

Ensuring Network Integrity, Continuity and Process Enforcement with Route Analytics

*Layer 3 Visibility for Mission-Critical
Network Services*



Packet Design



Introduction

For government, public service, national security and military agencies, maintaining network continuity is no casual matter. The sensitivity of these organizations' missions means that there is no room for network outages, application degradation, or network integrity breaches. The good news is that these agencies have built highly redundant IP networks to ensure application traffic delivery. The bad news: these massive and complex networks lack layer 3 network management visibility. This makes it much harder to ensure network integrity and continuity in the face of failures, threats, hostile incidents, and even the less sinister but more common culprits of human error, unenforced operational procedures and time-consuming network troubleshooting procedures.

Route analytics technology, with its real-time, network-wide understanding of the operational routing topology and the traffic flowing across all network paths and links, provides the missing visibility. With route analytics, network engineers can easily monitor for network topology and traffic problems, strengthen change management processes, proactively uncover network vulnerabilities, and accurately simulate failure scenarios and resulting network behavior. Effective use of route analytics as part of the network management process can help ensure the network will always be ready to pass the test for mission-critical requirements.

The Inherent Unpredictability of IP Networks

A major reason that IP became the de facto worldwide standard for data communications networks is its automated resiliency based on intelligent IP routing protocols that control the traffic routing topology. But while IP's distributed routing intelligence makes it efficient and resilient, it also makes IP network behavior unpredictable and harder to manage. IP routing protocols automatically calculate traffic routes or paths from any point to any other point in the network based on the latest known state of network elements. Any change to those elements causes the routing topology to be recalculated dynamically. While this means highly reliable traffic delivery with low administrative overhead, it also creates endless variability in the active routing topology. Not only can a large network be in any one of millions of possible active routing topology states, but application traffic patterns are by nature unpredictable. Network problems – router software bugs, misconfigurations and human error – only add to that unpredictability. And that's just counting normal, day-to-day challenges. Add the potential for natural disasters and hostile incidents that can cause large-scale disruptions too, and it is a formidable technical challenge to ensure network integrity and continuity.

These same issues show up in the time-consuming process of linking network causes to application problems. For example, when a user reports an application performance problem that doesn't stem from an obvious hardware failure, pinpointing the root cause can be quite difficult because in a large, complex network, IT engineers have no way to know the route the traffic took through the network, the relevant links servicing the traffic, whether those links were congested at the time of the problem, or even which devices were servicing the traffic.



Today's Network Management—Many Points of View, No Big Picture

Network management's purpose is to overcome the complexity inherent in a large network and provide better visibility to network operations and engineering. The overarching architectural principle of today's network management is to gather information on a vast number of different "points" in the network (e.g., routers, switches, security devices, servers), then correlate various point data to infer service conditions. The key mechanism for doing this is the Simple Network Management Protocol (SNMP). The main data gathered is:

- Device health: uptime, current status, CPU and memory utilization
- Fault indicators: up/down status, uptime, dropped packets, errors
- Traffic information: interface utilization: bytes in/out, packets in/out, configuration
- Service utilization information: utilization per class of service, threshold violations

While having this point data is critical – for example, an interface or device that fails, runs out of memory, or is congested with traffic can have a direct impact on application delivery – the sum of all this point data is much less than the whole picture. Just knowing that an interface is full of traffic doesn't tell you *why* it is full. Where is the traffic coming from and going to? Is the traffic usually on this interface, or was there a change in the network or elsewhere that caused it to shift to this interface? If so, from where, when and for how long? Without answers to these questions, there is no real understanding of the logical operation of the network as a whole.

While there are correlation algorithms for deducing certain types of network conditions, SNMP was never built to understand a network's logical behavior. SNMP's key limitation is that it is too periodic – polling cycles from 30 seconds to several minutes simply cannot produce an accurate portrait of the network's routing state, with its rapid and high-volume state changes. Even speeding up the polling cycle – say, to every five seconds – would still miss many routing state changes, and anyway would generate so much management traffic overhead as to be impractical.

The Impact of Lack of Visibility into Logical Network Operations

The lack of understanding of IP networks' logical operations significantly impacts an IT organization's ability to maintain a high state of network readiness to support critical applications.

Network Integrity Risks: Without a clear understanding of how the network as a whole is operating from a logical and routing point of view, it is easy to miss important information when assessing security and integrity risks. Three major risk points emerge from lack of logical network visibility:

- *Lack of an Accurate Network Map for Security Engineering:* A clear understanding of the active routing topology is essential when designing and deploying security choke points to provide defense against external attacks. Without an accurate, always-updated routing map, security engineers must rely on outdated, manually assembled, static maps of the network, educated guesses and assumptions which can often prove to be disastrously wrong. For example, in one network, engineers discovered that a WAN link they thought had been decommissioned was still active months later.



- *Back-Door Internet Routes:* When securing a network, it is critical to know where the Internet peerings are located. More than once, due to the complexity associated with the sheer size of many government agencies, along with human error or miscommunication, Internet peerings have been created outside the purview of central IT department oversight. Without a way to reliably see where all BGP Internet peerings are, and what routes are being advertised across them, it is very hard to plan the overall defense of a network against Internet-based attacks.
- *Improperly Filtered Peerings with Other Organizations:* Internet peerings aren't the only backdoor into a secure agency's network. Peerings with other government agencies and government contractors create a potential integrity risk for the network because of the connections that reside elsewhere in those third-party networks. Without routing visibility through those peerings, it is possible for an organization with a lower security classification to leak routes into a more secure network and thus compromise its integrity.

Route and Flow Analysis—A New Frontier for Network Visibility

The answer to the need for better insight into network-wide service delivery lies in a new trend in network management: the combination of router-based traffic-flow measurement and analysis with route analytics technology. The value of this pairing comes from analyzing the network based not on point devices and interfaces, but on end-to-end traffic flows and the precise paths or routes that those flows travel as they move across the network.

Router-based traffic-flow monitoring has existed for years, but became practical only recently when deployed in ASIC-based routers. The most popular traffic-flow monitoring mechanism today is Cisco's Netflow, while a forthcoming IETF standard called IPFIX will create a vendor-neutral format for traffic-flow collection. In either case, a router samples packets on its interfaces, collects statistics on a per-flow basis, and forwards collected flow data to a server, where it is analyzed. Understanding end-to-end flows provides much greater insight into service delivery than simply collecting generic interface bandwidth statistics via SNMP, because a flow can be correlated to a service. For example, a packet stream from an application server to a user can be correlated to a particular traffic flow based on the source IP address of the server, the destination address of the user, the Class of Service marking, and possibly the TCP or UDP port numbers.

For all its benefits, router-based traffic-flow collection and analysis still generates management traffic overhead, and thus is impractical to turn on at every interface in the network. This leaves IT in the same old visibility bind with regard to network behavior. Network engineers can tell which end devices are exchanging traffic across the network, but still have no idea how specific flows, or all traffic in aggregate, is traversing the network. As a result, traffic-flow information on its own fails to provide the visibility needed to ensure predictably high performance.

However, when combined with real-time information about network-wide routing state, traffic-flow information becomes extremely powerful. The challenge, unmet until recently, is how to collect all routing-state changes.

Route analytics technology, which has been adopted by many large enterprises, service providers and government agencies in recent years, is the answer to understanding routing state in IP networks. A



route analytics device – typically a network appliance running specialized software – acts like a router: it listens to routing protocol updates sent by all routers in the network and computes the network-wide routing state in real-time, just as all the “real” routers do. While the route analytics device itself is passive, never advertising itself as a place to send traffic, it provides real-time visibility, always up-to-date routing-state knowledge, and a completely accurate historical record of all past routing changes. It knows every route or “path” that any traffic takes at any point in time – hence the name “path-based” network management.

The understanding of network-wide routing topology and the full detail of routing changes provide the basis for many useful analyses of the routing control plane. However, it is when combined with traffic-flow data that the full power of route analytics information emerges. Collecting all traffic flows and mapping them to the precise routes they traverse over time makes it possible to create an integrated, accurate map of all routing and traffic for the entire network core. Based on its always-current map and its continuously recorded history of routing and traffic changes, route analytics provides a number of unique network management capabilities:

- **Live network topology monitoring:** Route analytics can monitor a network for real-time changes in its routing state and thus detect critical problems in network behavior.
- **Network modeling:** Since it has a completely accurate picture of routing and traffic behavior on all links and paths in the network, route analytics can be used to simulate routing and traffic changes and show exactly how the entire network would behave as a result.
- **Proactive routing topology audits:** Route analytics can be used to perform thorough audits of a network’s routing topology to find weak points, vulnerabilities, unintended configurations and potential failure points.
- **Traffic trending and capacity planning:** Engineers can easily pinpoint trends on network-wide traffic or a subset of traffic by link, CoS, or application group.

Applying Route Analytics to Ensure Network Service Delivery

Integrated path and flow-based network management provides a new and far more useful picture of network and application behavior. While it helps IT ensure day-to-day application traffic delivery and speeds troubleshooting, it is also extremely useful for ensuring network integrity and continuity and for strengthening process controls.

Ensuring Network Integrity and Continuity: Route analytics can be used in a variety of ways to ensure that networks are designed for and maintain ongoing compliance with integrity and continuity requirements. For example:

- **An always-accurate map:** Route analytics provides network and security engineers an accurate, up-to-the-moment map of the entire routed network topology that can be used to ensure that all appropriate traffic routes are covered by defense security measures.
- **Eliminating routing vulnerabilities:** Using route analytics’ insights, engineers can ensure that backdoor routes and improperly filtered peerings with other organizations are detected and shut down.



- Monitoring overall reachability: Engineers can set thresholds for overall routed prefixes and be alerted if the volume of prefixes crosses a threshold, indicating either a pervasive loss of communication or the injection of potentially damaging new routes into the network.
- Monitoring loss of BGP peering, redundancy and reachability: Often major inter-organizational connections are routed using the BGP protocol. Route analytics can monitor BGP peerings to detect if they go down, and can also monitor BGP Autonomous Systems for loss of redundancy or reachability.
- Alerting on changes to critical traffic routes: In any network, certain traffic routes/paths are particularly sensitive. Engineers can monitor and be alerted to changes to these routes in real-time.
- Simulating disaster recovery scenarios: Route analytics' network modeling capabilities can be used to simulate various failure scenarios to see if the network as configured will behave as desired. For example, engineers can model failures on critical WAN links and observe how traffic will reroute and whether various classes of service will breach their engineered QoS prioritization parameters – and potentially drop traffic.
- Routing and traffic continuity analyses: Using route analytics' proactive routing audit, traffic trending and capacity planning capabilities, engineers can comprehensively analyze their networks to find vulnerable points such as single points of failure. They can also obtain an inventory of Equal Cost Multi-Path routes that may be implemented to ensure fast recovery from routing adjacency failures, and project traffic trends to detect anomalous traffic behaviors or see whether any links will become congested in the near term.

Enforcing Process Controls: One of the most important aspects of process control in secure IT environments is change management. Not only do engineers need to ensure that their networks can deliver a complex, changing matrix of application traffic at various service levels, but also that the network's behavior will continue to meet critical continuity and disaster recovery requirements as it grows and changes. With route analytics, engineers can model routing and traffic changes and see exactly how the entire network would behave after the change. This provides the perfect complement to device-oriented change management controls on OS versions and command syntax. For example, with route analytics an engineer can model a change of Expedited Forwarding (EF) traffic for a group of users based on the projected rollout of a sensitive new application. The simulated new traffic will be overlaid not on an abstract model, but on the traffic and routing matrix as it actually exists in the network at a time (e.g., peak usage period) chosen by IT engineers. The new traffic and routing picture will then show whether EF or any other traffic class is affected on any link in the core IP network. If not, engineers can proceed with confidence in the rollout, knowing that the network will continue to support existing application requirements. Additional failure analysis simulations can ensure that the new application rollout will not compromise network continuity in the face of disasters or threats.

Increasing Network Quality: Aside from its other benefits, troubleshooting also gets much faster with route analytics, since engineers can see the route/path that a particular service traveled across the network at the time a problem occurred, then analyze all links to see if key applications or CoS were breaching their volume thresholds. If there was congestion, further analysis can show whether a routing issue caused traffic to shift, or, if additional, unexpected traffic was present, where it originated, its destination and the route that included the problem link. Even if a routing or traffic problem isn't the root cause, knowing the precise path provides the most accurate possible starting point for examining devices and interfaces involved in servicing application traffic. Ultimately, faster troubleshooting leads to more problems solved that otherwise could continue to lurk under IT's radar.



and impact current network quality, as well as create further risks in the face of disaster and threat scenarios.

Conclusion

IT departments are under increasing pressure to help organizations achieve better top- and bottom-line results through applications that automate business processes. To succeed, network managers need the same level of automation and productivity in managing the network's delivery of those applications. Ultimately, no matter what other network management infrastructure has been deployed in organizations with large, complex networks, there is no substitute for having full visibility into the network's operational behavior. Packet Design's routing and traffic analysis solutions provide this visibility. Jim Metzler of Kubernan states:

“Variability in how a network delivers application traffic across its multiple paths over time can undermine the fundamental assumptions that IT organizations count on to support many other aspects of application delivery. Organizations with large, complex IP networks need visibility into the operational architecture and dynamic behavior of those networks. Packet Design's route analytics solutions provide the combination of real-time, network-wide routing and traffic-flow visibility that IT needs to manage networks' inherent variability and deliver consistent application performance.”

IT organizations tasked with maintaining applications that support government, national security, and critical public services in particular need this visibility, since failure of IT services doesn't just affect a top or bottom line but potentially the health and safety of a large number of people. As a result, route analytics is being deployed as a critical network management tool by many government, military, public safety and utility agencies around the world.

Packet Design is the pioneer and industry leader in route analytics. Nearly 300 global enterprises, government agencies and service providers have deployed Route Explorer and Traffic Explorer routing and traffic analysis solutions. For more information on Packet Design solutions, please visit our website at <http://www.packetdesign.com> or email us at info@packetdesign.com.