DYNAMIC NETWORK MODELING

Joel Obstfeld – jobstfeld@juniper.net
Director, SW Engineering
CTO Team, Platform Systems Division

NETWORK MODELING OVERVIEW

NETWORK MODELING OBJECTIVES

Data networks are dynamic entities

- Networks ‘evolve’ over time
- Introduction of new services and technologies
- New hardware triggered changes to original designs

Networks may initially reflect the planned architecture but tend to diverge

Network modeling enables us to

- Better understand the network and to be able test various assumptions
- Investigate new designs, services, impact of new technologies

Sometimes we’re even surprised by what we find

WHY MODEL NETWORKS?

Network Capacity planning amidst traffic uncertainty
- Sensitivity to traffic changes, projected growth
- Identify where to most effectively add capacity

Disaster planning and recovery
- Single failures are ‘simple’
- Multiple failures introduce much more complexity, unforeseen behaviours

Optimizations
- Benefits (& risks) of merging or evolving current networks to new designs
- Engineering the network for services
- Route-reflector, Content-Cache placement efficiency

Quantifying the benefits of technology choices
- The efficiency of Statistical multiplexing vs. Circuit switching
- Costs associated with the choices

Multi-layer Topologies – interaction of Layer-1 and Layer-3 environments
- Fibre routing, SLRG
- Explicit diverse path planning

Routing Metric & Policy Engineering
- Planning IGP metrics to influence traffic paths vs trying it in the live network
- BGP policy engineering

BUILD A COMPLETE VIRTUAL REPRESENTATION OF THE DYNAMIC NETWORK

Move beyond just nodes & lines on a screen

Applying solutions from virtualization/data-centers to networking modeling

Augment network modeling with virtual machine and virtualized networking tools to go beyond the diagram

- Build high-fidelity models of operational networks
- Demonstrate new service architectures such as Seamless MPLS
- In-depth investigation of the ‘what-if’ scenarios
- Enable operations teams to validate planned changes
- Pre-flight checks before config push
- Replay events from the live network, e.g. BGP routing disasters
- Architecture change phasing plans

FULL VIRTUAL ROUTER APPLIANCES

- Full-featured virtual router running JUNOS – VJX1000
- SW-based Forwarding plane
- Details available from http://www.juniper.net/us/en/products-services/software/junos-platform/junosphere/#overview

- IPv4/IPv6 Unicast / Multicast
- Routing: OSPF, BGP, RIPv2, Static routes, IS-IS,…)
- Multicast: IGMPv3, PIM, SDP, DVRMRP, Source Specific)
- MPLS: Layer 2 VPN (VPLS), Layer 3 VPN, LDP, RSVP
- Encapsulations: Ethernet (MAC and tagged), PPPoE
- NAT/Stateful Firewall Filters/Intrusion Detection
- Tunneling: GRE, IP in IP
- COS
- User Authentication/Access: RADIUS, RSA SecureID, LDAP
- J-Web, CLI

QUINTESSENTIAL CASE STUDY

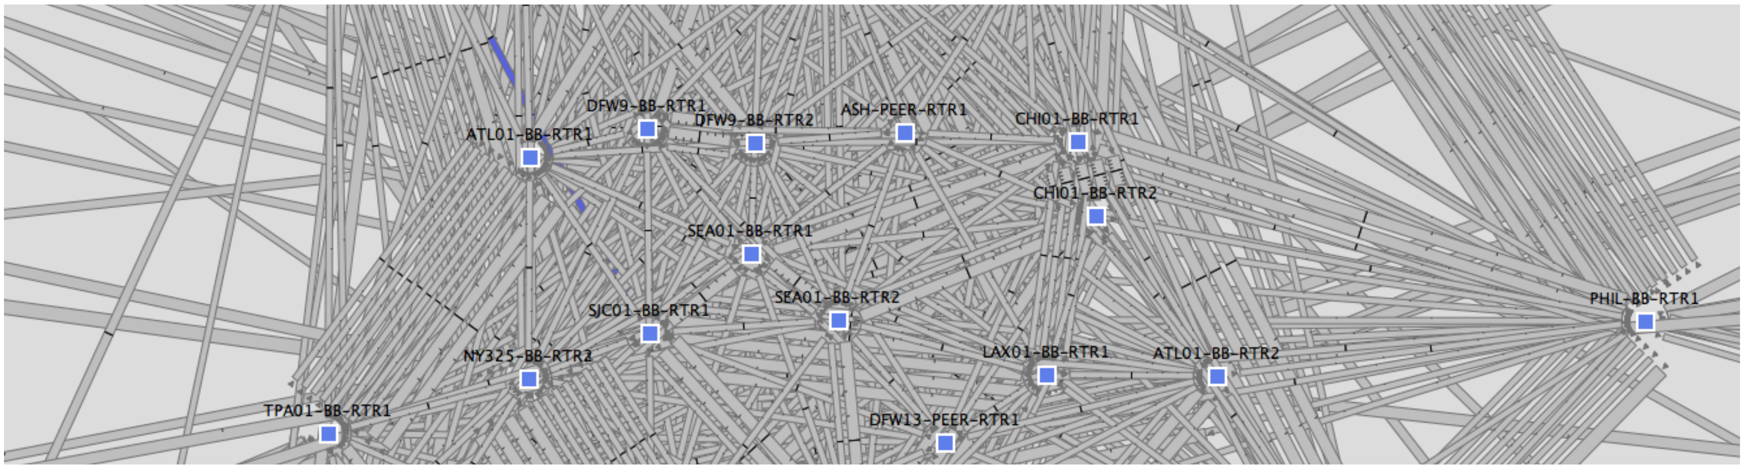

A large Service Provider "A" expands its network via acquisition of another Service Provider "B"

SP A instigates plans to merge networks, save costs, drive innovation

Years later, networks are still distinct, migration still being studied

Why? **Too risky and no way to feasibly simulate a migration**

COMPLEXITY – SIMPLIFIED

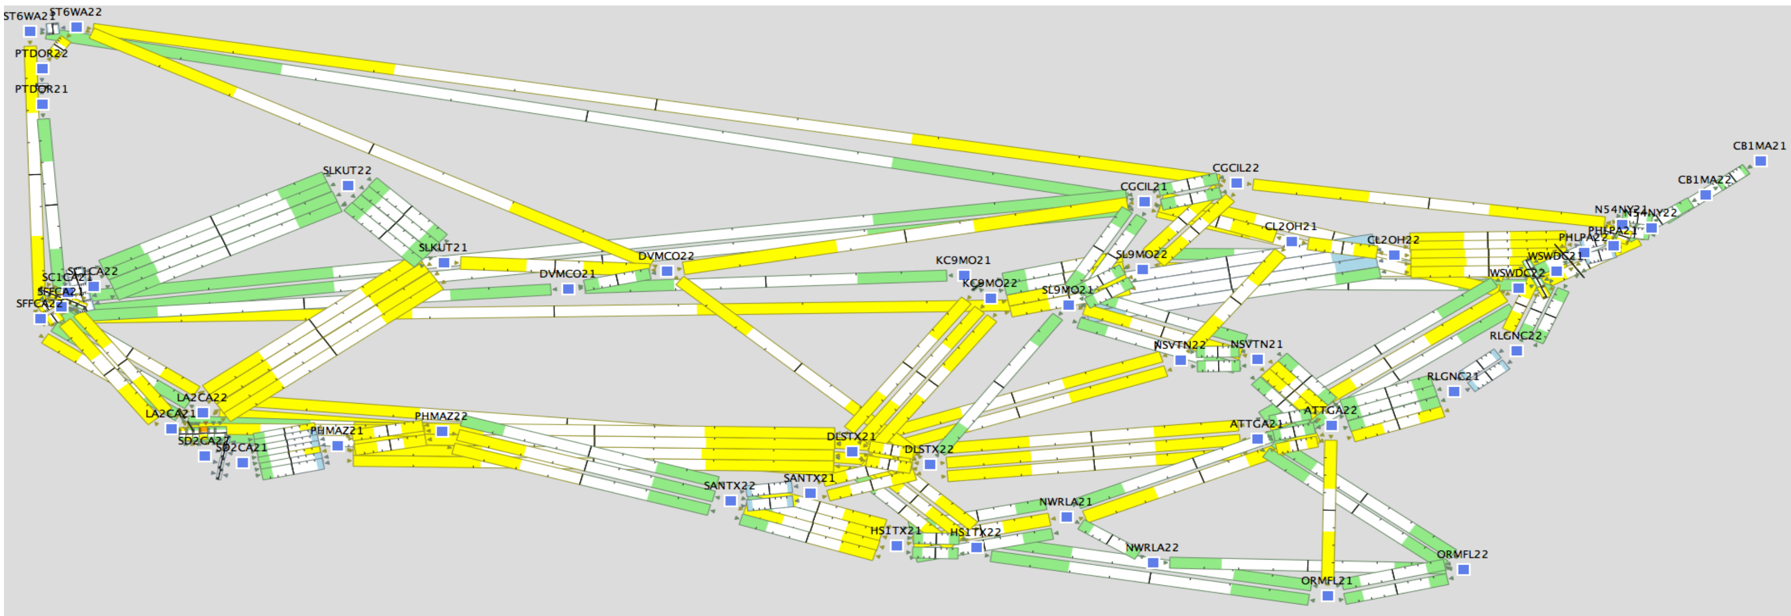

- Model current network state
- Simulate the interconnects
- Work through the migration steps – line-by-line configuration changes
- Identify the failure points in the migration process

Dynamic modeling can help reduce the risk and improve confidence that the migration will work

LEAST POWER COST SELECTION

LEAST POWER COST SELECTION
INTRODUCTION

Least Power Cost Path Selection is an approach to determining how to 'place' traffic demands across a network in a cost-effective manner

It also provides a framework for assessing the impact of introducing next-generation network technologies from a power-budget perspective

Ability to take advantage of differences in cost of power, power-source types, time of day variations, that IGPs don't currently take into account

SIMPLE POWER PROFILING

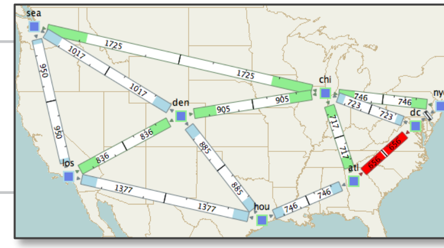

Offline network models, derived or obtained from actual SP topologies

- An information base of power-draw (Watts) of each node based on its computed inventory
  - ‘Common equipment’ plus each installed FPC & PIC
- Resulting in the cost (Watts) per Gb for data passing through a node

Expected Peak & Average Power per node

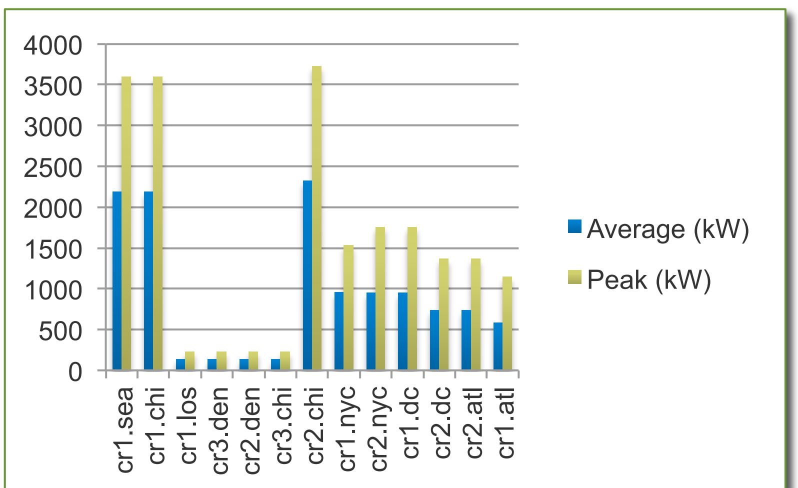

Watts per 1Gig demand through the node

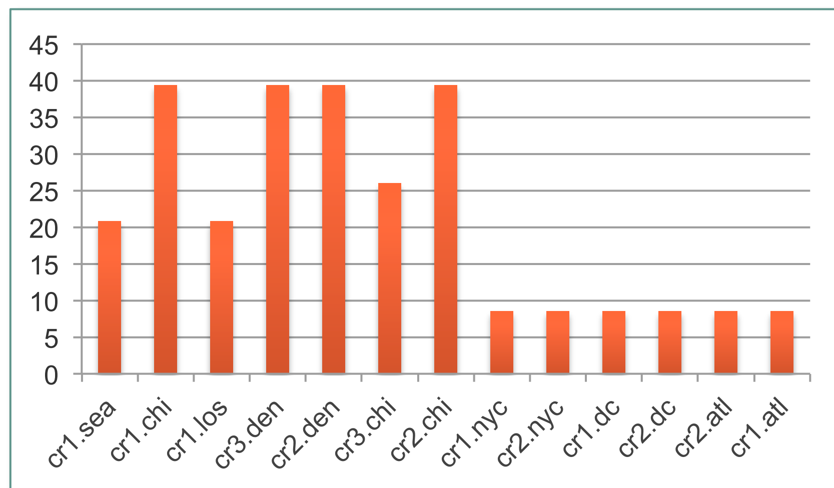

JUNIPER NETWORKS

GENERATE INVENTORY AND POWER ANALYSIS

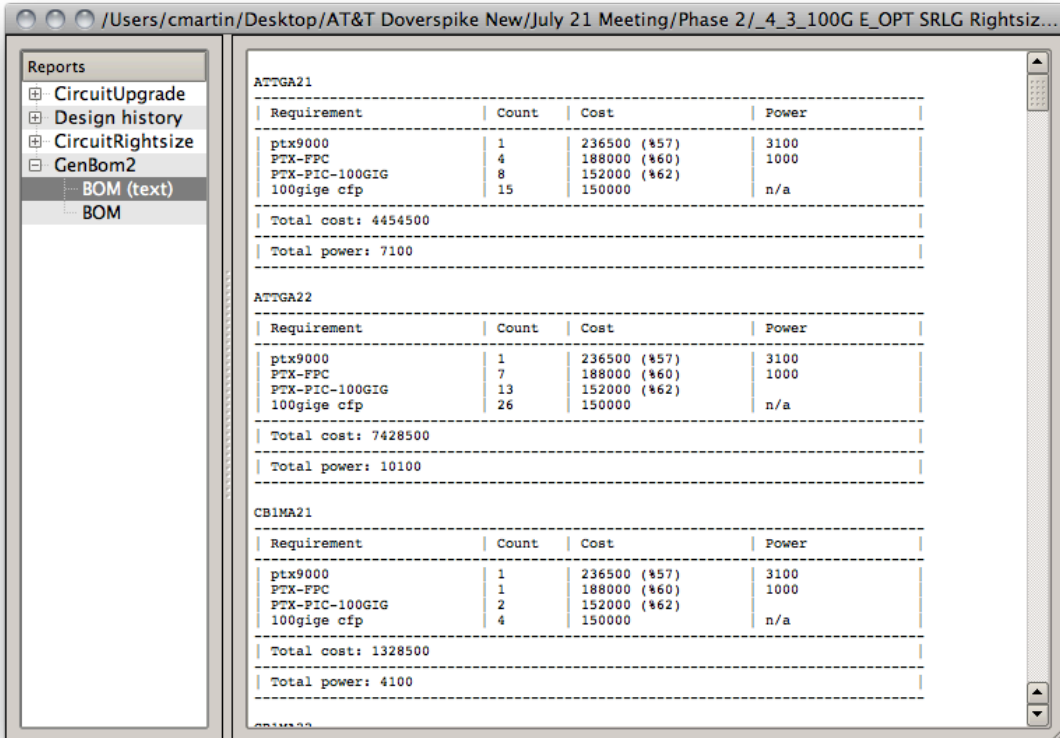

Requirement	Count	Cost	Power
ptx9000	1	236500 (%57)	3100
PTX-FPC	4	188000 (%60)	1000
PTX-PIC-100GIG	8	152000 (%62)	
100gige cfp	15	150000	n/a

Total cost: 4454500

Total power: 7100

ATTGA22

Requirement	Count	Cost	Power
ptx9000	1	236500 (%57)	3100
PTX-FPC	7	188000 (%60)	1000
PTX-PIC-100GIG	13	152000 (%62)	
100gige cfp	26	150000	n/a

Total cost: 7428500

Total power: 10100

CB1MA21

Requirement	Count	Cost	Power
ptx9000	1	236500 (%57)	3100
PTX-FPC	1	188000 (%60)	1000
PTX-PIC-100GIG	2	152000 (%62)	
100gige cfp	4	150000	n/a

Total cost: 1328500

Total power: 4100

Populated power and hardware data into MATE that easily allows for the generation of inventory and power requirements

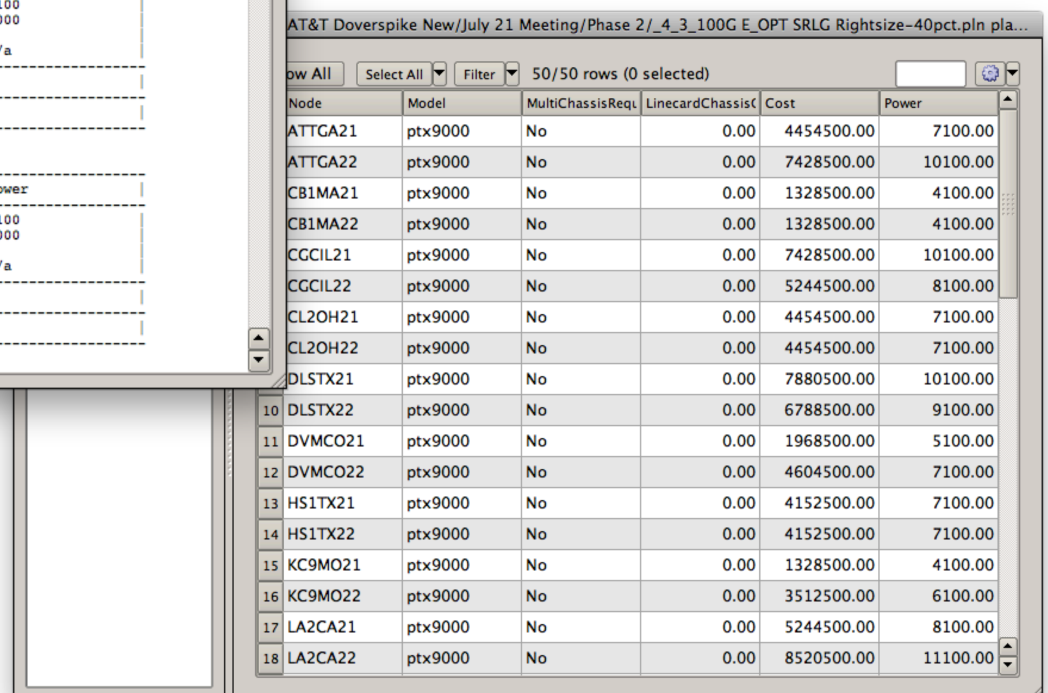

	Node	Model	MultiChassisRequ	LinecardChassis(	Cost	Power
	ATTGA21	ptx9000	No	0.00	4454500.00	7100.00
	ATTGA22	ptx9000	No	0.00	7428500.00	10100.00
	CB1MA21	ptx9000	No	0.00	1328500.00	4100.00
	CB1MA22	ptx9000	No	0.00	1328500.00	4100.00
	CGCIL21	ptx9000	No	0.00	7428500.00	10100.00
	CGCIL22	ptx9000	No	0.00	5244500.00	8100.00
	CL2OH21	ptx9000	No	0.00	4454500.00	7100.00
	CL2OH22	ptx9000	No	0.00	4454500.00	7100.00
	DLSTX21	ptx9000	No	0.00	7880500.00	10100.00
10	DLSTX22	ptx9000	No	0.00	6788500.00	9100.00
11	DVMCO21	ptx9000	No	0.00	1968500.00	5100.00
12	DVMCO22	ptx9000	No	0.00	4604500.00	7100.00
13	HS1TX21	ptx9000	No	0.00	4152500.00	7100.00
14	HS1TX22	ptx9000	No	0.00	4152500.00	7100.00
15	KC9MO21	ptx9000	No	0.00	1328500.00	4100.00
16	KC9MO22	ptx9000	No	0.00	3512500.00	6100.00
17	LA2CA21	ptx9000	No	0.00	5244500.00	8100.00
18	LA2CA22	ptx9000	No	0.00	8520500.00	11100.00

POWER ANALYSIS

We compute full power usage details, for each chassis and all its active components

- We can then compute the power cost for each demand in the network
- (Future) TE option to route LSPs across the least power cost path

/Users/cmartin/Desktop/AT&T Doverspike New/July 21 Meeting/Phase 2/100G E_OPT SRLG Rightsize-40pct.pln plan info

Reports
- CircuitUpgrade
- Design history
- CircuitRightsize
- Simulation Analysis
- PowerAnalyzer
  - Node Power (none) (587)
  - Site Power (none) (50)
  - Network Power (none) (1)

Show All | Select All | Filter | 587/587 rows (1 selected)

	Node	Slot	Hardware	Component Draw (w)	Total Draw (w)	Maximum Draw (w)	Annual kWh	Annual Carbon Footprint (lbs)	Energy Consu	Expected Therma	Max Thermal Out
1	LA2CA22	chassis	commons	3644.00	7964 (total)	24750 (total)	69804.46	90536.4	39.8	27180.9	84471.0
2	ATTGA22	chassis	commons	3644.00	7424 (total)	24750 (total)	65071.36	84397.6	37.1	25337.9	84471.0
3	CGCIL21	chassis	commons	3644.00	7424 (total)	24750 (total)	65071.36	84397.6	37.1	25337.9	84471.0
4	DLSTX21	chassis	commons	3644.00	7424 (total)	24750 (total)	65071.36	84397.6	37.1	25337.9	84471.0
5	SFFCA21	chassis	commons	3644.00	7424 (total)	24750 (total)	65071.36	84397.6	37.1	25337.9	84471.0
6	DLSTX22	chassis	commons	3644.00	6884 (total)	24750 (total)	60338.26	78258.7	34.4	23494.9	84471.0
7	SL9MO22	chassis	commons	3644.00	6884 (total)	24750 (total)	60338.26	78258.7	34.4	23494.9	84471.0
8	WSWD...	chassis	commons	3644.00	6884 (total)	24750 (total)	60338.26	78258.7	34.4	23494.9	84471.0
9	CGCIL22	chassis	commons	3644.00	6344 (total)	24750 (total)	55605.16	72119.9	31.7	21651.9	84471.0
10	LA2CA21	chassis	commons	3644.00	6344 (total)	24750 (total)	55605.16	72119.9	31.7	21651.9	84471.0
11	PHLPA22	chassis	commons	3644.00	6344 (total)	24750 (total)	55605.16	72119.9	31.7	21651.9	84471.0
12	SL9MO21	chassis	commons	3644.00	6344 (total)	24750 (total)	55605.16	72119.9	31.7	21651.9	84471.0
13	WSWD...	chassis	commons	3644.00	6344 (total)	24750 (total)	55605.16	72119.9	31.7	21651.9	84471.0
14	ATTGA21	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
15	CL2OH...	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
16	CL2OH...	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
17	DVMC...	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
18	HS1TX21	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
19	HS1TX22	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
20	N54NY...	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0
21	NSVTN...	chassis	commons	3644.00	5804 (total)	24750 (total)	50872.06	65981.1	29.0	19808.9	84471.0

WHAT CAN WE DO WITH THESE MODELS?

The previous slide demonstrates an efficient inventory & ROI tool for introducing NG technology

- e.g. Inserting PTX will pay for itself in X months but …

We can incorporate other factors

- Facilities costs e.g. PECO vs. Co-Lo, etc.
- And of course … the Traffic Matrix

Facility Costs

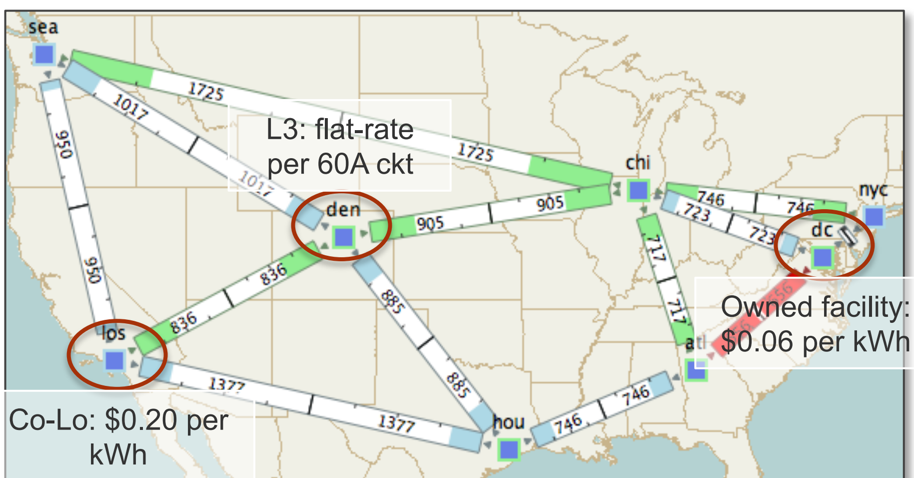

Traffic Matrix

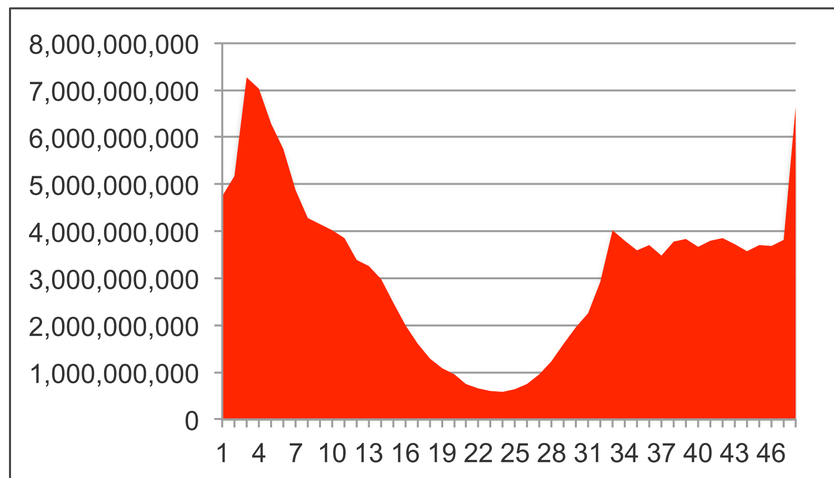

WE CAN COMPUTE A DEMAND DRIVEN POWER UTILIZATION

The cost of individual demands can be quantified as well as the effects of demand aggregation

- In figure #1 the Nodes transited for a demand from Seattle to NYC

  pe1.sea, cr1.sea, cr1.chi, cr2.chi, cr1.nyc, pe1.nyc

  - Path cost = 159.89 W/Gb, B/W of demand at time X = 13Gb
  - Path cost of demand = 2.079KW

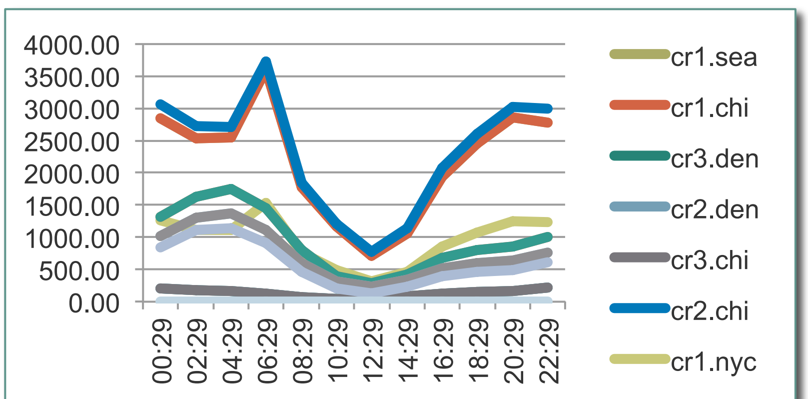

Time of Day, Demand driven Power Utilization

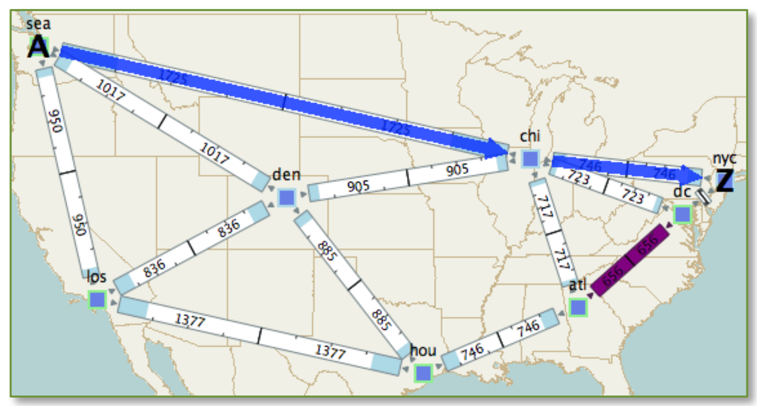

Figure 1: Seattle to NYC demand

JUNIPER NETWORKS

WHAT'S NEXT

Combining power consumption information & routing analytics enabling more informed decisions about:

- Taking advantage of cheaper power sources, time-of-day pricing, lower carbon tonnage etc.
- Route high B/W demands more efficiently &/or react to natural disasters
- Automated provisioning of paths – TE explicit paths, for example

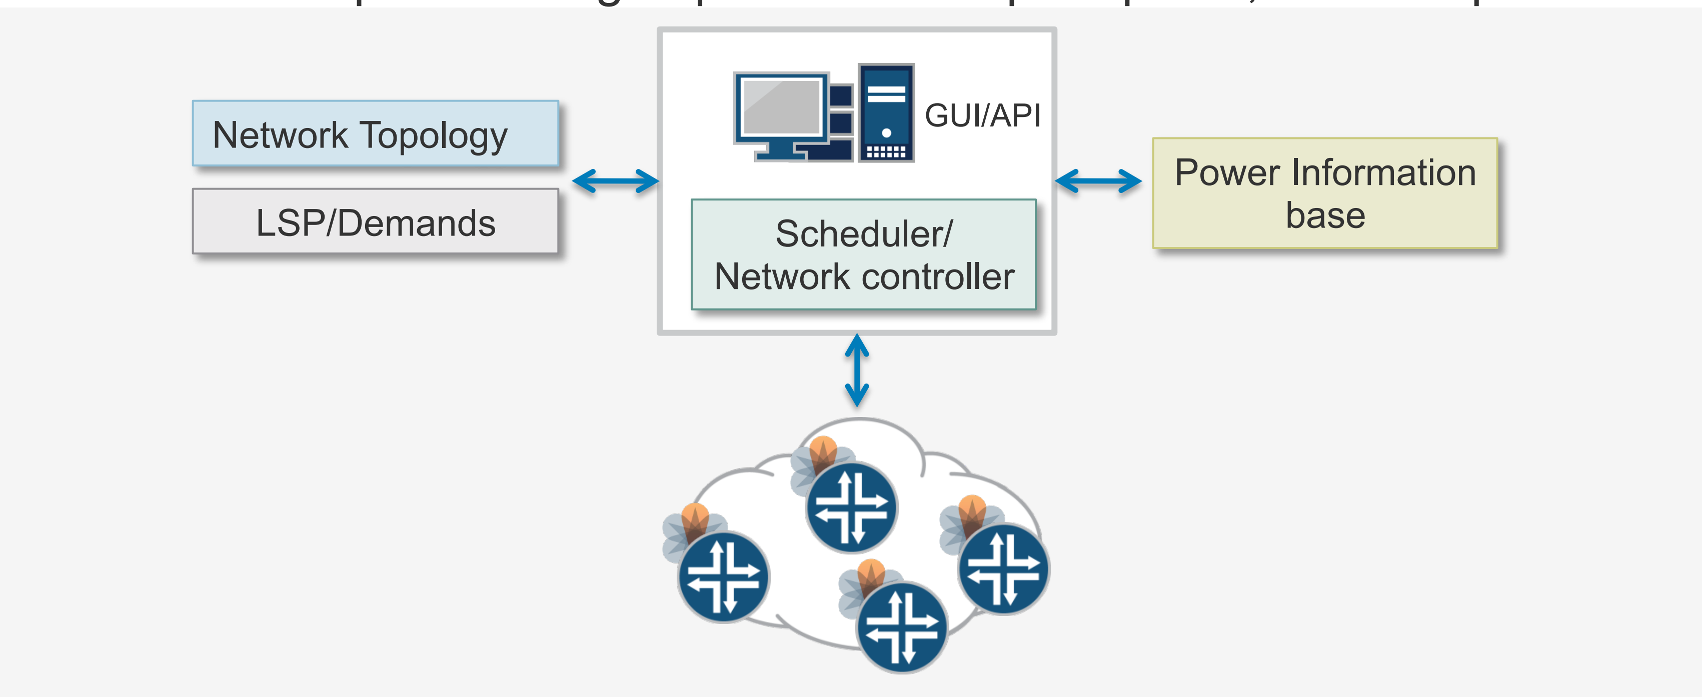

CONSTRUCTING NETWORK MODELS

NECESSARY DATA FOR SUCCESSFUL MODELS

Diagrams are good – data is better

A successful approach requires:

- A topology (OSPF or IS-IS LSDB where possible)
- The link types in the network (speeds)
- Any layer 0/1 protection schemes in use
- A traffic matrix (demands from node to node)
- Time of day traffic behaviour
- Expected growth forecasts (ie, 40% YOY)

In addition, we can build a more accurate model if we know

- Any engineering guidelines for things like latency bounds, QoS handling, etc
- The layer 1 topology
  - Resulting SRLGs, Lambda contention, Span lengths
- The current hardware inventory of the network
- Any measured utilization on interfaces (MRTG for example)

LOADING THE NETWORK

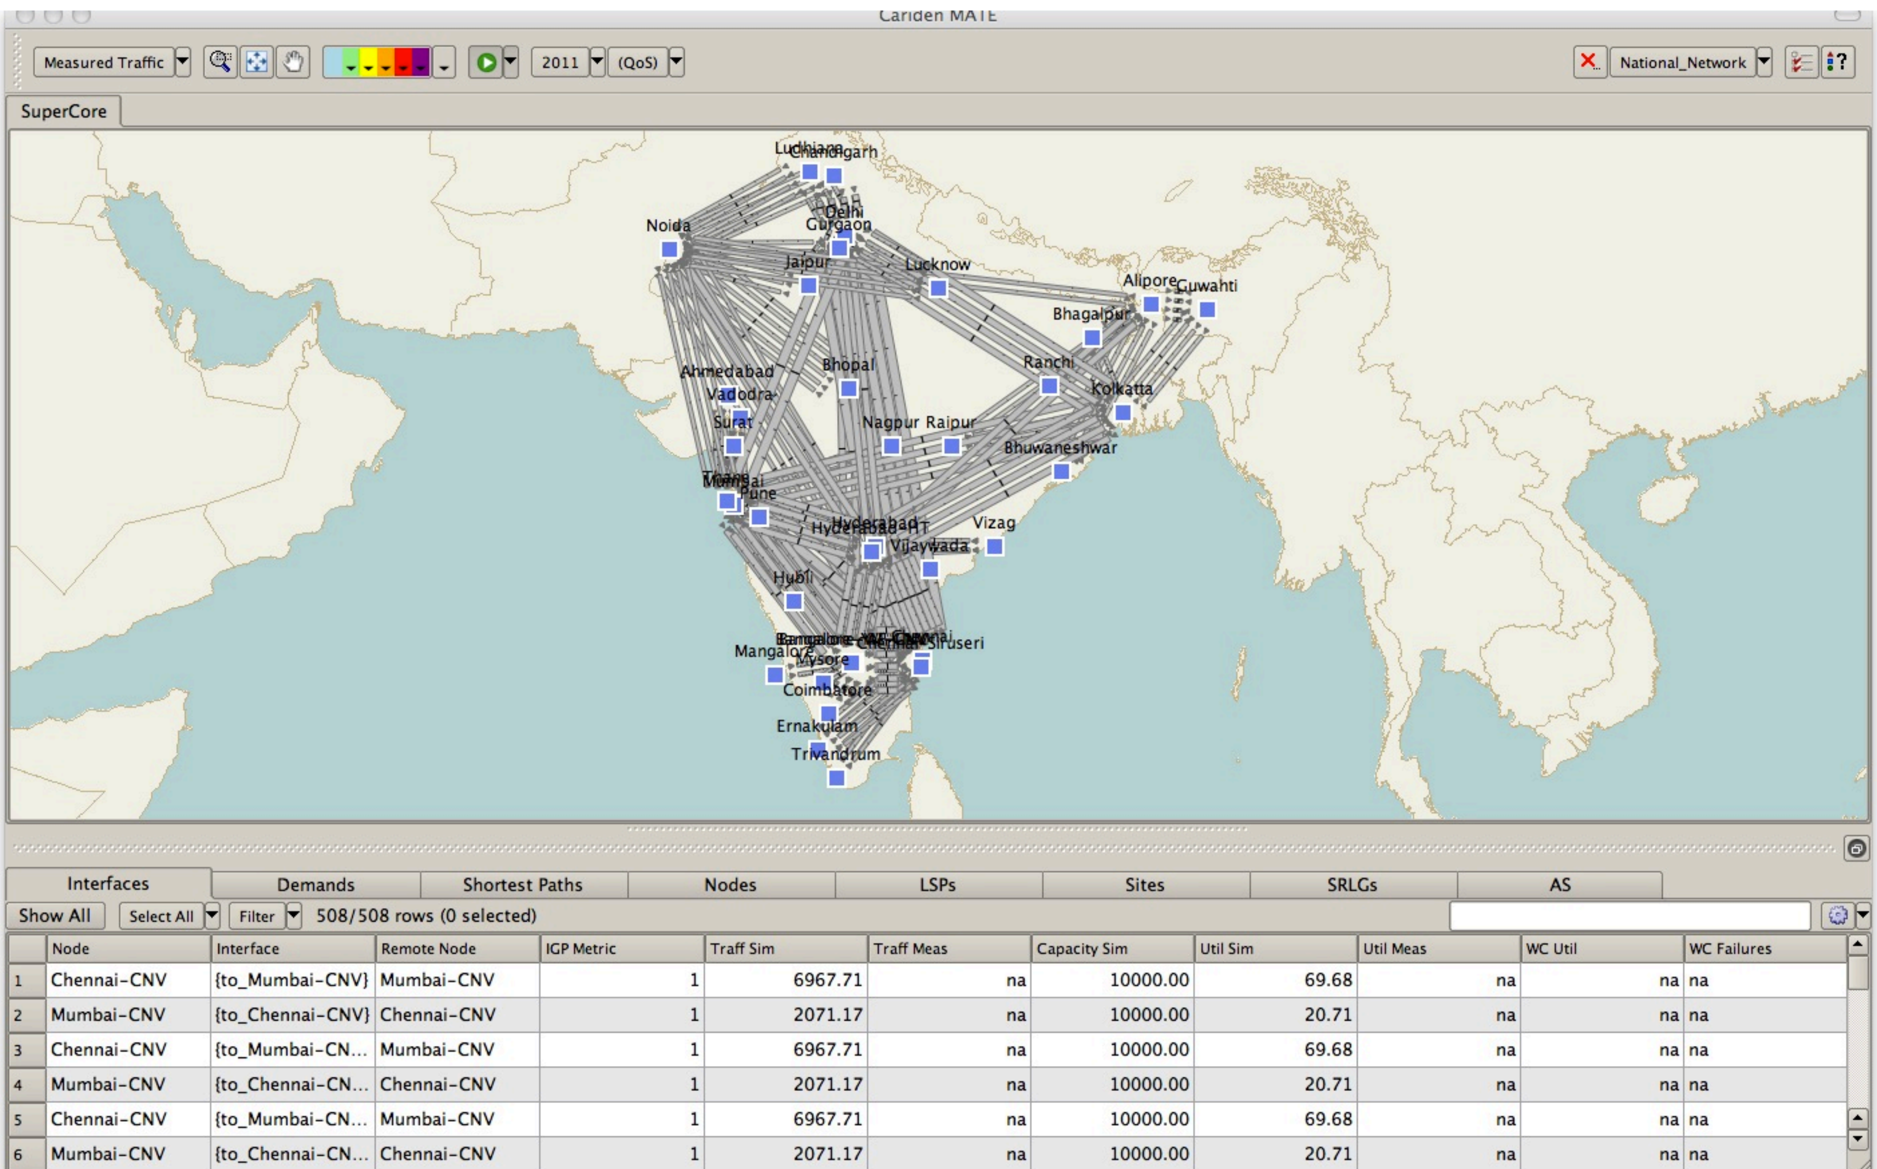

TRAFFIC DEMANDS APPLIED

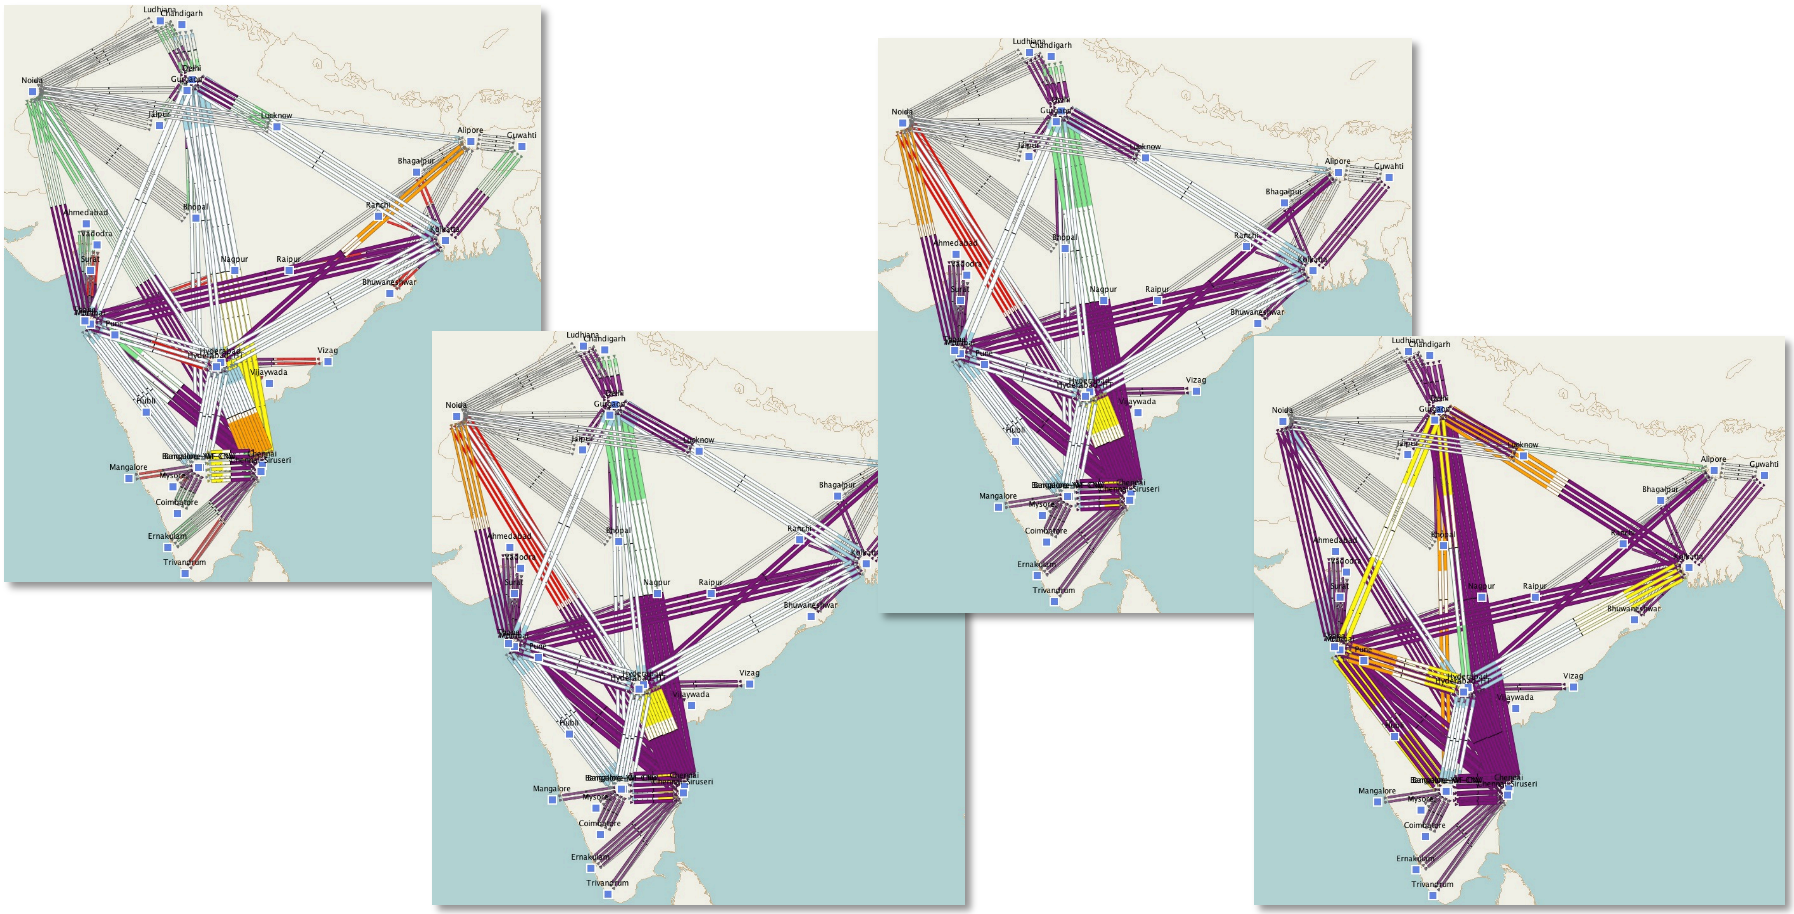

2012 projection → 2015 projection

JUNIPER NETWORKS

AND WHEN THINGS GO WRONG…

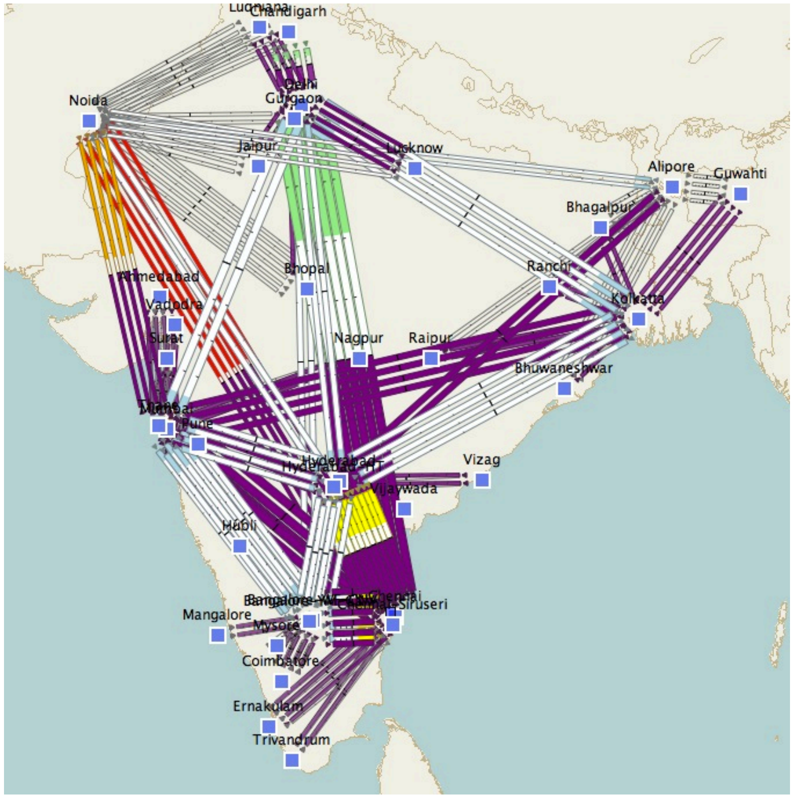

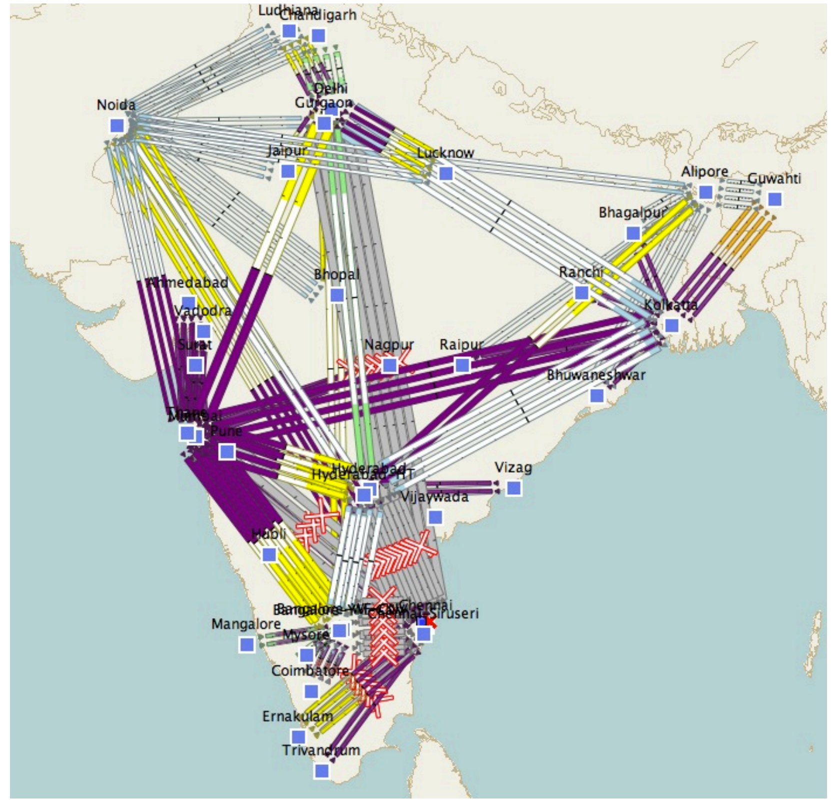

Core node failure

Traffic flows rerouted

FULLY VIRTUALIZED NETWORK MODEL

VIRTUALISE THE MODEL

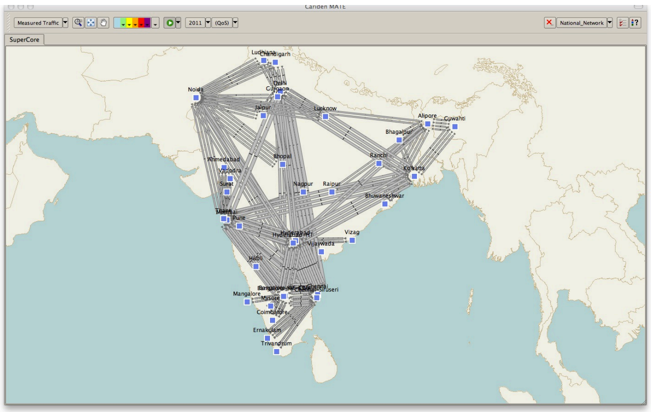

Build diagram

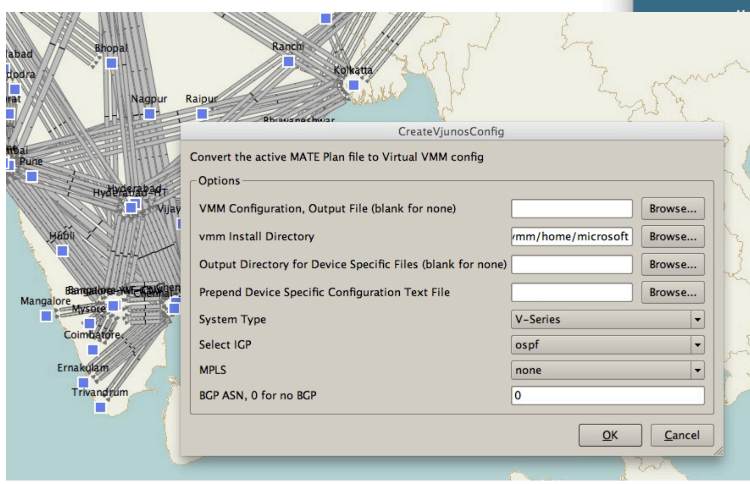

Generate the configuration files

Log in to platform

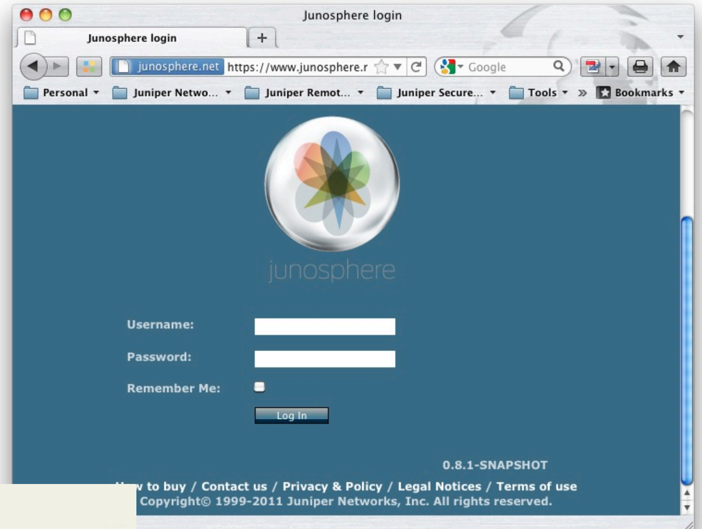

JUNIPER NETWORKS

STARTUP

Upload configuration files

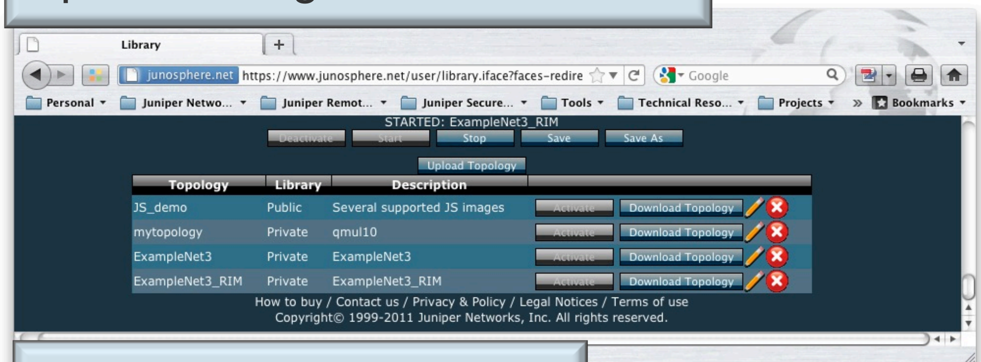

Start Virtual Routers

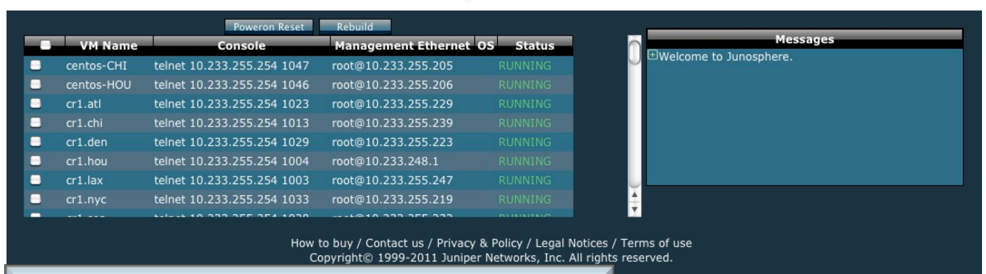

Connect to Virtual Routers

Log in and use

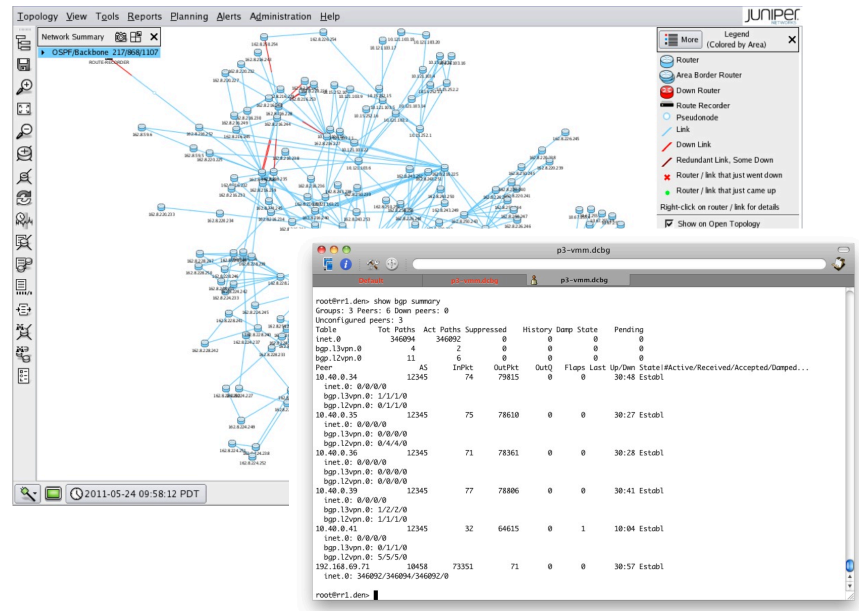

JUNIPER NETWORKS

PLATFORM COMPONENTS

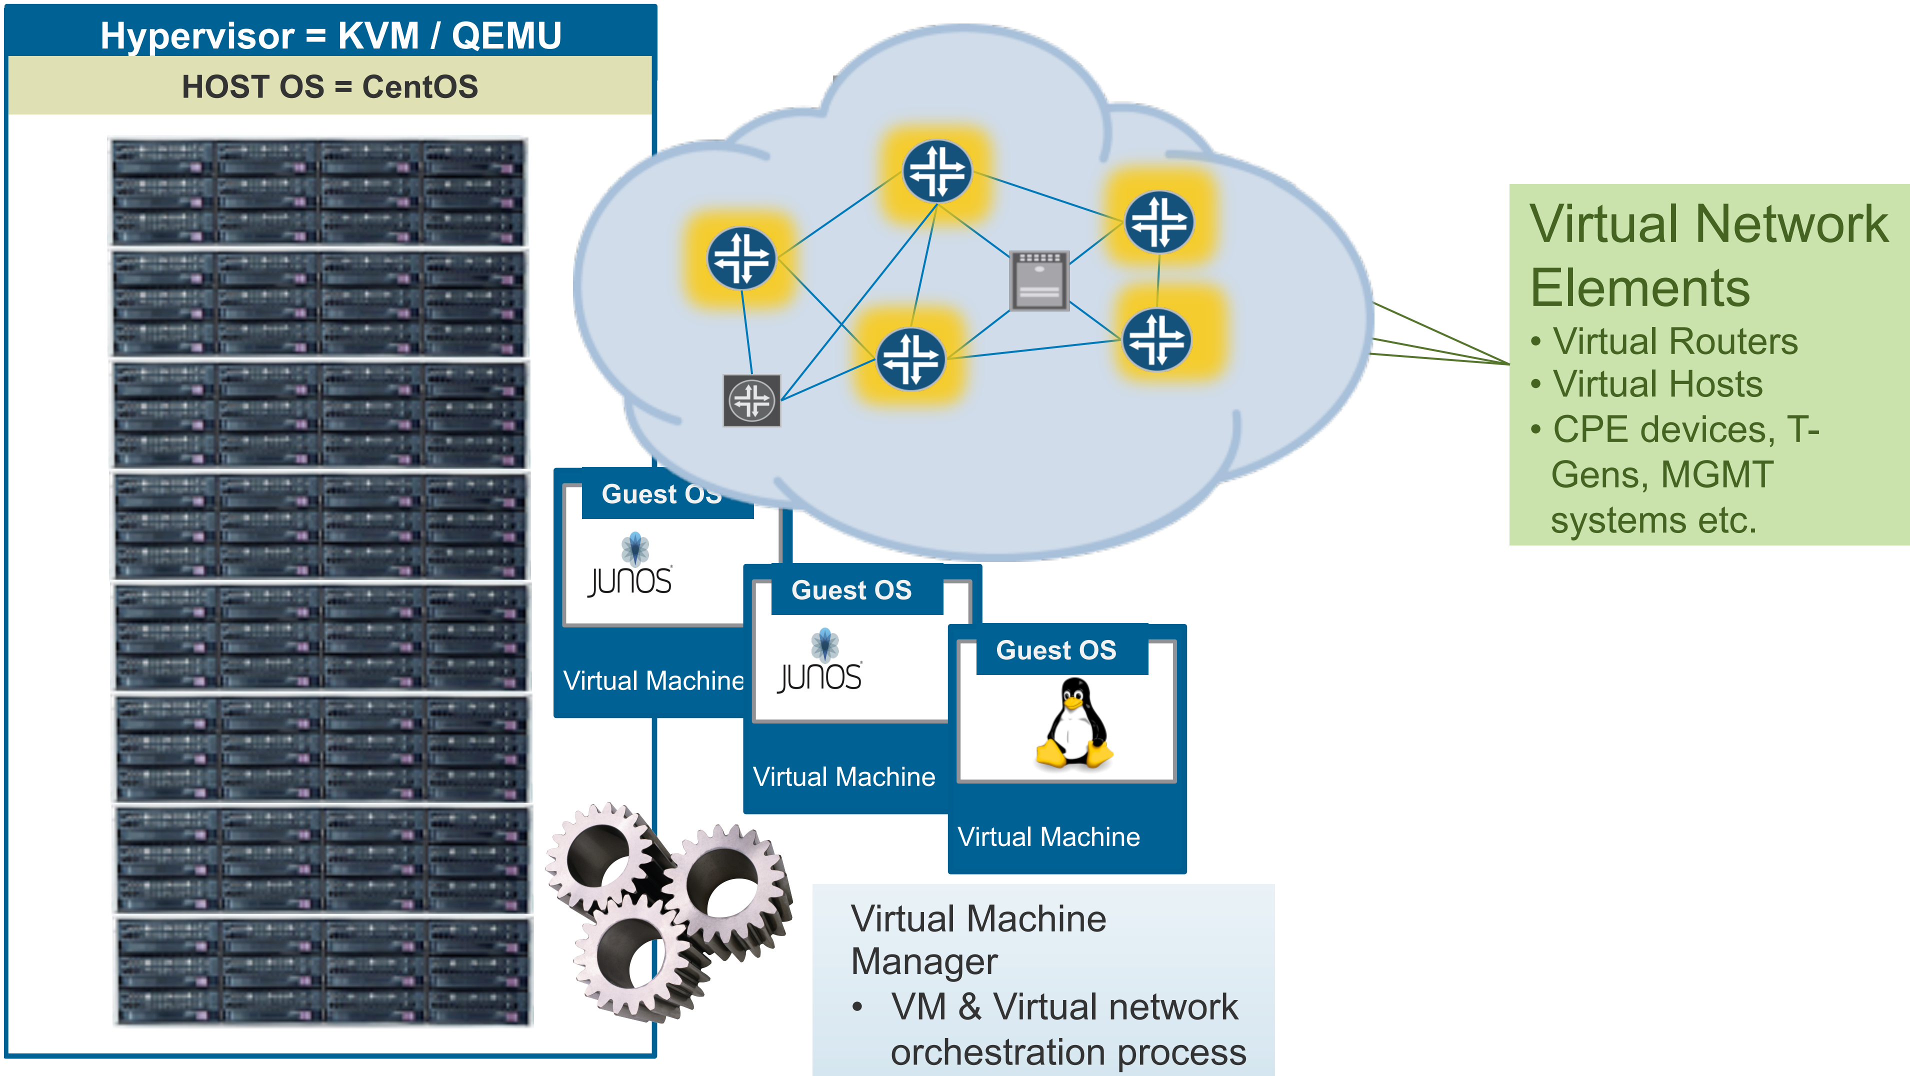

VMM EXAMPLE

Open-source tools in development with University of Adelaide

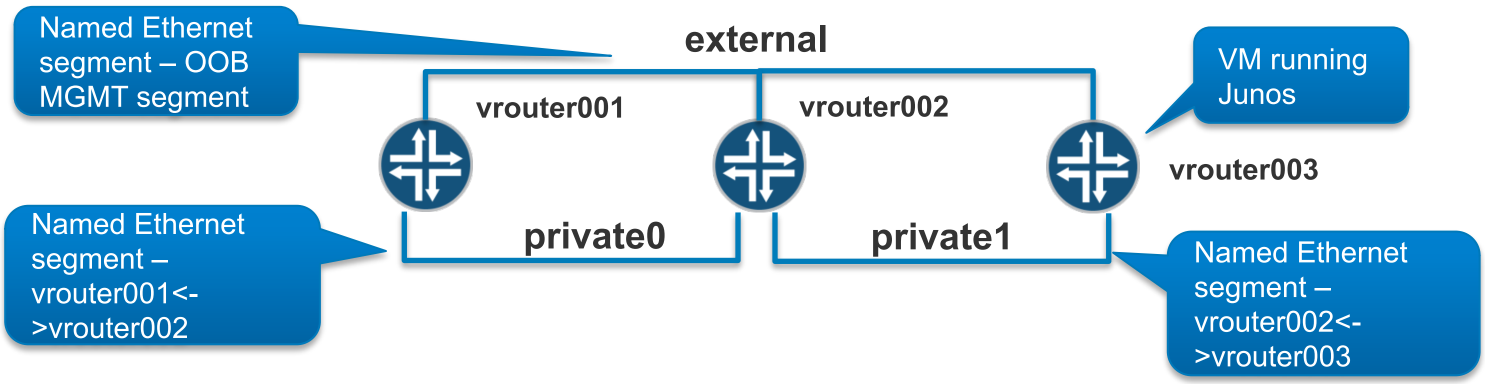


```
// description – bind global definitions
#include "/vmm/bin/common.defs"
config "config" {
        display "NULL";
          vm "vrouter001" {
                    // description – hostname of set on VM
                    hostname "vrouter001";
                    // description – Operating system image to load –
                    VJX_DISK
                    // description – ge 0/0/0 interface to the outside world
                    interface "em0" { EXTERNAL;};
                    // description - ge 0/0/1 link between vrouter001 and vrouter002
                    interface "em1" { bridge "private0"; };
                    // description - configuration file to load on the router
                    install "ENV(HOME)/active/configset/vrouter001.conf" "/root/junos.conf";
                    };
```

VMM EXAMPLE

```
    vm "vrouter002" {
            hostname "vrouter002" ;
            VJX_DISK
            //description - interface to the outside world
            interface "em0" { EXTERNAL;};
            // description - link between vrouter002 and vrouter001
            interface "em1" { bridge "private0"; };
            // description - link between vrouter002 and vrouter003
            interface "em2" { bridge "private1"; };
            // description - configuration file to load on the route
            install "ENV(HOME)/active/configset/vrouter002.conf" "/r
junos.conf";
            };
    vm "vrouter003" {
            hostname "vrouter003" ;
            VJX_DISK
            //description - interface to the outside world
            interface "em0" { EXTERNAL;};
            // description - link between vrouter003 and vrouter002
            interface "em1" { bridge "private1"; };
            // description - configuration file to load on the router
            install "ENV(HOME)/active/configset/vrouter002.conf" "/root/
junos.conf";
            };
    PRIVATE_BRIDGES
};
```

Connect to OOB MGMT segment

Named Ethernet segment – vrouter001<->vrouter002

Named Ethernet segment – vrouter002<->vrouter003

Named Ethernet segment – vrouter003<->vrouter002

Per-VM Junos configuration file

Macro call to enable Ethernet bridge segments

LOOKING INSIDE THE VIRTUAL WORLD

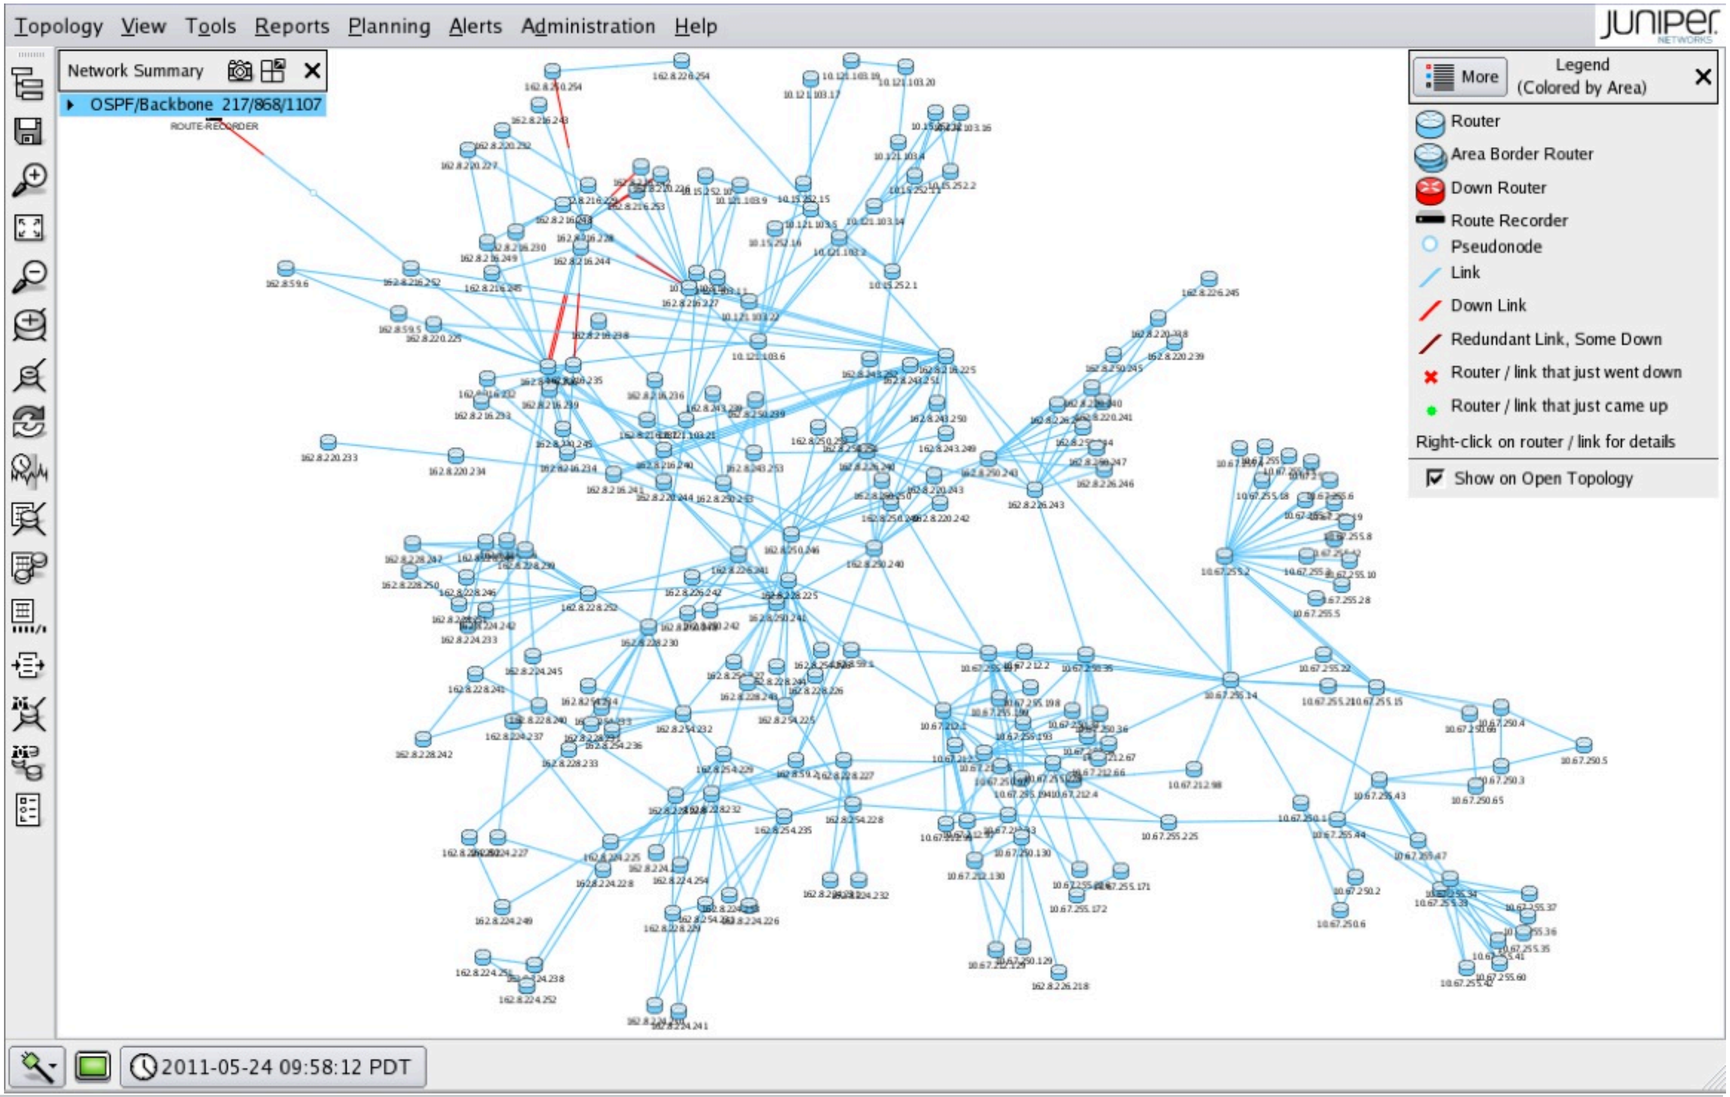

MANAGEMENT TOOLS

Can take tools that you use in you physical network and deploy into the virtual – VMs don't know that they're not physical devices

- MRTG monitoring SNMP stats and generating graphs

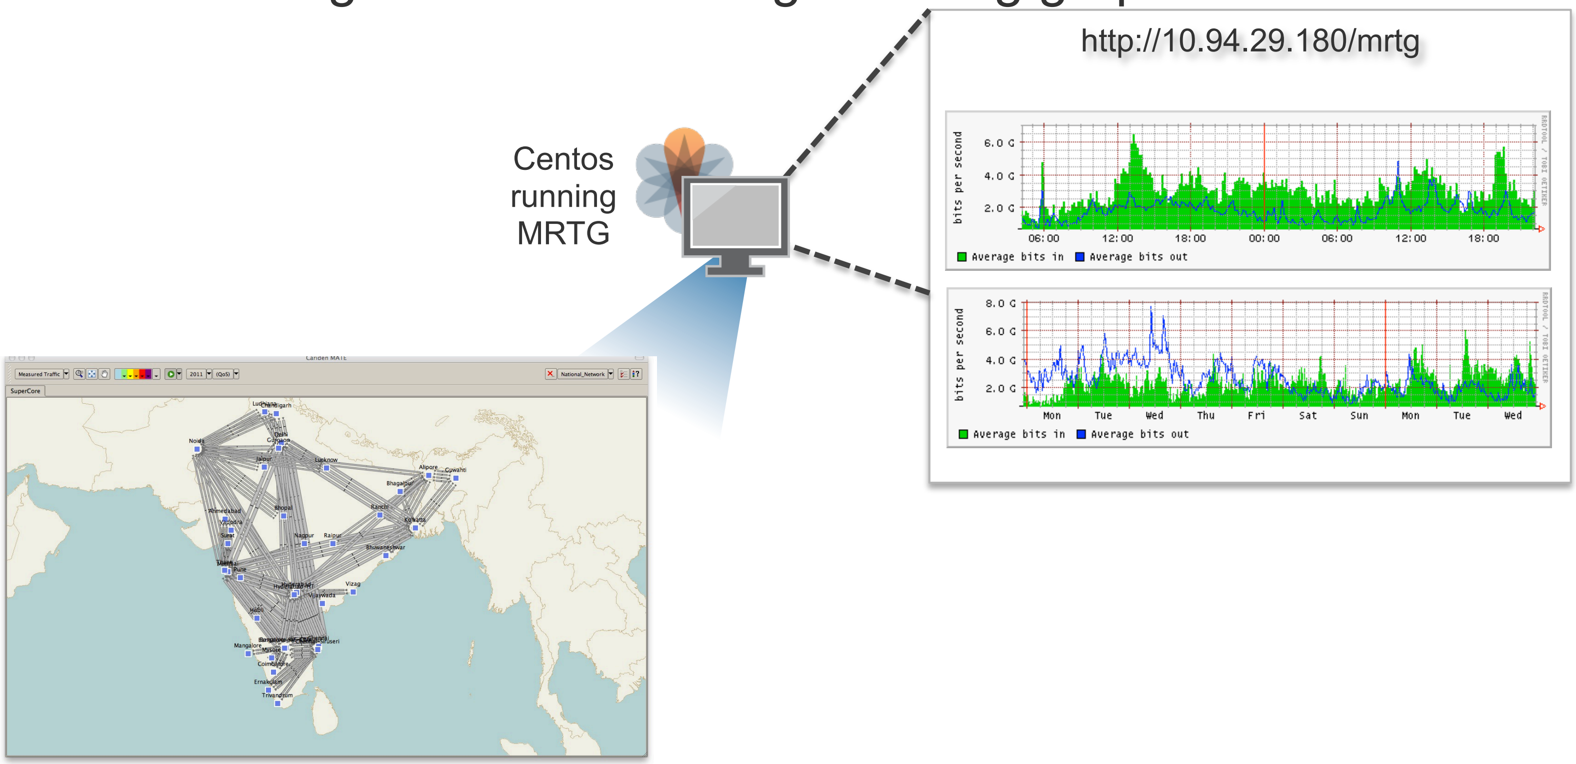

SPIRENT ROUTER TESTER INTEGRATION

Spirent Test Center Virtual is a software package that extends & complements the capabilities of Spirent Test Center – a virtual traffic-generator for the virtual world

Network operators can:

- Run real-world traffic and control-plane events over virtualized networks

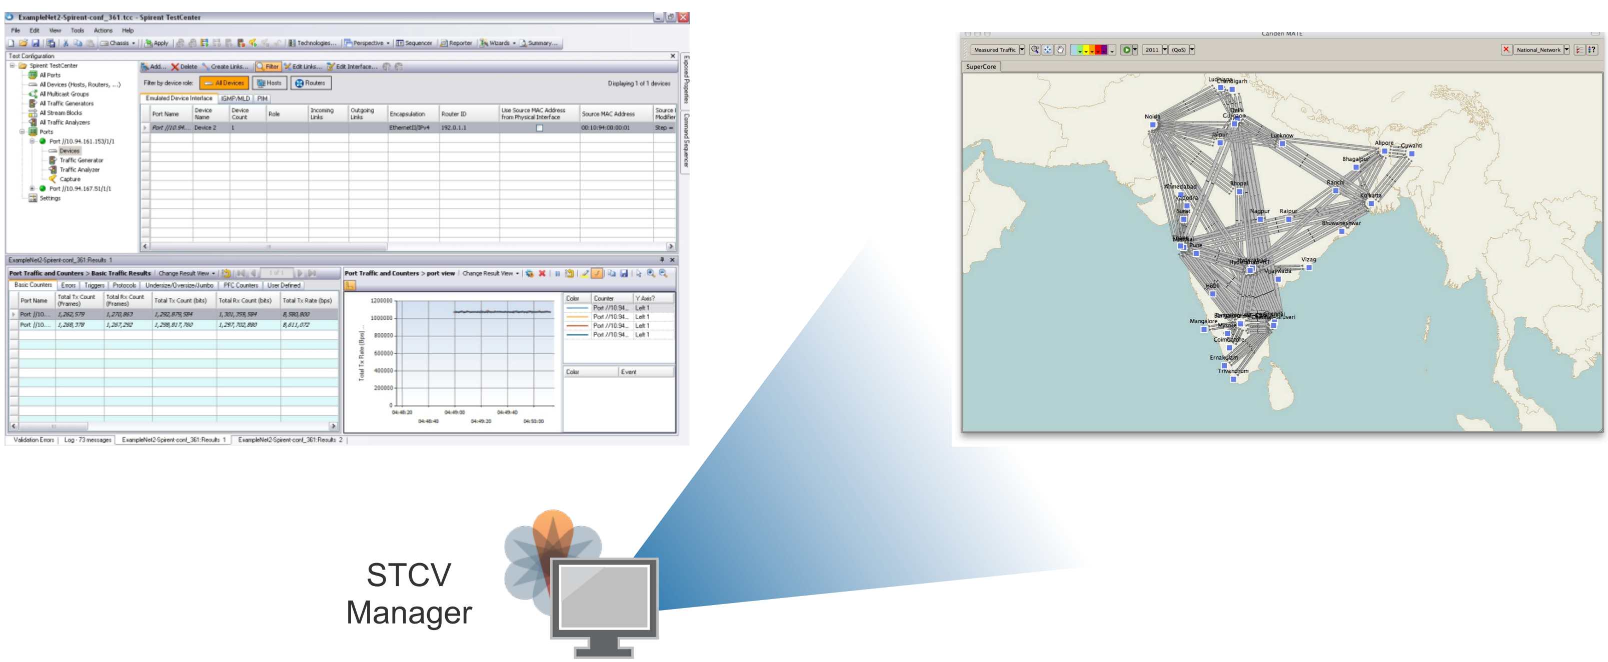

VIRTUALIZED NETWORKS IN OPERATION

NANOG 51/52/53 – RPKI Workshop

- http://www.nanog.org/meetings/nanog52/presentations/Sunday/110612.nanog-lab-agenda.pdf

RIPE62 – RPKI Workshop

- http://ripe62.ripe.net/programme/meeting-plan/tutorials

Queen Mary University of London – Networking labs

- Paper at SIGCOMM/ACM Toronto
- http://edusigcomm.info.ucl.ac.be/pmwiki/uploads/Workshop2011/20110311002/sigcom2011_VindyaWijeratne.pdf

Loughborough University – Networking labs

Universitat Wien – Networking labs

Boston University – Network Security

Roma Tre University – Network software development

- Paper at IEEE Network Operations and Management Symposium, Hawaii
- http://www.ieee-noms.org/cfp.html

Internet Institute Japan – Internet routing research

Cambridge University, Systems Research Group – Internet routing research

Networks from 50+ International Service Providers

SUMMARY

Fully virtualized networks are a powerful tool

- Interconnect with existing physical network labs
- Node-accurate ‘what-if’ testing ground
- Ability to provide operations teams with a ‘safe’ environment to learn via ‘break & recover’ methods
- Reduce the risk of complex configuration changes, service migrations etc.
- Build and operate planned networks before any equipment is physically deployed
- Development platform for programmable networks & orchestration
- Platform for research, development and collaboration

THANK YOU

everywhere