WHITE PAPER

WHERE IS ALL MY TRAFFIC GOING?

by Mark Roberts and Brad Hale

solarwinds

Where is All My Traffic Going?

INTRODUCTION

Over the last few years, there has been a dramatic shift in how IT services are delivered to organizations. What was once solely implemented and managed on-premises is now shifting to a hybrid infrastructure or entirely to an external cloud provider. These changes are being made to reduce costs, reduce infrastructure management, or a variety of other reasons. Regardless of the motivation, the level of direct control you have diminishes as you move further into the cloud.

Even if you maintain your own datacentre, it is likely that you are using some applications or services, such as Office 365® or Salesforce®, that are delivered over the internet. Performance of these SaaS applications is heavily dependent on the WAN, which, in turn, is provided by third-party service providers. When connectivity is interrupted or performance degrades, troubleshooting and identifying the problem outside of your network becomes very difficult. The problem resides somewhere in the "invisible space" between your end-user and the application service (or SaaS) provider.

INTERNET ROUTING PRIMER

The intent of this paper is not to go into the gory details of internet routing protocols, but to show how you lose control and visibility once your data leaves your corporate network.

Let's compare mail delivery to packet delivery over the internet. When a letter arrives to be sorted, the street address is ignored, as the first priority is getting it to the correct region. This is similar to IP addresses in networking. When a source machine generates a packet for a destination, such as when a user on a web browser requests a news website showing today's news, that source machine performs a Domain Name Service (DNS) lookup to convert a textual address (for example, www.bbc.co.uk) into an IP address for a server for that page.

Once the source computer has this IP back from DNS, it will look at its own network settings and determine that this is a packet destined for a location outside of its local network and forward it on to its default gateway. The default gateway, sometimes called the default router or gateway router, is the device responsible for connecting the local subnet to the rest of the world. The receiving router will look at the destination IP address and determine which router to pass the packet along to by looking at its own routing table. This process is repeated until the last router is able to send the packet to the final destination.

In the figure below, you will see an example of traffic that has to be passed all the way up to a Tier 1 provider, whereby the journey can be described as:

- 1. Source PC wants to perform an HTTP Get request for a web page on a destination server.
- 2. The browser application generates the request and passes this to the network layer on the PC.
- 3. The PC performs a DNS lookup and is given an IP of 212.58.226.75.
- 4. The PC determines this is not a local IP address and sends the packets to its default gateway router, usually at your office.
- 5. This device has no knowledge of the 212.*.*.* network and sends the packets to its default gateway router.
- 6. Rinse and repeat until the packets hit a Tier 1 provider, which knows that the packets can be sent to a peering partner (more on this below), so it sends the packets to a peer exchange partner.
- 7. That router knows that 212.58.*.* traffic should be passed to a specific Tier 2 provider router, as it has received BGP information that it is responsible for that network.
- 8. The journey down the path has now begun, with more accurate routing performed as it moves through the IP/subnet definition.
- 9. Finally, the service provider router will receive the packet and, based on its routing table, will know where the end point exists. At the last hop, the end point's layer 2 information will be retrieved from the ARP table or by broadcasting an ARP Who Has request.
- 10. The web server receives the packets and generates the necessary HTTP response back to the client for whom the process is repeated.

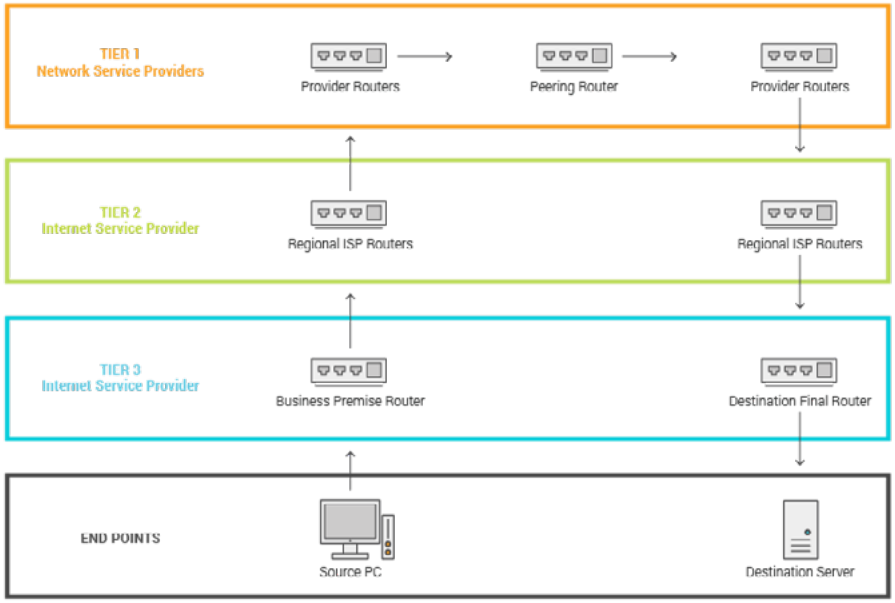

Internet Service Provider Transit and Peering

Packets traverse the internet based on paths that are identified in routing tables, which in turn are defined by routing protocols such as BGP (Border Gateway Protocol). Internet service providers (ISPs) make money by charging other companies to transit traffic across their network. The more traffic that transits, the more money the ISP makes. Traffic routing between providers is managed through peering agreements and internet exchange points (IXP).

Internet Exchange Points

An internet exchange point is a physical network access point through which major network providers connect their networks and exchange traffic. IXPs were created to minimize the part of an ISP's network traffic that had to go through an upstream (higher tier) provider. This creates a more efficient method for routing, as large volumes of traffic can be processed within the high capacity environment. Member ISPs can pass traffic to each other within that LAN without having to pass transit up the chain and over multiple hops. This route "shortcutting" lowers costs, improves fault tolerance, and improves performance by reducing the number of hops in the path.

Peering

Similar to IXPs, peering is a method in which ISPs and large organizations create contractual agreements to exchange traffic between themselves. Some agreements are based on cost related to the volume of traffic exchanged, while others are based on mutual benefit. Either way, routing rules can be determined by cost or benefit.

Shorter, Better Paths

While there are few differentiators between ISPs these days, if your business consumes a lot of resources from certain providers such as Amazon®, Azure®, or CDNs, then the peering agreements ISPs have can have a significant impact on network performance. Say you use Salesforce for your CRM platform. If you had an ISP with peering agreements with Salesforce, the path to the end point would be shorter and higher performing than an ISP who did not. This gets more complex, as the longer the journey the packets take, the more networks are traversed, meaning more transit links and peered networks are involved.

Some companies will choose their ISP purely because they have strong transit links into internet backbone Tier 1 providers and/or short, high-quality paths to common business-critical resources.

Longer, Cheaper Paths

Due to the cost model and peering agreements ISPs have, it's not outside the realm of possibility for ISPs to deliberately inject routes into their network to pass traffic through financially less expensive routes. This will reduce their cost exposure, but the resulting path may be more expensive in terms of latency or quality than if a more intelligent (and correctly announced) path was applied.

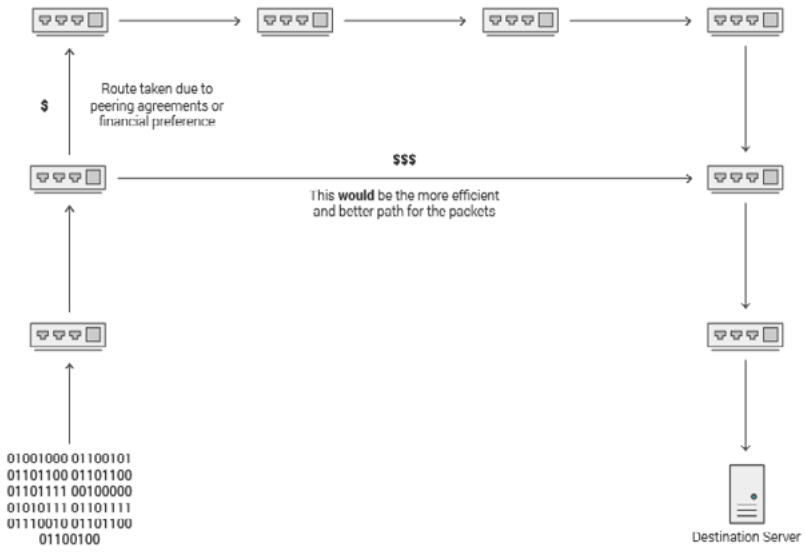

Routing Changes/Bad Routes/Multipath Routing

Internet routing affects the levels of performance you receive, and is more related to things that change infrequently, if they change at all. There are many situations in which internet routing impacts your service levels to end points, as it is designed to provide economy, multiple paths, fault tolerance, and dynamic changes. Multiple paths provide redundancy and greater capacity, and over 80% of paths to destinations on the internet are multipath. As a result, the path taken on one transaction may not be the same path taken on the next.

WHEN IT'S IN YOUR NETWORK, IT'S UNDER YOUR CONTROL

When an organization runs its own datacentre, they have complete control over application delivery and performance. Users directly access services within the corporate-owned and -managed IT infrastructure. This provides the ability to monitor and manage access, routes, and devices to ensure peak performance. Troubleshooting is achievable at all levels, as you have the capabilities to capture the necessary data and information.

Tools such as packet capture, Ping, SNMP polling, flow analysis, and syslog and trap event capture, along with an effective monitoring solution, provide network engineers with just about everything they need for day-to-day monitoring of an internal network. You can easily capture data, set thresholds, and create alerts that provide insight on when and where issues may exist.

But what happens when you run a hybrid environment that relies heavily on cloud applications such as Salesforce, Gmail®, or Google®? Suddenly, the visibility and control that you previously had no longer exists. And even worse, perhaps your IT department no longer has the resources or the technical skillset to properly diagnose issues. Just because you do not have an internal datacentre doesn't mean you can stop "watching" the network.

Let's take a simple example: as a business, you have selected a WAN network provider to connect your regional offices, shops, and home office workers. Whether or not the WAN network is connected to your ISP network connection, you have network connectivity external to your building. You have chosen your ISP because of pricing, quality of connectivity, SLAs, customer service, or other reasons. It is unlikely that most organisations have control over who and how the ISP exchanges traffic with (routes, priorities, etc.).

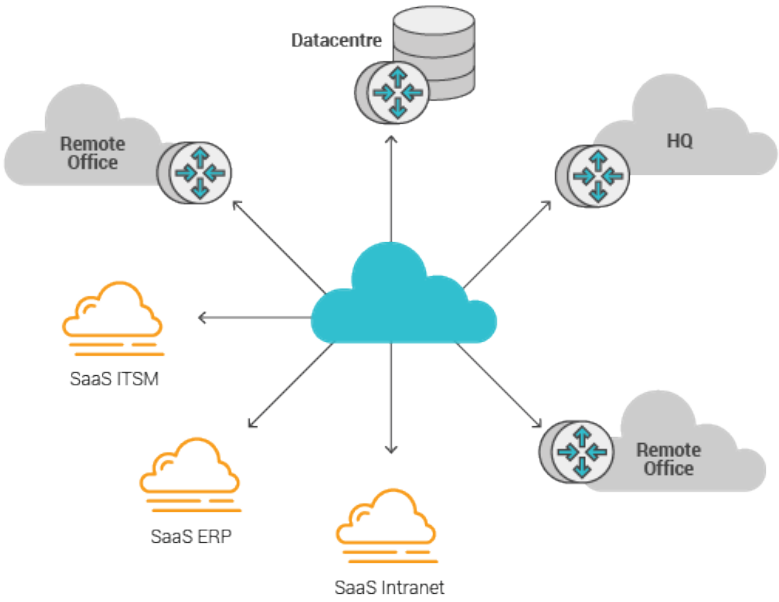

When the CTO comes along asking why a certain service is being reported as slow for users, knowledge beyond your own silo is necessary to ensure job security. One of the tools that has been historically used for monitoring path performance is Traceroute. The original objective of Traceroute was to provide visibility into the hops along a path and their latency. Unfortunately, as paths have become more dynamic and the need for security has increased, Traceroute has been rendered obsolete. Chris O'Brien has written an excellent white paper on the limitations of Traceroute (The Short Falls of Traceroute in Modern Multipath Networks), and I urge you to read it.

When the issue is related to network performance, knowing what paths the traffic is on allows a strong level of understanding for root cause and potential causal performance degradation.

HOP-BY-HOP NETWORK PATH VISIBILITY

Not all paths will perform the same, but, unfortunately, we lose visibility once the packet leaves our network. Immediate and regular monitoring of the paths your external traffic takes will allow you to answer questions such as:

- Which devices are transiting the packets?
- Who owns these devices?
- When changes occur, does this impact or improve performance?
- Is multipath routing is being used, and if it is, are each of the paths performing well?

SOLARWINDS NETPATH

With the release of version 12 of SolarWinds® Network Performance Monitor (NPM), the patent-pending NetPath™ feature became available. Research by SolarWinds found that a combination of a number of probing techniques and protocols could provide a visual solution to the limitations of Traceroute, thus giving birth to NetPath. Additional capabilities and information were included, providing a more powerful solution to the challenge of route identification and the performance delivered not just by the full source to destination, but by each hop.

A probe agent is used as the source. It is then configured with a destination and network port for which it will perform regular probing and performance capture of the path, whether to your own infrastructure, cloud-based, or hybrid networks.

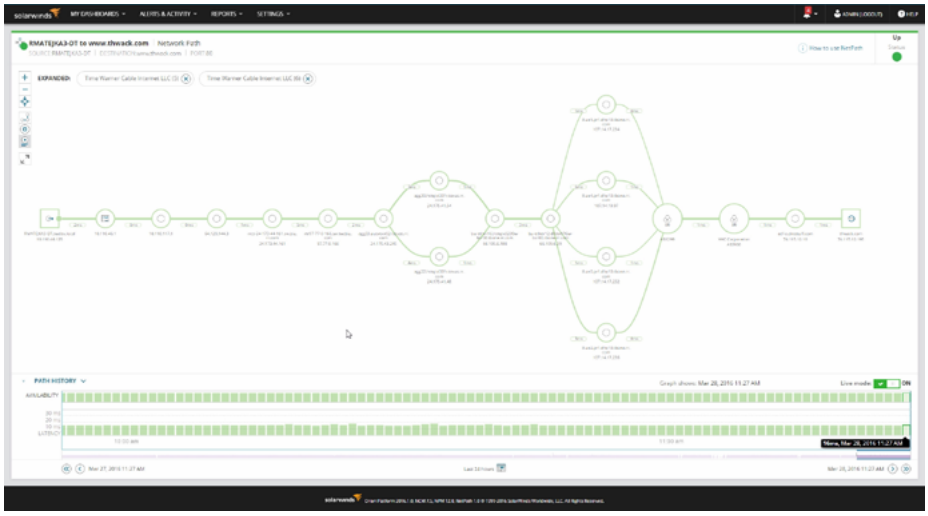

Where NPM is already monitoring routing devices in your network, you can correlate device metrics such as CPU, memory, and interface traffic utilisation, as these can be causes of performance issues. If you have Network Configuration Manager (NCM) and/or NetFlow Traffic Analyzer (NTA), relevant data will be presented from those solutions as well. For example, NCM will show configuration changes that occurred around the time of performance issues that breached thresholds.

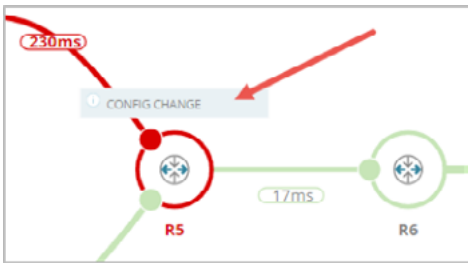

SolarWinds built NetPath to probe for information from each device along the path, providing the ability to see which network domain is causing the performance issues. Additionally, NetPath is able to provide details on who owns each device in your path and provide you with contact details to report your issue.

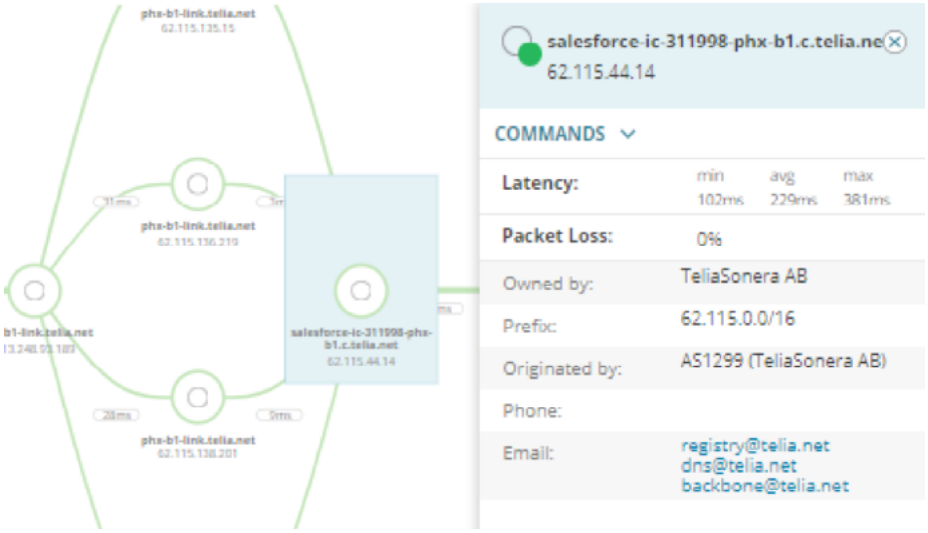

Should you contact them? Of course. You are an indirect customer of theirs and have every right to let them know that you have detected issues affecting routing performance across their network. In fact, many customers and ISPs have reported that visibility and knowledge of detected problems on their network are important to them.

Download a free 30-day trial of SolarWinds Network Performance Monitor and see what NetPath can do for you.

ABOUT THE AUTHORS

Mark Roberts is the Technical Director of Prosperon Networks, a U.K.-based reseller and provider of consultancy and training services for SolarWinds. A regular blog poster and contributor to THWACK®, as a proud THWACK MVP member, he is keen to share his knowledge and ideas with those who are responsible for delivering IT service to users.

Brad Hale is the Product Marketing Principal for SolarWinds network management products. With over 25 years in software, semiconductors, and systems, Brad brings an extensive background in product management, product marketing, and strategy.