ThousandEyes for Application Delivery

White Paper

ThousandEyes

ThousandEyes is
now part of Cisco.

Summary

The rise of mobile applications, the shift from on-premises to Software-as-a-Service (SaaS), and the reliance on third-party services has increased the complexity application delivery. Online Operations and Site Reliability teams in firms of all types are now responsible for many parts of the application delivery chain. When things go wrong, no matter where, they're expected to react quickly to minimize the impact on their customers. The challenge for application providers is that users may often experience service degradation when the application stack seems to be working perfectly fine, making troubleshooting extremely difficult.

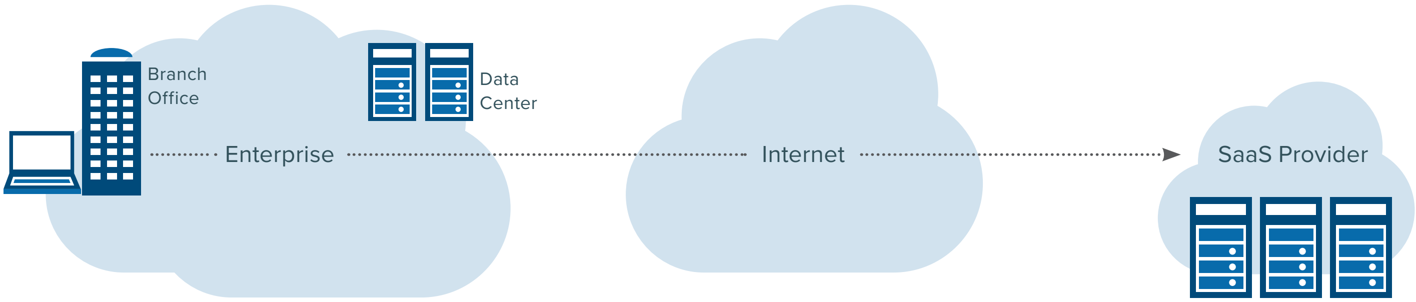

Figure 1: Application delivery through different network segments.

Challenges in Delivering Online Applications

Service degradation may be due to problems anywhere in the application delivery chain: 1) inside the user's local network or the customer's corporate network, 2) in the Internet, or 3) in the application provider's data center. Traditional performance monitoring solutions that rely on server-side instrumentation are ineffective in troubleshooting these different network segments between the end user and the application, each of which can negatively impact user experience. To troubleshoot problems quickly, application providers need a new solution that can find problems anywhere along these different network segments and can correlate network issues with application performance.

The primary challenge that online operations teams face today is lack of deep visibility beyond their application stack. Users may still complain about the application being slow or unavailable when nothing seems wrong with the application as seen by the provider. This is typical when things go wrong in the application delivery. Several things outside the data center environment may be impacting the user experience. A third-party DNS service the application relies on could be misbehaving, a CDN used by the provider could be adding longer delays, or network problems somewhere along the path could be negatively impacting the application. Furthermore, the source of the problem might be within the enterprise or local network using the application, making it almost impossible to troubleshoot. There are a host of things that can go wrong. Application providers need to track a variety of different metrics.

Key Metrics for Application Delivery

Key metrics to track application performance at the web application layer include:

- **Time-to-first-byte:** This metric is also called response time and is a good approximation of the time it takes for the server to start sending data to the user. The download process of each component follows a sequence of steps including: 1) DNS lookup, 2) TCP connect, 3) SSL negotiation, 4) HTTP GET (or POST). Receive time is the time from the DNS lookup till the first byte of data received by the client.

Users may often experience service degradation when the application stack seems to be working perfectly fine, making troubleshooting extremely difficult.

- **Page Load Time:** The time it takes for the browser to trigger the load event for the page. It includes the time to download the HTML, parse the markup, process CSS, render the page and download all images and files required in the page.
- **Transaction Time:** The time to perform key user actions in a script that simulates user actions in a group of pages (e.g. log in to a site and click on certain links).

Key metrics to track at the network layer are:

- **Latency:** The round-trip time (RTT) between the client and server is referred to as network latency. Latency depends mostly on the physical distance between the endpoints, routing behavior, as well as the congestion state of the network. Higher latencies cause higher application response times, as well as lower TCP (Transport Control Protocol) throughputs, making the application seem slow.
- **Packet Loss:** When packets are dropped anywhere along the path, the result is degradation of service and lower TCP throughputs, slowing down the overall transfer of objects from server to client. In extreme situations, the application may become completely unavailable due to packet loss.
- **Capacity and Available Bandwidth:** Capacity between two endpoints is the maximum data rate that can be achieved in the absence of any cross-traffic. Cross-traffic will use a portion of the capacity, and the remaining is the Available Bandwidth, which determines how fast TCP can send messages through the connection.
- **Routing availability:** In the Internet, networks exchange routing information using the Border Gateway Protocol (BGP). BGP allows independent neighboring networks to talk to each other to decide what traffic they will exchange. BGP changes (e.g. misconfigurations) can render an entire network unreachable or induce severe performance degradation on applications.

While each of these network metrics are important, one of these metrics may have a disproportionate impact, depending on the type of application (e.g. if the application involves a lot of data transfer, lack of enough available bandwidth can severely cripple user experience).

Connecting the Dots with ThousandEyes

ThousandEyes provides a unique solution that allows SaaS providers to characterize and troubleshoot the different elements of application delivery using measurements from lightweight agents. Application providers can conduct measurements using Cloud Agents spread across the globe or easy-to-install Enterprise Agents in their own data centers or customer environments.

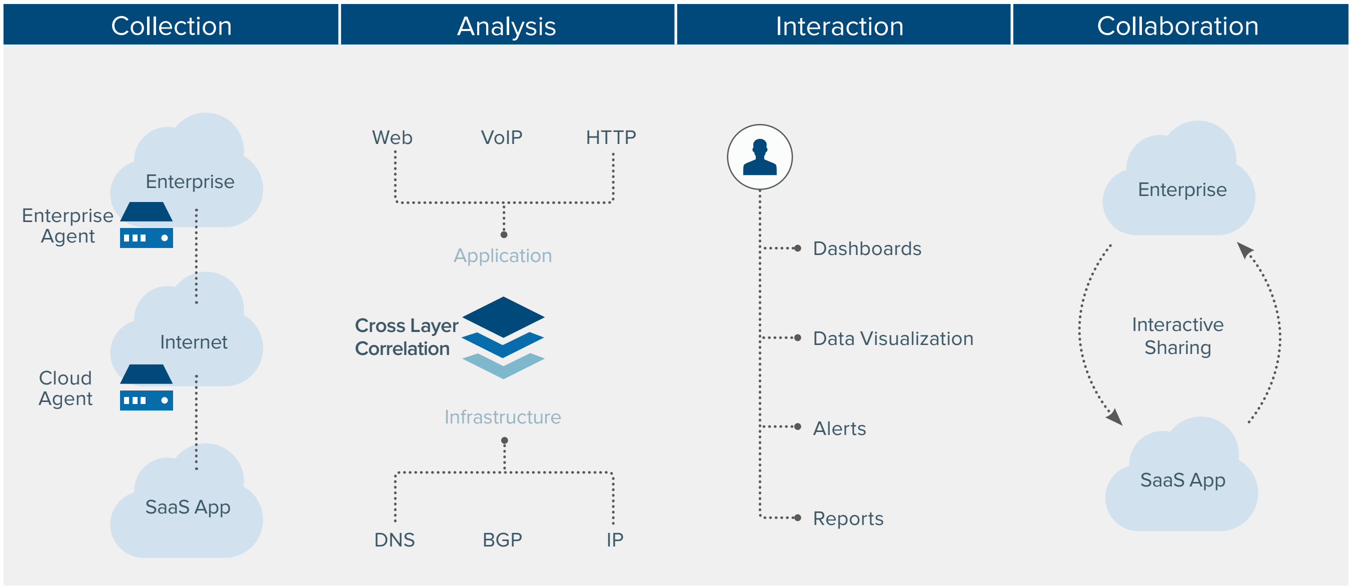

ThousandEyes goes beyond monitoring and provides actionable network intelligence.

ThousandEyes provides deep visibility into each layer of application delivery and a connecting thread between these layers, making it possible to jump from layer to layer to understand where the problem is. Online Operations teams can immediately identify which component of the application delivery (e.g. DNS, HTTP, Network) is broken and dive deeper into that specific component to visualize the problems and interact with the data. Plus, ThousandEyes clearly shows large scale Internet outages, learned from data in all customer tests, to speed up the process of root cause analysis.

Finally, using interactive sharing provides a way for others to see the same data and interact with it as well, enabling SaaS providers to collaborate with their ISPs, DNS providers, CDN providers and even customers to resolve problems faster. Figure 1 shows the flow involving, collection, analysis, interaction and collaboration.

Troubleshooting Infrastructure Problems with Cloud Agents

Data collected from Cloud Agents is very useful to identify problems with third-party DNS providers and CDN providers, as well as network issues within and close to their data centers. The more Cloud Agents that detect a problem, the higher the likelihood that customers are being impacted by it.

Cloud Agents are useful to conduct periodic testing to ensure application health, and detect issues closer to data centers, such as problems with upstream ISPs. Let’s look at a real example of troubleshooting with Cloud Agents:

- **HTTP Server Availability Problems:** Figure 2 shows an example of availability problems detected from Cloud Agents in Europe. Further examination reveals that all of these agents were failing in the TCP Connect phase and hence unable to establish a connection to the web server.

> “ThousandEyes has become one of the key tools in our operations center for troubleshooting network and application performance problems.”
>
> **Ivan Batanov**
> *VP Engineering*
> *ServiceNow*
>
> servicenow

- **Packet Loss Within an ISP Network:** Figure 3 visualizes the network layer at the exact same time. One can see several hops in the network dropping packets, indicating severe network issues impacting the application. Further analysis indicates the problems localized to different locations within the same ISP.
- **DNS-Based Load Balancing:** Figure 3 also shows a load balanced service across multiple data centers and web servers. Performance across locations can be compared to ensure optimal routes and proper geographic distribution. This can be especially helpful for environments with global data centers, a CDN edge or that use Anycast.
- **BGP Activity Contributing to Loss:** To understand if lossy behavior has anything to do with Internet routing, we jump to the BGP layer where there is a corresponding spike in BGP routing activity, shown in Figure 4. A closer examination reveals that BGP routes moved from one provider to another. The provider that the routes moved away from was the one dropping packets, thus indicating an inconsistency that often develops when routes change at the BGP layer but packets are still being forwarded along older paths.

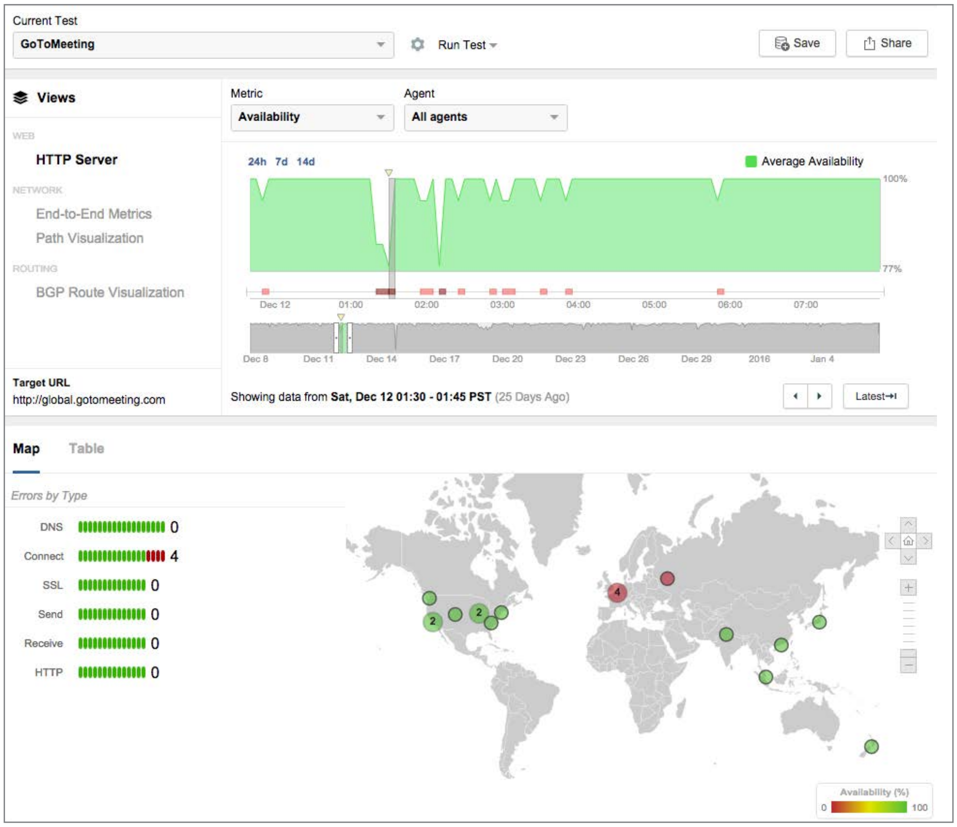

Current Test: GoToMeeting

Metric: Availability

Agent: All agents

Target URL
http://global.gotomeeting.com

Showing data from **Sat, Dec 12 01:30 - 01:45 PST** (25 Days Ago)

Errors by Type

DNS	0
Connect	4
SSL	0
Send	0
Receive	0
HTTP	0

Figure 2: HTTP Server layer visualization indicating agents impacted (in red) in the world map, and step where errors are occurring (TCP connection establishment).

In the example above, the Cloud Agents helped troubleshoot an application availability issue caused by lossy network behavior due to a BGP routing change. Even when BGP is stable, problems could still occur at various network hops and the Cloud Agents help keep a close watch on the upstream ISPs that the application providers rely on. Knowing exactly where the problem lies can enable application providers to change their routing policies or DNS load balancing to route around the problem.

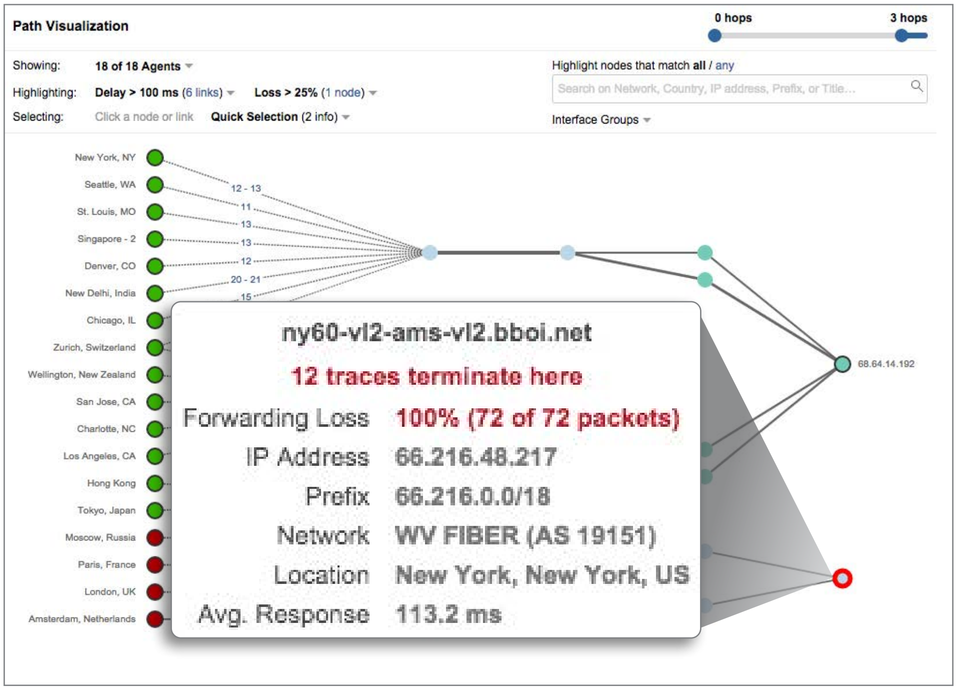

Path Visualization

Showing: 18 of 18 Agents

Highlighting: Delay > 100 ms (6 links) Loss > 25% (1 node)

Selecting: Click a node or link Quick Selection (2 info)

Highlight nodes that match all / any

Search on Network, Country, IP address, Prefix, or Title...

Interface Groups

ny60-vl2-ams-vl2.bboi.net

12 traces terminate here

Forwarding Loss	100% (72 of 72 packets)
IP Address	66.216.48.217
Prefix	66.216.0.0/18
Network	WV FIBER (AS 19151)
Location	New York, New York, US
Avg. Response	113.2 ms

Figure 3: Path visualization highlighting where packets are being dropped.

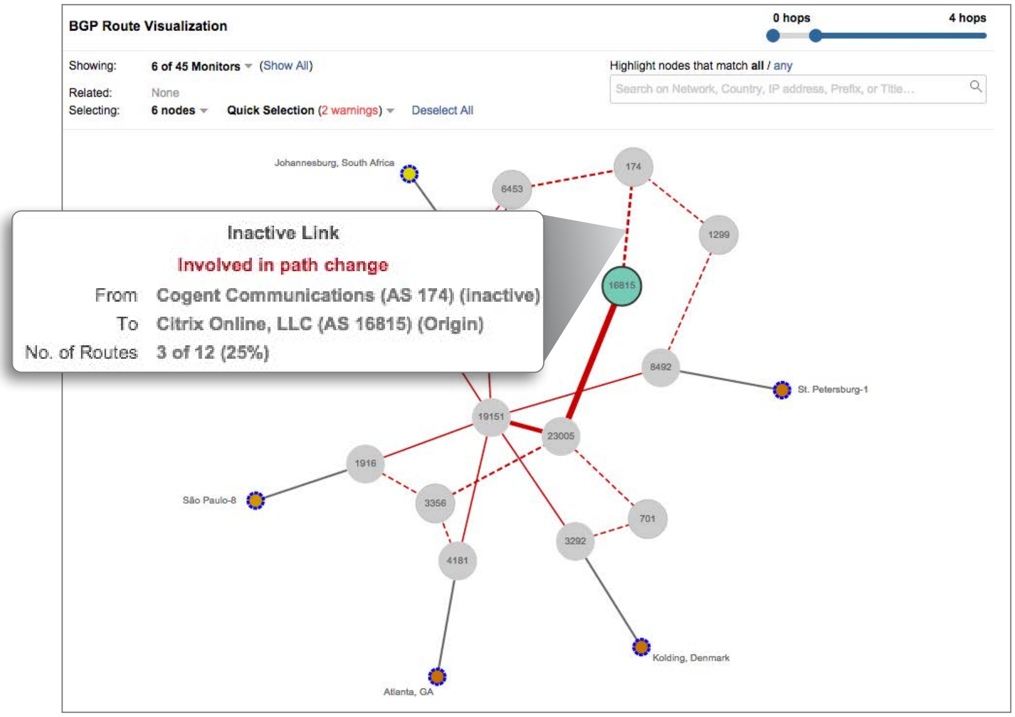

BGP Route Visualization

Showing: 6 of 45 Monitors (Show All)

Related: None

Selecting: 6 nodes Quick Selection (2 warnings) Deselect All

Highlight nodes that match all / any

Search on Network, Country, IP address, Prefix, or Title...

Inactive Link

Involved in path change

From	Cogent Communications (AS 174) (Inactive)
To	Citrix Online, LLC (AS 16815) (Origin)
No. of Routes	3 of 12 (25%)

Figure 4: BGP layer showing routes shifting from one provider to another.

Using Enterprise Agents for Data Center and Customer Environments

Enterprise Agents can provide additional vantage points to assess performance. As applications become more distributed, many application providers need visibility between data centers to understand communication between different application components and micro-services. And for SaaS providers, Enterprise Agents can offer a view from within end-customer environments that characterizes the underlying network paths and behavior. This is really important when responding to customer issues that cannot be reproduced using Cloud Agents.

ThousandEyes Enterprise Agents can be easily deployed as a virtual appliance or a Linux package to conduct measurements from the data center or the enterprise customer environment. Once installed, Online Operations teams can then run a series of tests from the application layer to the network layer and not only measure the application performance but also understand network issues contributing to the service degradation. Typical network issues that are caught by Enterprise Agents include suboptimal routing, DNS problems, proxy bottlenecks, capacity bottlenecks and MPLS misconfigurations. Sometimes, Enterprise Agents can also help catch problems with DNS providers and CDNs, or even the application provider's own network, especially if it is wide and complex, making it difficult to cover entirely through Cloud Agents.

Conclusion

With ThousandEyes, Online Operations teams can troubleshoot application delivery problems in hybrid environments in a matter of minutes. Cloud Agents are an easy way to get started with ThousandEyes since they don't require any special deployment effort except test configuration. Enterprise Agents can be deployed at customer sites and be configured to perform periodic tests to the application provider's application. ThousandEyes can reduce the MTTR (Mean Time to Resolution) of infrastructure problems from hours and days to a few minutes.

Please visit www.thousandeyes.com for more information.

ThousandEyes

201 Mission Street, 17th Floor
San Francisco, CA 94105
(415) 513-4526

www.thousandeyes.com

About ThousandEyes

ThousandEyes is a network intelligence platform that delivers visibility into every network your organization relies on, enabling you to resolve issues faster, improve application delivery and run your business smoothly.