

# *Route Analytics*

---

*Discover, Resolve and Prevent  
the hardest problems in IP Networking*

# IP Networks Are Far from Predictable

---

- Unpredictable behavior still plagues IP networks:
  - Brown-outs, intermittent disruptions and performance degradations, mysterious problems
- Most of the events in IP networks leave no audit trail, and often go unexplained
- Unpredictable and unexplainable behavior is a major challenge:
  - Disrupts application performance and availability
  - Impacts user and IT productivity
  - Reduces confidence in IT



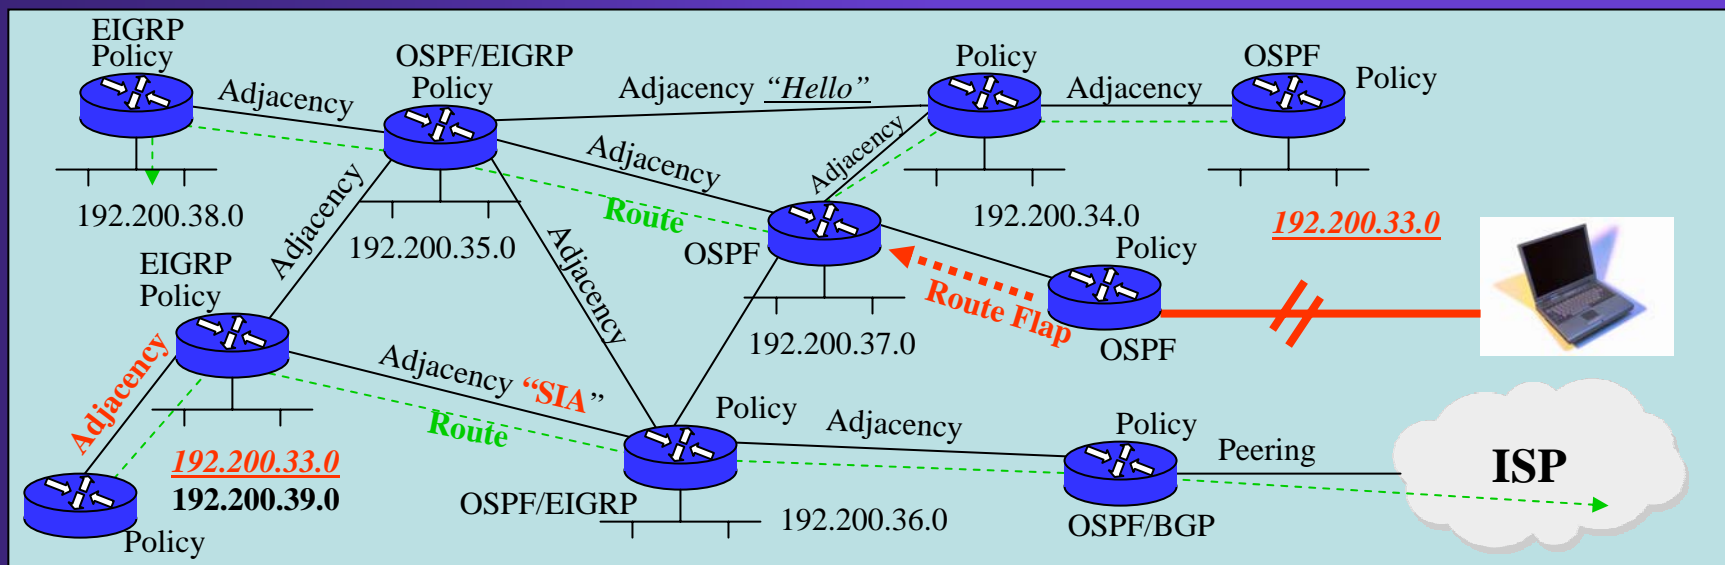
Users



IT

# Key Problem: Traditional Network Management Focused on Devices

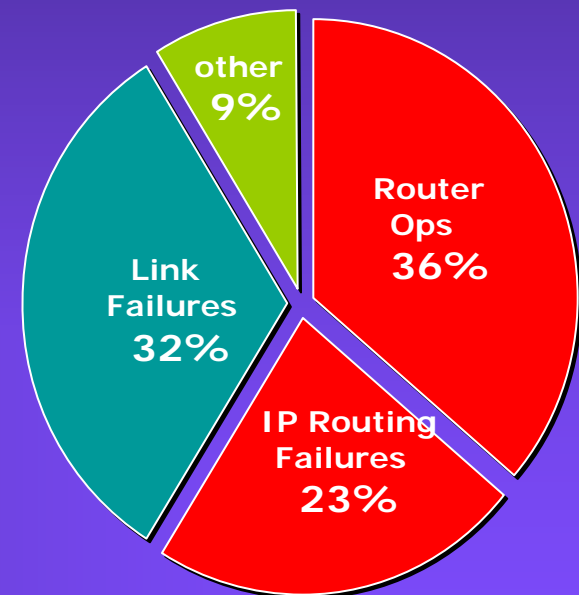
- IP Networks not simply a collection of device elements
  - Up/down, CPU, memory, interface up/down/stats, counters
- Complex network-layer ***interactions*** between devices
  - Routers, switches, DNS/DHCP, hosts
- 1000x volume of ***logical elements***
  - Protocols, protocol events, links, adjacencies, routes, prefixes



# Root Cause of Unpredictable Behavior Going Unmanaged

- Dynamic network-layer problems:
  - Inter-device and multi-device
  - **Logical**, rather than physical in nature
  - Network-wide rather than local
  - Stem from the **accumulation** of:
    - Misconfigurations
    - Software bugs
    - Rapid rate of network change due to dynamic IP routing protocols
  - **Often latent** until exposed by adverse or unforeseen events
  - Often exhibit symptoms only **intermittently**
  - Generate high volumes of indicators **too fast for SNMP** polling to detect
    - 1000's of routing events per second

Causes of Downtime in IP Networks



Sources: University of Michigan, Sprint

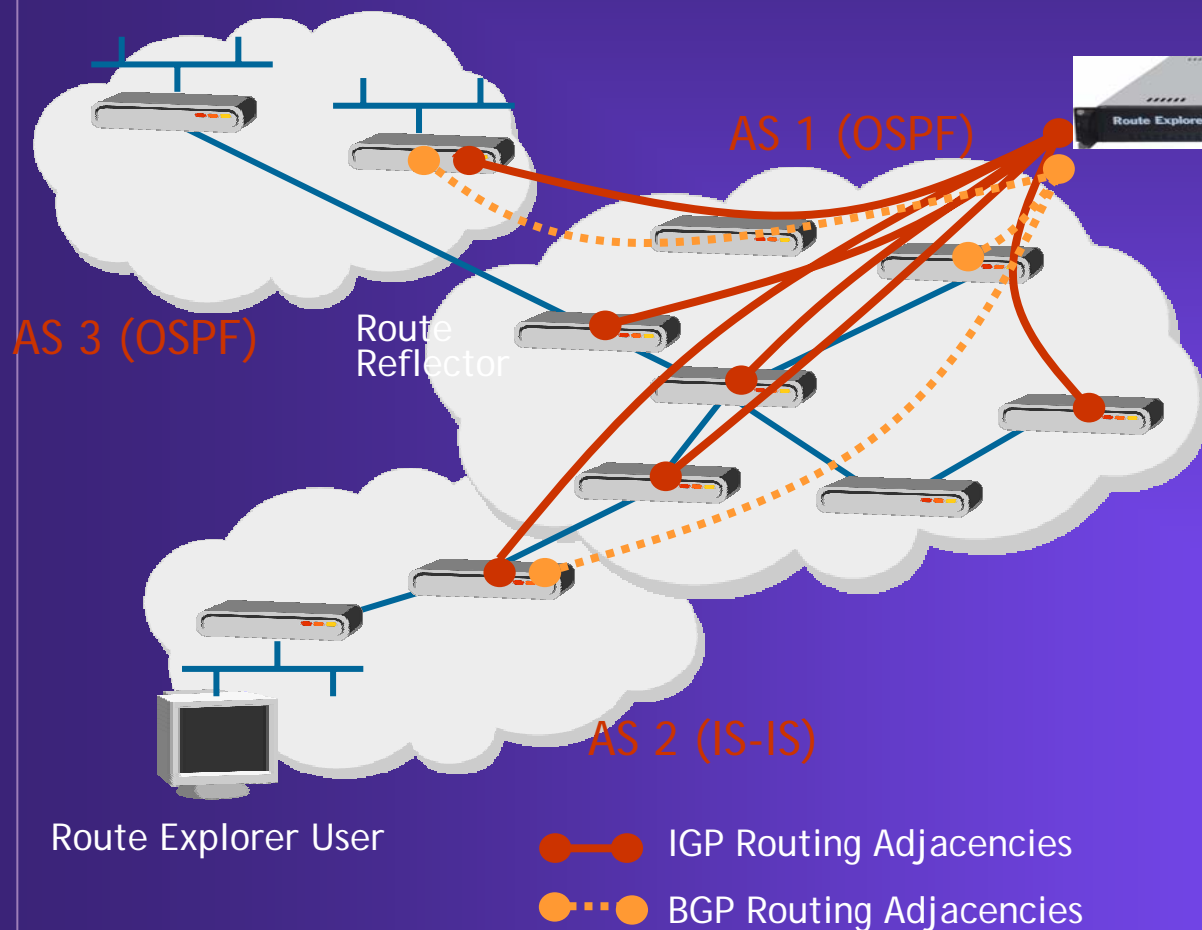
# Route Analytics Goals

---

- *Complete concurrent monitoring* and analysis of complex networks with multiple routing protocols and multiple ASes
- Generate *real-time IP network map* allowing visibility into how traffic is being routed for the first time
- *Alert generation* for potential critical problems so steps can be taken to prevent an outage
- *Historical report generation* summarizing key variables related to routes, reducing analysis time
- *Scenario planning* on “as running” network to validate operational plans -- before making the changes
- *Tracking of changes* as they’re deployed for validation

# Route Explorer™

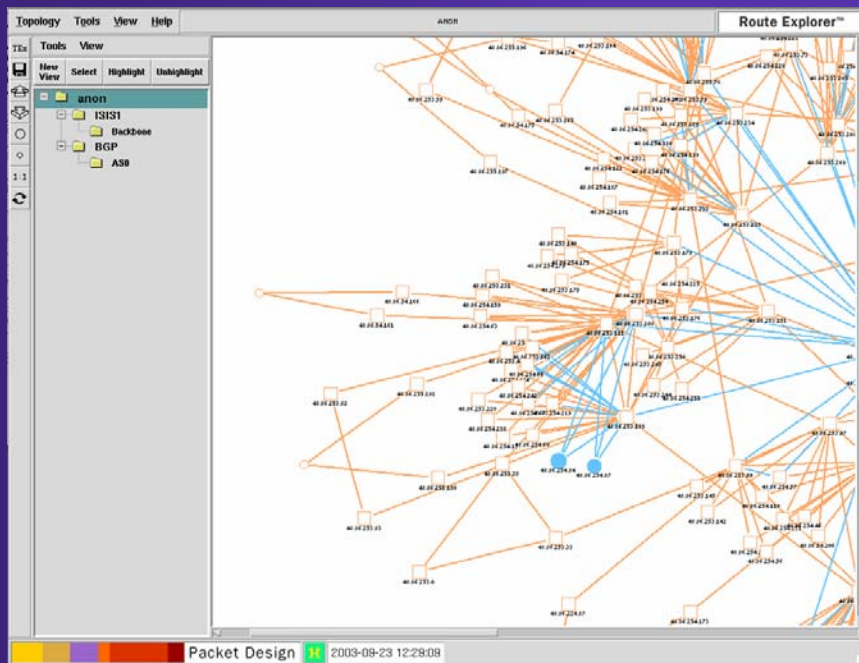
## Appliance-based Route Analytics



- Complete concurrent monitoring of multiple routing protocols – OSPF, IS-IS, EIGRP, BGP, mBGP (2547bis)
- A single appliance can monitor multi-AS networks
- Provides routing protocol-specific or network-wide viewing and analysis

# Unified View of Network in Real-Time

## *Unique Live View of the Routed Network*

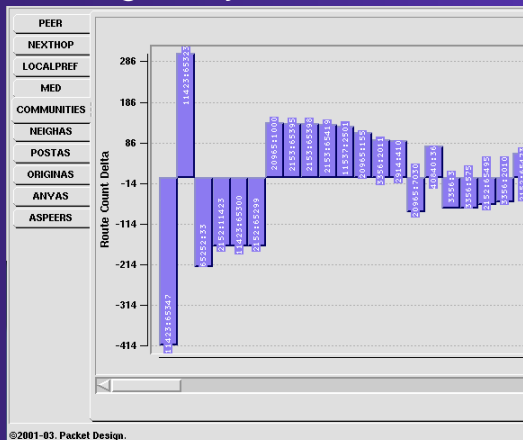


- Operates in the IP control plane – just like a router
- As changes are detected in the network, the topology map is instantly updated
- Detailed data can be easily accessed: link status, link metrics, new prefixes
- A specific source/destination can be highlighted for viewing of the active route between routers


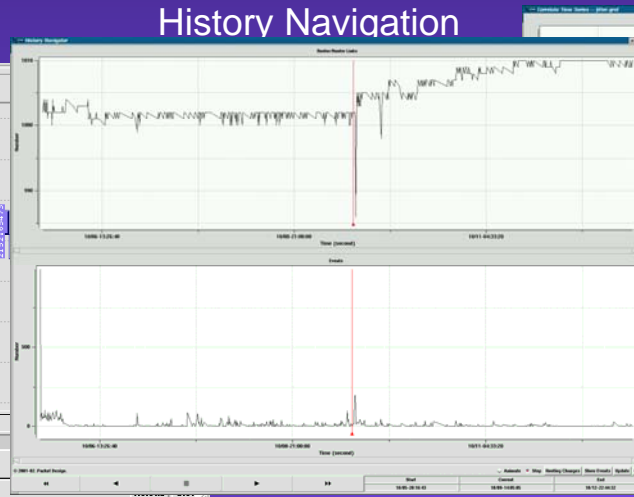
# Reduce Time to Analyze Network Problems

Correlated External Data

Routing Analysis



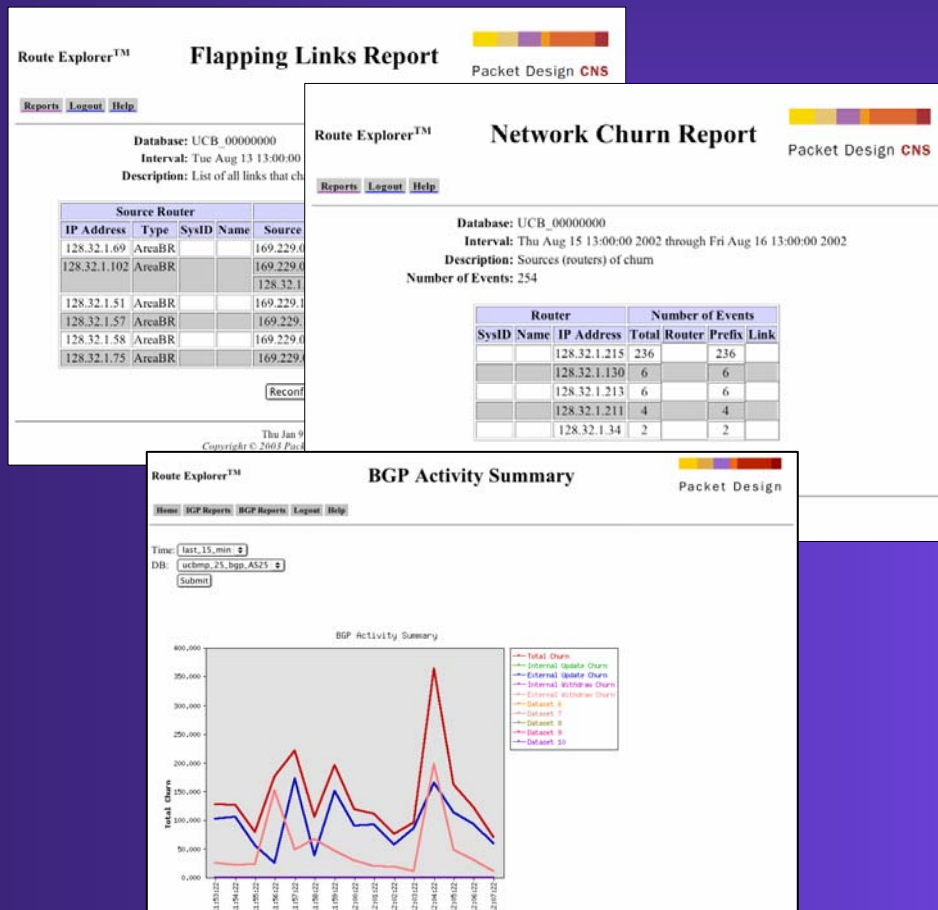
History Navigation



- Detects IP routing faults not found by SNMP-based management systems
- Historical routing events are easily replayed, at user-defined speed, to enable quick identification and diagnosis of problems
- 3<sup>rd</sup> party time-series data can be imported and correlated with routing events to enhance problem analysis

# Comprehensive Reports

*Proactive Analytics*



**Route Explorer™ Flapping Links Report**

Database: UCB\_00000000  
Interval: Tue Aug 13 13:00:00  
Description: List of all links that ch

IP Address	Type	SysID	Name	Source
128.32.1.69	AreaBR			169.229.0
128.32.1.102	AreaBR			169.229.0
128.32.1.51	AreaBR			169.229.1
128.32.1.57	AreaBR			169.229.0
128.32.1.58	AreaBR			169.229.0
128.32.1.75	AreaBR			169.229.0

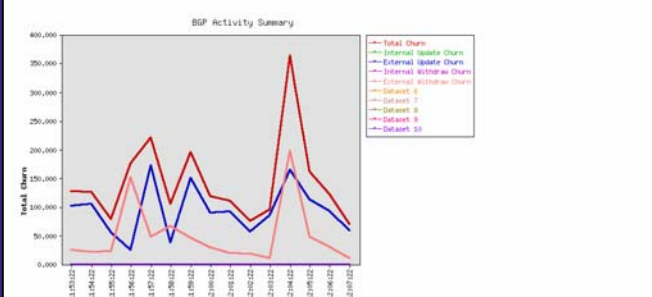
**Route Explorer™ Network Churn Report**

Database: UCB\_00000000  
Interval: Thu Aug 15 13:00:00 2002 through Fri Aug 16 13:00:00 2002  
Description: Sources (routers) of churn  
Number of Events: 254

Router		Number of Events	
SysID	IP Address	Total	Router Prefix Link
	128.32.1.215	236	236
	128.32.1.130	6	6
	128.32.1.213	6	6
	128.32.1.211	4	4
	128.32.1.34	2	2

**Route Explorer™ BGP Activity Summary**

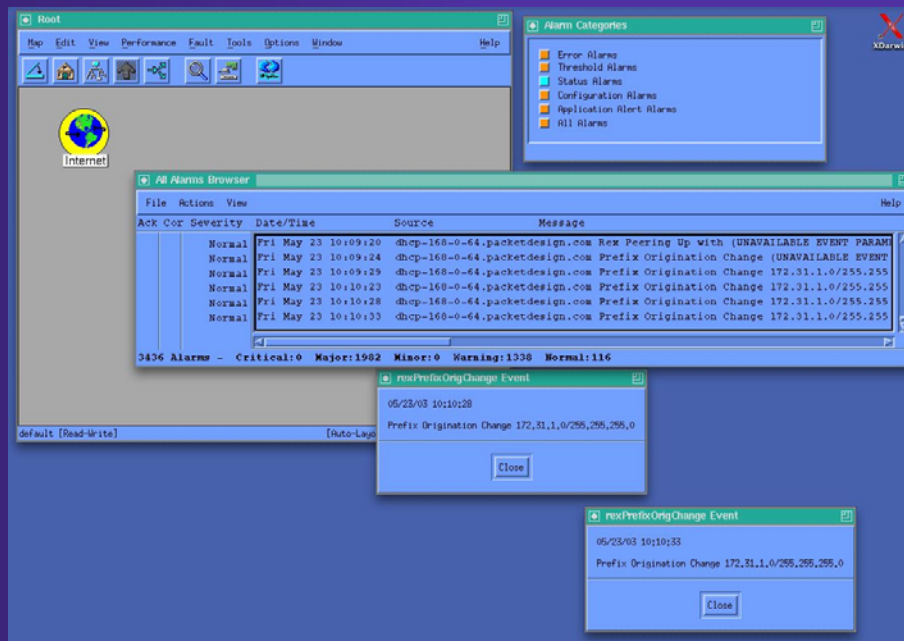
Time: last 15 min  
DB: ucdbmp\_25\_bgp\_ASYT



The BGP Activity Summary graph shows Total Churn on the y-axis (0 to 400,000) and time on the x-axis (from 11:00:00 to 12:00:00). The legend includes: Total Churn, Internal Update Churn, External Update Churn, Internal Withdrawal Churn, External Withdrawal Churn, and Datacenter 6 through 10. A significant peak in Total Churn is visible around 11:45:00.

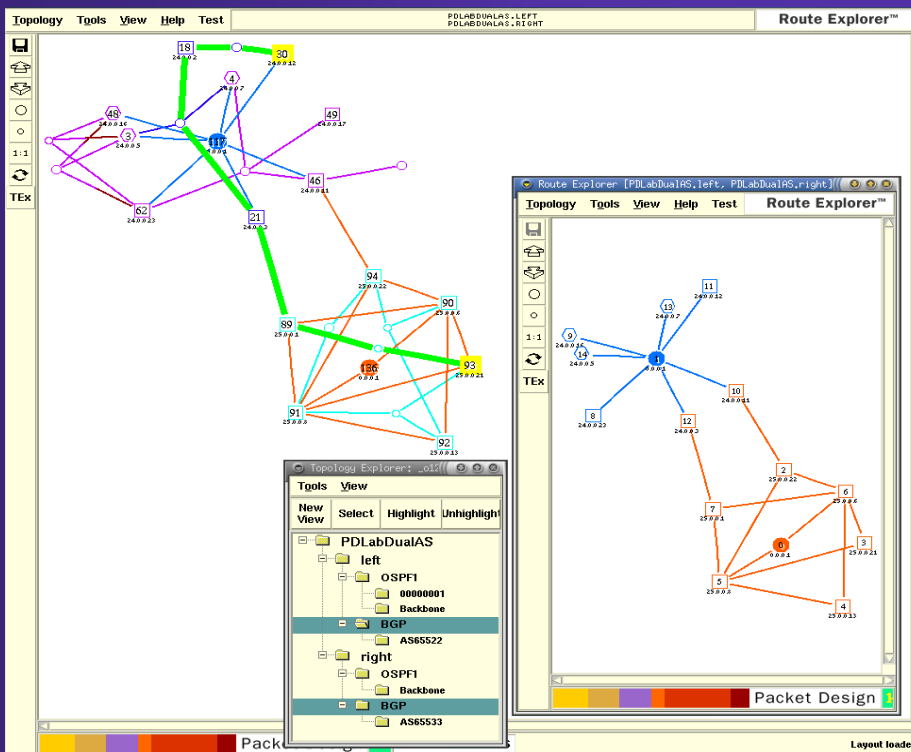
- Predefined reports provide detailed routing activity data and higher-level trend information; examples
  - Flapping links
  - Link metric changes
  - New prefixes and routers
- Web-based reports can be generated for any time period recorded in the database

# Real-Time Alerts Signal Potentially Critical Problems



- Alerts are sent based upon user-defined thresholds allowing monitoring of select routing attributes
- Types of alerts generated include
  - Route flaps
  - Excessive routing events
  - Router adjacency loss
- Alerts can be sent as
  - SNMP traps to user's NMS
  - Recorded in Syslog for later analysis

# Know How Changes Will Affect Your Network



- Route planning and analysis based upon actual network
- Perform impact analysis of possible network failures and configuration changes
- Evaluate the effects of maintenance activities on the network – before, during and after, avoiding potential problems

# Scenario 1: See which exit routers your traffic is taking

---

- In large multi-AS enterprises, connectivity from business units to business partners and customers is crucial. For example:
  - Vendor support web sites: Cisco TAC
  - Web services: Salesforce.com
  - General business productivity: Google.com
- In ISP networks, traffic imbalance at the peering or transit points can mean loss of earnings, or an unreachable Internet prefix can mean lost customers
- Route Explorer can show you exactly which BGP border routers your traffic is taking to get to their destination

# Verifying Exit Routers

Topology Tools View Help Route Explorer™

HighlightByExitRtrDialog

Dst (Prefix / DNS Name): www.cisco.com

DNS Name "www.cisco.com" is resolved to IPAddr 198.133.219.2

OK Unhighlight Cancel

**Show exits from all routers to a prefix**

Packet Design 2003-10-19 11:37:22

Topology Tools View Help UCE / AUG / 08 Route Explorer™

**Show exit and path from a particular router to a prefix**

Find Path To Prefix

Src Rtr (ex: 18.1.0.0 ): 169.229.128.170

Dst (Prefix / DNS Name): www.packetdesign.com

DNS Name is resolved to IPAddr 66.39.57.162

OK Unhighlight Cancel

Packet Design 2003-10-05 15:37:02

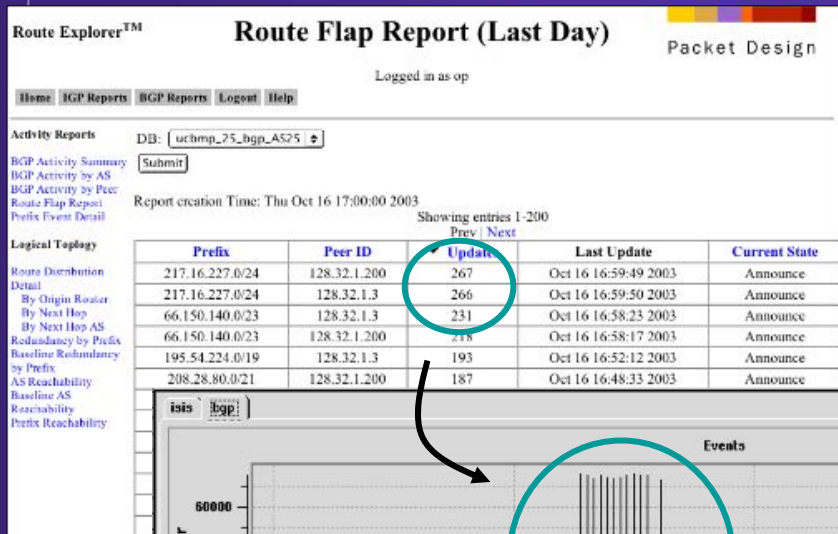
# Scenario 2: Detecting a BGP MED attribute “oscillation”

---

- In large BGP networks with route reflection, MED oscillations are potentially service disrupting:
  - Can lead to dropped packets
  - Undesirable router CPU load
- Oscillation causes a severe route flap
  - arises when different MEDs advertised by EBGP peers for the same prefix cause multiple route reflectors to repeatedly advertise them into the IBGP mesh
- Nearly impossible to isolate without Route Explorer
  - BGP analysis features lead to identification within minutes

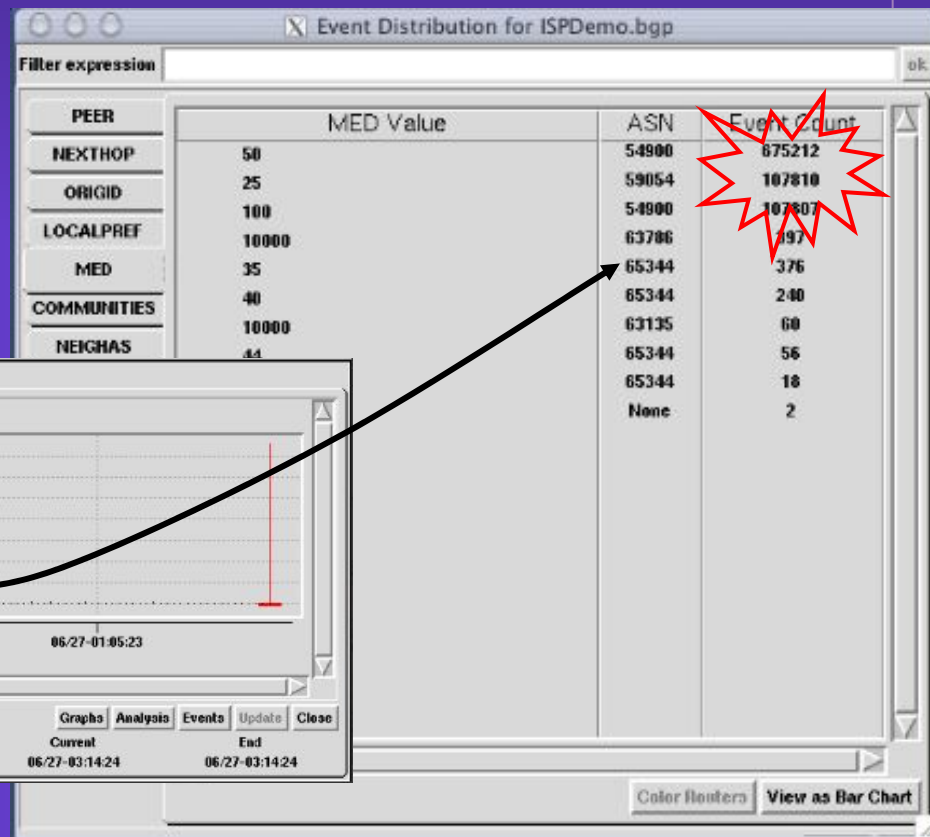
# Diagnosing a MED Oscillation

## 1. Daily report shows a route flap



## 2. History Monitor confirms it

## 3. Event Analysis identifies MED oscillation



# Scenario 3: Monitoring Redundant Routes


---

- How do you know when route redundancy has been reduced to only one link?
  - Monitor SNMP traps – time consuming; down link may be missed if its flapping
  - Monitor application performance – if performance is being affected, at least one redundant link has already failed
- The most efficient way to identify a downed redundant link is by monitoring at layer 3; what you need is
  - An alert to let you know a link has failed
  - Reports that help you quickly pinpoint where a failure has occurred

# Route Explorer – Helping you keep redundant routes up

**Route Explorer™**

## Redundancy by Prefix



Packet Design

Home IGP Reports BGP Reports Logout Help

---

DB:

Showing entries 25901-25950

[Prev](#) | [Next](#)

Prefix	# Next Hops	Next Hop(s)	Next Hop AS(es)
6.10.0.0/15	1	128.32.0.90	11423
6.9.0.0/20	1	128.32.0.90	11423
6.8.0.0/20	1	128.32.0.90	11423
6.5.0.0/19	1	128.32.0.90	11423
6.4.0.0/16	1	128.32.0.90	11423
6.3.0.0/18	1	128.32.0.90	11423
6.2.0.0/22	1	128.32.0.90	11423
6.1.0.0/16	1	128.32.0.90	11423
221.240.0.0/12	2	128.32.0.70 128.32.0.90	11423 11423
221.232.0.0/14	2	128.32.0.70 128.32.0.90	11423 11423
		128.32.0.70	11423

- Monitoring
  - Recording of all events, including prefix changes and withdrawn routes
  
- Alerts
  - User-defined “watch list” can alert when redundancy is lost
  - Alerts can be sent to SNMP NMS or Syslog
  
- Reports
  - Summary in easy to read format to expedite analysis

# A True Customer Case Study: 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, ...
  - Route explorer provides the tools for this

# Challenges

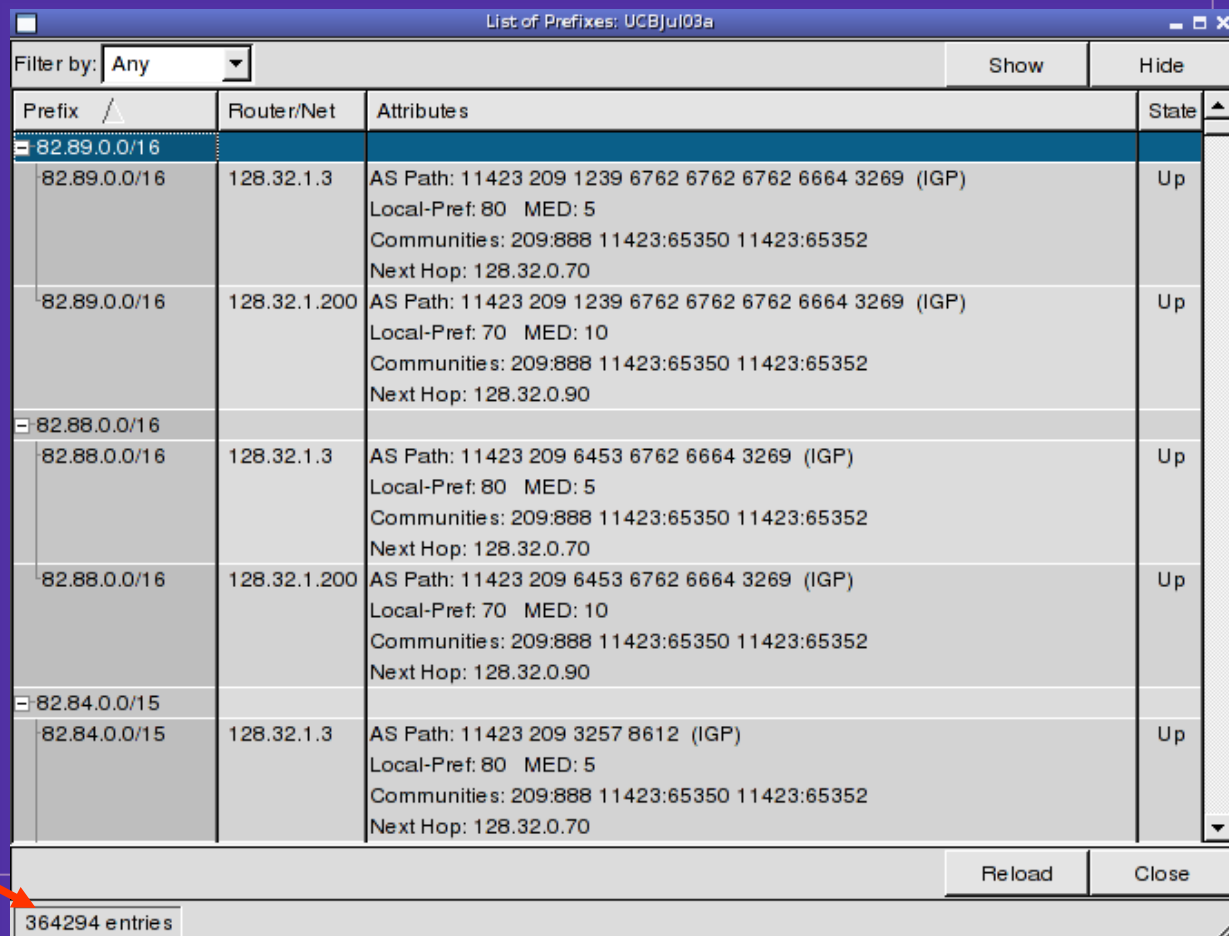
---

- Who can look at “show ip bgp ...” and tell what are the leaked routes? At which router do you issue this command?
- BGP maintains a very large state
  - 150K prefixes
  - 300K routes at a dual-homed AS
  - 1500K routes in a tier-1 ISP
- BGP sends very large number of messages
  - Simple session resets can send 100,000s of messages
- How can one tell which routes are the newly leaked routes?

# Central View of BGP Routes

- Like a route server, a central prefix list can show all the routes, but a nice GUI doesn't solve the job

Too many



Prefix	Router/Net	Attributes	State
82.89.0.0/16	128.32.1.3	AS Path: 11423 209 1239 6762 6762 6762 6664 3269 (IGP) Local-Pref: 80 MED: 5 Communities: 209:888 11423:65350 11423:65352 Next Hop: 128.32.0.70	Up
82.89.0.0/16	128.32.1.200	AS Path: 11423 209 1239 6762 6762 6762 6664 3269 (IGP) Local-Pref: 70 MED: 10 Communities: 209:888 11423:65350 11423:65352 Next Hop: 128.32.0.90	Up
82.88.0.0/16	128.32.1.3	AS Path: 11423 209 6453 6762 6664 3269 (IGP) Local-Pref: 80 MED: 5 Communities: 209:888 11423:65350 11423:65352 Next Hop: 128.32.0.70	Up
82.88.0.0/16	128.32.1.200	AS Path: 11423 209 6453 6762 6664 3269 (IGP) Local-Pref: 70 MED: 10 Communities: 209:888 11423:65350 11423:65352 Next Hop: 128.32.0.90	Up
82.84.0.0/15	128.32.1.3	AS Path: 11423 209 3257 8612 (IGP) Local-Pref: 80 MED: 5 Communities: 209:888 11423:65350 11423:65352 Next Hop: 128.32.0.70	Up

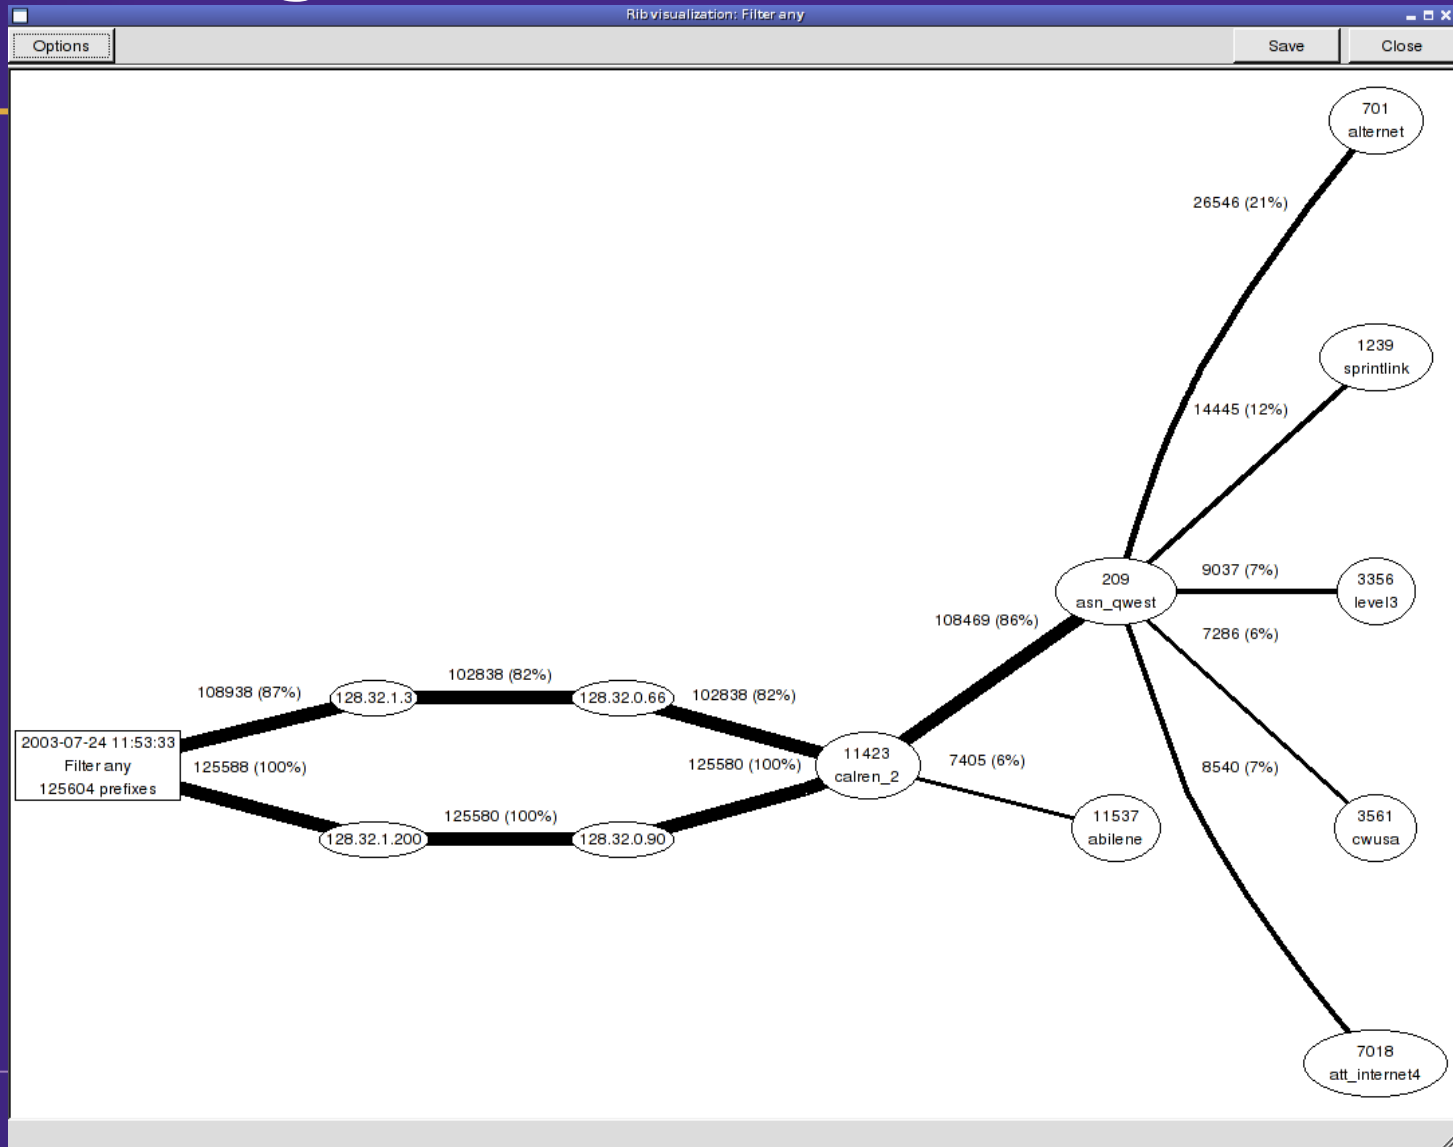
364294 entries

# Advanced BGP Analytics

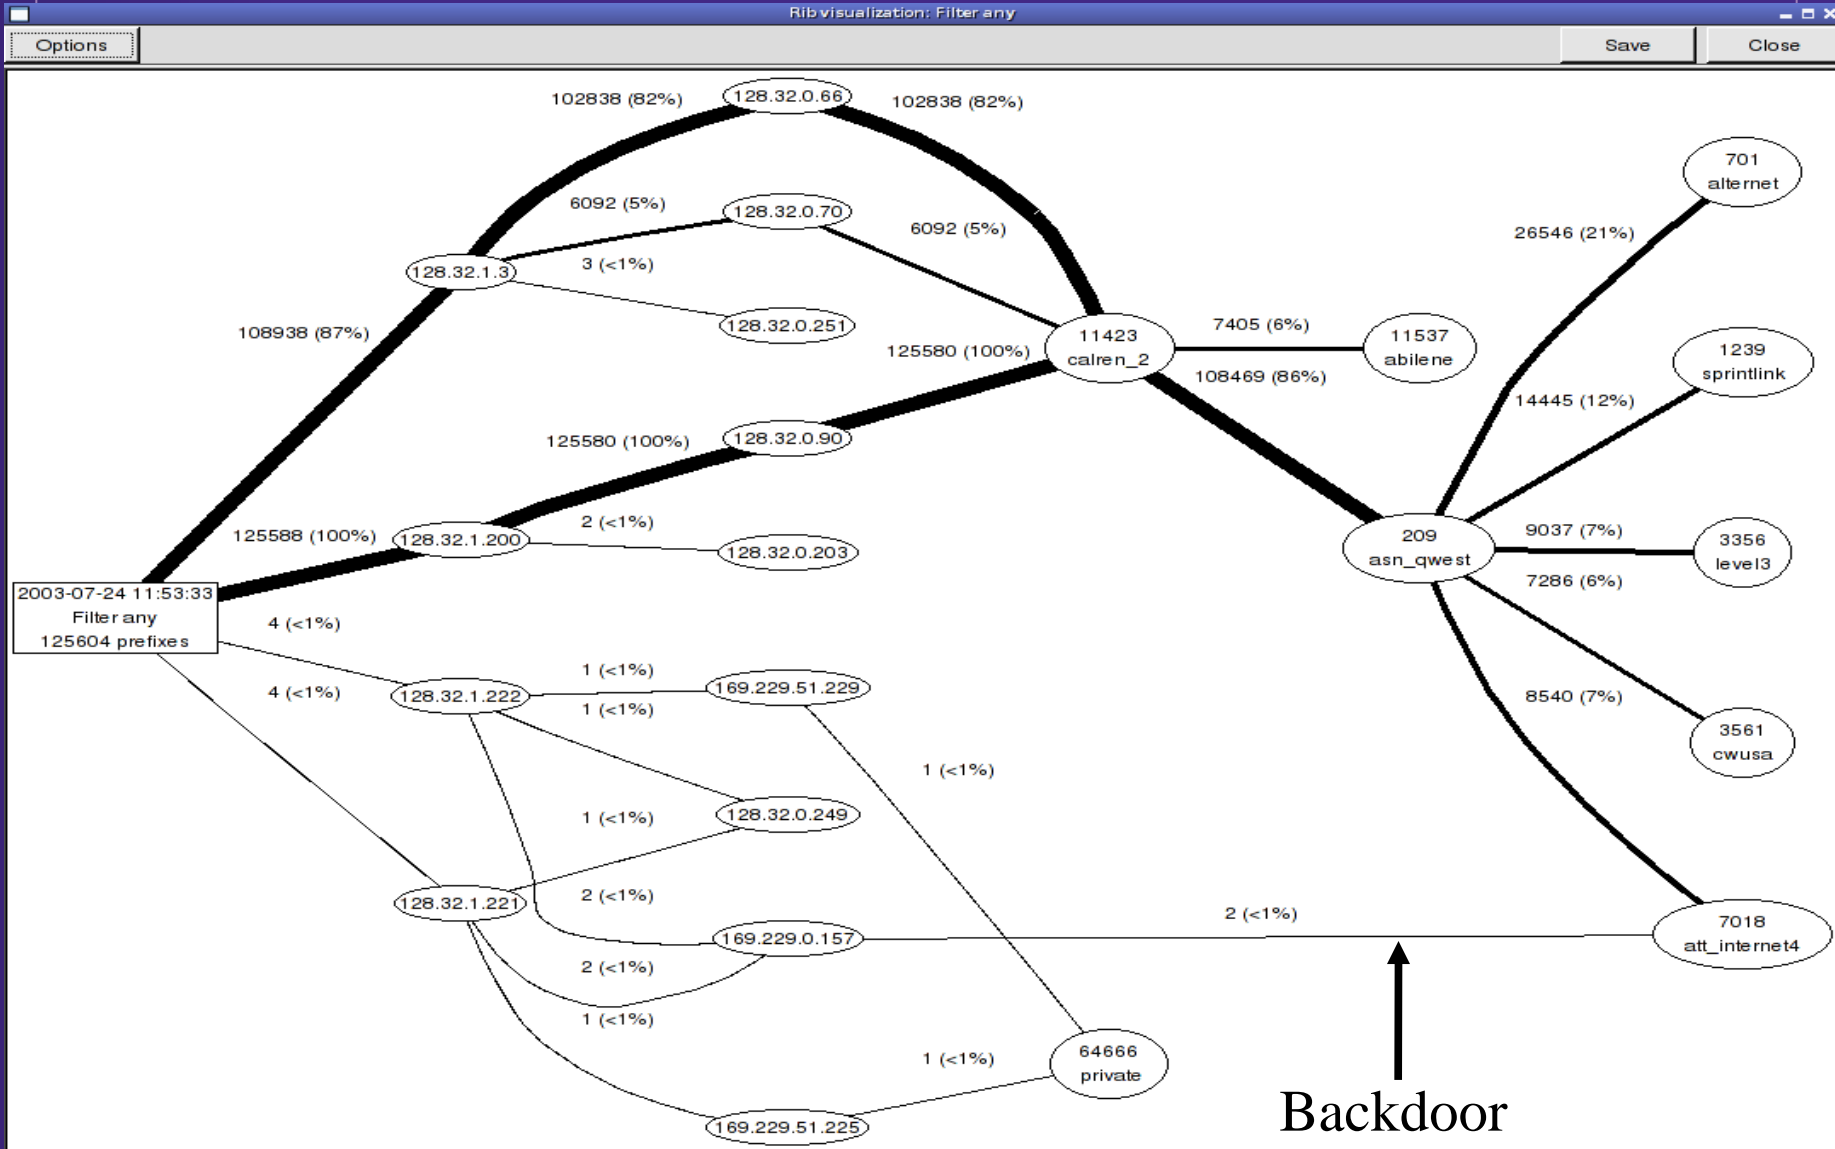
---

- Routing table *visualization*
- Per BGP attribute routing table breakdown and visualization
- *Before and After analysis* can highlight what is different between a known state of the network and the state in question
  - This was the feature used in detecting the customer that injected the extra routes
- Route table baselining can separate the typical route announcements from the new ones
- Root cause analysis can process 100,000s of BGP message to determine the cause of these events

# Routing Table Visualization



# Detailed Visualization Exposes Backdoor Routes



# Two Backdoor Routes

---

- Hard to notice in a “show ip bgp” output...
- In UCB, this was due to some department having a special arrangement to reach certain AT&T customer
  - Often backdoor routes are not allowed by administrators
- Backdoor peerings can have severe impact
  - What if AT&T started sending full routes
  - That router does not have enough memory for full routing tables
  - Better have prefix access lists on this backdoor peering for additional safety

# Attribute Break Down of BGP Routes

- Attribute breakdown in conjunction with visualization is very effective for understanding policies
  - Also for verifying policies

Rib Browser for: BGP/AS25 [ any ]

Filter by: Any Analyze Matching Analyze Excluding

	MED	AS	Route Count
Peer			
Nexthop	10	11423	125580
Originator	5	11423	108920
Local Pref	0	Internal	11
<b>MED</b>	30	11423	10
Communities	0	7018	
Neighbor AS	20	Internal	
2nd Hop AS	101	Internal	
Origin AS	1	64666	
Any AS	30	Internal	
AS Peers	120	Internal	1

10 entries


2003-07-24 11:53:33 Reload Close

Show Routes  
Visualize  
Filter Analysis

Color Routers View as Bar Chart

# Visualization of Community 2152:65297

- 2152:65297 is CENIC tag for Los Nettos Routes

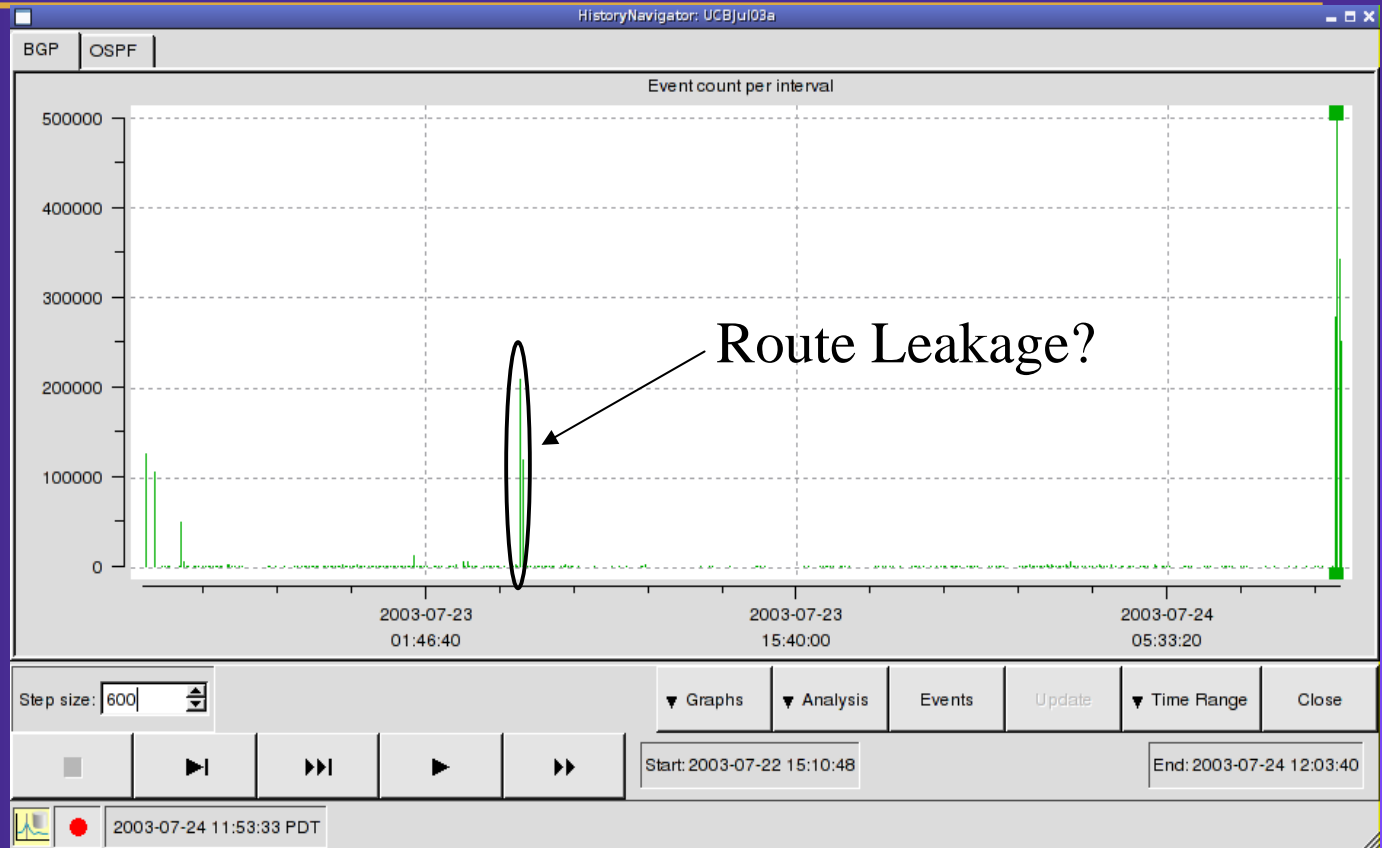


?

- Why is KDDI tagged? Is KDDI using precious resources that were meant for Los Nettos?

# Leaking Routes

- When an AS leaks routes, BGP sends many messages
- 1000 times more chatty than OSPF in Berkeley



# BGP Messages

- This is because BGP talks at the prefix-level
  - A million messages does not mean a million different things happened. BGP cannot say the actual incident, e.g. peering down or flaky router or peer leaking routes.

All events: UCBJul03a/BGP					
Time	Router	Operation	Neighbor/ Prefix	Attributes	Area or AS
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.93.192.0/19	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.100.192.0/19	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.100.203.0/24	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.100.205.0/24	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.116.192.0/19	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25
2003-07-23 05:05:18.305069	128.32.1.3	Withdraw	212.116.205.0/24	AS Path: 11423 3356 701 8895 25233 (IGP) Local-Pref: 80 MED: 5 Communities: 3356:3 3356:86 3356:575 3356:666 3356:2010 11423:65350 11423:65351 40840:36 Next Hop: 128.32.0.66	UCBJul03a.BGP/AS25

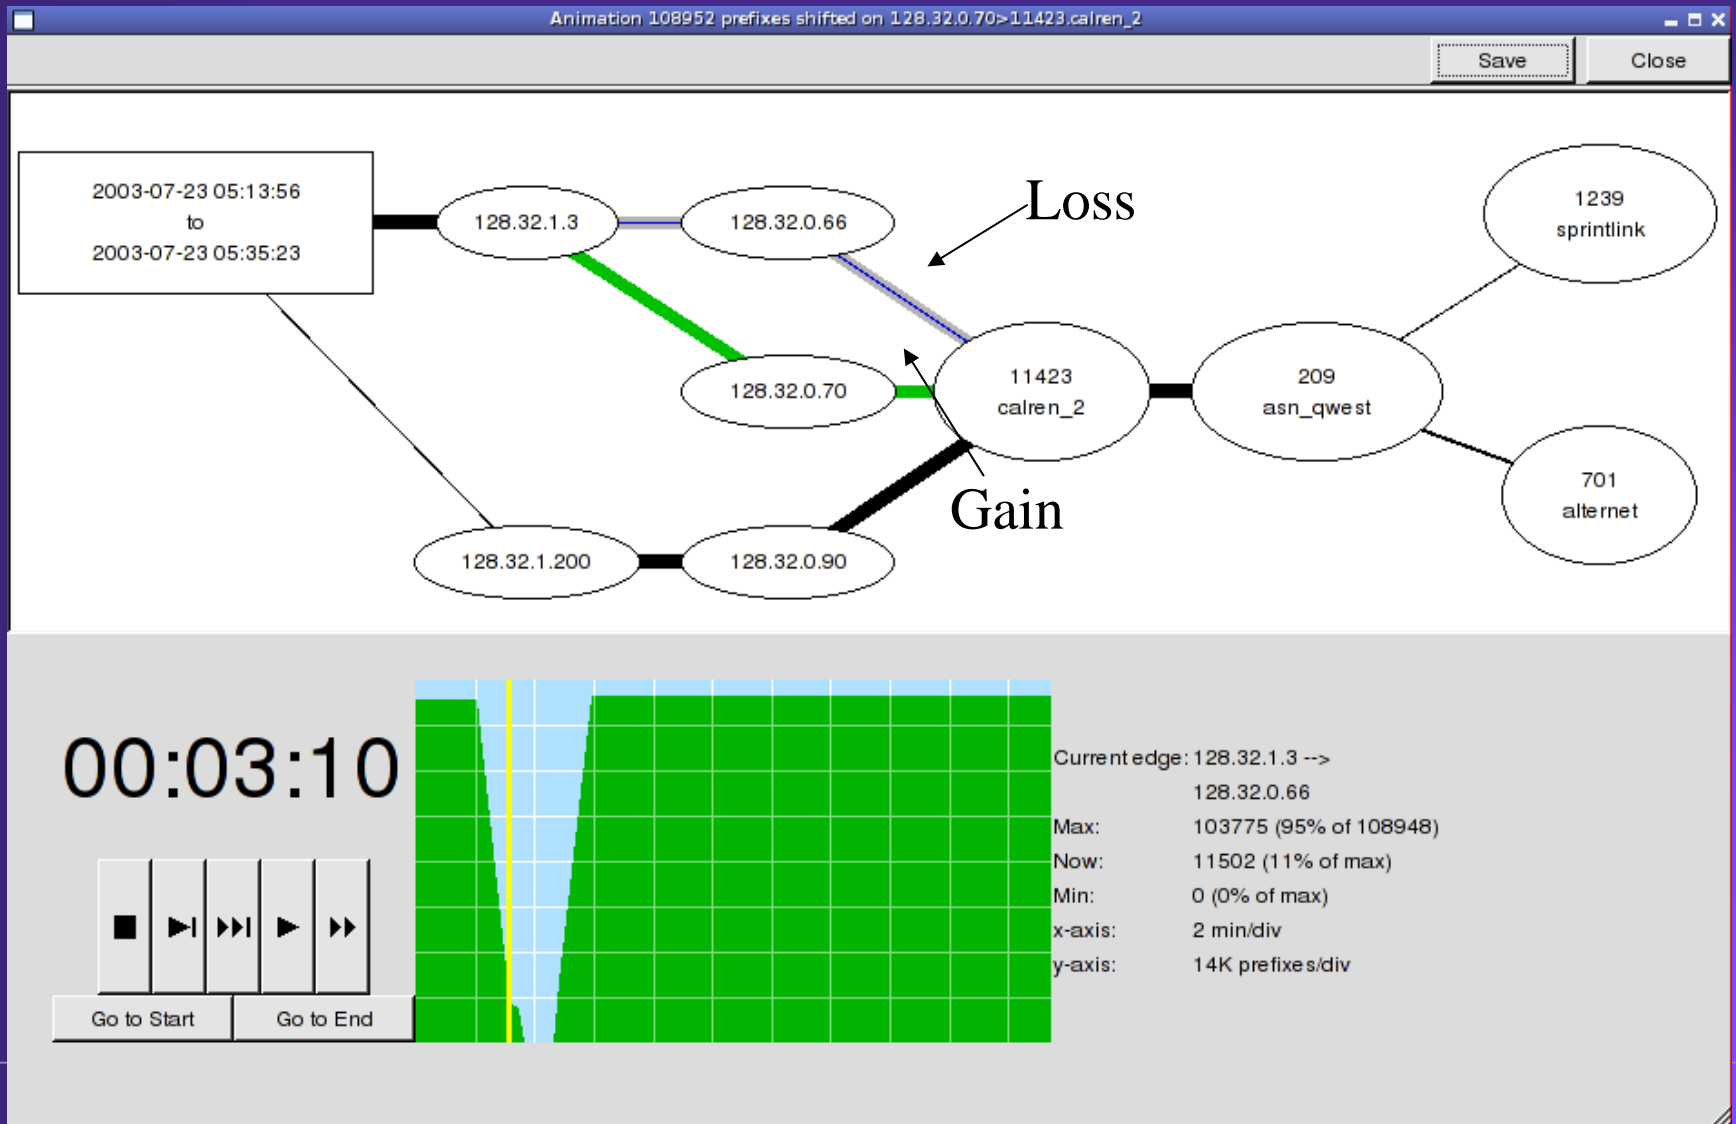
334241 entries

# BGP Root Cause Analysis

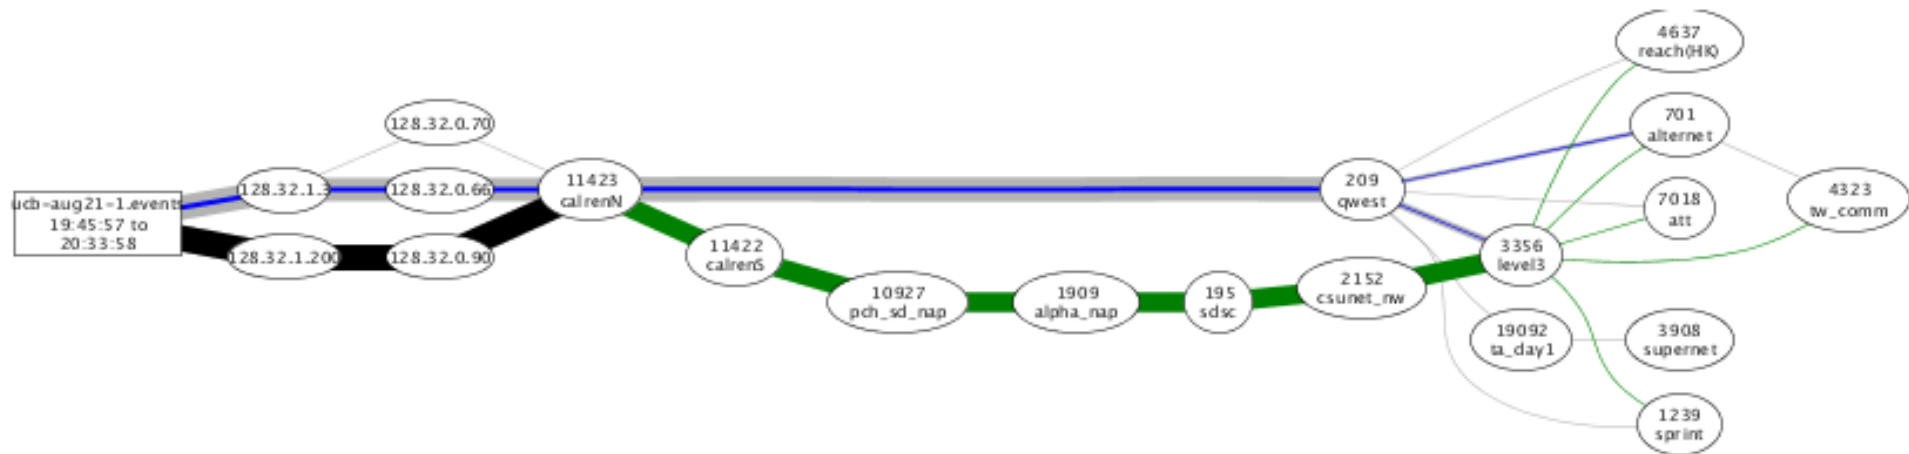
Root Cause Analysis Results: 20030723 05:05:13 to 20030723 05:37:49	
Description	View Details
108952 prefixes shifted on 128.32.0.70>11423.calren_2	Animation 330709 Events 108952 Prefixes
Close	
1 entry	

- 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

# 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

# Meta Coverage of Route Analytics

---

"... new route analytics products present better insight ... and enable enhanced relationship identification for mapping of services to infrastructure components."

"Superior business service management and optimized configurations are possible with proper inclusion of these products in management tool portfolios."

*Glen O'Donnell*

*"Route Analytics Enrich Technology Relationships"*

*Meta Delta, 4 February 2004*

# NW Fusion RA Coverage

---

“Route analytics will become a must-have for enterprise network managers...”

Denise Dubie

“Route Analytics to evolve beyond niche,”  
NW Fusion, 20 April 2004.

# One Big Award... So Far

---



Network Management  
Category

“Route Explorer stands out in the network management category because it gives enterprise network managers *visibility into their routed networks that has heretofore been unavailable*,” said Steve Steinke, Network Magazine editor-in-chief. “Packet Design has *proved that innovation in the network IT space is alive and well*.”



Networking &  
Management Tools

Finalist, eWeek Excellence award,  
Networking & Management Tools  
Category