Packet Design

Real-World Uses of Route Analytics

Greg Hooten
Senior Routing Engineer
Packet Design Inc.
gahooten@packetdesign.com

Packet Design, Inc.

"Harnessing the Intelligence of IP"

- Established in 2003
- Headquarters in Palo Alto, CA
- Fourth major release of Route Explorer™ product
- Pioneer and leader in Route Analytics for IP networks
- More than 100 customers worldwide
- Demo station 2 at Apricot

Cisco Systems | Technology Developer Partner

hp invent

ENTERPRISE MANAGEMENT ASSOCIATES — IMPACT BRIEF

HP's Strategic Partnership with Packet Design Signals a Landmark in Network Troubleshooting

2004 Finalist eWEEK

NetworkMagazine PRODUCT OF THE YEAR 2003

Packet Design

Why Route Analytics?

IP Routing Lacks Real-Time Management Visibility

- Routing is too dynamic to measure via SNMP polling intervals
- No existing network management tools provide visibility into IP routing and Layer 3 topology
  - SNMP only sees interfaces and links
  - Application performance management measures end to end statistics
  - Planning tools are primarily offline
  - Traffic analysis measures usage per link and end to end

Causes of Downtime in IP Networks

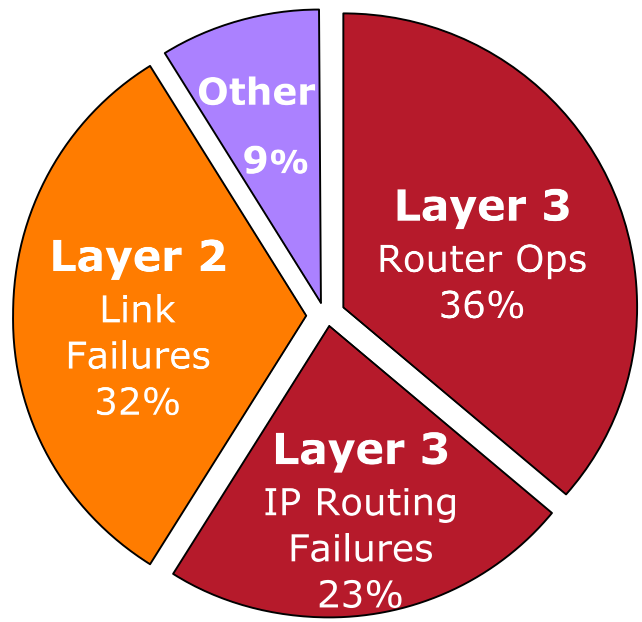

Sources: University of Michigan, Sprint

Packet Design

Route Analytics Defined

- Peers passively with routers
- Records all routing updates in real-time
- Stores complete history
- Computes topology across protocols, Areas, AS
- Real-time, network-wide analysis of IP routing
- Monitoring, Reporting, Alerting, Troubleshooting
- What-if analysis and modeling

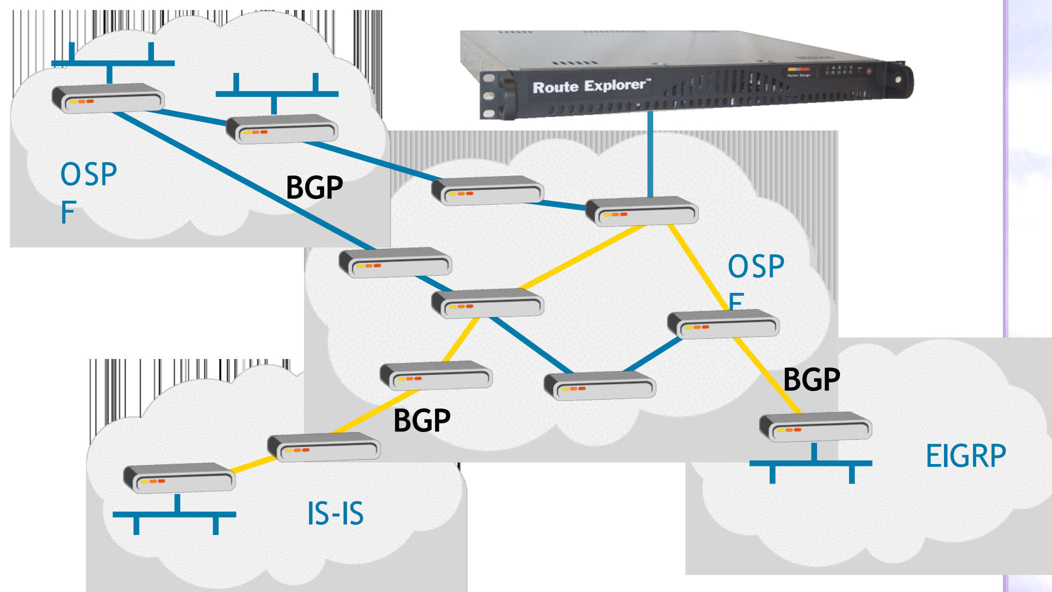

- BGP
- OSPF
- IS-IS
- EIGRP
- MPLS L3 VPNs

Network Map

- Green, blue and magenta are different OSPF areas
- Purple are BGP speakers
- Red elements are down
- Yellow is a path between two routers

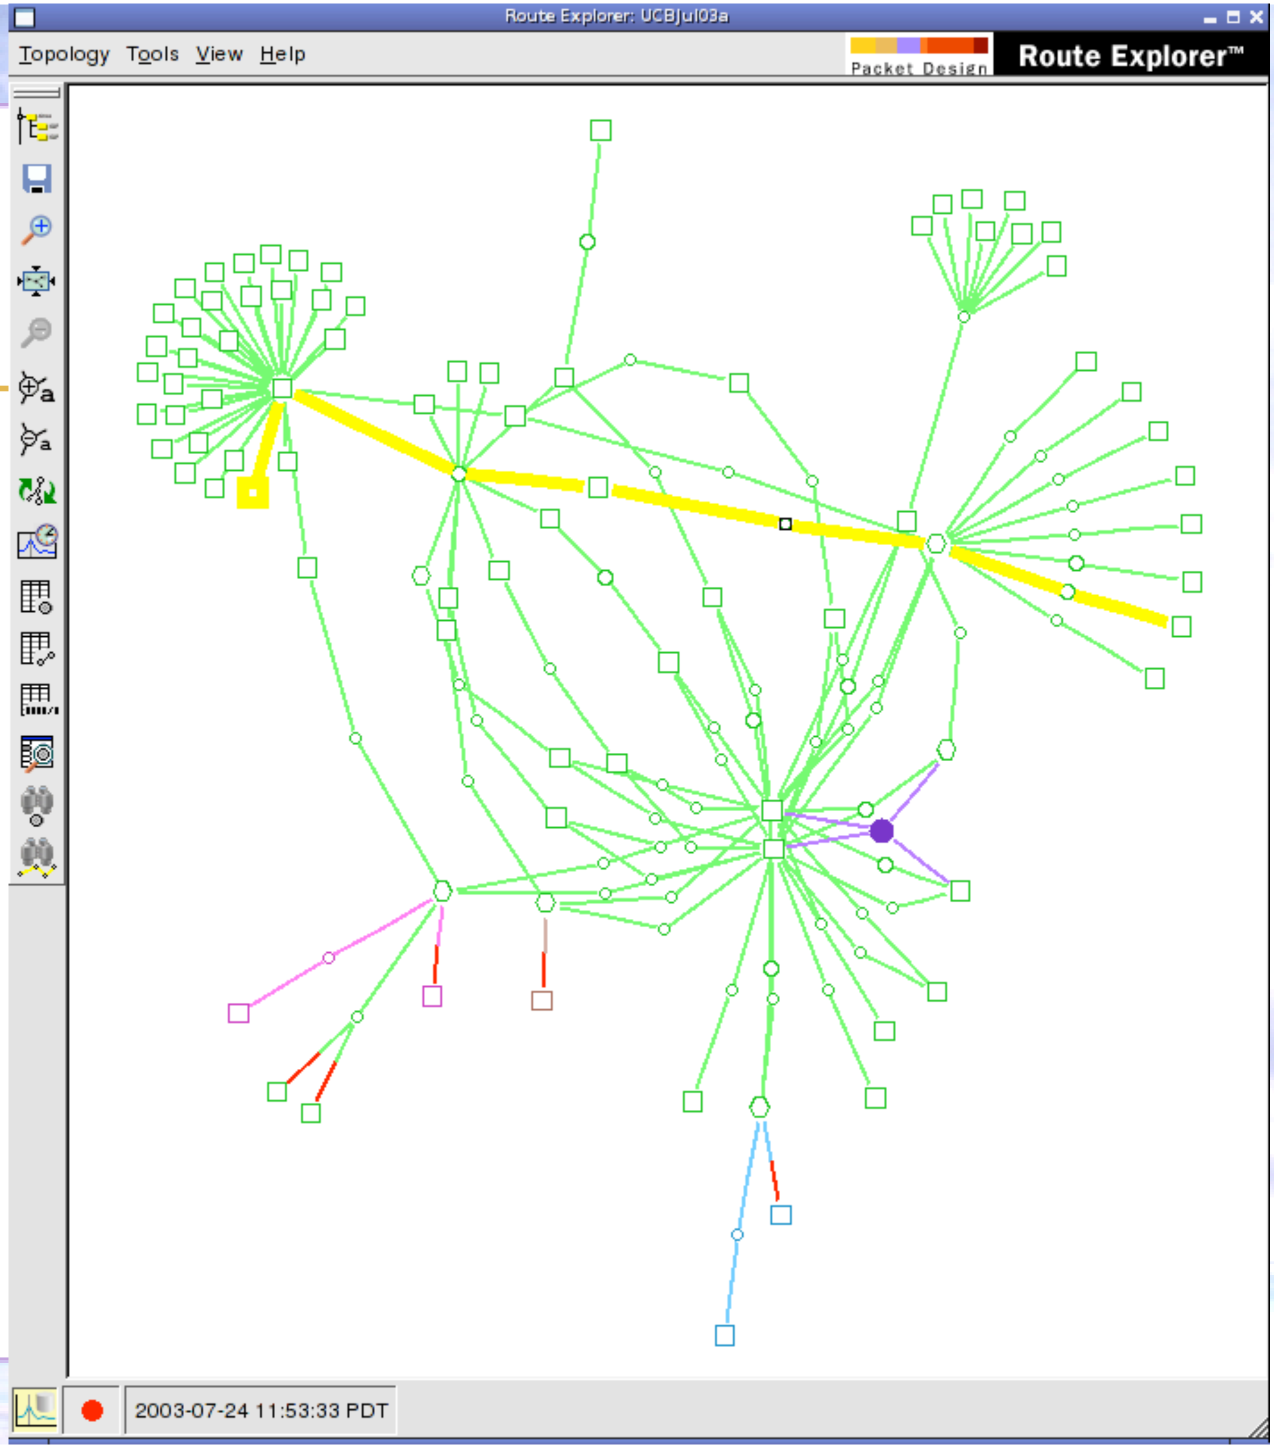

Routers, Links, Prefixes

Packet Design

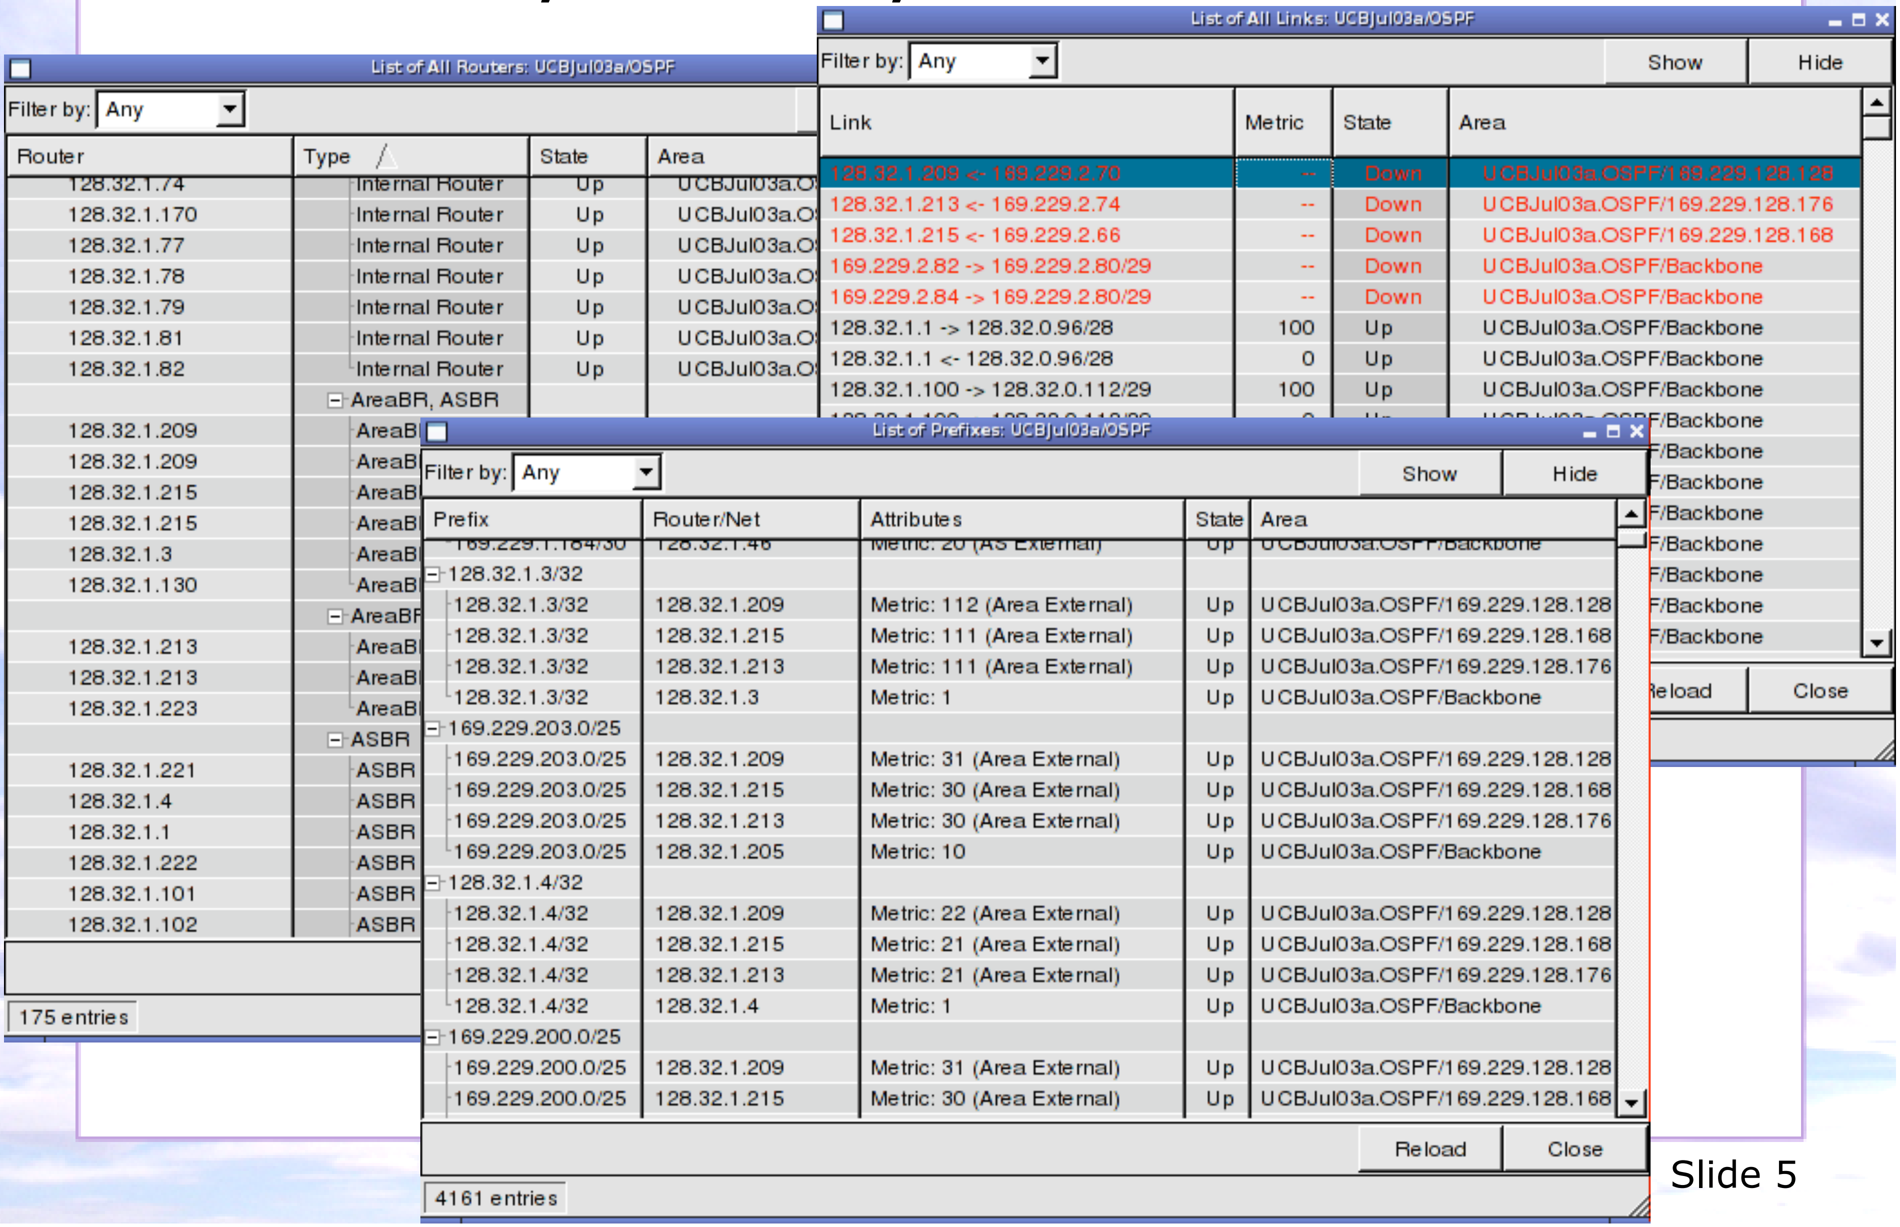

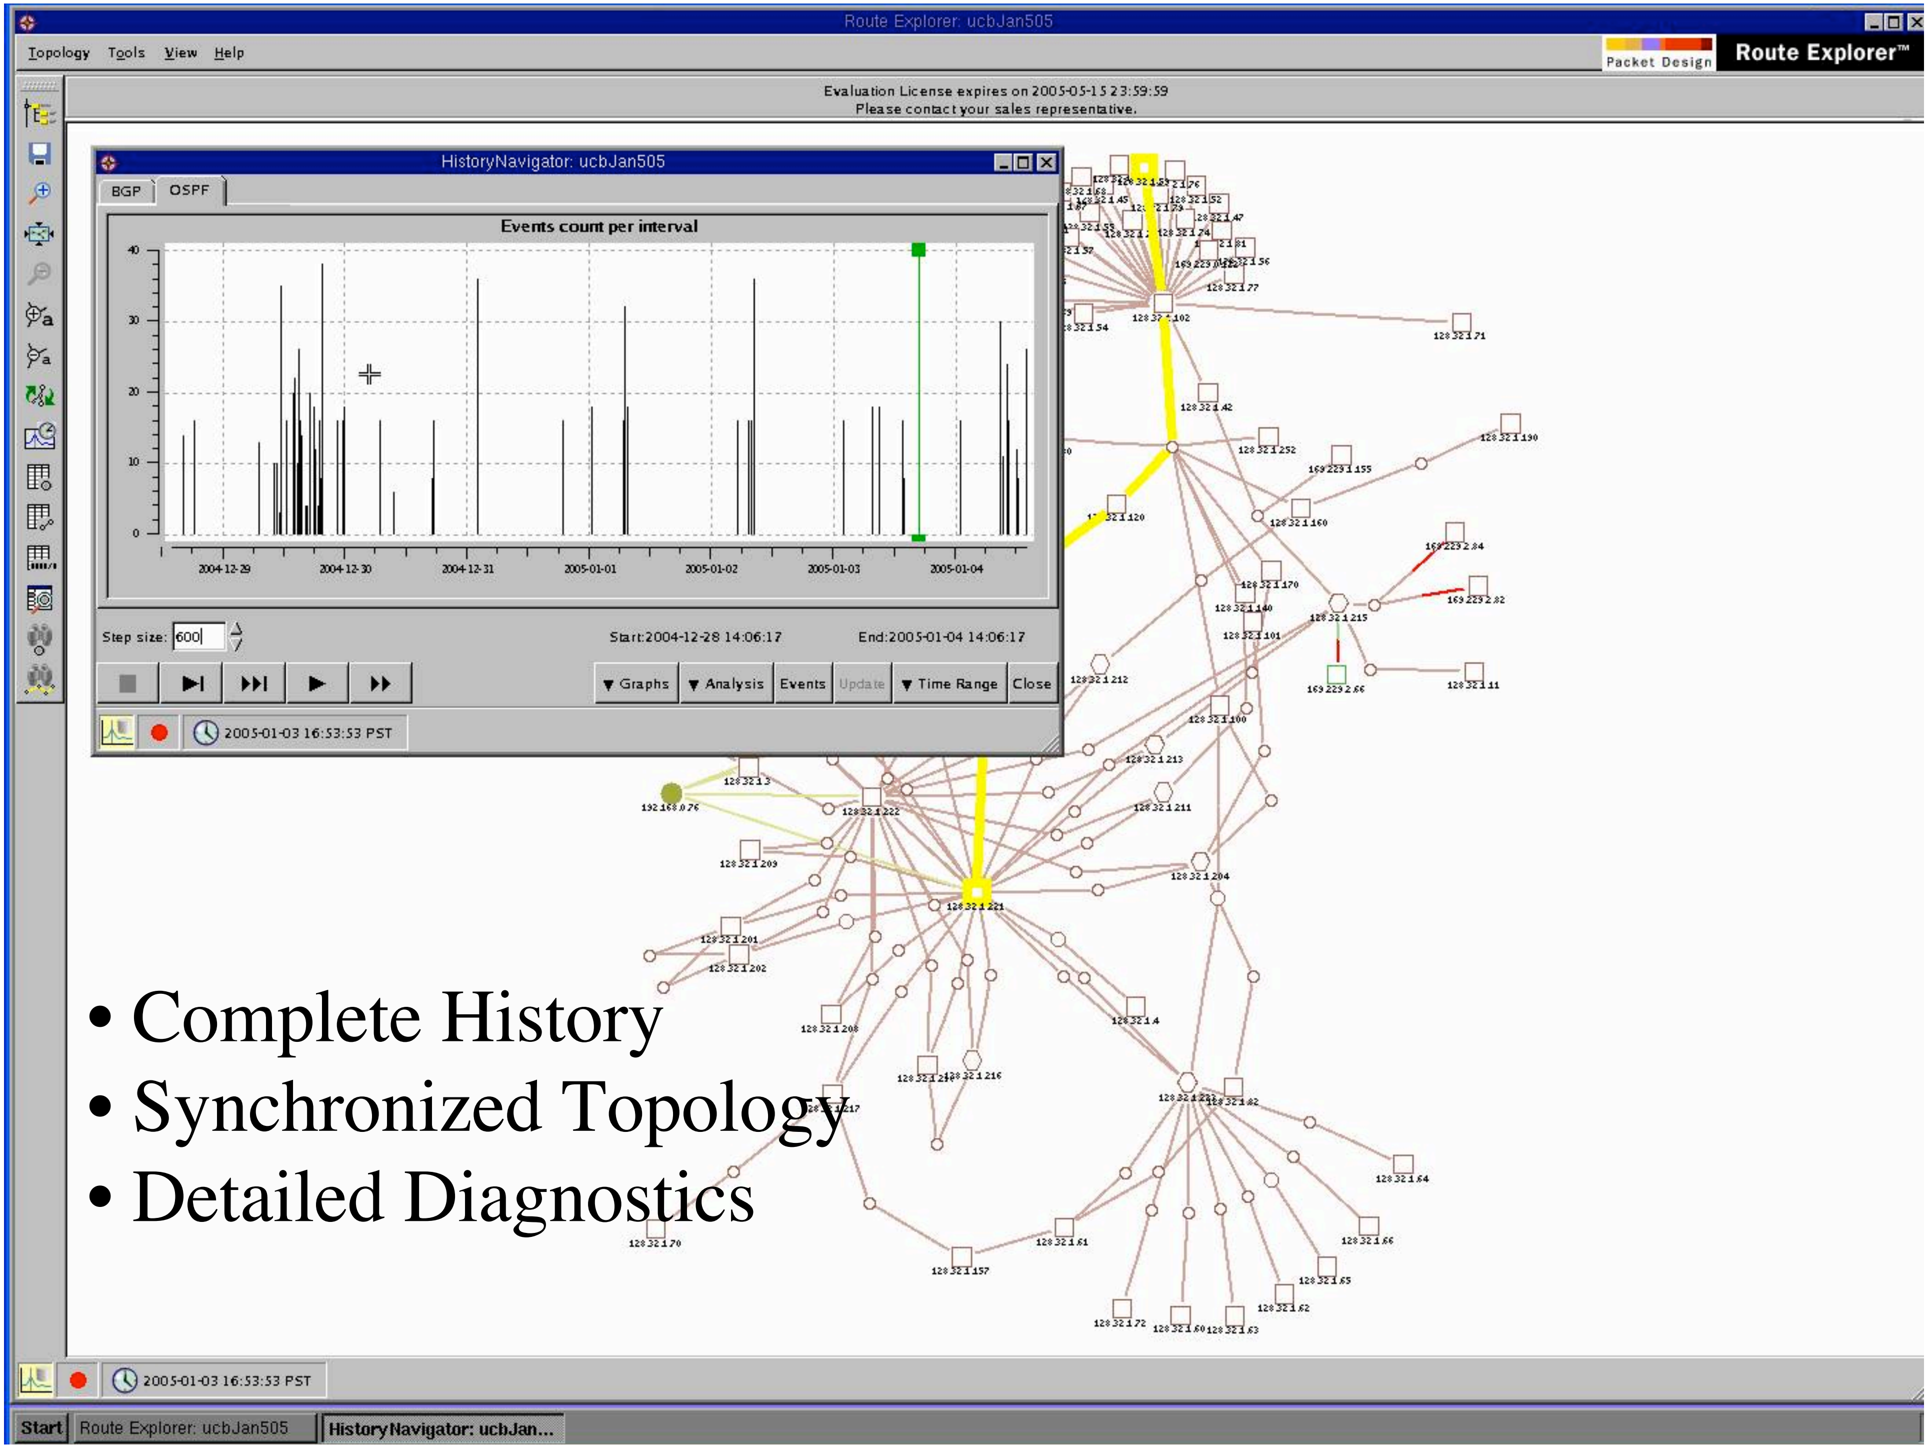

- Complete History
- Synchronized Topology
- Detailed Diagnostics

Who's Using Route Analytics Today?

- Large, meshed routed networks
- Networks with multiple Internet peerings
  - large enterprises
  - global service providers
  - government/military

Some Real-World Use Cases

- Highlighted:
  - MPLS VPN Per-Customer Layer 3 Monitoring
  - Routing Convergence/Propagation Measurement
  - Peering BGP Root Cause Analysis
- Other Use Cases
  - Change planning and validation
  - Failure scenario analysis
  - Customer forensics
  - Inter-protocol analyses

MPLS VPN Route Analytics Monitors:

- Per-Customer VPN Reachability
  - "Are all the customer VPN, BGP, IGP (OSPF, IS-IS, EIGRP) routes functioning properly?"
- Per-Customer VPN Privacy
  - Fundamental selling point of MPLS VPNs
  - "Are all customer VPN Route Distinguishers properly configured?"
- Per-Customer VPN Policy
  - "Is the customer's desired routing architecture (Full-mesh, hub and spoke, partial mesh, etc.) working?"

Provider's Concerns

- Customers can leak extra routes and consume precious resources
- Provider needs to monitor the customer's routes and routing activity
- Provider can misconfigure
  - Route Distinguishers: used to distinguish customer routes
  - Route Targets: used to determine intra-VPN routing policy
  - Distributing one customer's routes to another
  - VPN becomes no longer “private”.
- Provider needs to monitor PEs, RDs and RTs associated with each customer and flag any deviation

Baselining Customer VPNs

- For each route, determine how long the route was available over the last week
- If the route was available 80% of the time or longer, then include the route in the baseline
- Similarly for PEs, RDs, RTs
  - i.e. if a PE is announcing at least one route for a customer for 80% of the week, then include that PE in the baseline for that customer

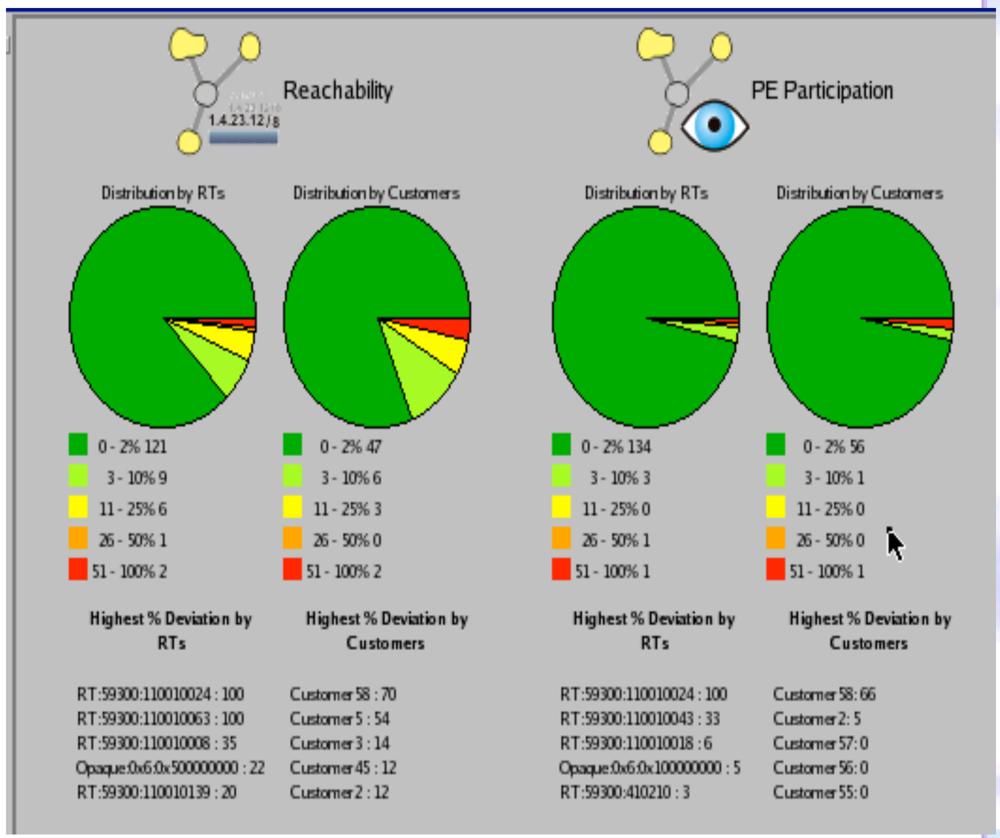

Route Deviation per Customer

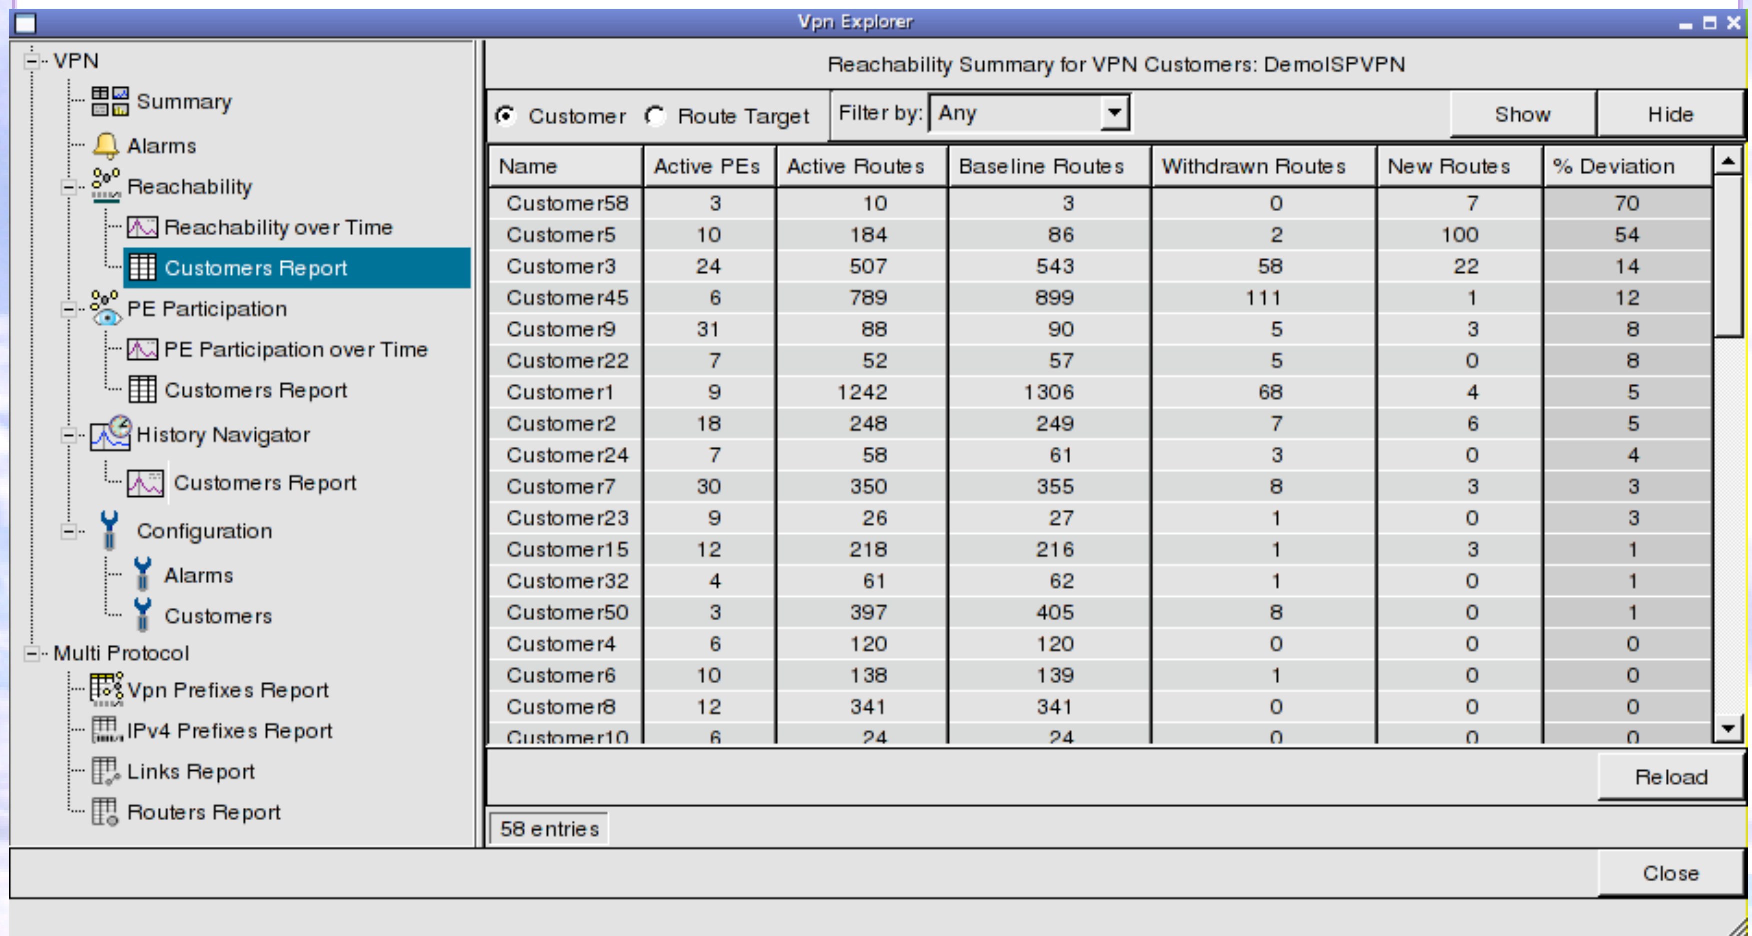

Name	Active PEs	Active Routes	Baseline Routes	Withdrawn Routes	New Routes	% Deviation
Customer58	3	10	3	0	7	70
Customer5	10	184	86	2	100	54
Customer3	24	507	543	58	22	14
Customer45	6	789	899	111	1	12
Customer9	31	88	90	5	3	8
Customer22	7	52	57	5	0	8
Customer1	9	1242	1306	68	4	5
Customer2	18	248	249	7	6	5
Customer24	7	58	61	3	0	4
Customer7	30	350	355	8	3	3
Customer23	9	26	27	1	0	3
Customer15	12	218	216	1	3	1
Customer32	4	61	62	1	0	1
Customer50	3	397	405	8	0	1
Customer4	6	120	120	0	0	0
Customer6	10	138	139	1	0	0
Customer8	12	341	341	0	0	0
Customer10	6	24	24	0	0	0

Packet Design

Breakdown Customer5's Deviation to its PEs

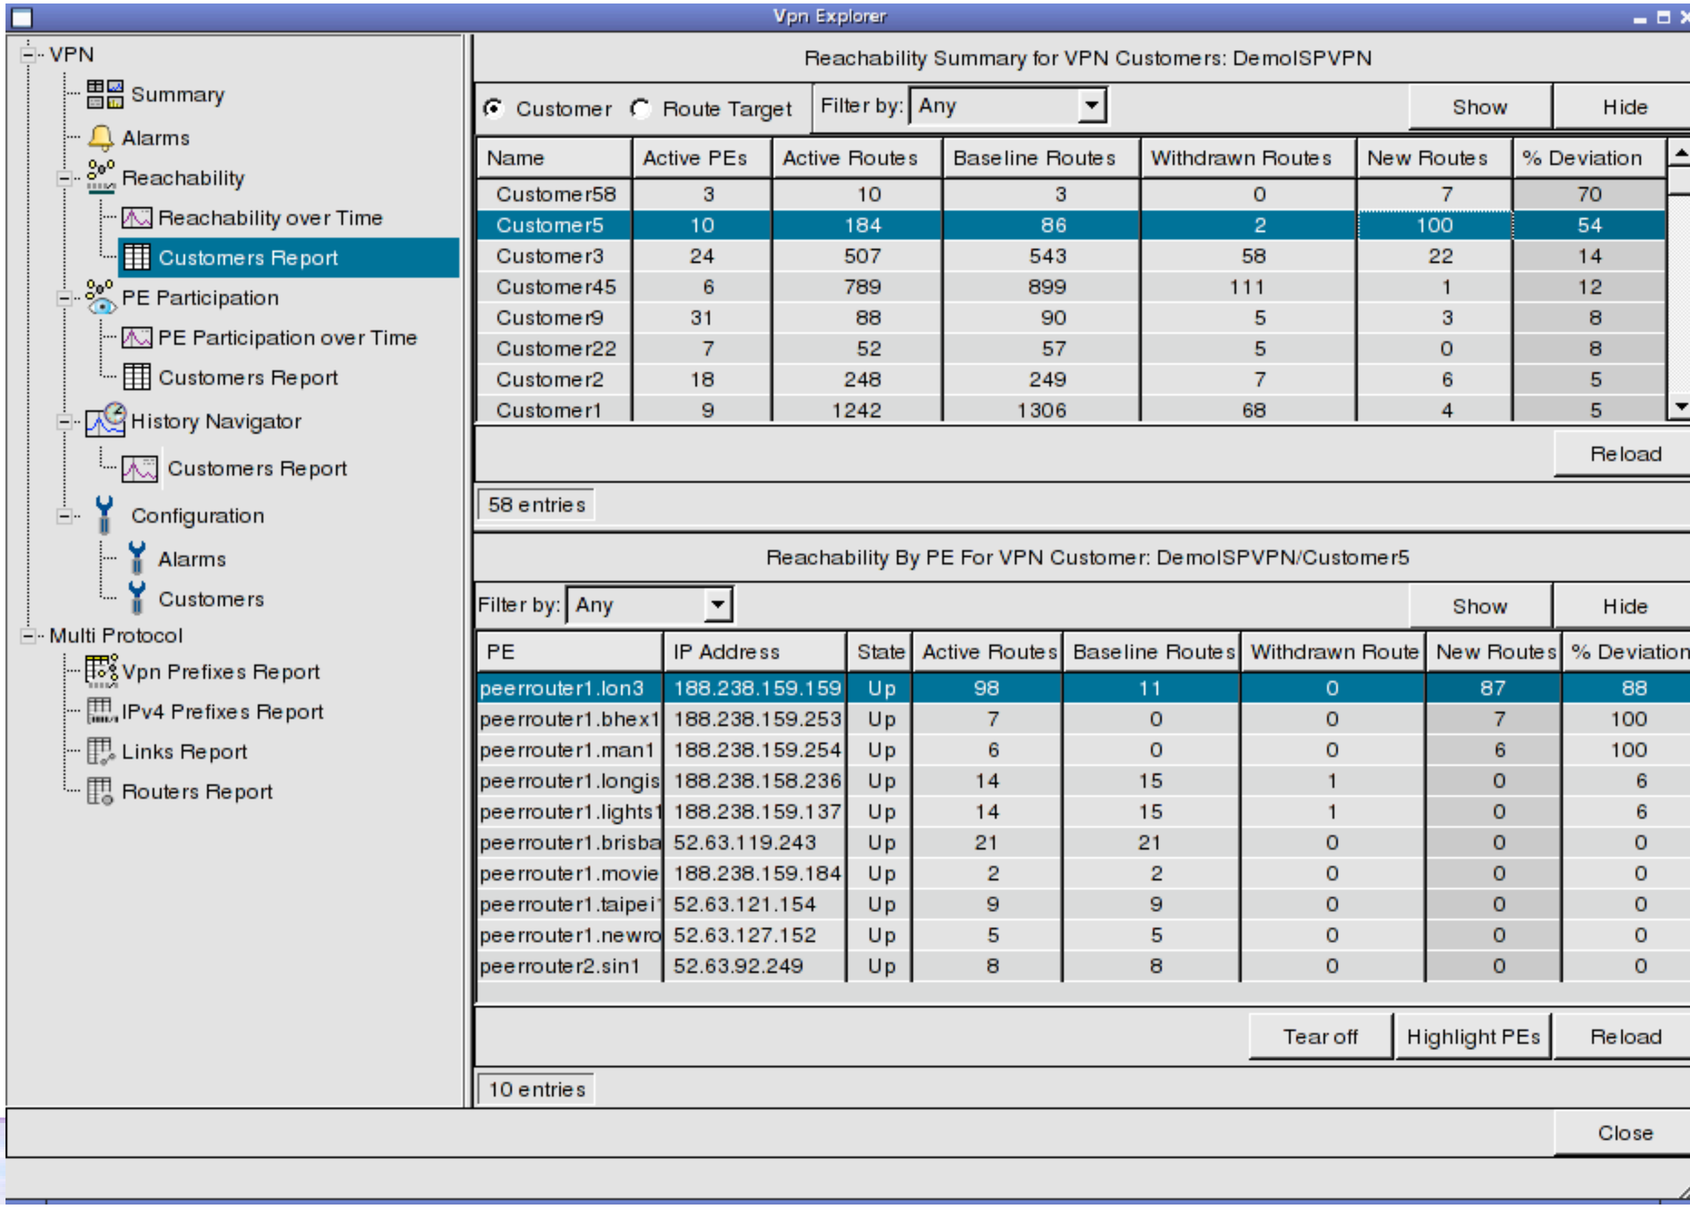

Name	Active PEs	Active Routes	Baseline Routes	Withdrawn Routes	New Routes	% Deviation
Customer58	3	10	3	0	7	70
Customer5	10	184	86	2	100	54
Customer3	24	507	543	58	22	14
Customer45	6	789	899	111	1	12
Customer9	31	88	90	5	3	8
Customer22	7	52	57	5	0	8
Customer2	18	248	249	7	6	5
Customer1	9	1242	1306	68	4	5

PE	IP Address	State	Active Routes	Baseline Routes	Withdrawn Route	New Routes	% Deviation
peerrouter1.lon3	188.238.159.159	Up	98	11	0	87	88
peerrouter1.bhex1	188.238.159.253	Up	7	0	0	7	100
peerrouter1.man1	188.238.159.254	Up	6	0	0	6	100
peerrouter1.longis	188.238.158.236	Up	14	15	1	0	6
peerrouter1.lights1	188.238.159.137	Up	14	15	1	0	6
peerrouter1.brisba	52.63.119.243	Up	21	21	0	0	0
peerrouter1.movie	188.238.159.184	Up	2	2	0	0	0
peerrouter1.taipei1	52.63.121.154	Up	9	9	0	0	0
peerrouter1.newro	52.63.127.152	Up	5	5	0	0	0
peerrouter2.sin1	52.63.92.249	Up	8	8	0	0	0

Packet Design

Two Questionable PEs

- It is not only that there are 100 new routes for Customer5 that is troubling
- 13 of these routes are coming from PEs that are not part of this customer's baseline PEs
  - Perhaps customer is growing in number of sites
  - Perhaps it is a result of misconfiguration and these routes belong to a different customer
  - It needs to be verified by an operator
- How long has this been going on?

Customer VPN instability

- Customer3's network is causing instability

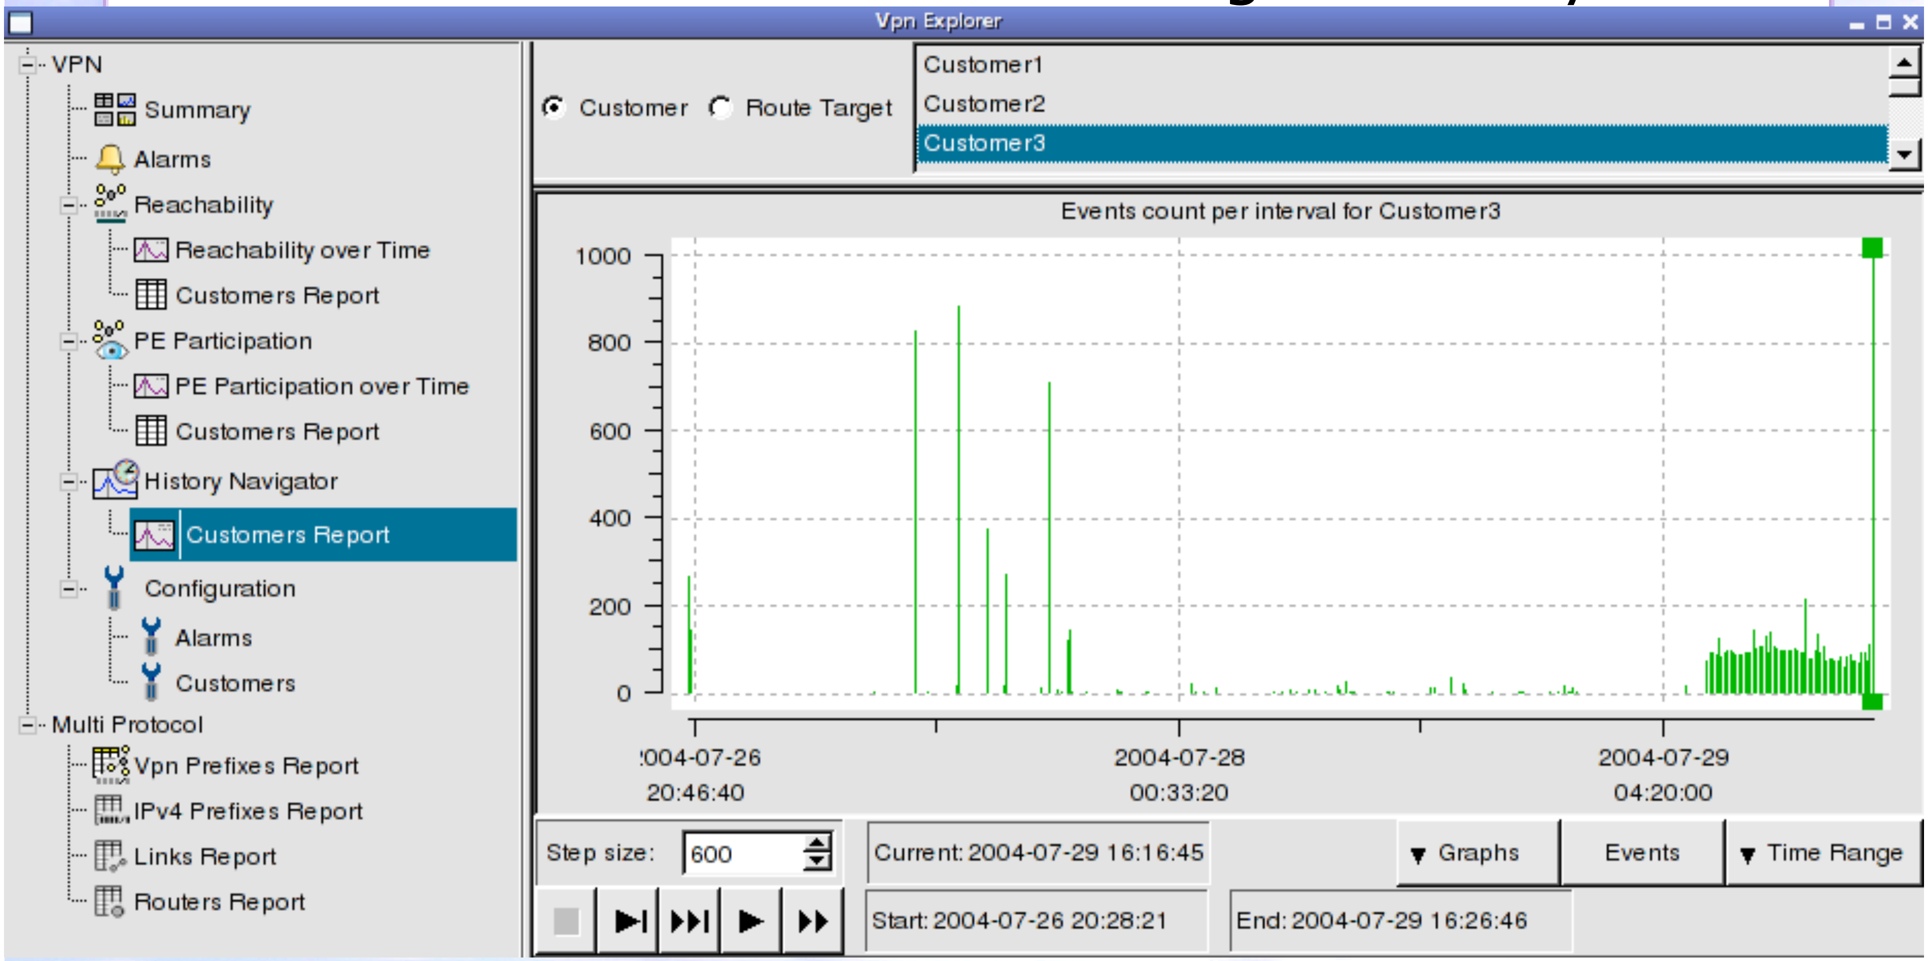

Route Convergence Measurement

- Challenge: Large financial WAN RFP required routing convergence SLA
- Solution Overview:
  - Deploy multiple route analytics appliances and routing event generator to measure synthetic route event propagation delay
  - Monitors routing “convergence” through measurement of routing event propagation times
  - Uses route analytics real-time recording of all routing updates

Convergence Measurement Architecture

- All systems are enabled with NTP and are synchronized to the same timezone for accurate time-stamping of routing events

- Routing event propagation time is calculated by extracting event timing data from Route Explorer's XML API

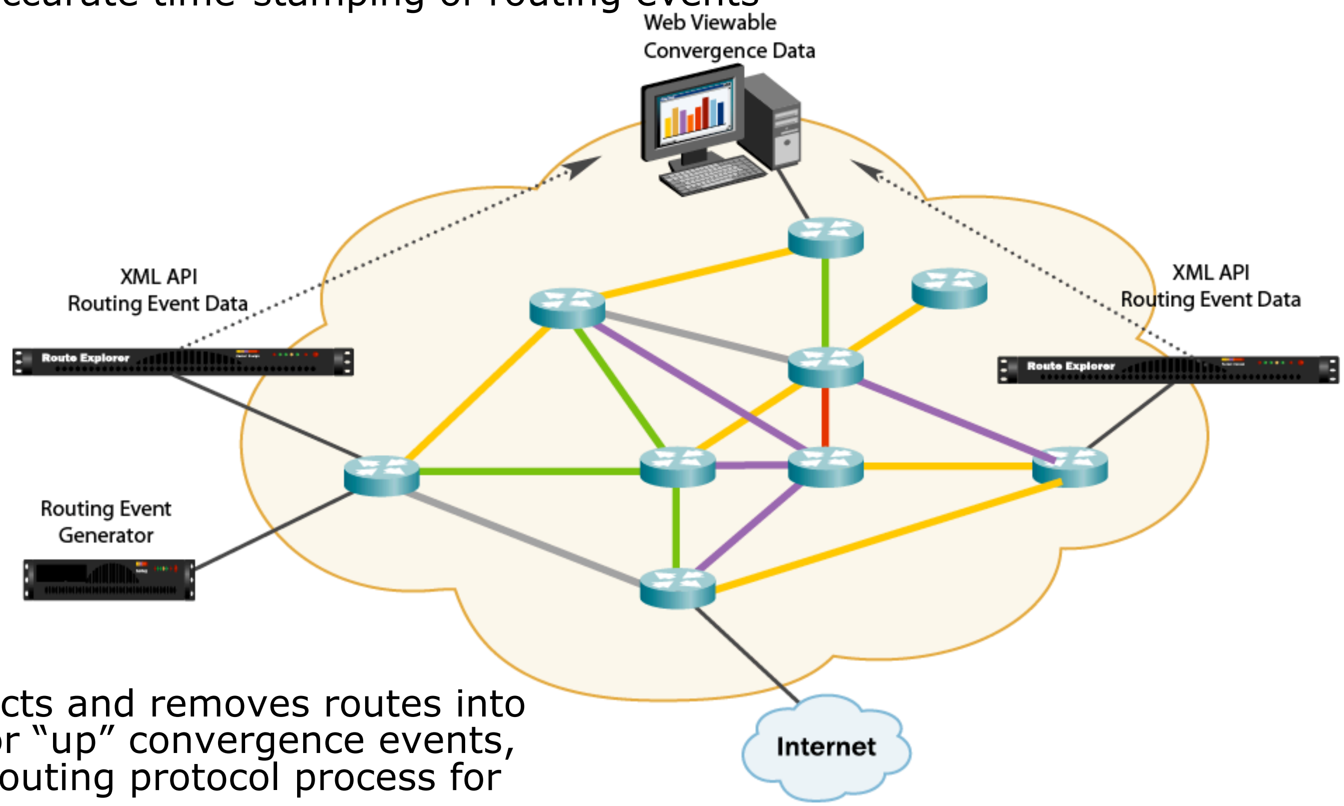

- Generator injects and removes routes into the network for "up" convergence events, and stops its routing protocol process for "down" events

BGP Root Cause Analysis

- Hard to diagnose BGP problems:
  - BGP was designed to facilitate routing, not diagnosis.
  - Deluge of data
    - Large number of routes and paths on a router, and the number of options for each route. Differences in a route may be subtle, but important, like path length or MED, but comparing them is not simple.
    - Most minor connectivity change produces hundreds of messages and a major peering loss can generate millions.
    - Many error-prone configuration knobs
- Route analytics addresses a number of scenarios

Peering Reset Between Two ISPs

- A tier-1 ISP's customer leaked thousands of routes
- Tier-1 ISP announced them to peers
- One peer had a prefix-limit configured and reset the session severing the communication between the two ISPs
- Tier-1 ISP needed to find what were the new routes, who was injecting them

Packet Design

Leaking Routes

- When an AS leaks routes, BGP sends many messages

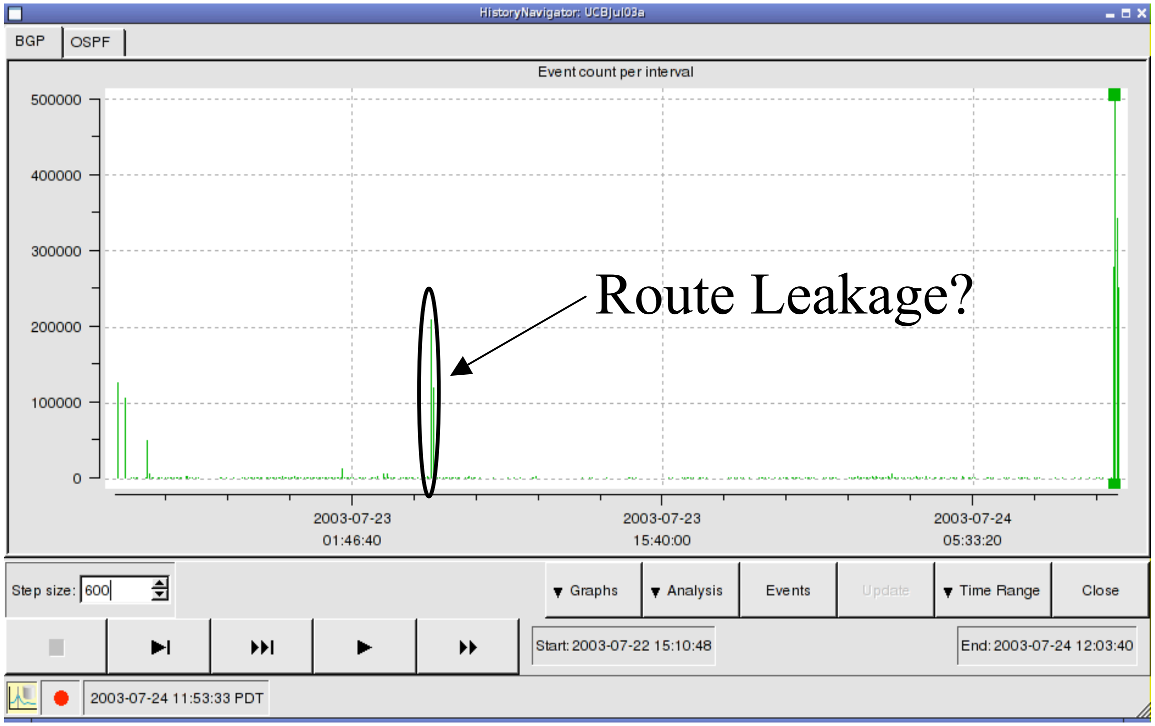

BGP Root Cause Analysis

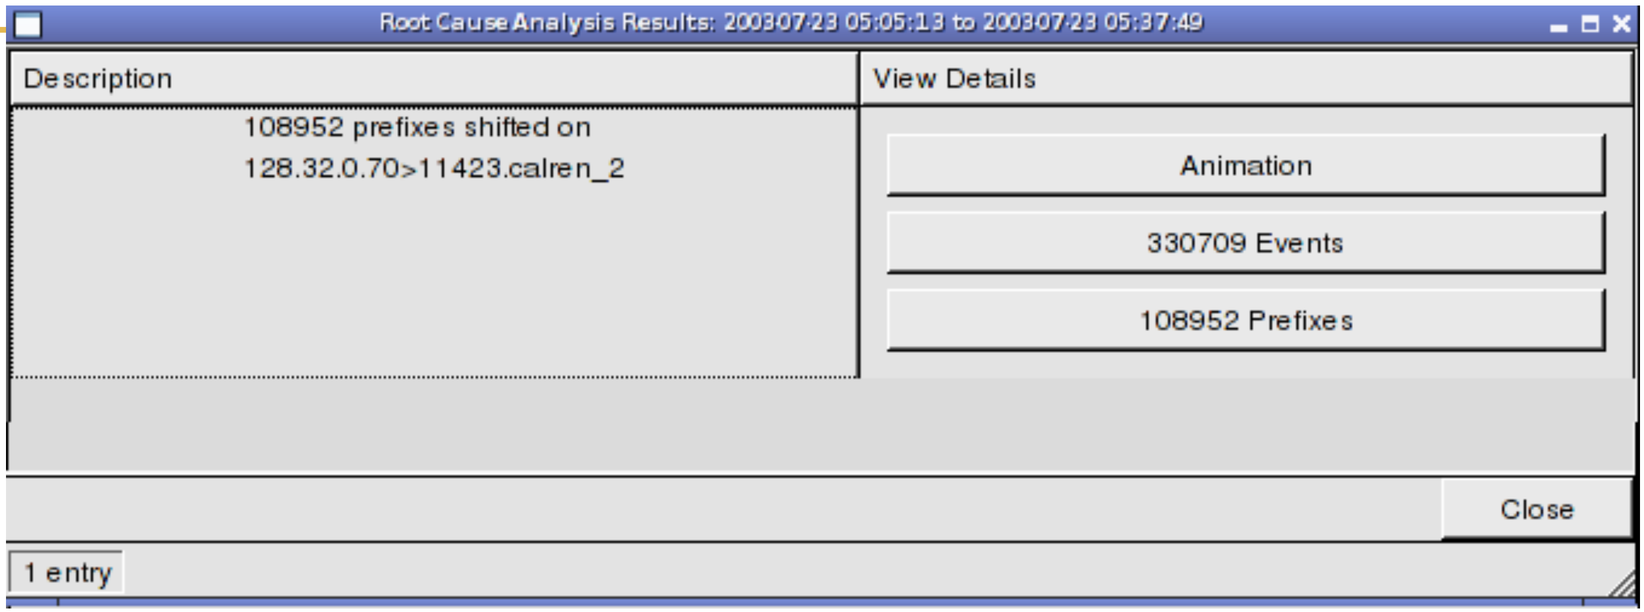

- RCA found only one thing happened: a peering was reset, prefixes shifted to another peering, and then back
  - BGP had to send 330709 announces/withdrawals

Packet Design

Visualization of Peering Reset

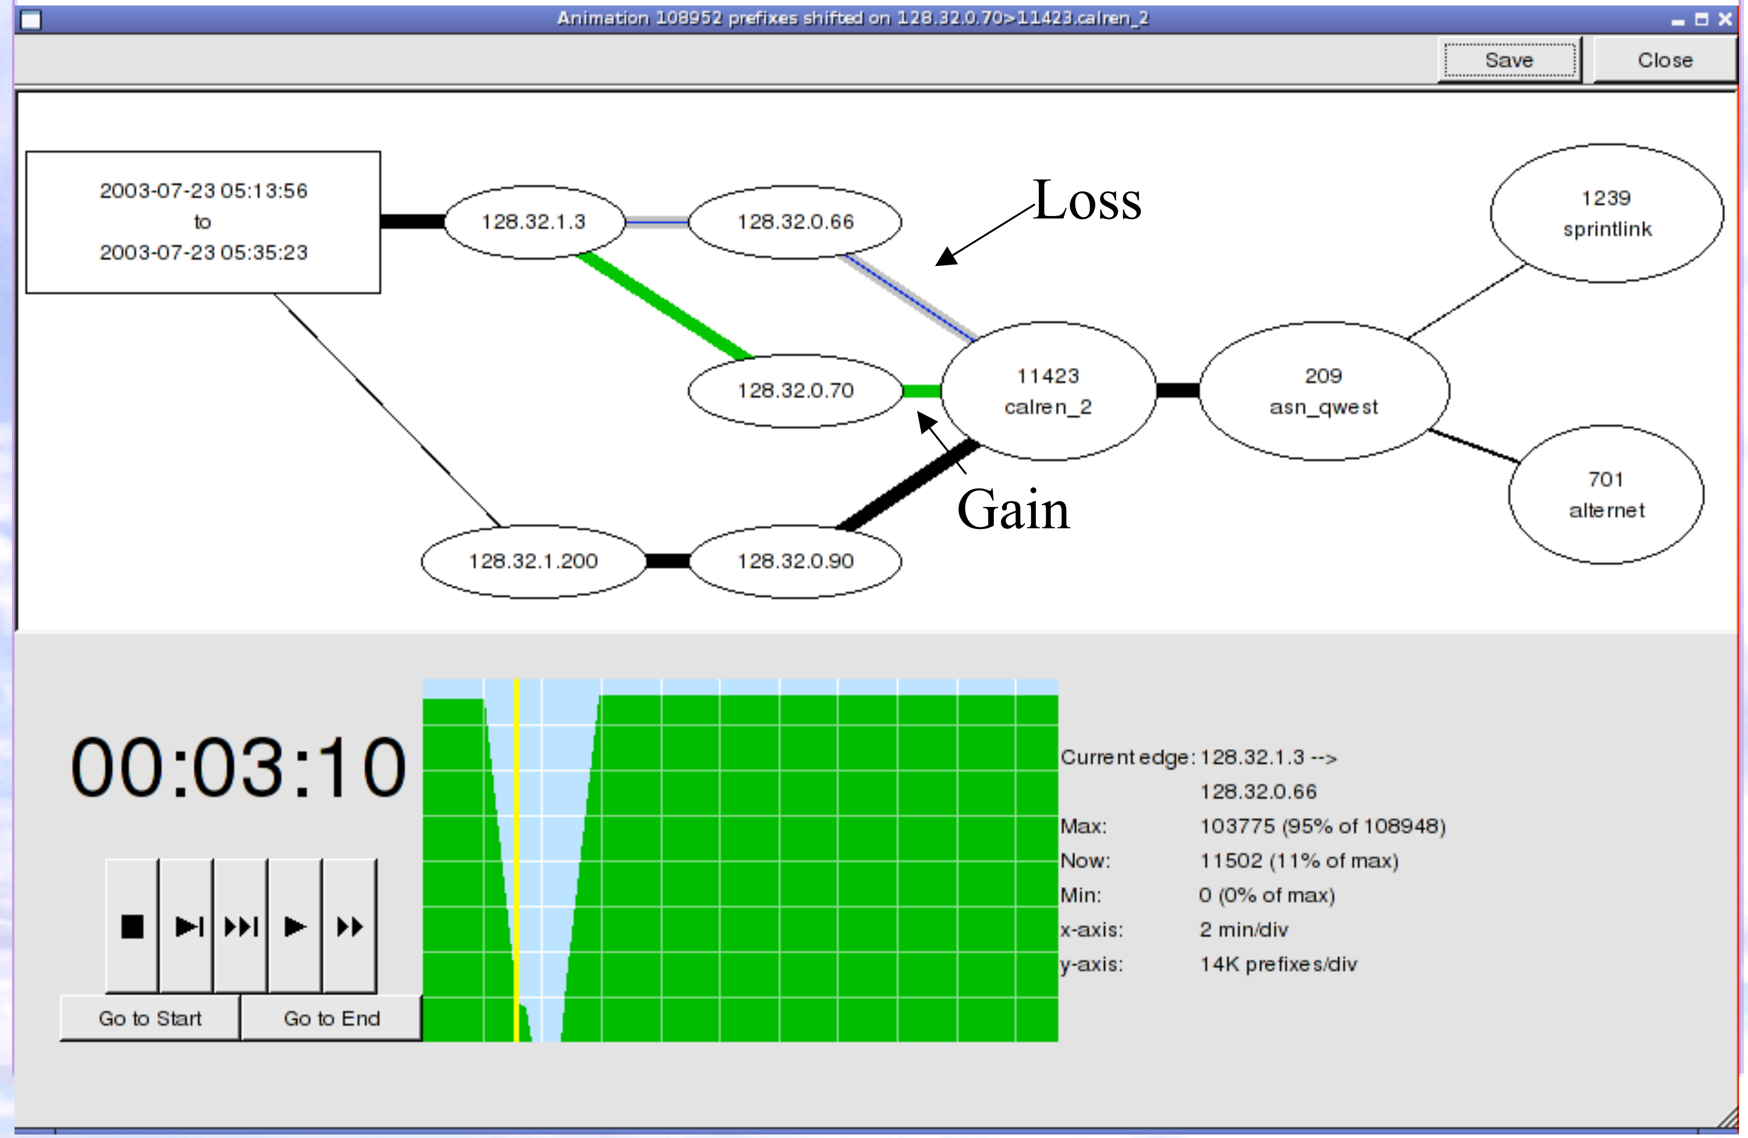

Route Leakage in an RCA Visualization

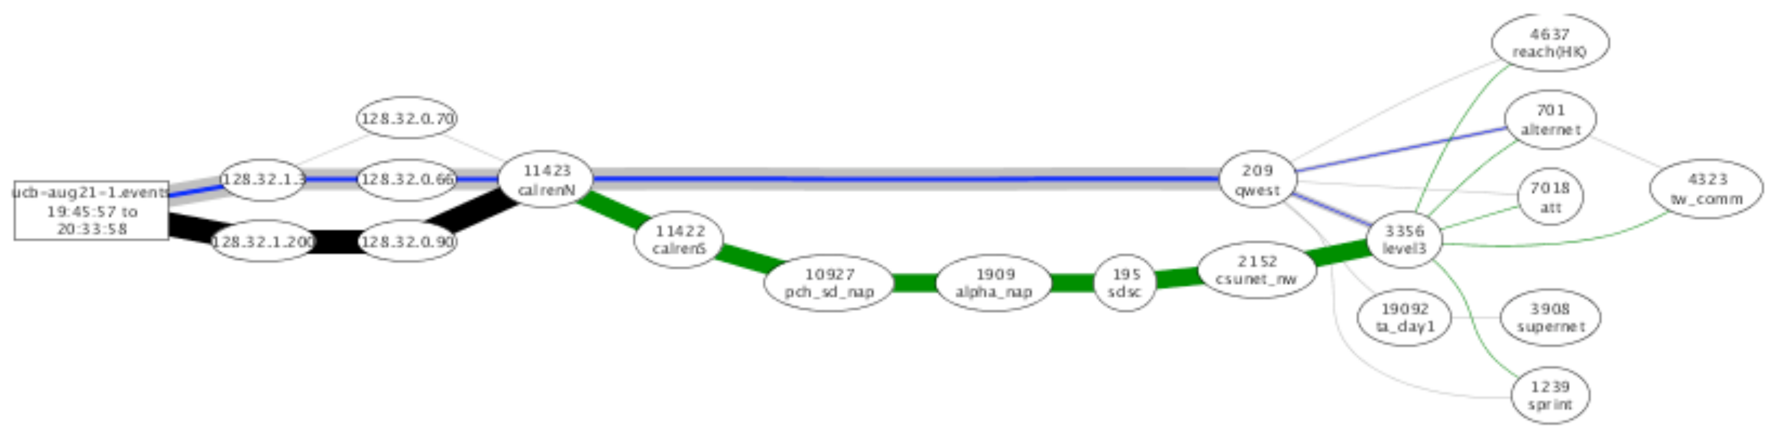

- About a million BGP events revealed
  - Peer leaking routes causing sub-optimal routes
  - Community mishap treats commercial routes as academic

Packet Design

Before and After Comparison

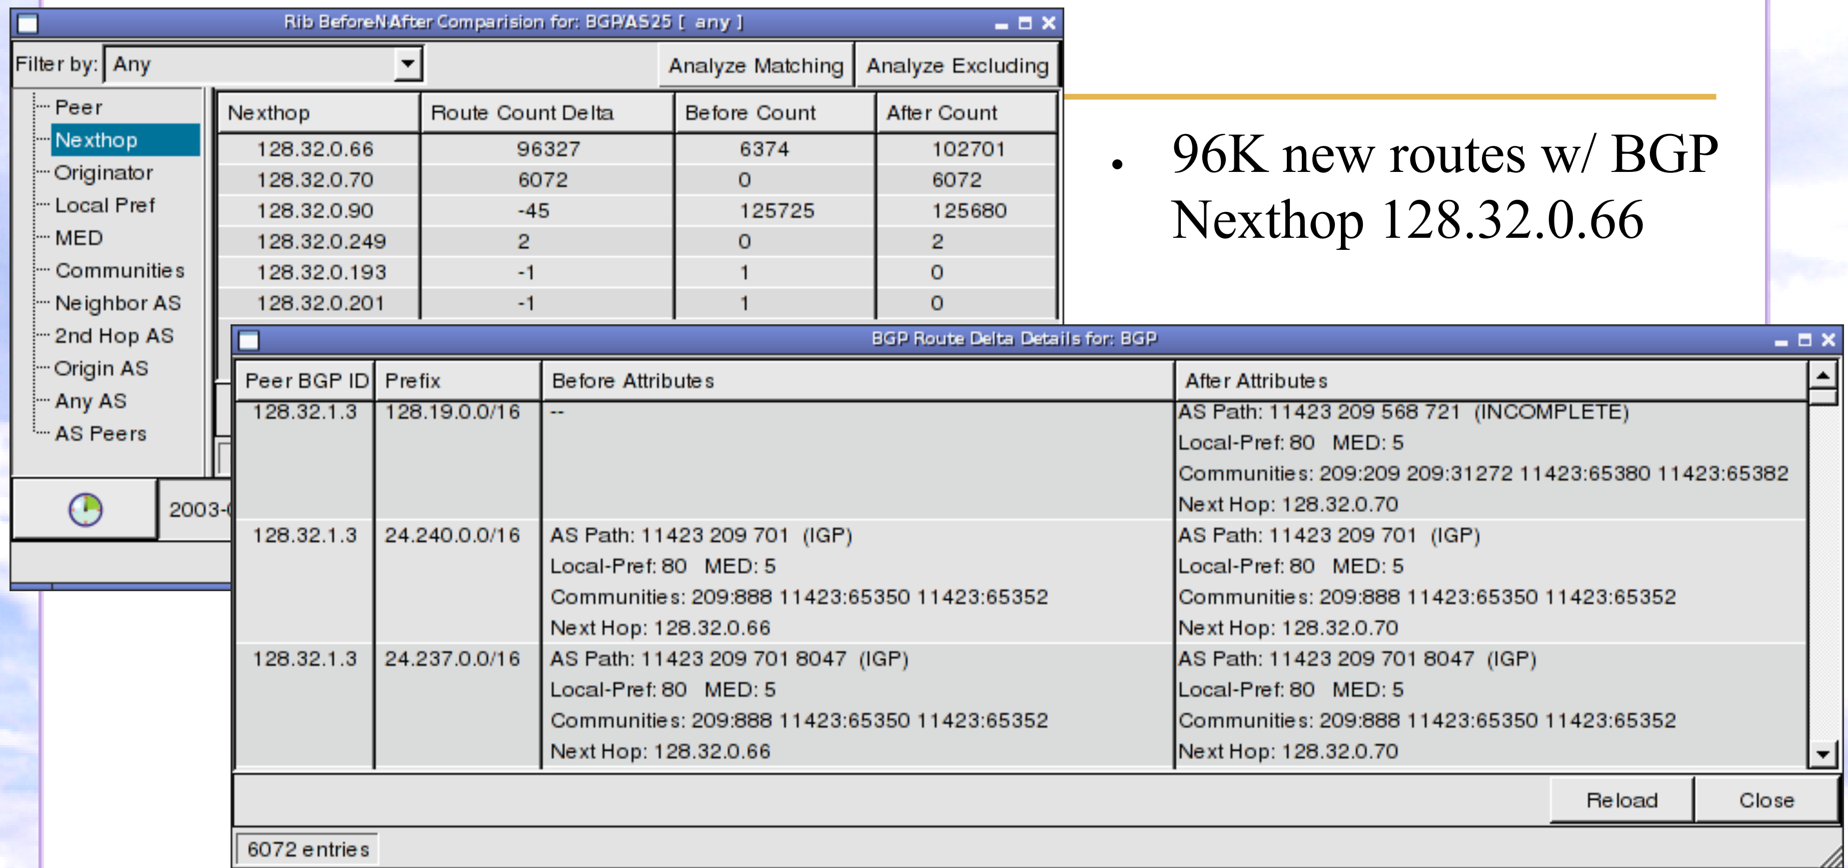

- 96K new routes w/ BGP Nexthop 128.32.0.66

- This was the feature used to determine the customer that leaked routes

MED Oscillations

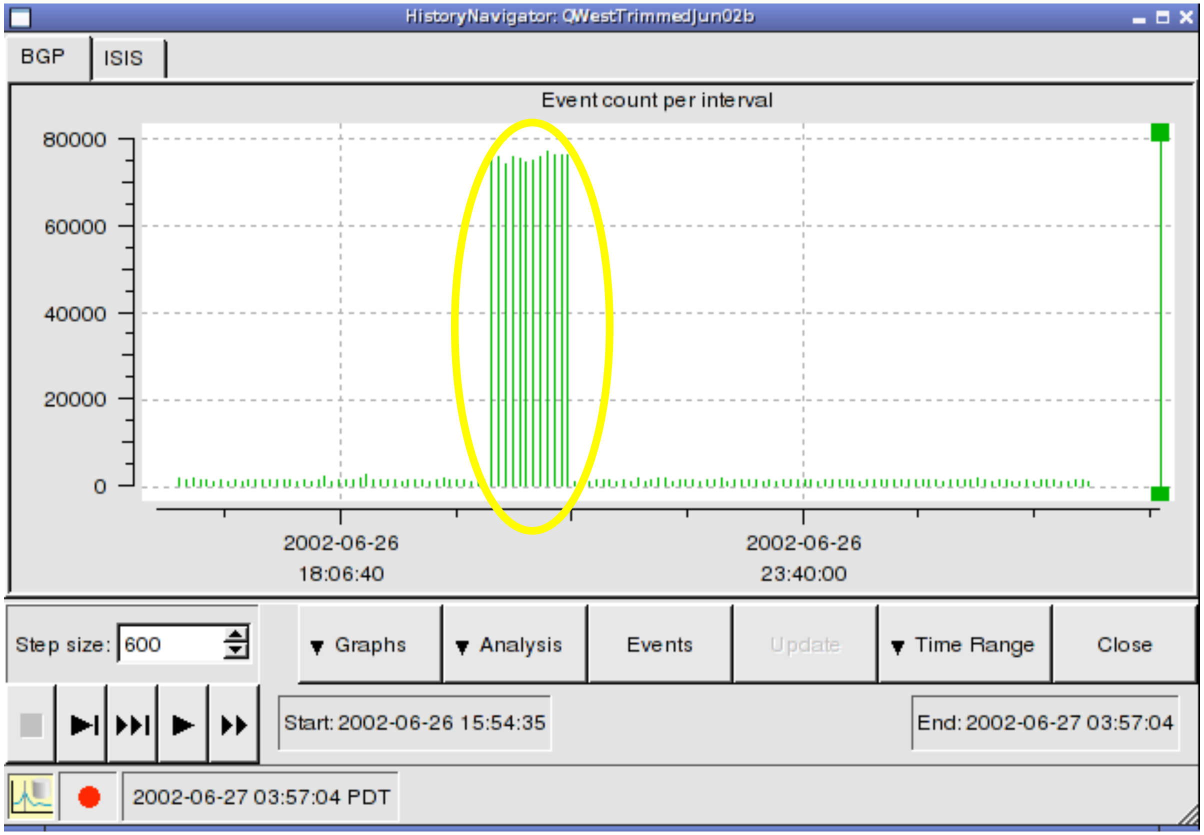

- MED oscillations caused BGP updates every 10s of microseconds
  - Expect very high router CPU utilizations across the network

MED Oscillations

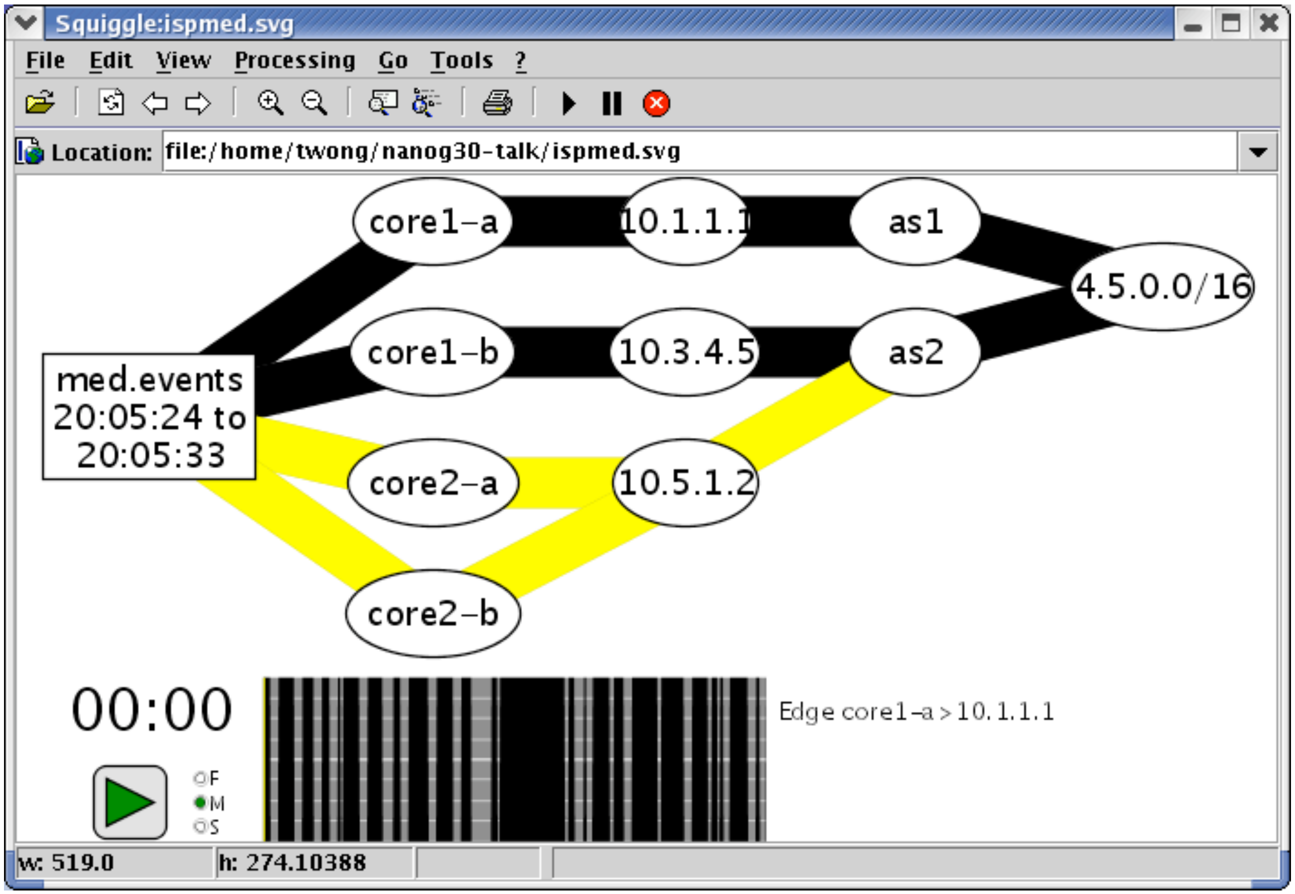

BGP Peering Flaps

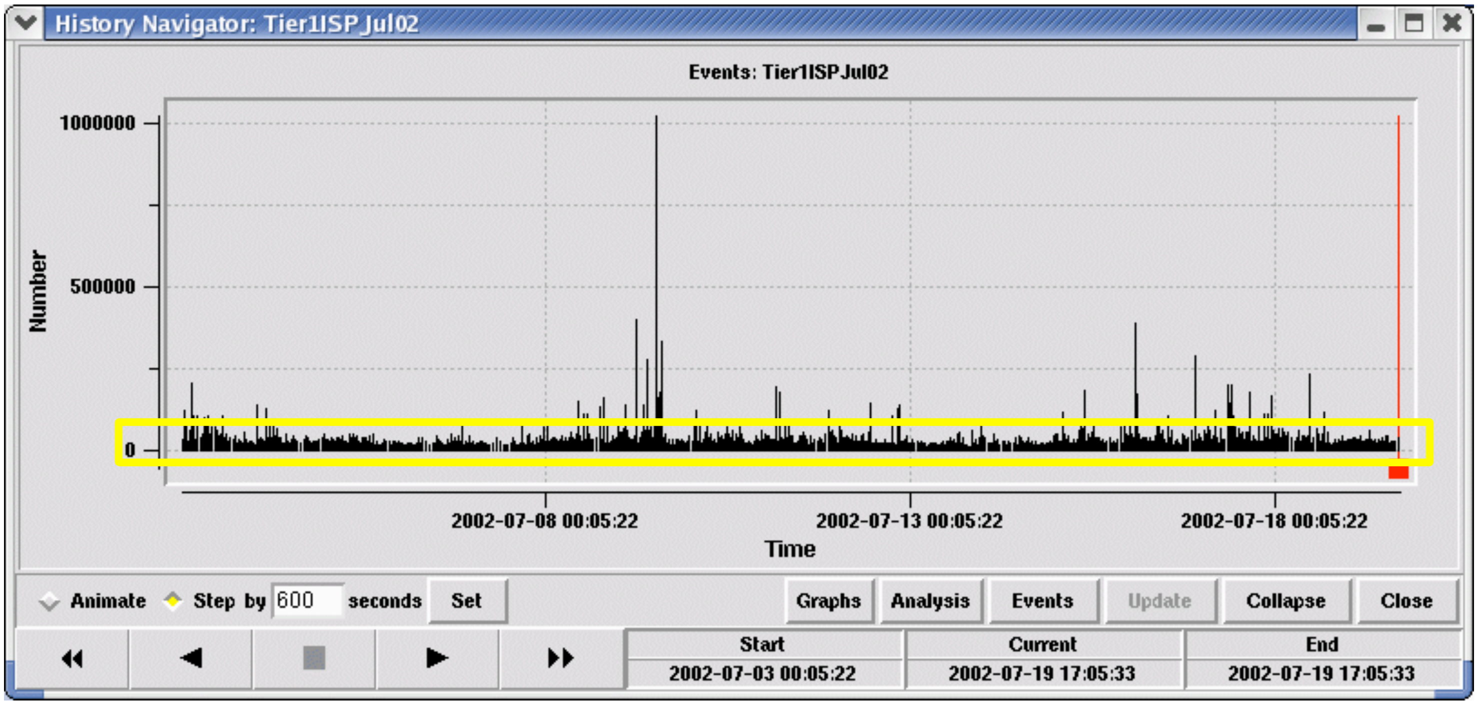

- Persistent route flapping once a minute looks like BGP noise here (yellow box)
  - This lasted for at least 1.5 months, and resulted in a very unhappy customer.

Persistent Customer Route Flaps

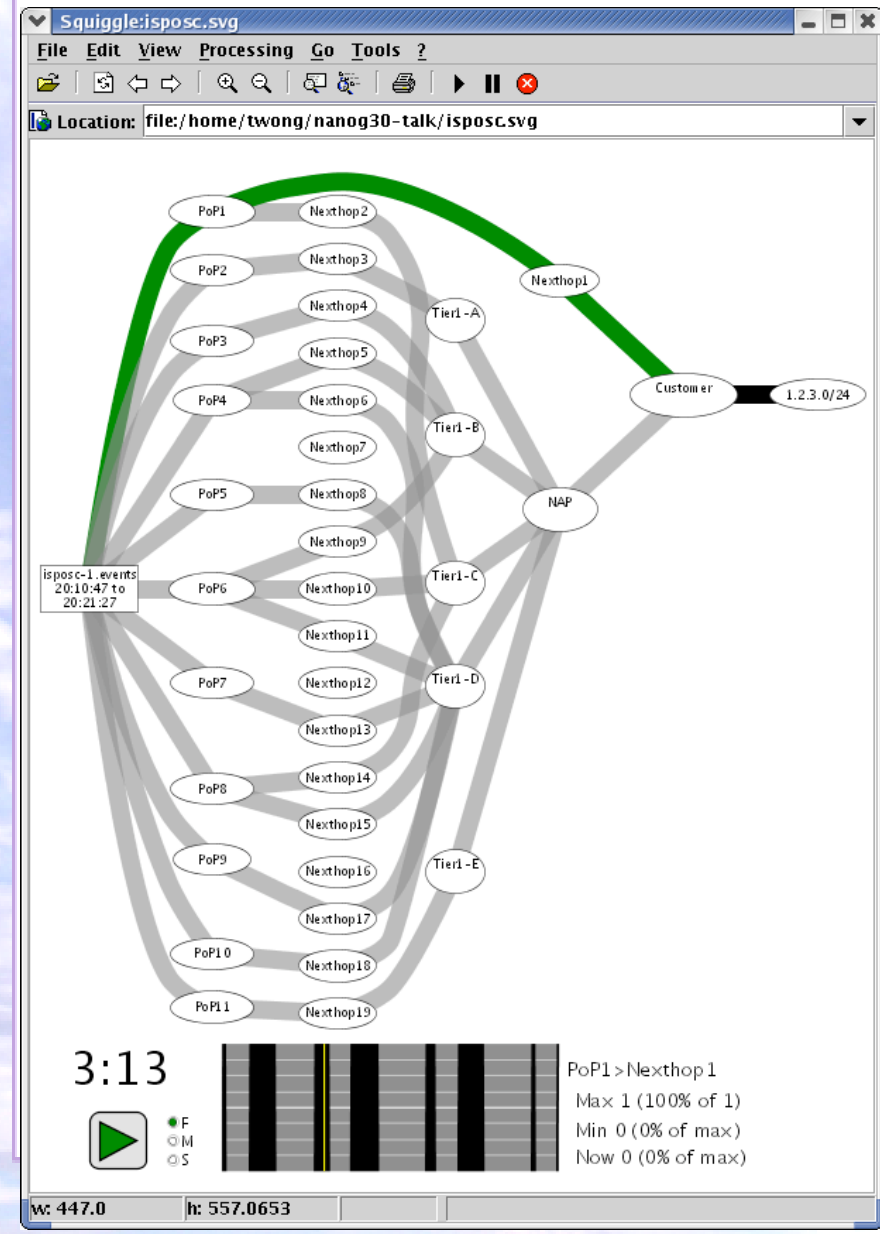

- Route is not flap damped
- When observed from peers of this ISP, the route is still a candidate for route flap damping
  - Customer is off the Internet

Route-Flow Fusion

- How it works:
  - Take limited number of flow exporters
  - Use IGP/BGP to project flows across all links
- The result:
  - Arrive at network-wide traffic/flow distribution
  - Correlated routing and traffic event/change analysis
- Applications:
  - Comprehensive peering and transit analysis
  - Modeling of network changes on actual routing and traffic to show precise impact
  - Know if network contributed to any traffic change
  - Optimize capacity planning, failure response
  - Service availability, customer forensics