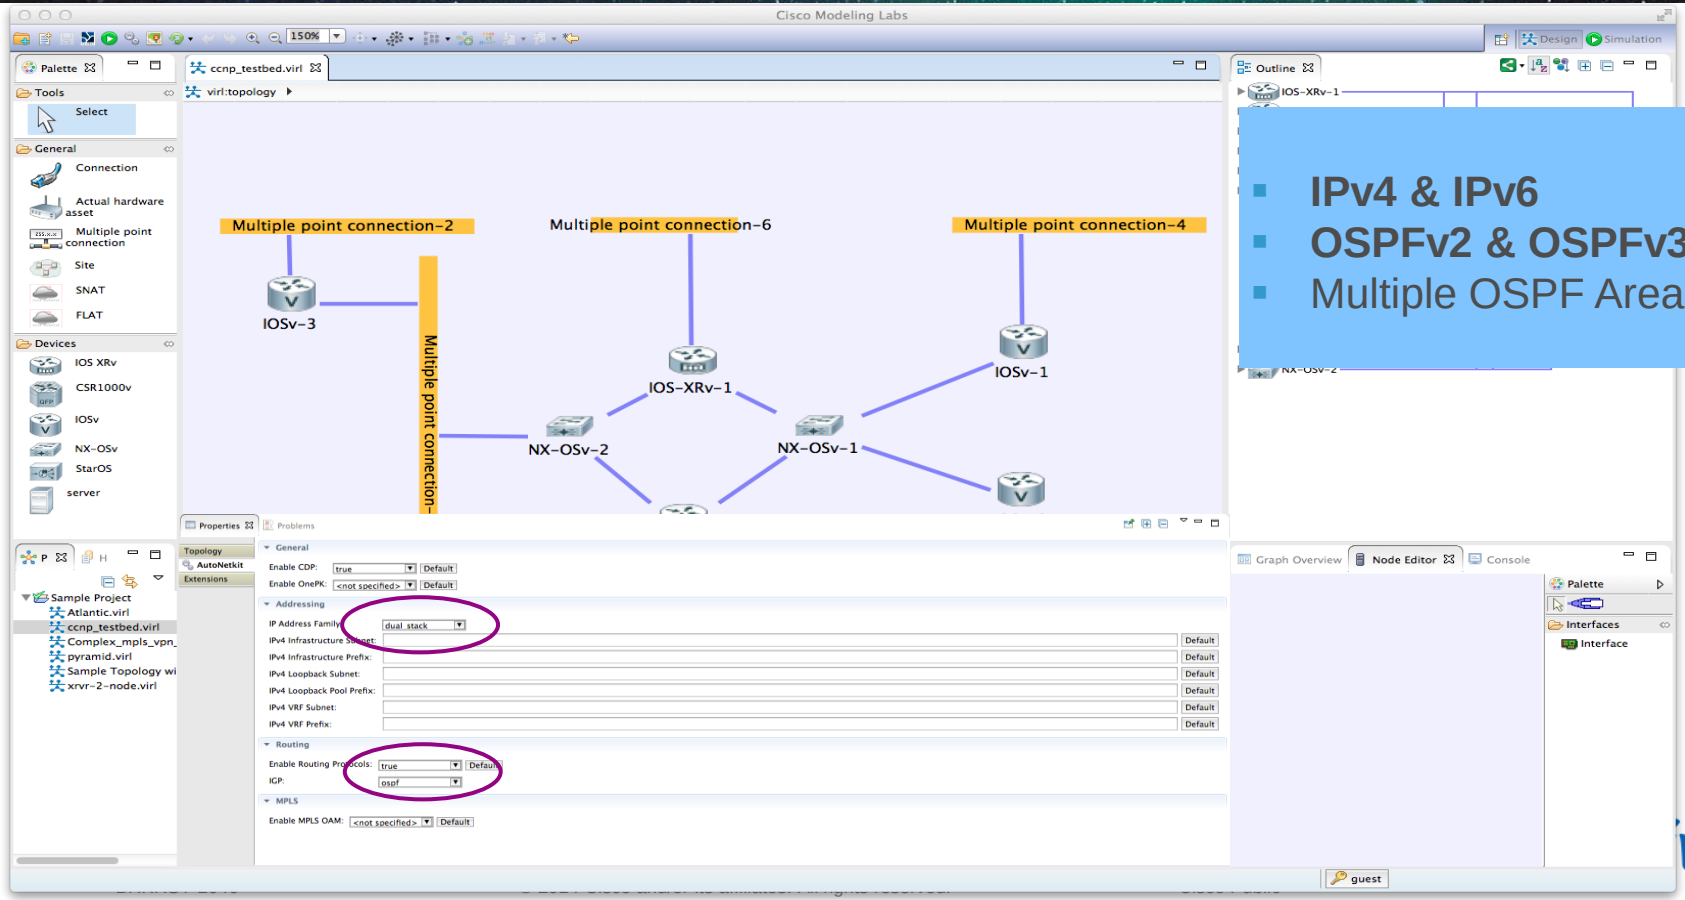

- **IPv4 & IPv6**
- **OSPFv2 & OSPFv3**
- Multiple OSPF Areas

Study Topology – Setting OSPF Areas

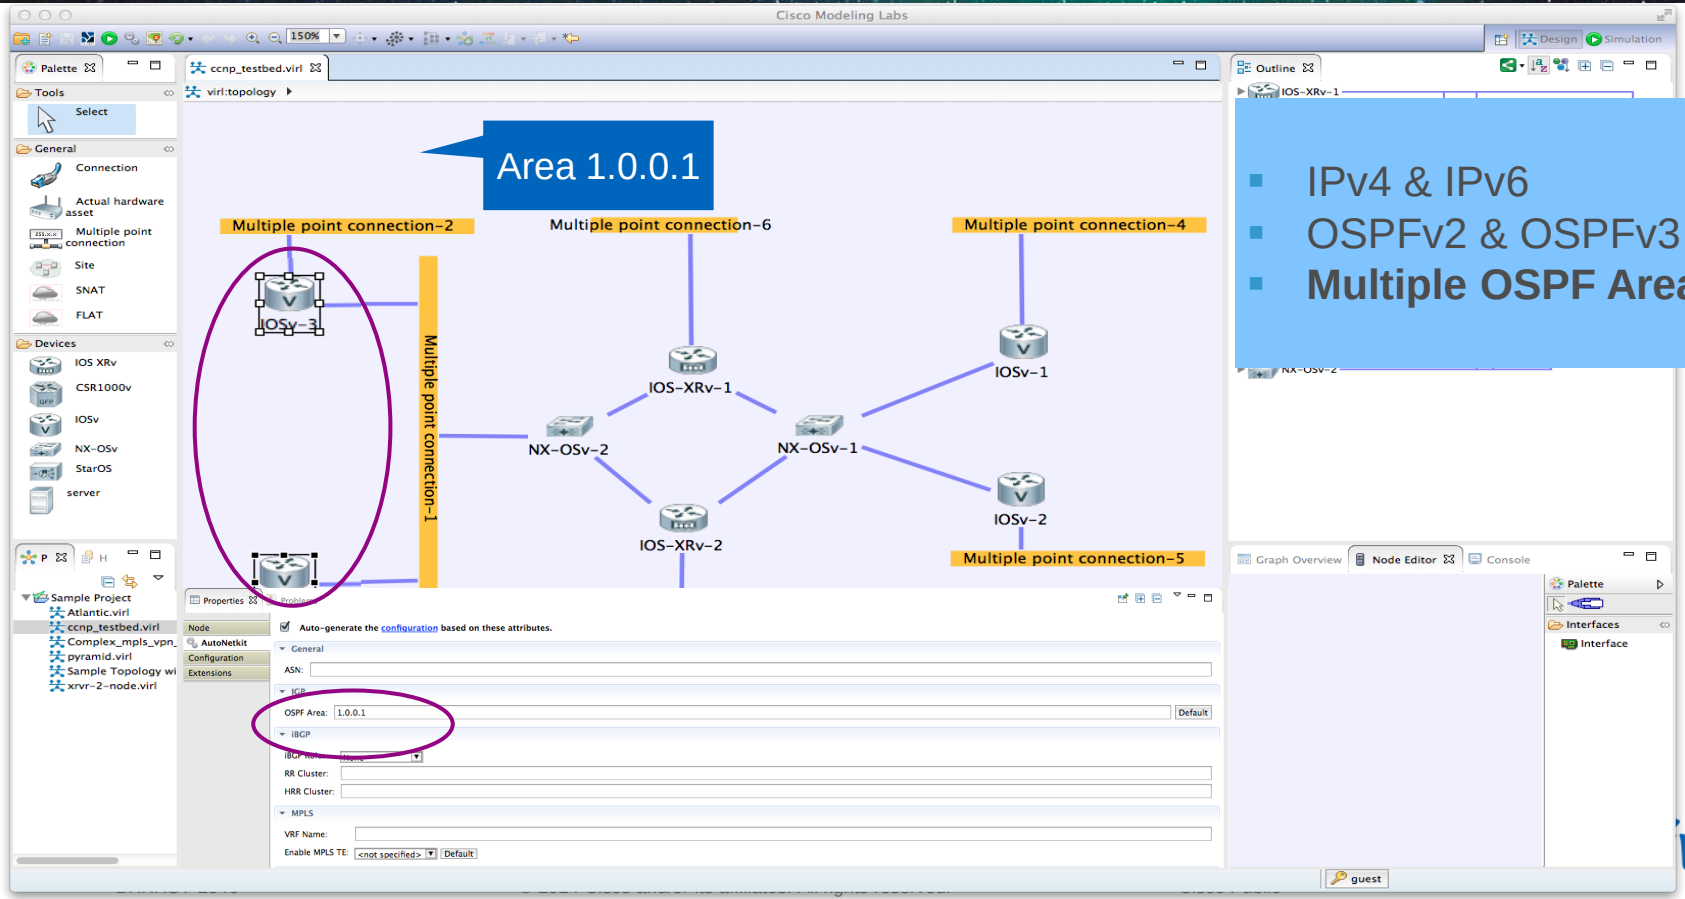

- IPv4 & IPv6
- OSPFv2 & OSPFv3
- **Multiple OSPF Areas**

Study Topology - Visualise

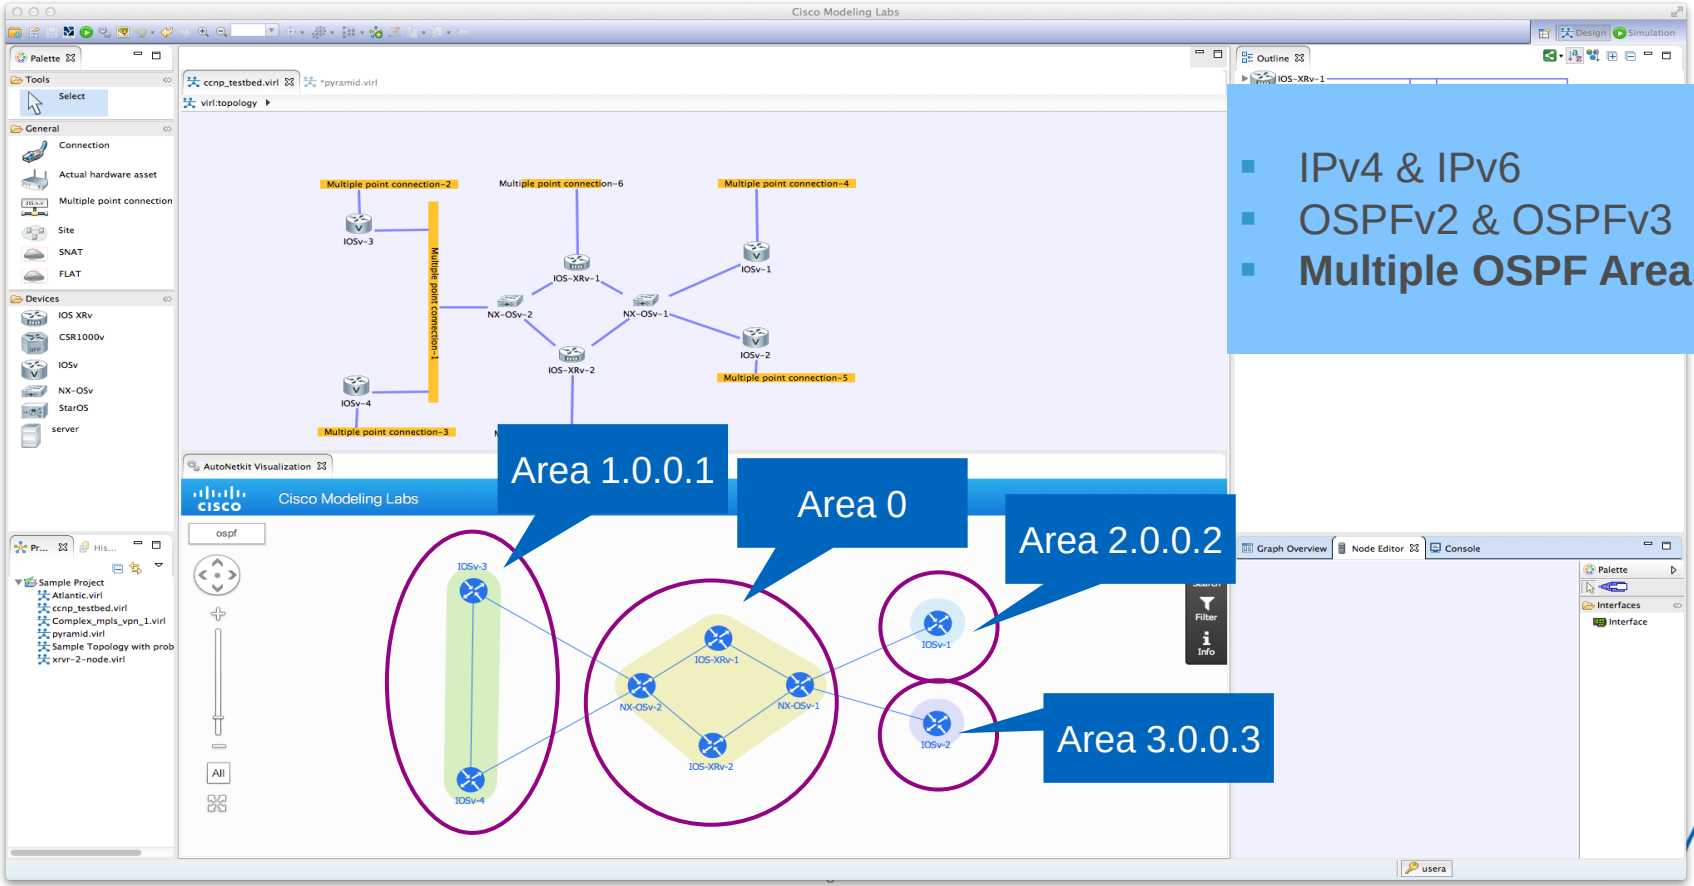

- IPv4 & IPv6
- OSPFv2 & OSPFv3
- **Multiple OSPF Areas**

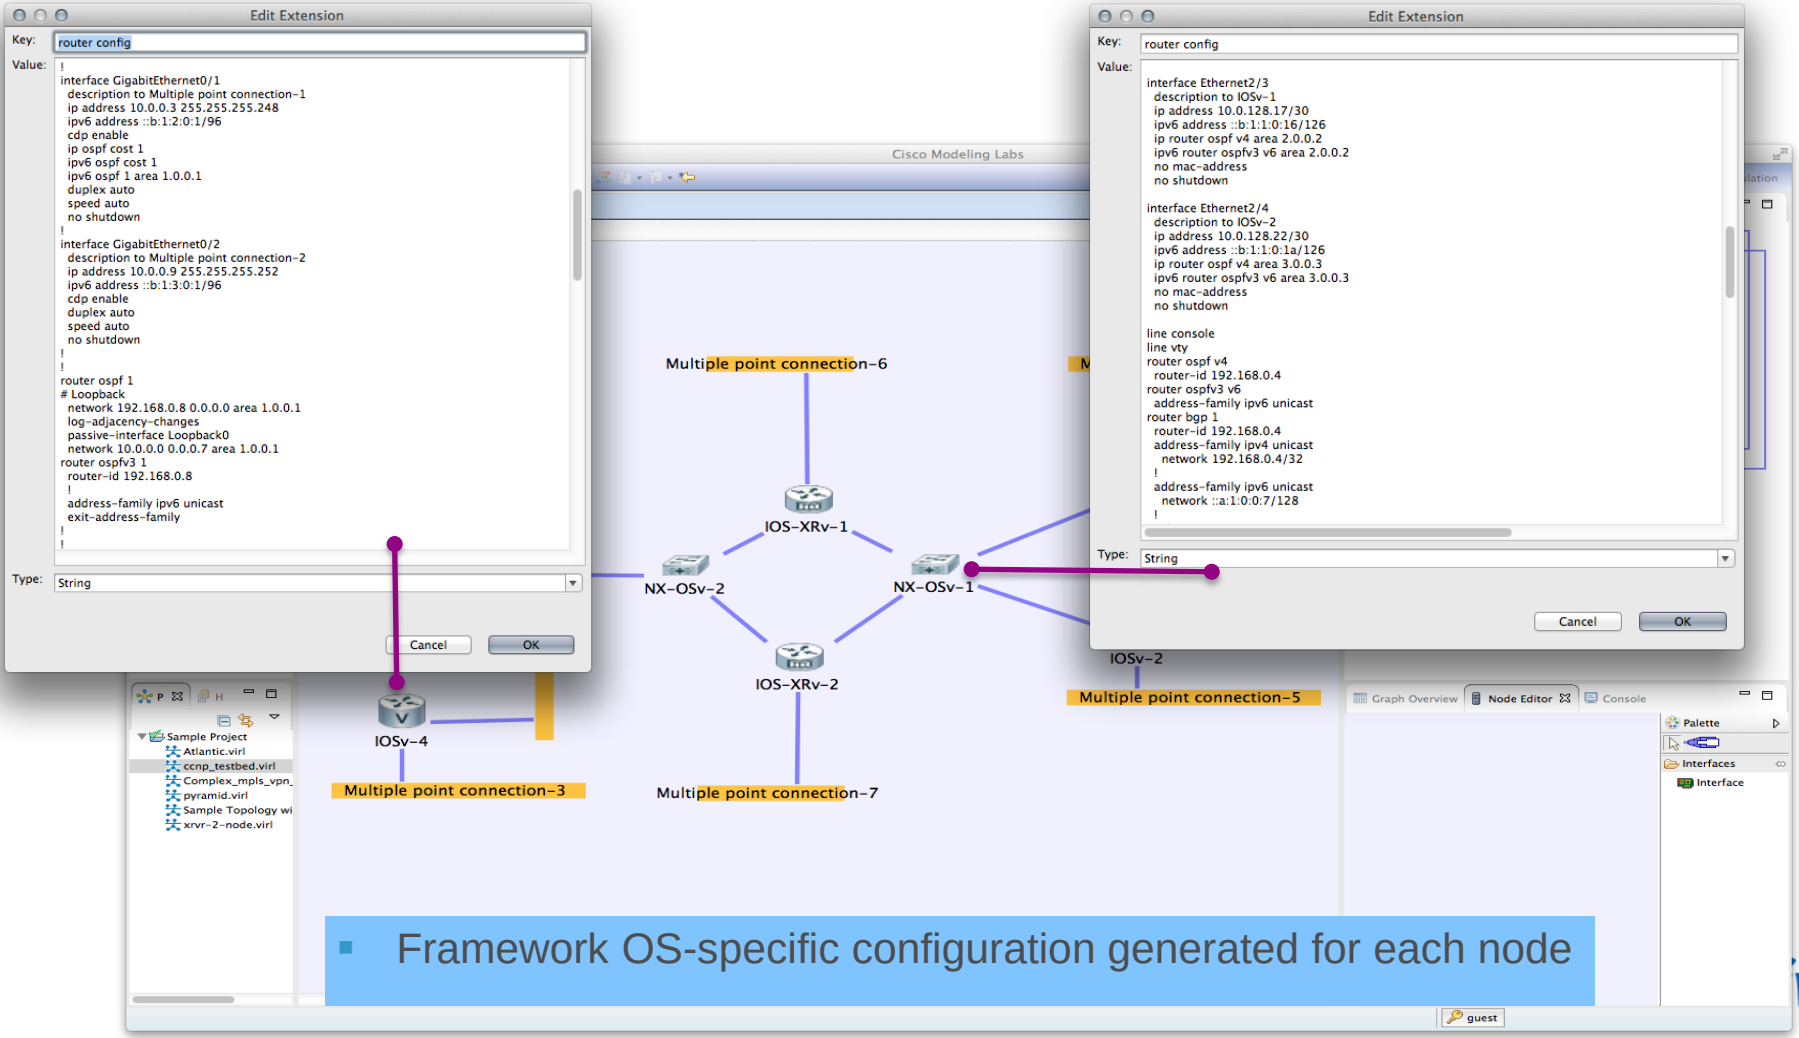

- Framework OS-specific configuration generated for each node

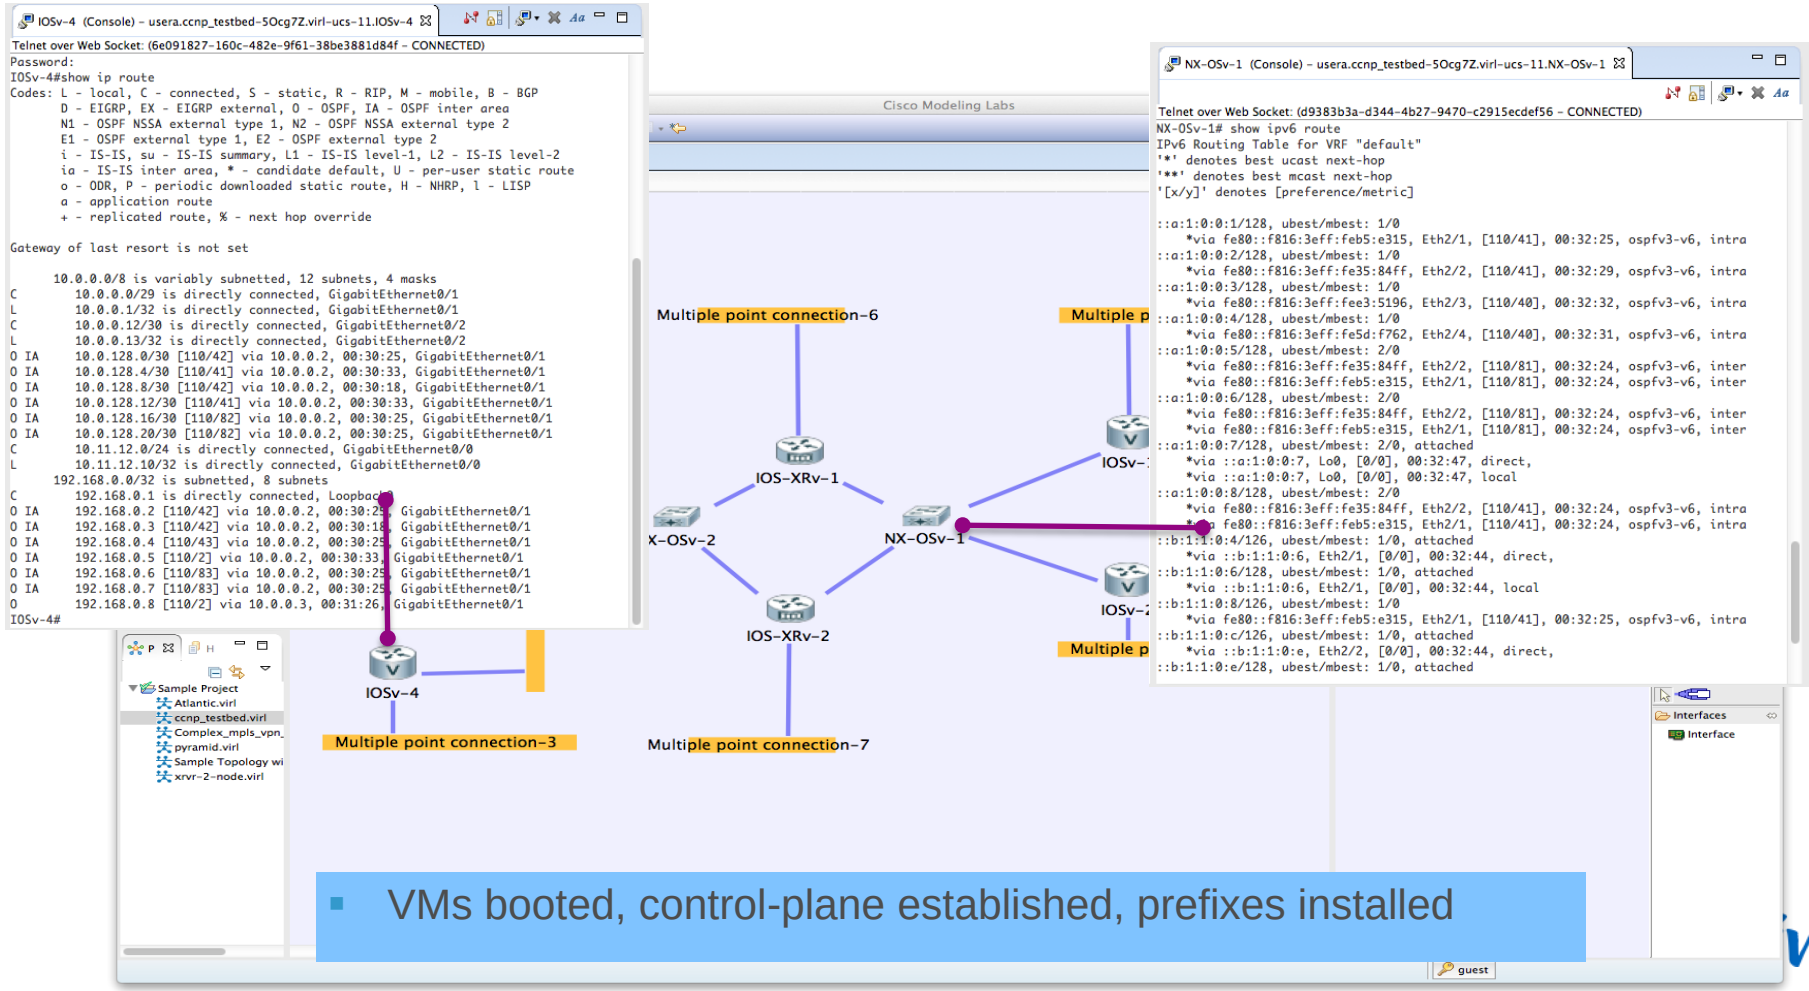

- VMs booted, control-plane established, prefixes installed

Study Topology

- NX-OSv-1 node shutdown
- IOS-XRv-1 sees OSPF adjacency terminate

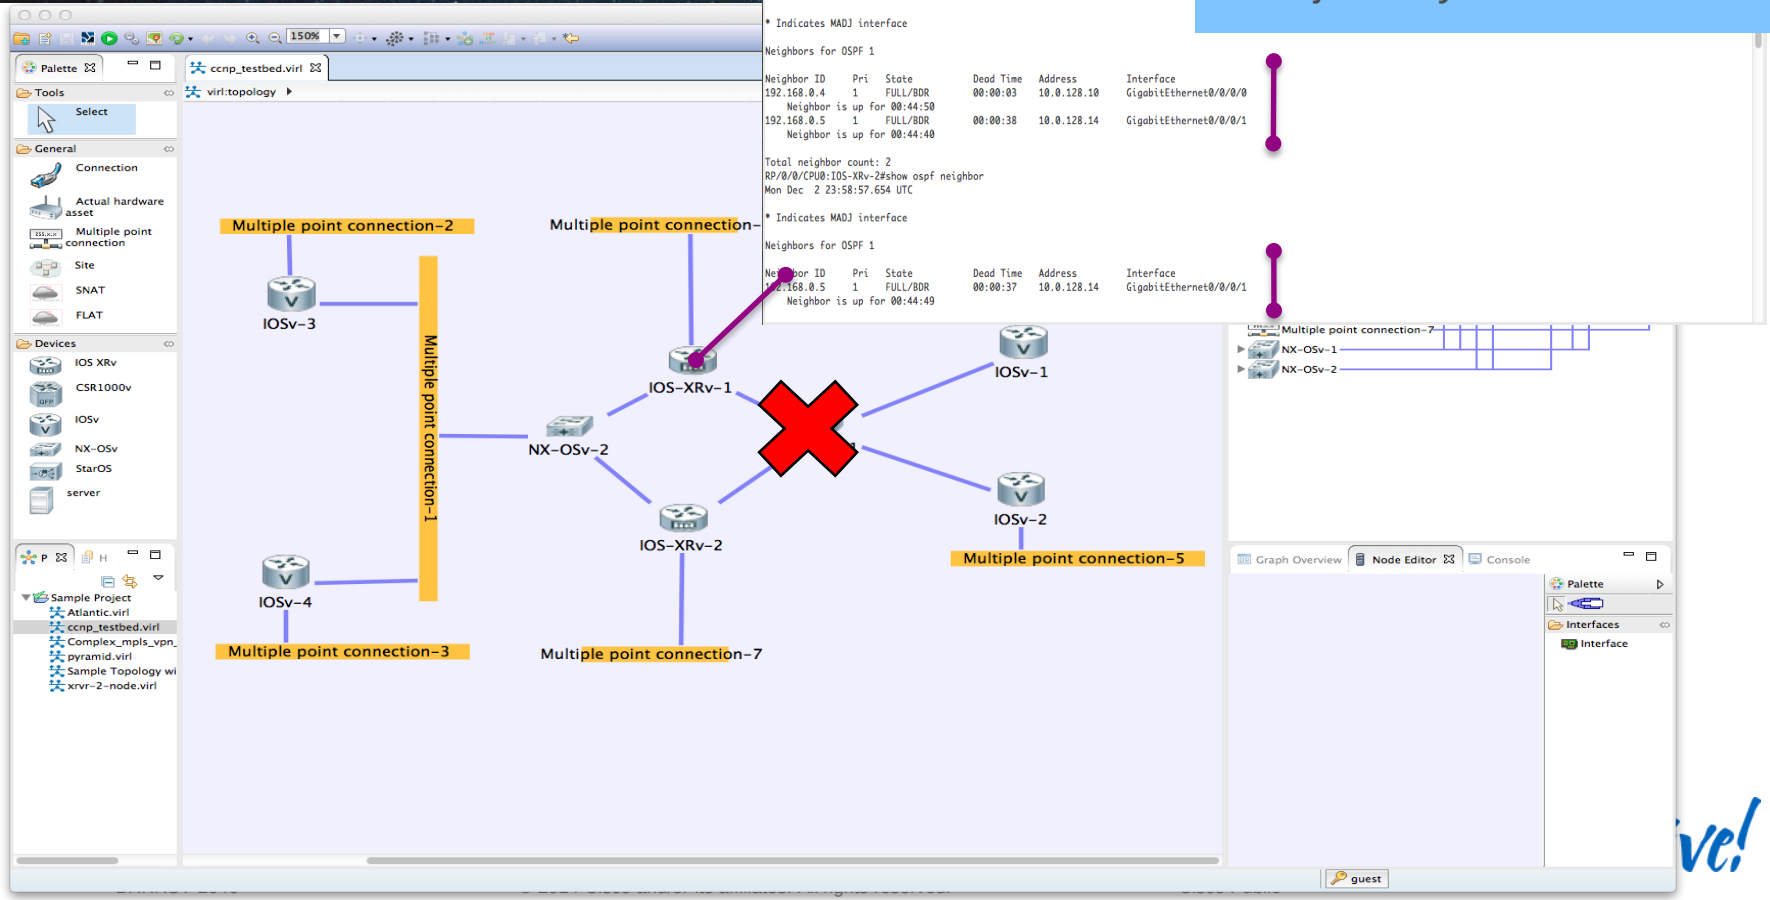

Study Topology

- Reconfigure for EIGRP

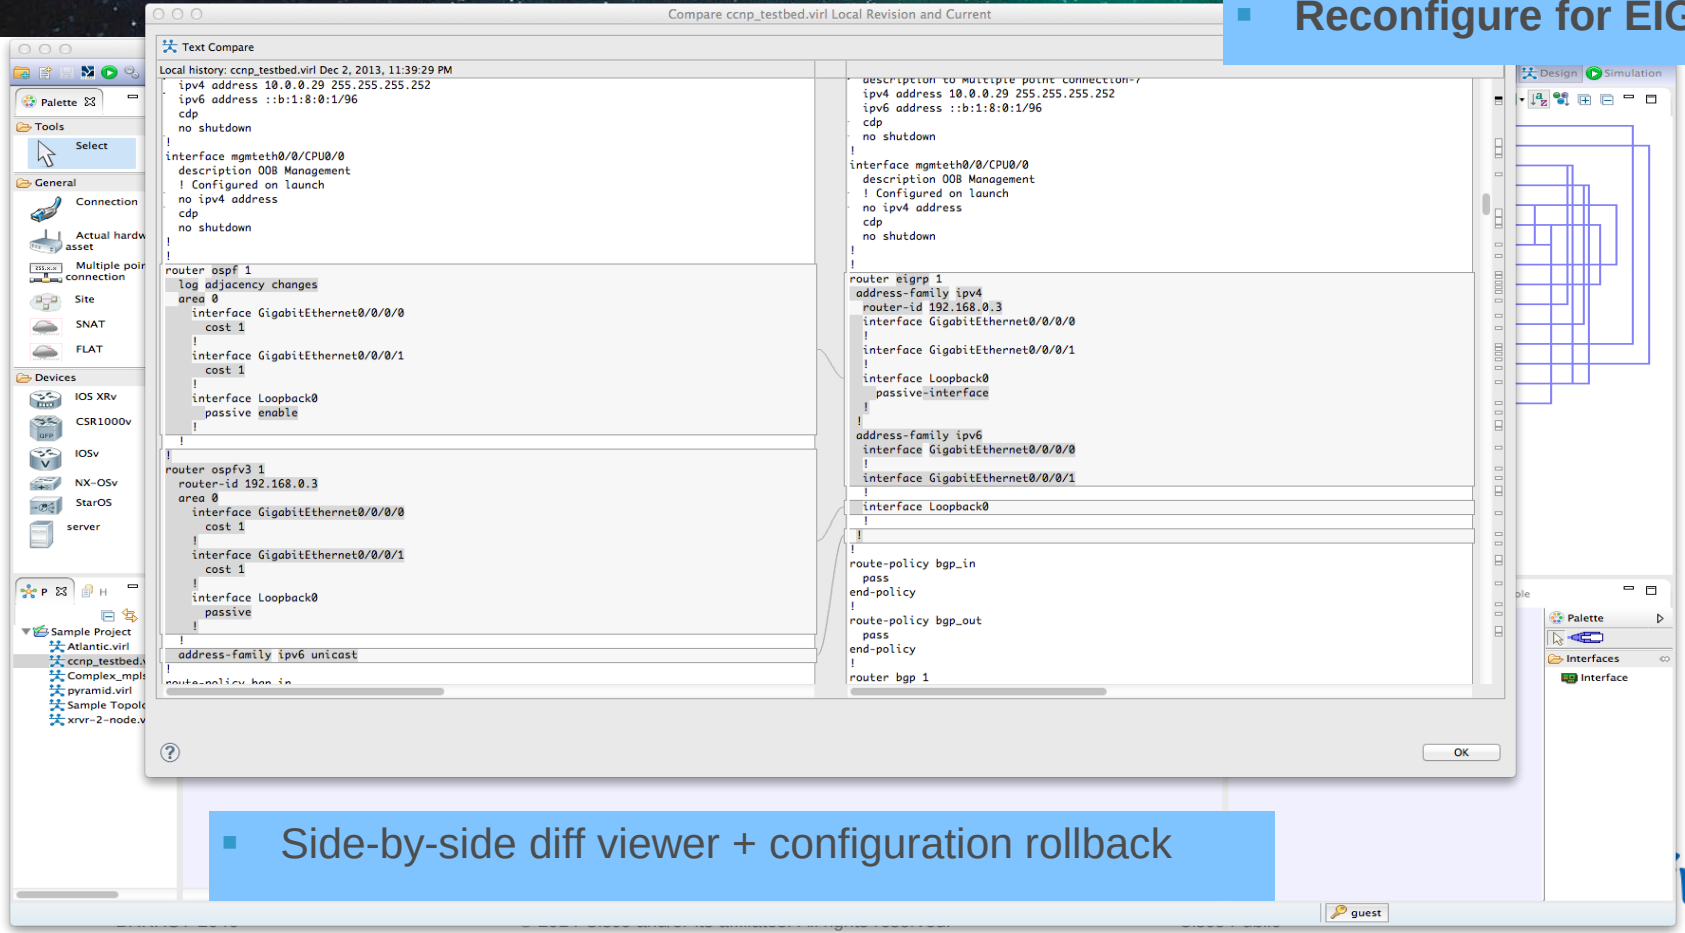

- Side-by-side diff viewer + configuration rollback

Topology Building – Physical to Virtual

Physical to Virtual

Networking functionality in Cisco Modelling Labs enables connectivity from Virtual Machines to the physical world

- Option 1 – Out-of-band management network connecting all VMs
- Option 2 – Connections into specific data interfaces on designated VMs

- .1q provides ‘trunking’ between VMs and physical devices…trunk to physical switch and ‘break-out’ to individual physical machines

Enables ‘hybrid’ test environments – 1 router, 99 VMs!

Connecting Physical to Virtual

- EIGRP enabled
- IPv4 only
- Two AS's – BGP between border peers

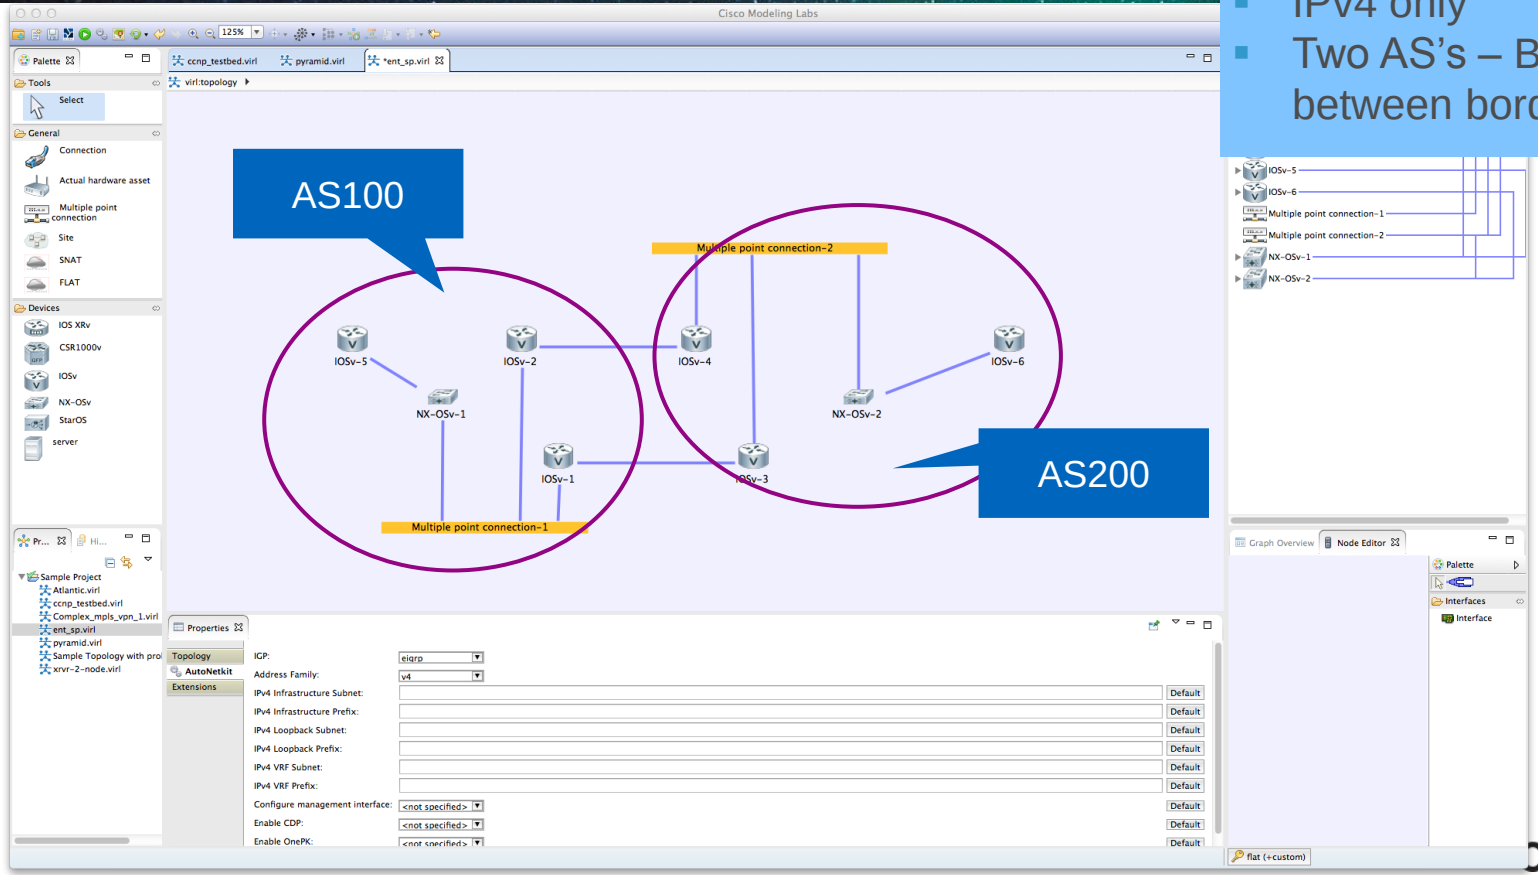

Connecting Physical to Virtual

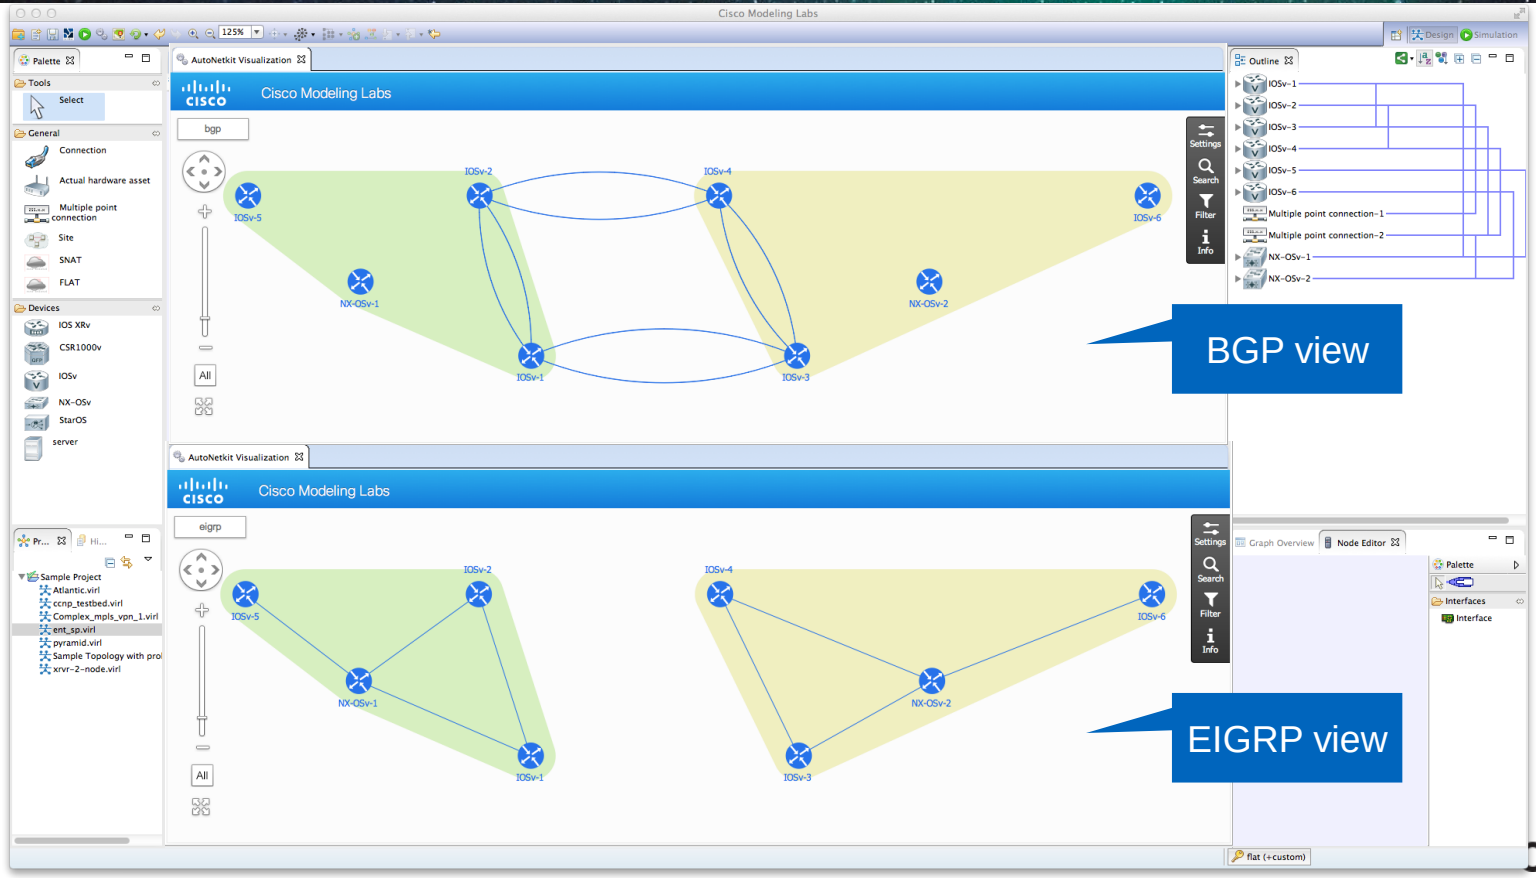

Connecting Physical to Virtual

- External physical Linux host connected to OOB mgmt network
- CACTI Graphing software installed
- Monitoring network as if nodes were ‘real’ routers

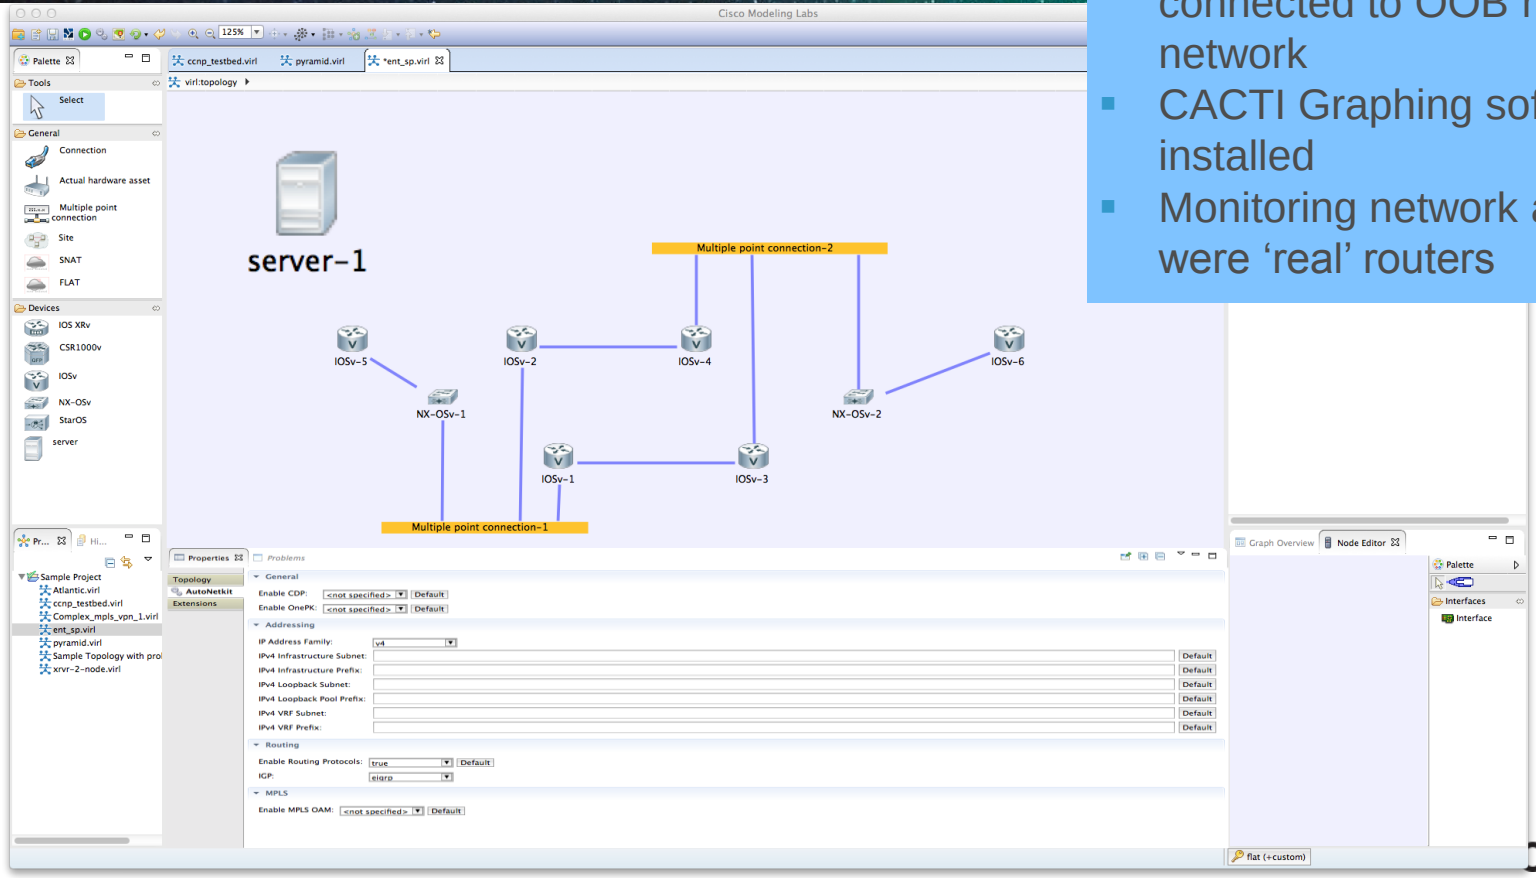

Connecting Physical to Virtual

- Graphing tool collecting stats from Cisco VMs, just as if they were physical devices

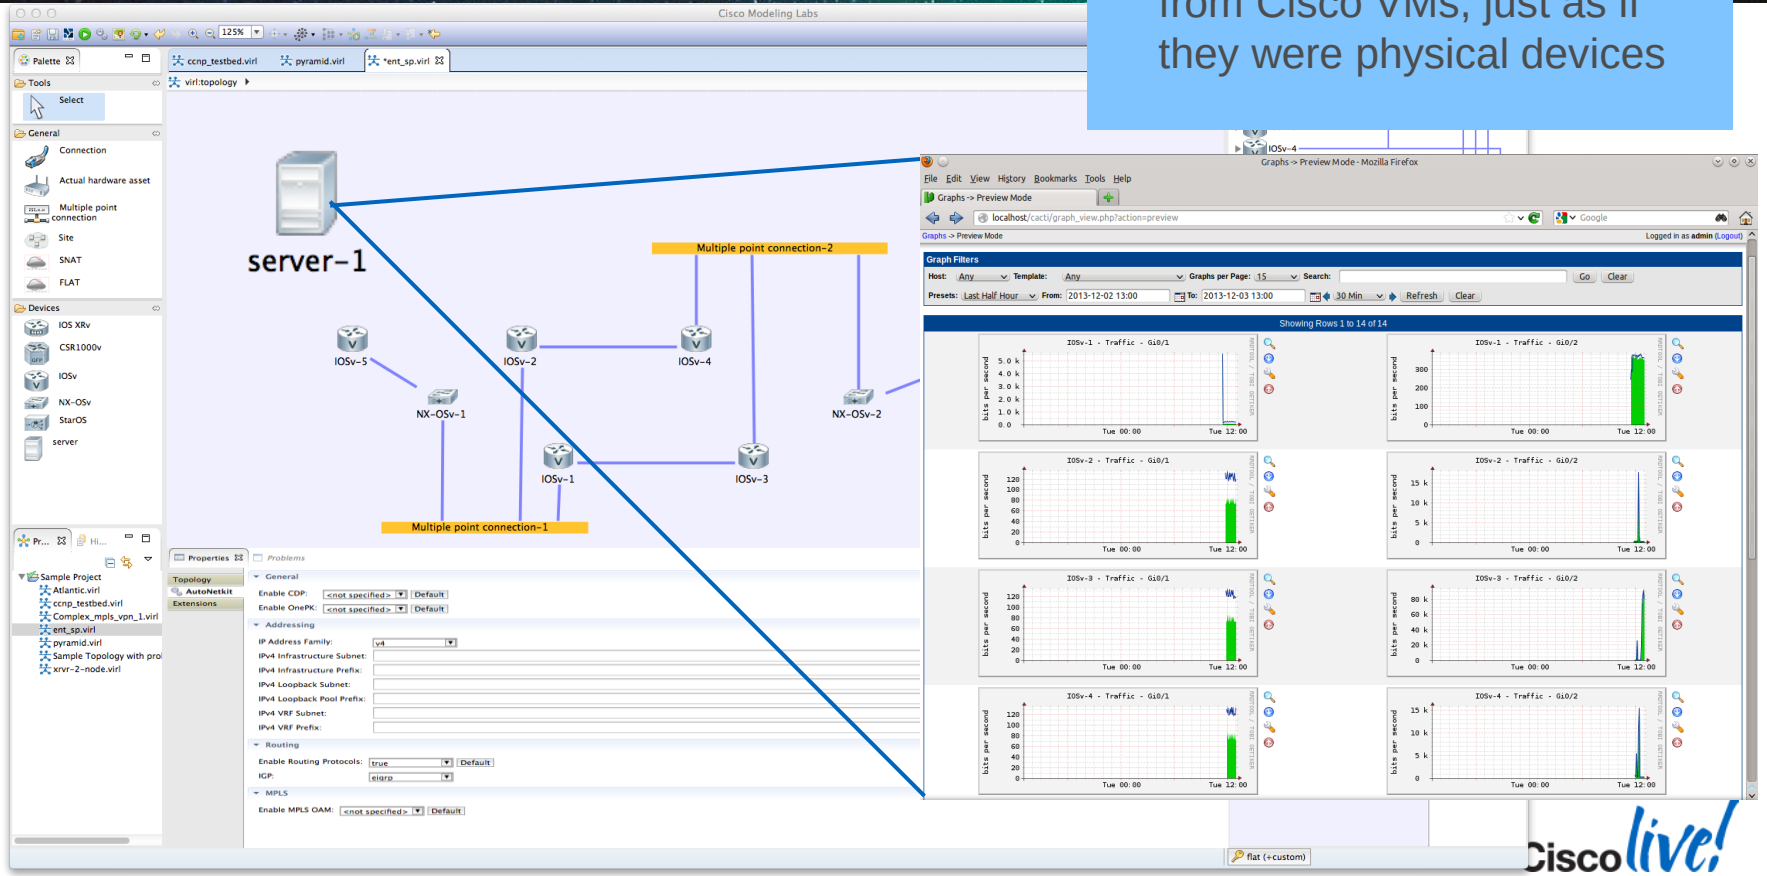

Topology Building – Cariden MATE Design Import

Cariden MATE Design Import

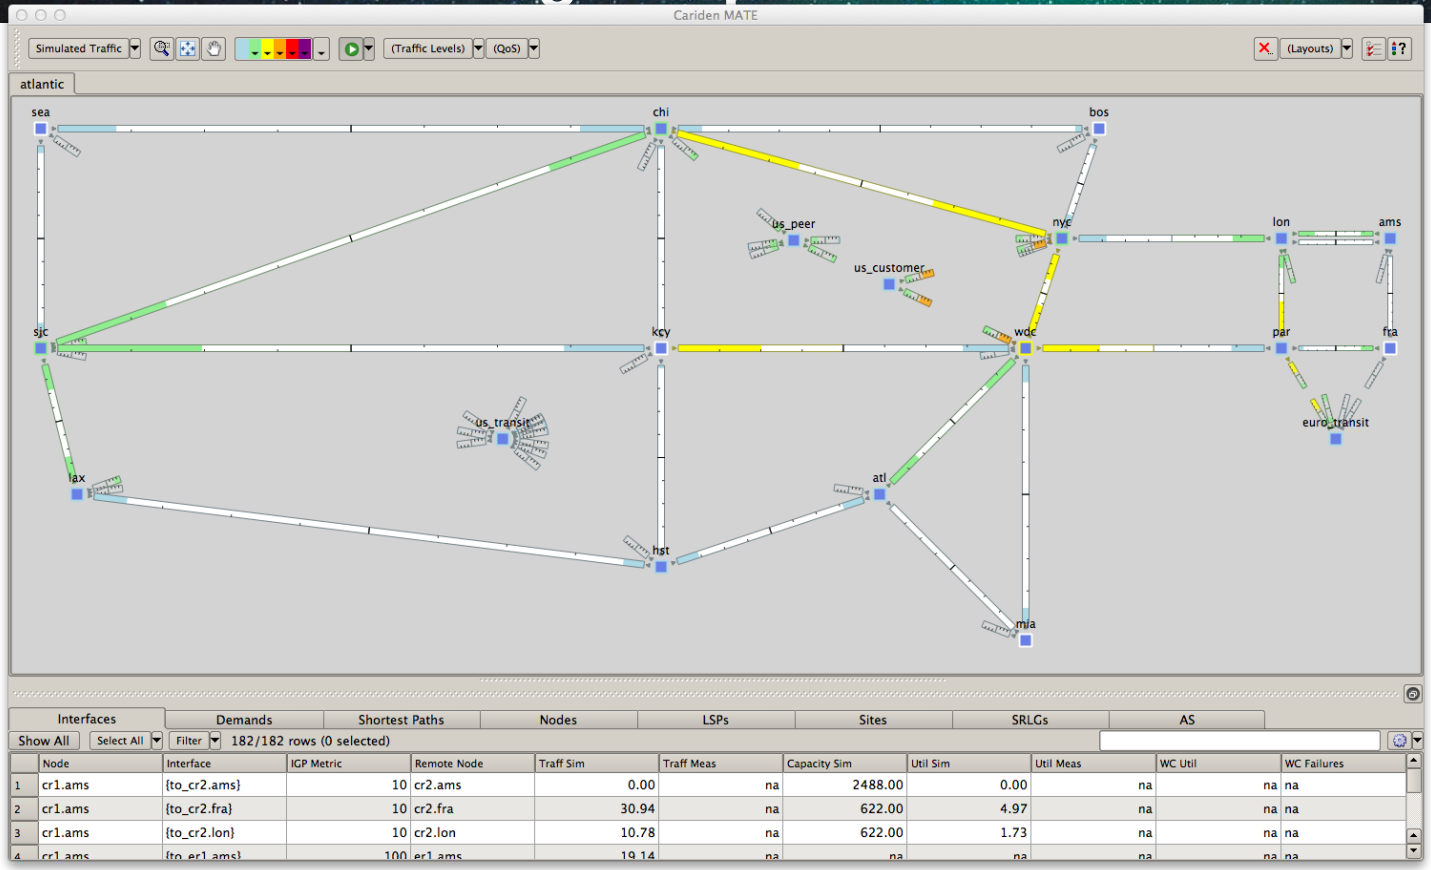

	Node	Interface	IGP Metric	Remote Node	Traff Sim	Traff Meas	Capacity Sim	Util Sim	Util Meas	WC Util	WC Failures
1	cr1.ams	{to_cr2.ams}	10	cr2.ams	0.00	na	2488.00	0.00	na	na	na
2	cr1.ams	{to_cr2.fra}	10	cr2.fra	30.94	na	622.00	4.97	na	na	na
3	cr1.ams	{to_cr2.lon}	10	cr2.lon	10.78	na	622.00	1.73	na	na	na
4	cr1.ams	{to_er1.ams}	100	er1.ams	19.14	na	na	na	na	na	na

Into CML Workbench

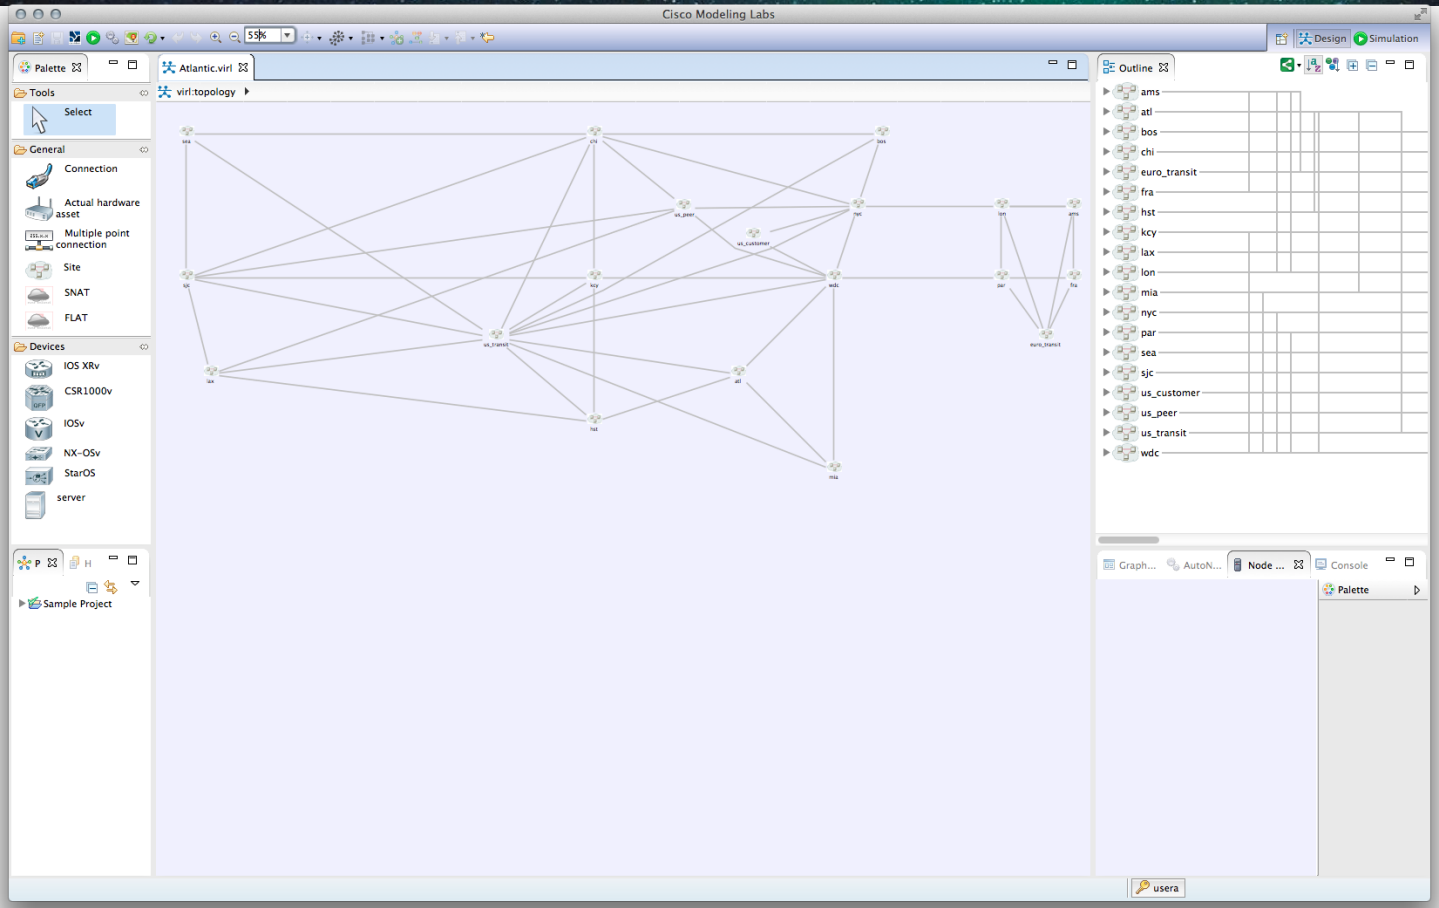

Cisco live!

To the AutoNetkit Visualiser

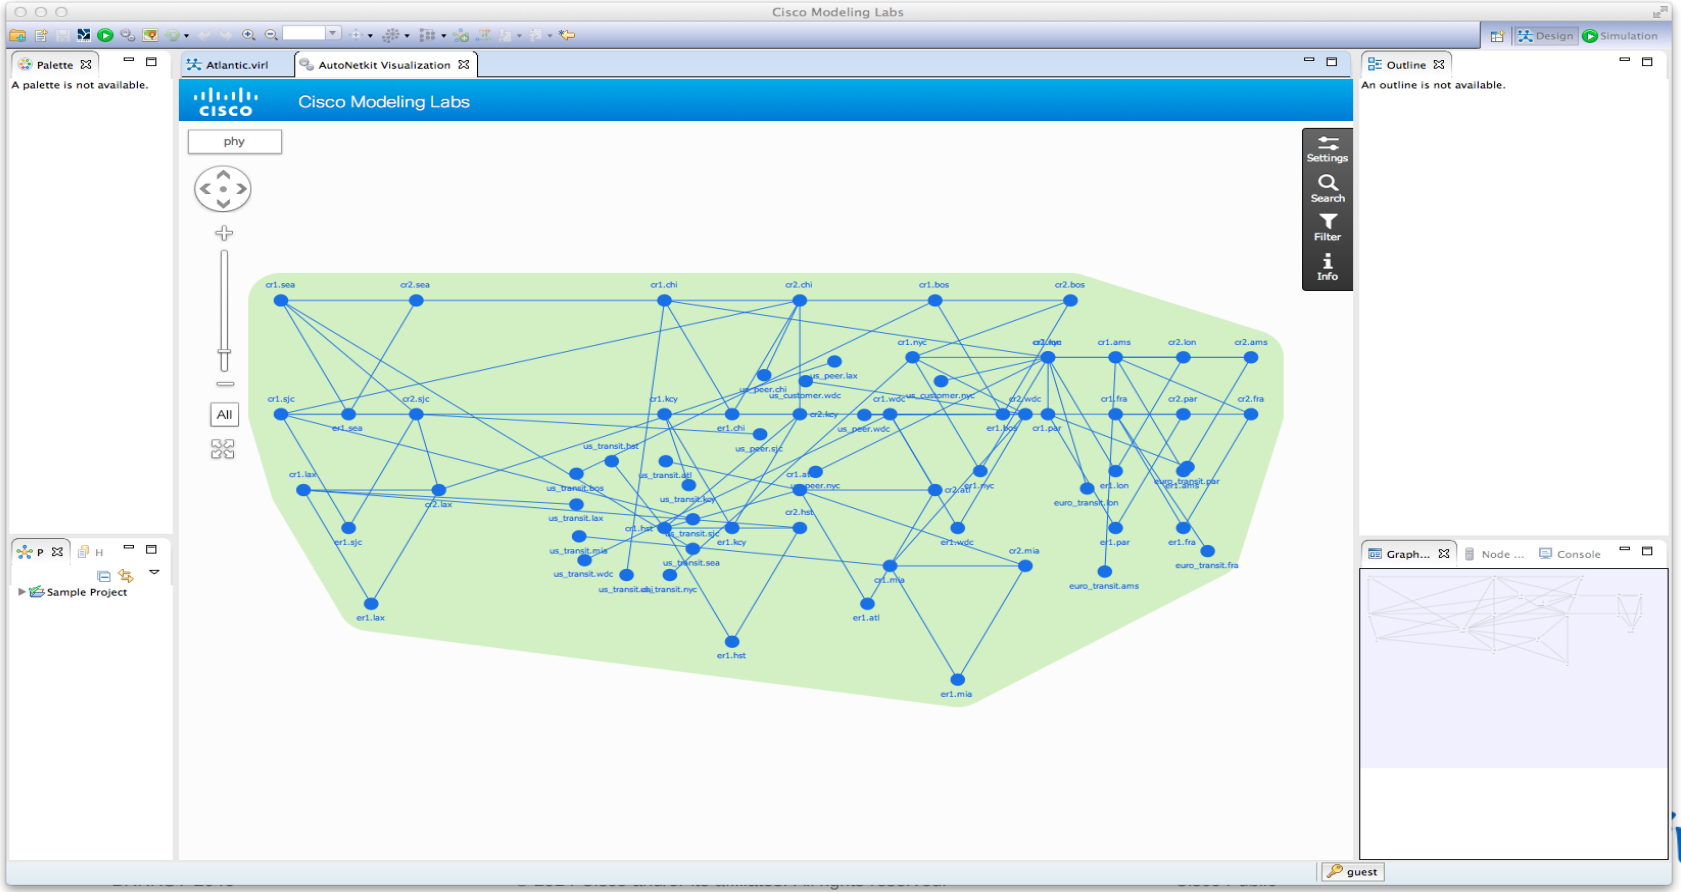

Cisco Modelling Labs in Action…Live Demo




Enterprise Case Study – Introduction of MPLS Traffic Engineering

Project Outline

ENT-B planned to deploy new CRS-1 routers and enable MPLS Traffic Engineering using Fast Re-route for the first time, controlling traffic flows from their edge locations to their Data-Centre

The need –

- Understand the configuration changes required to enable MPLS TE/FRR on IOS and IOS XR devices
- Implement Class-Based Tunnel selection profiles in IOS
- Observe changes to traffic flows under failure conditions
- Ensure zero (little) impact to traffic flows under failure conditions

The challenge –

- No lab equipment! Typically use the live network to test with…

Topology Requirements

- Full representation of network core and data-centre edge
  - 8 CRS-1
  - 17 IOS routers
  - 2 NX-OS router/switches
  - 55 interconnects

- Simulate traffic feeds from edge locations to data-centre to enable flow engineering

Simulation Topology

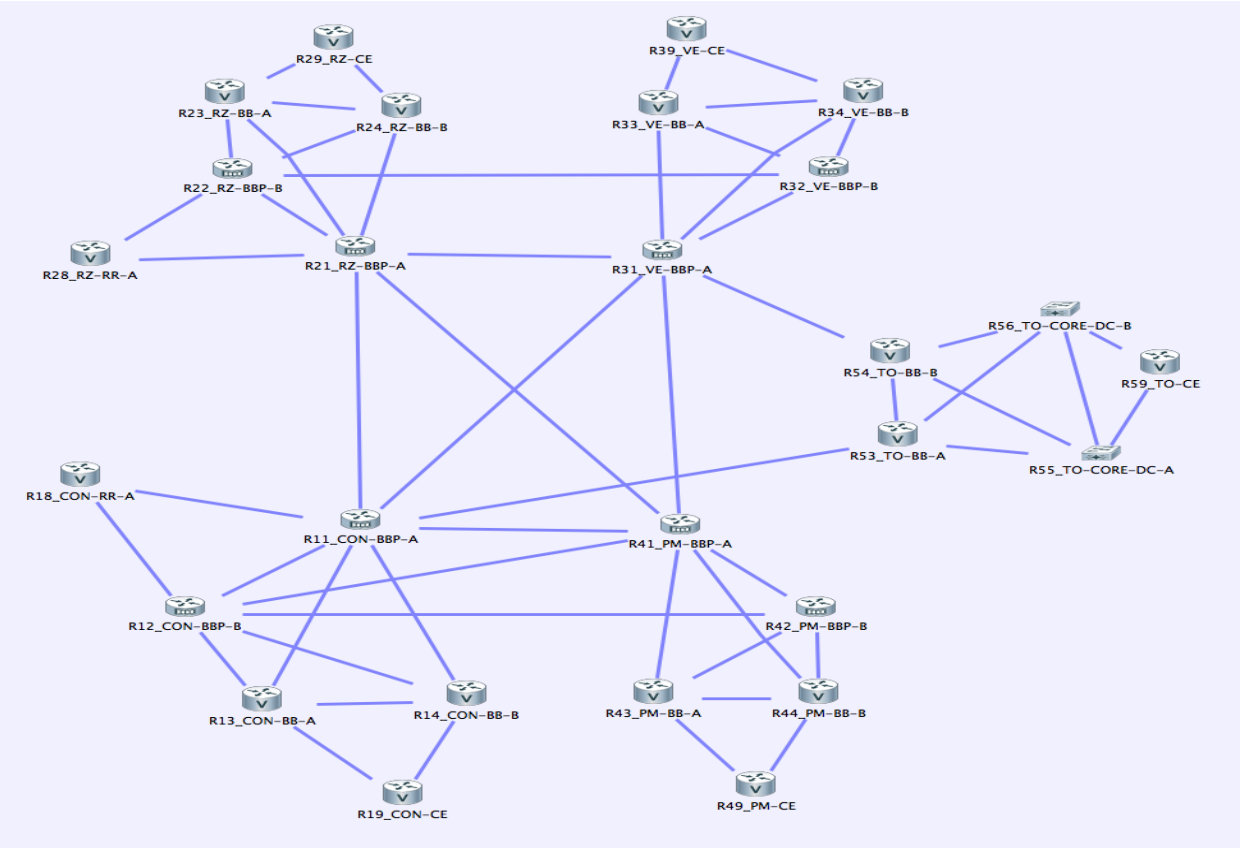

Benefits Brought to this Project

- Capability to simulate the production network has previously been impossible for the customer
- The simulation greatly increased the accuracy of configurations and design ahead of deployment with detailed analysis of MPLS TE/FRR implementation
- Operations team were able to spend time ‘hand’s on’, improving their confidence in the operation and behaviour prior to deployment
- “Surprises” in simulation are good, “surprises” in the live network…not so good

ISP Case Study – BGP Re-architecture

Project Outline

ISP-A planned to deploy new Peering & Transit nodes on their Internet Edge to replace legacy systems

The need –

- Perform the migration steps in the lab
- Observe changes to BGP routing from key locations within their network
- Understand affect on traffic flows between legacy PE and next-gen PE during the migration of BGP sessions
- Ensure no impact to customers

The challenge –

- Very limited lab equipment!

Topology Requirements

- Full representation of network core and Internet edge (12 CRS-3 and ASR9000-series routers running IOS XR)
- Use production BGP feeds
- Use production Route Policy Language configuration (over 12k lines)

High-Level Diagram

Simulation to highlight expected changes to the network-wide BGP table during this period, and to model traffic flows over PT1 <-> PT2 interconnect.

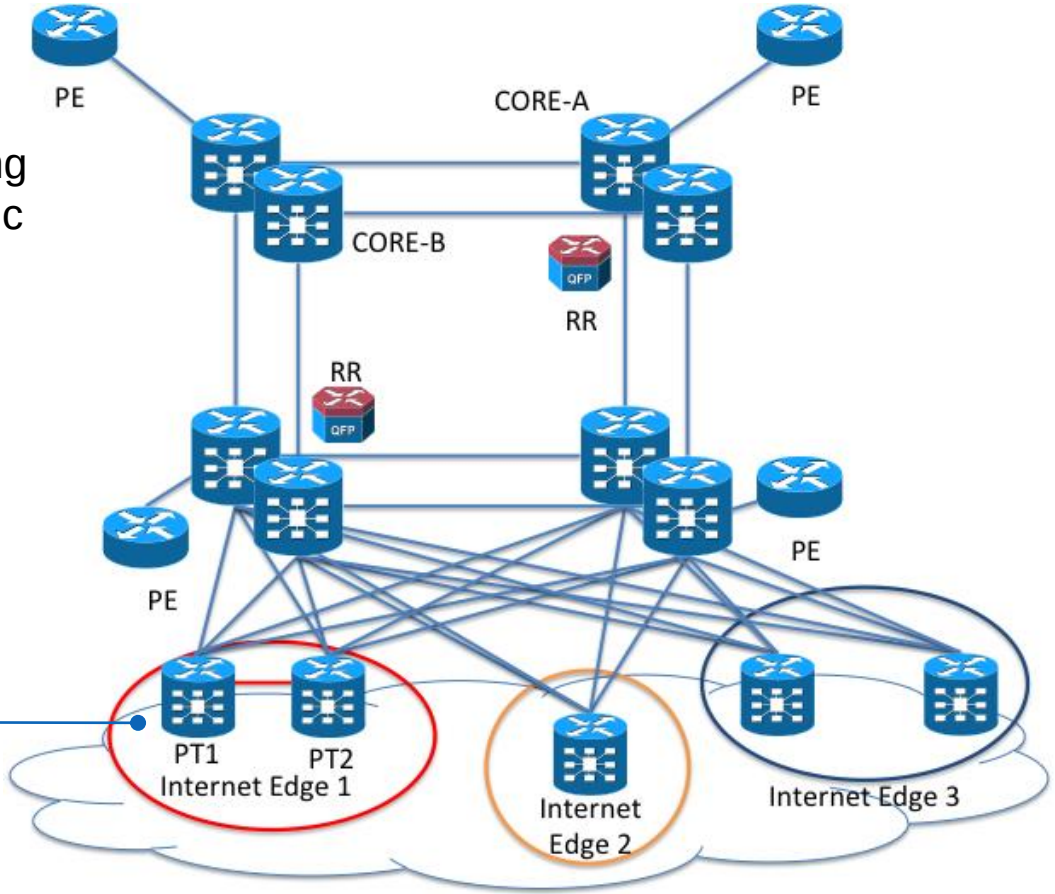

- BGP sessions to be moved from PT1 to PT2 over a 1 month period
- During transition both PT1 and PT2 will be the primary exit point

Simulation Topology

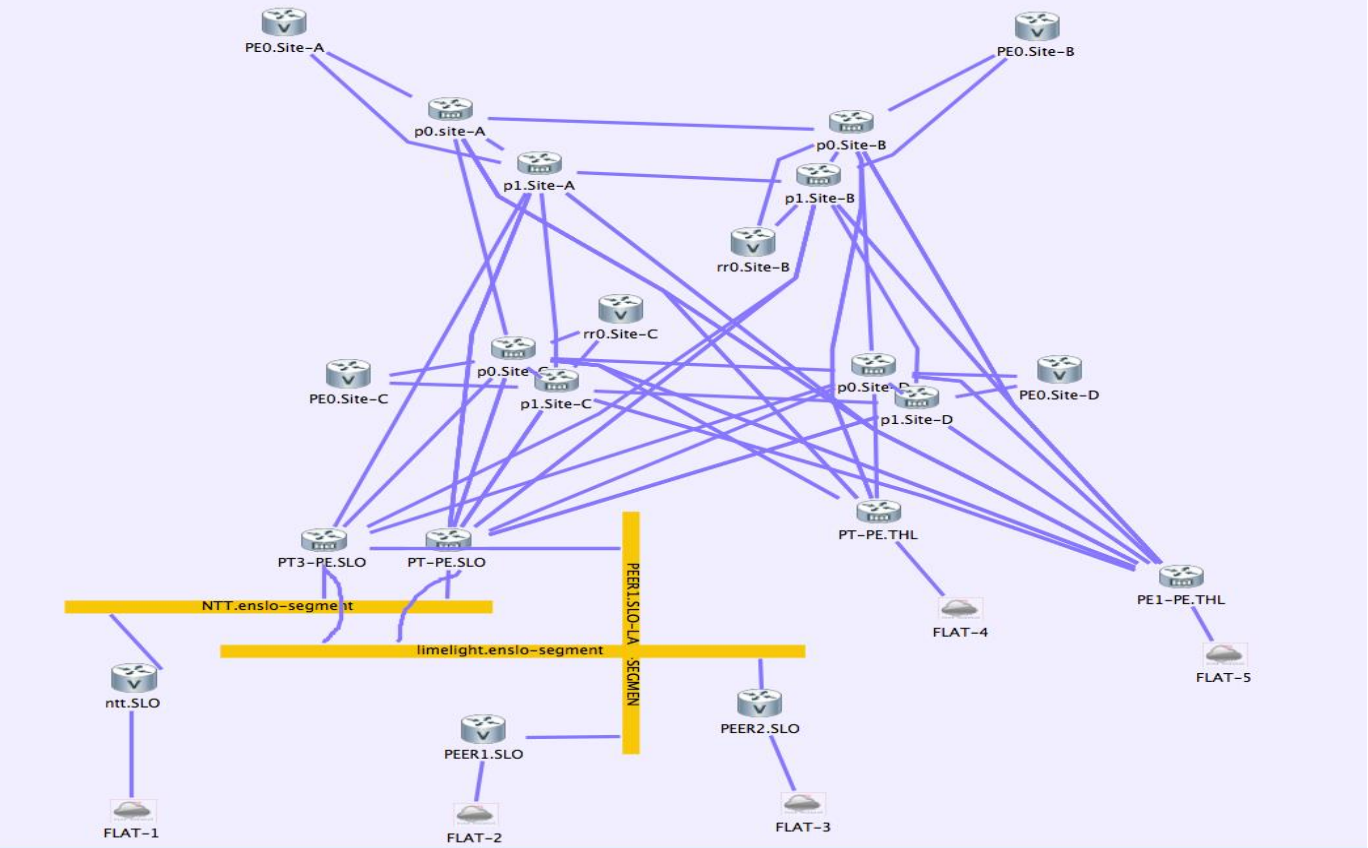

Highlights

- 1 hour to build topology foundations
  - 12 IOS XRv Nodes
  - 4 IOSv Nodes
  - 2 CSR1000v Nodes
  - 80 interconnects

- Traffic Generator connected into topology via network, replaying production BGP feeds (470k BGP prefixes, 700k paths)

- 1 Day to integrate production configuration into VIRL running configurations (inc 12,000 lines of RPL)

BGP Feed to Route Reflector

ACTIVE – guest.session-ygmEMb.virl.PT-PE1.THL (Cons | ACTIVE – guest.session-ygmEMb.virl.TRANSIT.SLO (Co | ACTIVE – guest.session-ygmEMb.virl.rr0.Site-B (Consol

Telnet: (10.52.212.251:17001 – CONNECTED)

```
**************************************************************************
RR0.SITE-B>en
Password:
RR0.SITE-B#sh ip bgp
BGP table version is 461004, local router ID is [redacted]
     Network          Next Hop            Metric LocPrf Weight Path
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 1.0.0.0/24       [redacted]128.219       15000     100      0 [redacted]419 15169 i
 *>i 1.0.4.0/24       [redacted]128.219       15000     100      0 [redacted]419 6939 7545 56203 i
 *>i 1.0.5.0/24       [redacted]128.219       15000     100      0 [redacted]419 6939 7545 56203 i
 *>i 1.0.6.0/24       [redacted]128.219       15000     100      0 [redacted]419 6939 4826 38803 56203 i
 *>i 1.0.7.0/24       [redacted]128.219       15000     100      0 [redacted]419 6939 4826 38803 56203 i
 *>i 1.0.20.0/23      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 *>i 1.0.22.0/23      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 *>i 1.0.24.0/23      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
     Network          [redacted]op            Metric LocPrf Weight [redacted]
 *>i 1.0.26.0/23      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 *>i 1.0.28.0/22      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 *>i 1.0.64.0/18      [redacted]128.219       15000     100      0 [redacted]419 6939 4725 7670 7670 7670 18144 i
 *>i 1.0.128.0/18     [redacted]128.219       15000     100      0 [redacted]419 3356 38040 9737 9737 i
 *>i 1.0.128.0/17     [redacted]128.219       15000     100      0 [redacted]419 38040 9737 9737 i
 *>i 1.0.160.0/19     [redacted]128.219       15000     100      0 [redacted]419 3257 2914 38040 9737 9737 i
 *>i 1.0.192.0/19     [redacted]128.219       15000     100      0 [redacted]419 38040 9737 9737 i
 *>i 1.0.192.0/18     [redacted]128.219       15000     100      0 [redacted]419 3257 6762 38040 9737 9737 i
 *>i 1.0.224.0/19     [redacted]128.219       15000     100      0 [redacted]419 3356 38040 9737 9737 i
 *>i 1.1.1.0/24       [redacted]128.219       15000     100      0 [redacted]19 15169 i
 *>i 1.1.64.0/19      [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 *>i 1.1.127.0/24     [redacted]128.219       15000     100      0 [redacted]419 3257 2516 2519 i
 --More--
```

BGP Prefix Scaling

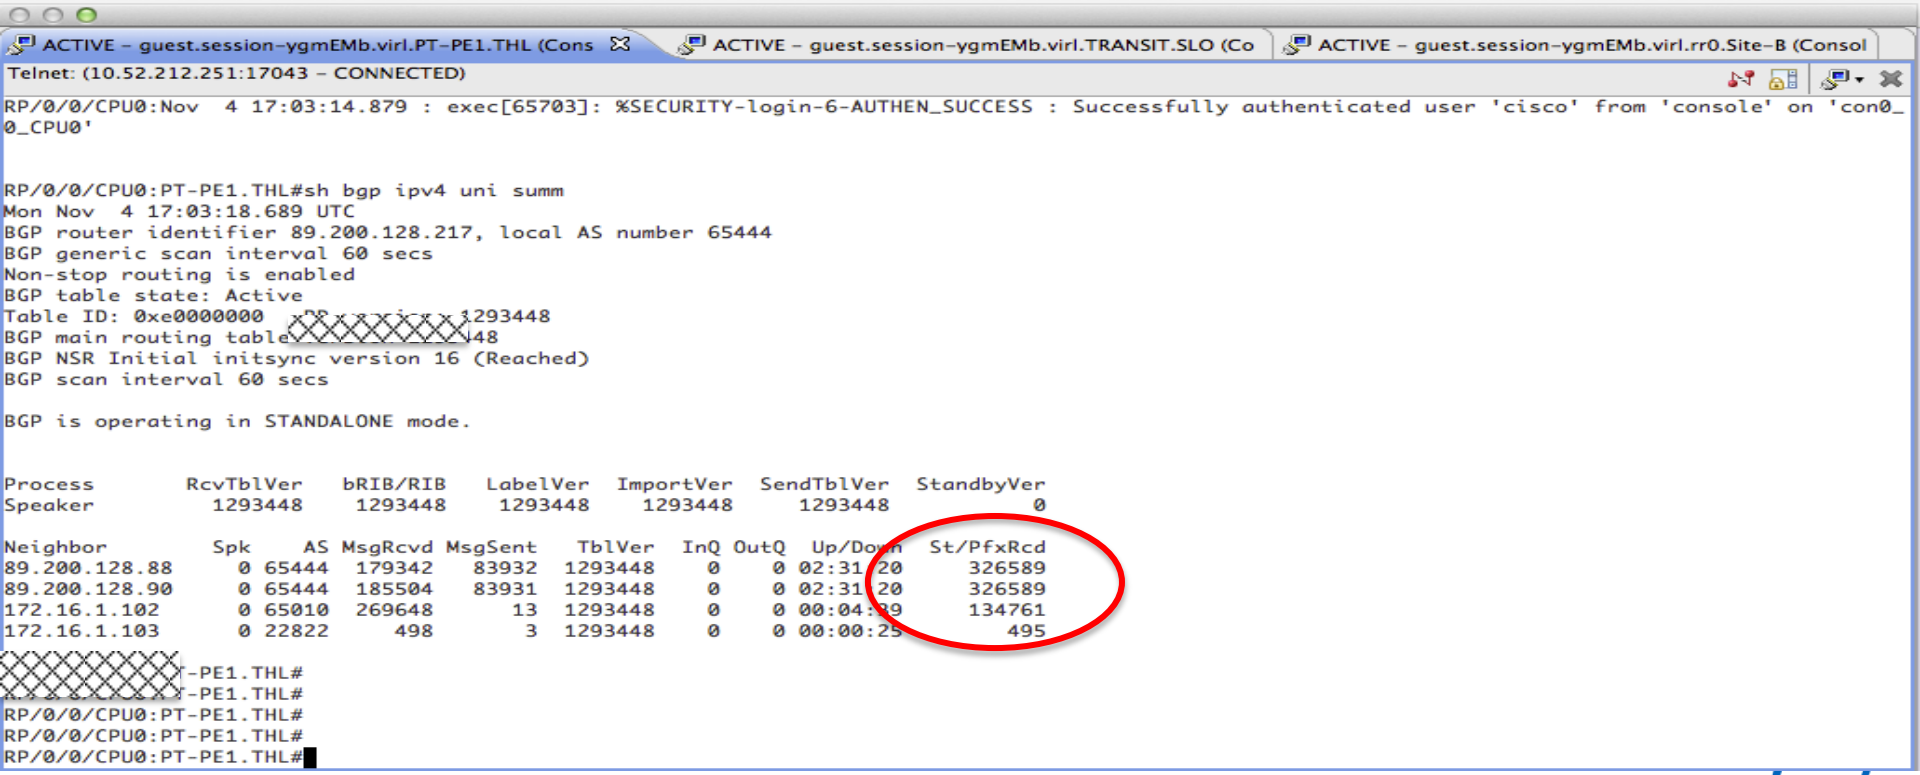

Using IOS XRv to Validate Configurations

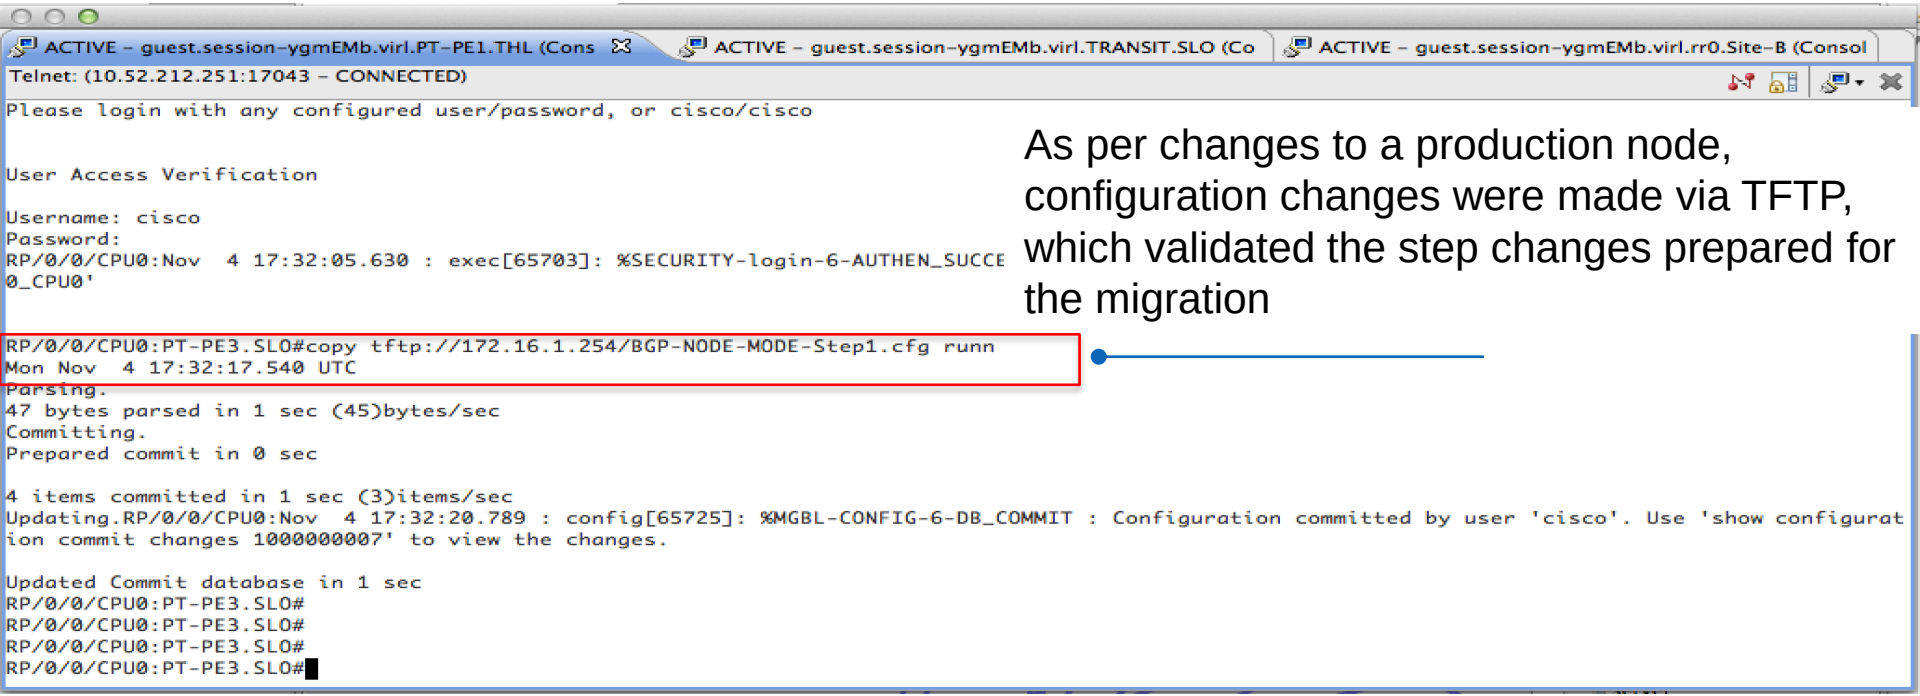

As per changes to a production node, configuration changes were made via TFTP, which validated the step changes prepared for the migration

Benefits Brought to this Project

- Far more comprehensive topology than their lab infrastructure could have provided
- The use of production IP addressing highlighted a number of key path changes
- The virtual topology was dedicated for this purpose, which allowed for complete control over the prefixes / protocols in use within the network
- The time to build the base configurations was drastically reduced compared with the time to build a physical alternative (combined with zero tear-down time – typically 0.5 days saved).

- Topology provides a robust foundation for future projects for the same customer network

Network Virtualisation – Because Breaking the Simulator is Free!

Cisco Modelling Labs

Many control plane manipulation/monitoring/management features in Cisco Network operating systems

What better place to mix them all up and set your network ablaze?

- For Training
- For Proof of Concept
- For configuration verification
- For bug-fix verification




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com


CISCO
TM