

Meeting the BGP Management Challenge with Route Analytics



Packet Design

Table of Contents

Introduction.....	3
The BGP Management Challenge	3
Common BGP Problems Affecting Internet Service	4
Customer-Affecting Disruptions	4
Performance Degradation.....	4
Backdoor Routes and Other Security Exposures	4
Customers Impact Service Provider Networks	4
Major Internet outages	5
Dark and “Murky” address space.....	5
Unpredictable and Intermittent Errors.....	5
Limits of Existing BGP Management Techniques.....	5
Route Explorer—Next Generation BGP Management	6
Route Explorer for BGP Routing.....	6
BGP Routing Table Visualization.....	7
BGP Routing Historical Analysis	7
Detailed RIB Analysis.....	8
BGP Root Cause Analysis.....	9
BGP Route Analysis in Production BGP Networks.....	10
Case Study 1: Diagnosing a Peering Reset	10
Case Study 2: Identifying Persistent ISP Customer Route Flaps.....	12
Case Study 3: Identifying MED Oscillations	13
Conclusion	14



Introduction

The Border Gateway Protocol (BGP) is a critical piece of IP network infrastructure today—it is the de facto inter-domain routing protocol in the Internet, and serves as the routing basis for new IP services such as MPLS VPNs. Yet, while implementations of BGP have improved over the last decade to provide greater stability and scalability, BGP’s manageability has not progressed significantly during that same time. Determining the root cause of BGP problems—involving hundreds of thousands to millions of BGP messages, is beyond human capacity. The lack of automated BGP diagnostic capabilities leads to undetected and hard to solve network problems, resulting in disrupted or suboptimal service, long Mean Time To Repair (MTTR), and reduced network engineering responsiveness and productivity. The impact on network performance can be devastating as increasingly stringent Service Level Agreement (SLA) requirements are being placed on mission-critical Internet and IP WAN services. Network engineering organizations need a more effective way to manage and troubleshoot BGP networks.

The first part of this technical brief discusses some of the problems encountered in managing BGP today, and the service impacts felt both in service provider and enterprise organizations due to these problems. The second part of this white paper introduces BGP route analytics, and explains how Route Explorer automates BGP management, allowing network engineers to understand, resolve and prevent BGP problems, and achieve higher service assurance levels for critical, BGP-based IP services.

The BGP Management Challenge

The Internet utilizes a two-tiered routing approach where Interior Gateway Protocols (IGPs) manage intra-domain routing, and BGP manages inter-domain routing for the Internet and for very large, multi-Autonomous System (AS) organizations. BGP is the glue that binds all ISPs and their large customers together to form the global Internet. BGP has also been extended to support new IP services such as MPLS VPN routing. Unfortunately, while BGP has been effectively designed to meet the high scalability requirements of its Internet and inter-AS routing applications, it is a very complex protocol that is difficult to manage, analyze and troubleshoot, for three major reasons:

- BGP’s chattiness: The volume of messages induced by a single root event can be very high. It is not unusual for a minor connectivity change to generate hundreds of messages, while a major peering loss can spawn literally millions of BGP messages.
- Large number of routes, paths and route options in BGP routing table: Much of the data represented in the routing information base (RIB) is repetitive, yet subtle differences may be very significant, such as path length or Multi-Exit Discriminator (MED) differences. Finding and comparing these differences in a routing table that contains hundreds of thousands of routes is difficult at best
- Many intricate and error-prone configuration options: The potential for misconfiguring routing attributes and policies, such as MEDs, local prefs, redistribution, filtering, aggregation or communities is significant.

BGP Route Analytics Technical Brief

Given the critical scope of BGP's operation, problems resulting from misconfigurations or other causes can have a serious impact.

Common BGP Problems Affecting Internet Service

BGP is primarily used in conjunction with Internet routing, and this is where BGP problems that go unnoticed or that are hard to solve have their heaviest impact. Following are some examples of common BGP-related problems.

Customer-Affecting Disruptions

Issues that affect customers, such as peering flaps that occur on a relatively localized scale and that only affect one customer, can go unnoticed and undealt-with for long periods of time, due to the relatively high ongoing "noise" level of BGP messages that mask smaller scale, but highly disruptive problems. In one case, a Tier 1's customer experienced persistent peering flaps—once a minute, for one and a half months, leading to a very unhappy customer.

Performance Degradation

BGP misconfigurations can cause suboptimal routing that slows performance and increases latency. Flapping routes and MED oscillations can cause continuous reconvergence, slowing down router performance due to increased load and impacting traffic. Intermittent and unexplained performance problems are one of the key issues for enterprise and service provider network engineering groups today.

Backdoor Routes and Other Security Exposures

Backdoor routes without proper security mechanisms in place, such as prefix access lists, can expose networks to unforeseen security holes and router performance problems. In some cases, backdoor routes may be exploited maliciously, but in other cases they may incur severe problems due to more innocent reasons—such as the backdoor peer announcing full Internet routes that may exhaust the memory of a router that was not anticipated to hold the Internet routing table. In addition, dark and murky address space can be used for spam and other attacks.

Customers Impact Service Provider Networks

Internet access customers can dramatically impact service provider networks. For example, misconfigurations can cause route leakage—such as injecting the Internet route table back into the Provider Edge (PE) router. Since PE routers are often configured with less memory than core routers, a large route leakage can cause the PE to run out of memory, causing performance degradation. In some cases, route leakages can cause inter-domain disruptions. In one case, a tier-1 ISP's customer leaked thousands of routes that were then announced to the ISP's peers. One peer had a prefix-limit configured and reset the BGP session, severing Internet communication between the two tier-1 ISPs.

BGP Route Analytics Technical Brief

Major Internet outages

Since BGP provides the intelligence for Internet-wide routing, it can have a huge impact on the Internet itself. An example is BGP routing black holes—when a network that should be routable by a particular BGP router, becomes un-routable due to some other router falsely announcing the same prefix. A routing black hole caused by a simple BGP misconfiguration which advertised an artificially short AS-PATH to all Internet routes caused a wide-spread Internet black-out in the 1990s. While global Internet outages are not common, these sorts of incidents illustrate the scope of disruption that is possible with BGP problems.

Dark and “Murky” address space

It has been estimated that nearly five percent of the global Internet is reachable from some network providers, but unreachable from other network providers. While intentional policy decisions to implement aggressive route filtering are involved in the majority of dark address space, router and policy misconfigurations play a significant role. Misconfiguration examples include leaking default or private routes into the Internet routing table, and utilizing canned routing configurations instead of context-appropriate configurations. In both cases, it is clear that even relatively innocent routing policy decisions can have major impacts that are not easily analyzed.

Unpredictable and Intermittent Errors

BGP does not stand alone as a routing protocol, but interacts with the IGP's operating inside each AS. These and other network-layer interactions increase the overall complexity of IP networks. As a result, BGP's complexity contributes to difficult to diagnose, network-layer problems that cause unpredictable behavior in IP networks and disruptions to key IP services.

Limits of Existing BGP Management Techniques

Today, management of a BGP network is tedious and manual, requiring expert intervention using router CLI commands, spreadsheets, and rudimentary analysis scripts to try to cope with BGP routing issues.

As mentioned earlier, BGP can generate a high volume of hard to decipher messages; in some cases up to a million or more messages per root event. BGP itself does not provide any interpretive information for the volumes of messages it generates, so it is up to the operator to figure out what they mean. However, in most problem cases such as those previously described, it would be very difficult or even impossible to look at the output of a “show ip bgp” command and discern:

- The difference between leaked routes and legitimate routes
- The cause of a large volume of messages—such as a downed peering, flaky router, MED oscillation, a peer leaking routes, etc.
- A single customer's peering flaps from the noise of millions of unrelated BGP messages

BGP Route Analytics Technical Brief

SNMP-based network management systems are useful for monitoring the status of individual devices, but provide virtually no information about complex network-layer problems such as those encountered in BGP routing. For an introduction to the concept of network-layer management and how it differs from SNMP management, please see the Packet Design white paper entitled: Network-Layer Management, which can be found at <http://www.packetdesign.com/>.

While the service impact of today's limited BGP management tools and techniques are obvious, what is often taken for granted is the wasted engineering and operations resources expended to maintain a basic level of order over BGP routing in a large network. Clearly, a next generation BGP management capability is required to allow network engineering to improve service assurance, as well as internal responsiveness and productivity.

Route Explorer—Next Generation BGP Management

The Route Explorer family is the industry's first route analysis solution offering an accurate and complete, real-time and historical view of network-wide IP routing. A single Route Explorer appliance provides network-wide IP route (Layer 3) visualization, monitoring, analysis and diagnostic capabilities by leveraging the information in the network's routing protocols and applying patent-pending route processing algorithms. Route Explorer enables network engineers to easily verify, monitor and optimize network operations, as well as detect, diagnose and resolve IP network problems faster than has been possible till now. Route Explorer supports an integrated view of routing operations across multiple areas and Autonomous Systems with support for EIGRP, OSPF, IS-IS, BGP and RFC 2547bis MPLS VPNs.

For an introductory explanation of how Route Explorer and route analytics works, please read the white paper: Route Analytics—Foundation of Modern Network Operations, which can be found at Packet Design's website, <http://www.packetdesign.com/>.

Route Explorer for BGP Routing

Packet Design's Route Explorer delivers breakthrough BGP route analytics for the largest service provider and enterprise networks. Route Explorer includes a suite of capabilities that dramatically aid the engineer in analyzing, diagnosing and resolving complex routing problems, including a BGP root cause analysis capability that can isolate the key events within a flood of BGP messages. Network managers are presented with actionable information, including:

- An accurate view of intra-domain and inter-domain Layer 3 topology at any point in time
- Tools to effectively analyze a virtually unlimited number of BGP events over any selected period of time
- Critical event notification derived from detailed analysis of large volumes of BGP messages
- BGP policy verification through snapshot-in-time visualizations of inter-AS routing
- Easily understood forensic analysis, including before-and-after comparisons, and animations of topological changes in routing over time
- The ability to easily isolate the few root cause events within large or prolonged BGP message spikes over any selected time period

BGP Route Analytics Technical Brief

Route Explorer for BGP offers the following capabilities that automate the difficult job of BGP routing analysis and troubleshooting.

BGP Routing Table Visualization

Route Explorer provides a static RIB (Routing Information Base) visualization tool that presents a graphical snapshot of BGP topology, including numerical counts and visual weighting of how many prefixes are carried by each routing branch. This view allows network engineers to see at-a-glance the behavior of BGP routing in the network, assess routing policies, enable evaluation of existing peerings and plan future peering or routing policies. Users can select how much detailed information they want to view, or reduce visual clutter by filtering on various parameters. An example of UC Berkeley's inter-domain BGP routing is shown in Figure 1.

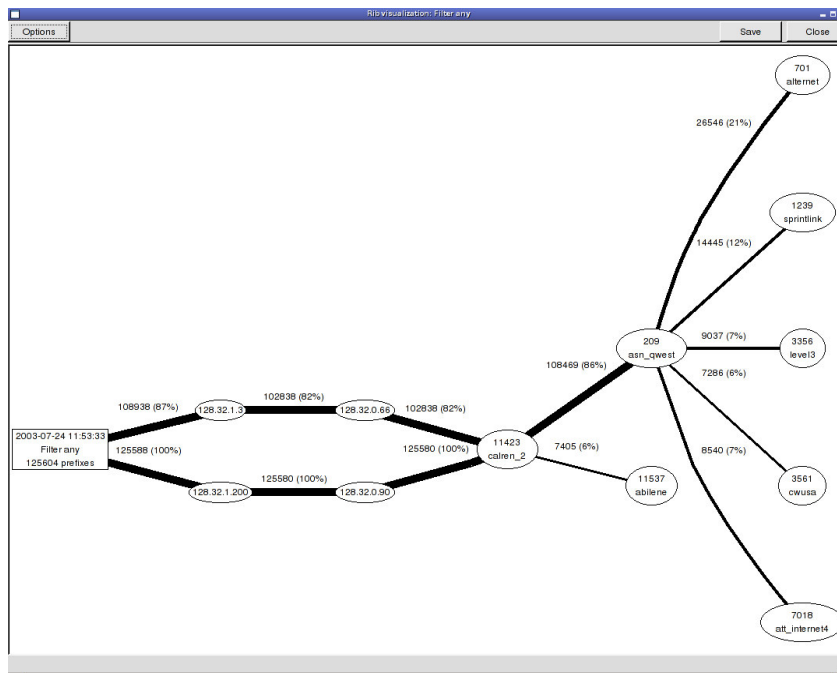


Figure 1: UC Berkeley's BGP RIB visualized by Route Explorer

BGP Routing Historical Analysis

One of the biggest challenges in BGP problem solving is assembling an accurate historical view of BGP routing activity. Route Explorer provides a history navigator that allows network administrators to view a histogram of BGP routing events. Since Route Explorer records all BGP routing updates from its peers, the history navigator provides an accurate picture of BGP activity in the network at any moment in time. The history navigator also serves as "home base" for selecting time periods for further detailed and automated analysis. Figure 2

BGP Route Analytics Technical Brief

shows a sample history navigator screen illustrating a suspicious spike of routing activity in UC Berkeley's BGP network that may indicate a potential problem.

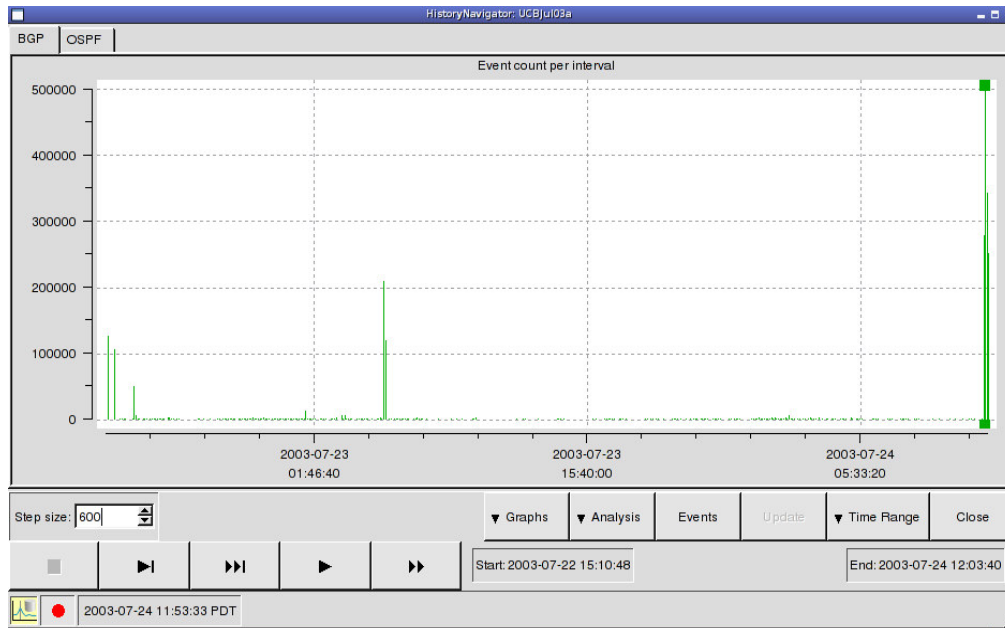


Figure 2: BGP History Navigator display for UC Berkeley

Detailed RIB Analysis

Detailed analysis of the state of the BGP RIB includes multi-level filtering on various BGP attributes for rapid problem isolation. In addition, a unique before-and-after analysis tool can easily highlight the difference between the state of the network at any two points in time, such as immediately prior to and following a spike in activity. Used in conjunction with the history navigator, these tools automate what has traditionally been a tedious process of gathering and parsing huge volumes of BGP events to understand the state of BGP routing at a given point in time. Figure 3 shows a RIB Analysis view of routes and their assigned MED values, which allows a network engineer to easily see if MED policies are correct for their entire network at any time in history. Right clicking on a value in the Route Count column brings up the option box shown, which allows the engineer to:

- View all the routes associated with the MED, along with its associated router, its up/down state, BGP attribute values, and AS/Area membership.
- Visualize the state of the topology with multi-AS route counts for the selected MED
- Further filter the routes associated with the selected MED, by other BGP attributes

BGP Route Analytics Technical Brief

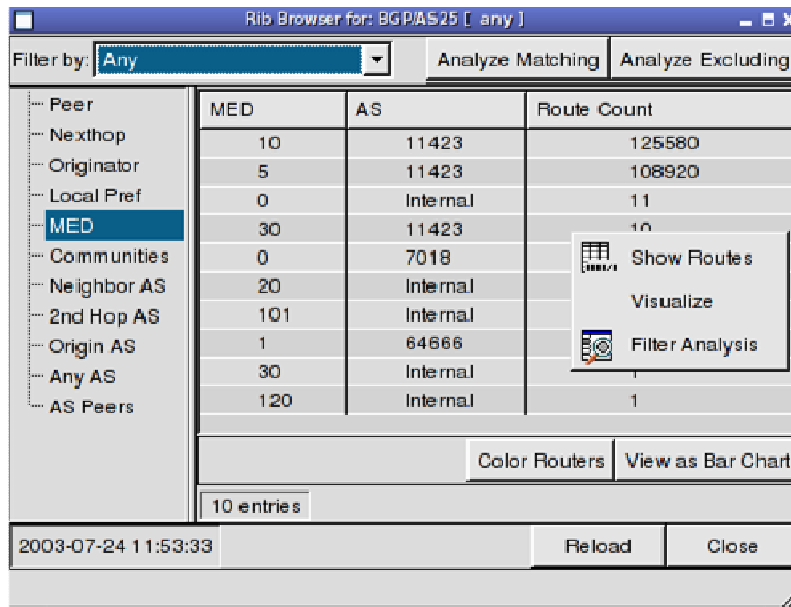


Figure 3: Detailed BGP RIB analysis with multi-level attribute filtering

BGP Root Cause Analysis

One of the most difficult tasks in managing a BGP network is identifying the critical, root cause event(s), or changes in the network, from the large volume of BGP messages that are a result of those event(s). Route Explorer provides a break-through root cause analysis capability that automates this task. By selecting any time period in Route Explorer's history navigator and launching the root cause analysis tool, the network administrator can see a summary of macro-level events and the number of BGP messages and prefixes associated with each event. This summary allows further exploration of the detailed messages and prefixes associated with each root cause event, as well as the option to view a dynamic animation, illustrating the macro BGP events on the affected network's topology in the selected timeframe.

Route Explorer's BGP root cause analysis capability provides network engineers with an intuitive, animated display of the changes to their BGP topology, helping them to visualize the "big picture" as to what happened, rather than spending hours manually analyzing the detailed router messages. The user can rewind and play the animation at various speeds to view how their multi-domain peering structures and routes changed over time. A multi-domain topology map, with graphic representation of route volume per peering provides insight into how external and next-hop peers have affected IBGP peers and overall routing behavior. Isolation of root-cause events such as peering flaps and MED (Multi Exit Discriminator) oscillations, or policy violations such as misconfigured community tags and unwanted back-door paths is performed in minutes rather than days. Topology animations highlighting critical events can be saved in SVG format and emailed between service providers, greatly helping to coordinate and speed inter-provider problem resolution.

BGP Route Analytics Technical Brief

BGP Route Analysis in Production BGP Networks

Route Explorer's BGP route analysis tools have been used to isolate and resolve difficult BGP problems in both enterprise and service provider environments. Following are three case studies from production BGP environments.

Case Study 1: Diagnosing a Peering Reset

UC Berkeley observed a spike of BGP messages as illustrated in Figure 4, and wanted to understand what had happened to their BGP network.

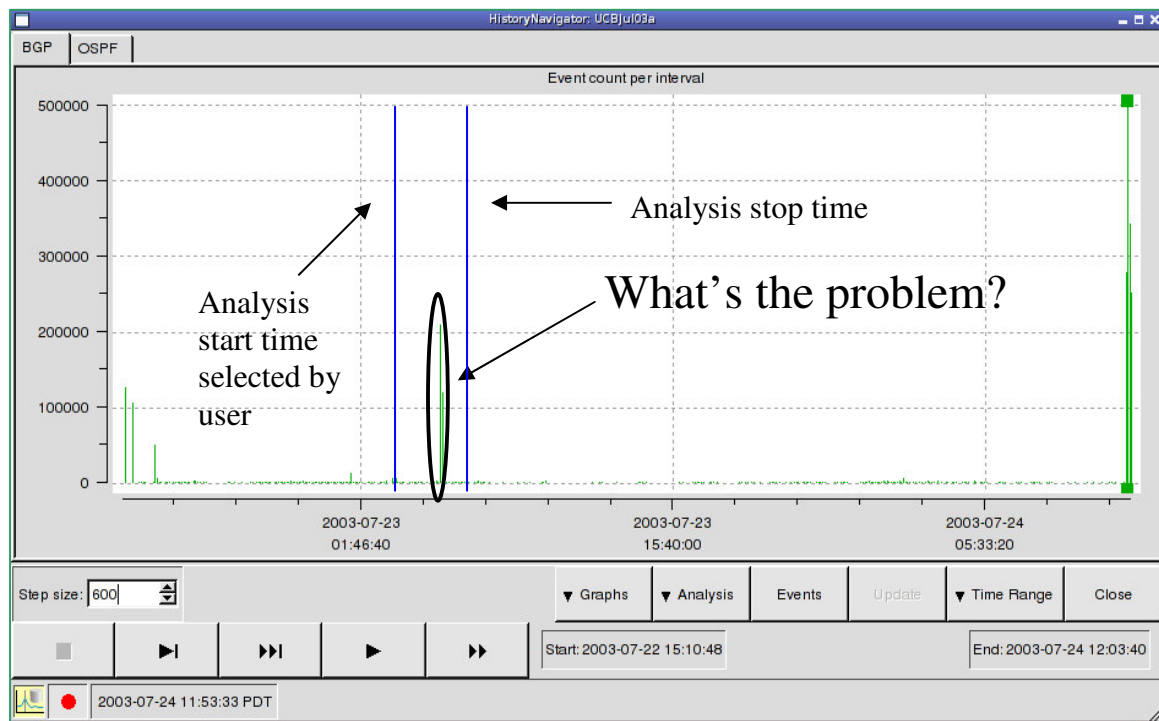


Figure 4: Suspicious spike of activity in UC Berkeley's BGP routing indicates a potential problem

Route Explorer's root cause analysis tool was able to determine that the 330K+ BGP messages occurring over the selected time period were the result of prefixes shifting between two routes. A peering had reset, shifting prefixes from one peer to another, followed by a shift of the prefixes back to the original peer. Figure 5 shows the root cause analysis of this spike of BGP messages.

BGP Route Analytics Technical Brief

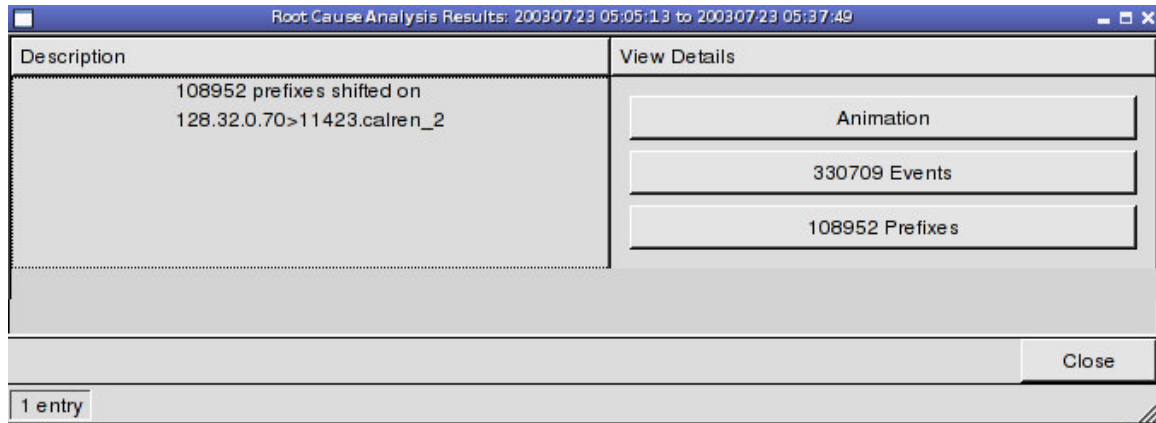


Figure 5: Root Cause Analysis of the UC Berkeley BGP event spike

A snapshot of the root cause event animation is shown in Figure 6, allowing network engineers to quickly and easily visualize what had happened. To see SVG animations of BGP root cause analyses, please visit: <http://www.packetdesign.com/products/applications.htm>

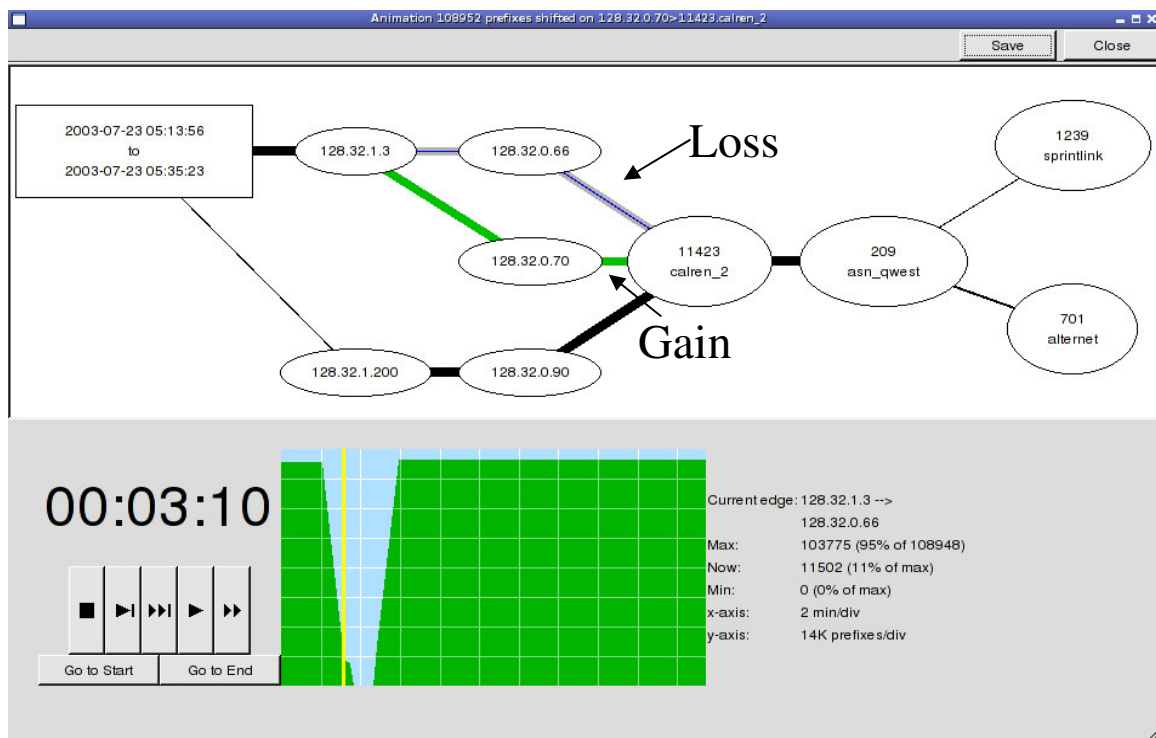


Figure 6: Root Cause Analysis animation of the reset peering at UC Berkeley

BGP Route Analytics Technical Brief

Case Study 2: Identifying Persistent ISP Customer Route Flaps

A tier 1 service provider's customer was complaining of poor performance and difficulty in accessing the Internet that had been going on for several weeks. No device issues were detected via SNMP management systems, so the service provider turned to Route Explorer to analyze the Layer 3 and routing conditions that might be affecting this particular customer's service. Initial analysis using Route Explorer's history navigator did not show any particular spike in routing activity of sufficient duration that explained the customer's persistent service difficulties. In fact, while the root cause of the problem was route flapping, the route flaps were hidden in the normal BGP message "noise" common on large, active networks, and they didn't show up as any noticeable spike in activity, as seen in Figure 7.

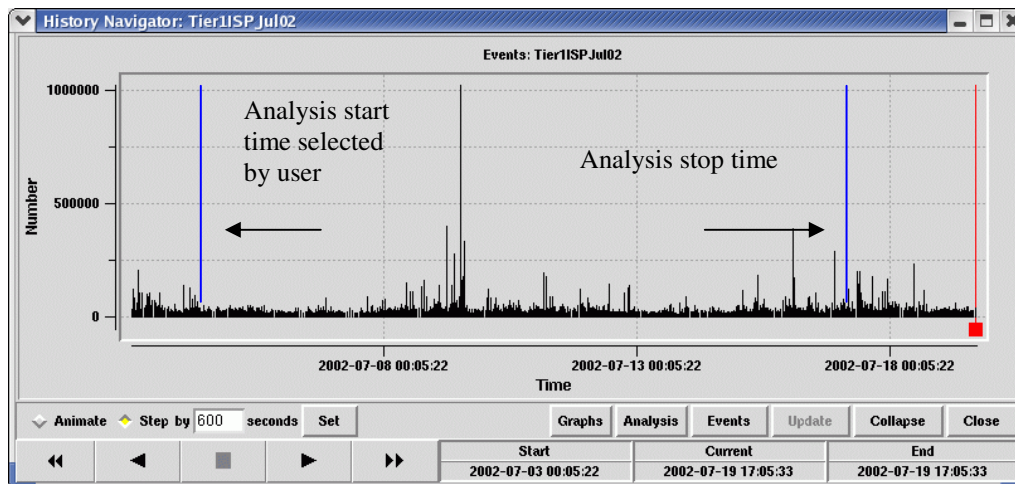


Figure 7: Tier 1 ISP BGP History Navigator display showing on-going BGP message activity level

Using Route Explorer's BGP root cause analysis tool and animation capabilities against the selected timeframe, the service provider was able to immediately identify that route flaps were the cause of the problem, as illustrated in Figure 8.

BGP Route Analytics Technical Brief

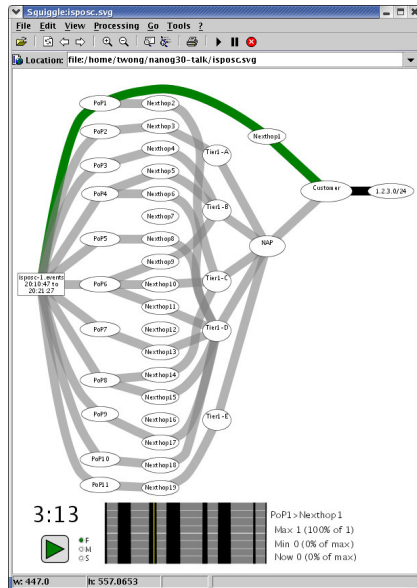


Figure 8: Route Explorer's BGP root cause analysis and animation tools reveal a customer's flapping routes

Case Study 3: Identifying MED Oscillations

A tier 1 service provider experienced router slowdowns in its network and observed very high router CPU utilization corresponding to the period of this slowdown. By selecting the relevant period of time in Route Explorer's history navigator (see Figure 9), they were able to analyze an unusual spike of BGP activity in the network as shown in Figure 10.

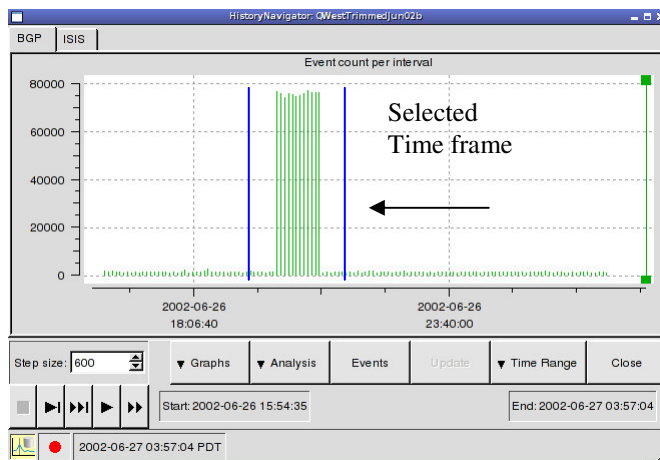


Figure 9: Unusual BGP activity spike observed by Route Explorer

BGP Route Analytics Technical Brief

By running Route Explorer's root cause analysis and animation tool, the service provider was able to discern that a MED oscillation was causing the problem, as illustrated in Figure 10:

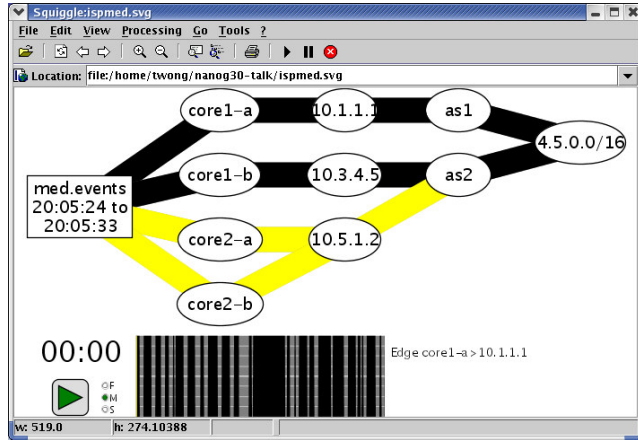


Figure 10: Root Cause Analysis animation shows the continuous MED oscillation that caused the router performance problem

Conclusion

BGP route analysis moves network engineering far beyond the tedious and manual troubleshooting techniques practiced today to a next-generation capability for preventing, understanding and rapidly solving BGP routing problems. Customer service level expectations from internal or outsourced IT departments and from service providers have dramatically increased, while lower prices are increasingly available in a fiercely competitive marketplace. Route Explorer provides the required next-generation management tools for IP networks, allowing network engineers and operators to maximize service assurance, responsiveness and productivity.

To learn more about Packet Design and Route Explorer, please:

- Email us at info@packetdesign.com
- Visit Packet Design's web site at <http://www.packetdesign.com>
- Call us at 650.739.1850



Packet Design

Corporate Headquarters

Packet Design Inc.
3400 Hillview Avenue, Building 3
Palo Alto, CA 94304
Phone: 650.739.1850
Fax: 650.739.0590
<http://www.packetdesign.com>

US Sales Offices

Western Region
Phone: 650.739.1886

Central Region
Phone: 469-737-5635

East/Federal Region
Phone: 978.779.8229

International Sales Offices

Europe, Middle East, Africa

Packet Design Ltd.
PO Box 3061
Wokingham, RG41 3GJ
United Kingdom
Phone: +44.118.977.5295
Fax: +44.118.977.5296