WHITEPAPER

Choosing the Right Technology for Remote Location Monitoring

IT has to rethink network performance and end-user experience when users and infrastructure are off-site. What's the best approach to make remote office technology work seamlessly?

Technology in the modern world is, above all, distributed. Cloud computing has pushed processing out of company-owned data centers into cloud data centers and SaaS providers. Depending on your favorite industry prediction source, there's a current 41% cloud adoption rate and an estimated 92% of workloads will be processed through a cloud data center in just a few years. Workloads and servers aren't in-house anymore. Users work from many locations on many devices.

For IT teams, these distributed apps, users and locations mean that there's physical and logical distance involved in supporting end users today. This distance may seem small when users don't run into problems, but it can seem huge when a user has an issue that IT can't easily troubleshoot. Those issues can easily be compounded when IT and the user aren't in the same place, and when there's no IT team at all at the remote office.

The complications of remote offices bring a lot of technology complications for IT. To make this new type of infrastructure work, you'll have to rethink the troubleshooting of the past. Let's talk about how IT should get started setting up new branch locations, including the choices involved around networks, cloud providers and metrics.

Choose Your Remote Location Technology

If you have the power and independence to pick your perfect mix of technology to support remote locations, that's great. What's more likely is that you'll be patching together the tools you'll use for these remote offices or sites. The technology needs for your remote offices will fall into a few top categories.

1. Cloud provider options

When it comes to cloud providers, location really matters. It will help quite a bit if you can match up the location of a cloud provider with the users who will be served by that cloud provider. This is especially relevant in the case of remote locations. Even if the branch office you're setting up or maintaining isn't far from your main office or data center, a different cloud provider may have a data center closer to that branch office. That smaller distance can have a big positive impact on network performance and end-user experience.

Choosing a public cloud provider has to do with location and also with the provider's networks. Those cloud providers depend on networks just like the rest of us, so do as much research as you can to see which provider might work best in each remote location. If more than one public cloud provider might be useful to you, do your homework on managing multicloud environments to balance complexity with performance.

2. Network options

Connecting remote locations to the main office, cloud or SaaS providers and to each other can potentially be quite a feat. There are various network technologies to choose, and you'll want to use the best tool for the job. Consider these common networks and whether they might be part of the network infrastructure supporting your remote sites.

WAN. Will you use an existing WAN or deploy a new one to remote offices or other locations? Either way, think about whether you'll require WAN optimization techniques or tools. You may also just skip to SD-WAN for some of the same features that WAN optimization provides.

SD-WAN. You can add SD-WAN products to your existing networks like broadband and MPLS to balance network loads and optimize performance. SD-WAN pairs multiple networks, then routes traffic depending on network health or IT-directed guidelines. Consider its limitations, though: SD-WAN can't see application context to route traffic accordingly, or guarantee redundancy.

MPLS. This type of network has been used to connect remote locations for ages, but it's expensive. Consider using it sparingly, since it provides good bandwidth and QoS. SD-WAN makes it easier to use MPLS only as needed.

Public internet. Plenty of businesses' branch offices or off-site locations use a broadband connection for user access. SD-WAN can make this more tenable if it's paired with MPLS.

WiFi. Business users depend on WiFi these days. This connection should be part of your monitoring scope along with the other networks, especially if you're providing WiFi to customers in stores or guests in hotels. You should compare wired and WiFi performance regularly to make sure user experience is good for both connections.

VPN. Businesses rely on VPN connections of several secure varieties: SSL and IPsec. A virtual private network (VPN) is essential when users are accessing company information off-site. Though they're secure, VPNs are still under IT control.

Remember, these are just the networks under your control. Now that so much of the application delivery path is outside of the traditional on-premises infrastructure, there's a lot you can't see. The cloud or SaaS provider networks are now essentially part of your network infrastructure. Those networks will fill in part of the application delivery path, so they add unknown hops in your user's network path. This is where monitoring comes in.

3. Monitoring options

How do you plan to track performance at each of your remote locations? Monitoring tools are particularly important when there are no IT resources at the remote locations. We're biased, of course, but here at AppNeta we think we've gotten the right combination of features down to monitor the entire user's network path. Being able to see the entire network path, even into cloud and SaaS providers, gives our IT team users an edge in finding problems and avoiding "Is it the app or the network?" questions.

One essential component we offer for remote locations is synthetic monitoring. That allows you to continuously see what users are seeing from their location. Synthetic monitoring uses scripting to emulate end-user paths through an application. That way, actual data gets loaded, and you can see the login page, not just the app's main page. Synthetics are especially important now because apps don't often go down entirely—they're more likely to become very slow. In addition, this approach lets you skip laptop agents, which are clunky to update and maintain over time.

Understanding User Experience

End-user experience at remote offices can be hard to quantify. Tracking help desk tickets and support calls can help you get a picture over time of where issues lie. Also pay attention to those ghost issues that are impossible to recreate but pop up now and then.

Consider these relevant metrics to get a good idea of how end users at remote locations are experiencing applications:

Response time: For cloud applications, response times can tell you pretty clearly where users might be experiencing frustration. Today's users don't tolerate much slowness, if any, when they're trying to open an app.

Throughput: This number reflects how well the service provider can handle your demands. For bandwidth-heavy apps, lack of throughput can lead to bad end-user experience.

Number of impacted users: Ideally, you won't need to know this number all that often. If you're hit by multiple issues and need to prioritize with service providers, it'll help to know how many users are affected by each issue.

Invocation time: This number helps quantify response time by looking at the time it takes the app to perform transactions through servlets, XML, APIs and more.

Know Your Remote Location Data

Ideally, you'll have a shiny new remote office or new hospital building to wire and set up. It's also possible you'll be patching together old hardware with new laptops, or some other combination of technology tools. Either way, do your best to get a baseline of network performance at that office or location, either before it's up and running or before you start monitoring in earnest. This baseline will be very useful when problems occur and you need to start pinpointing the cause.

When it comes to numbers and reports for the technology running your remote locations, there will be a few important areas to consider. Metrics is one, and service-level agreements (SLAs) is another. The two are linked—metrics should be featured in the SLA so you know what to expect from the provider during an outage or slowdown.

Ideally, you'll have a shiny new remote office or new hospital building to wire and set up. It's also possible you'll be patching together old hardware with new laptops, or some other combination of technology tools. Either way, do your best to get a baseline of network performance at that office or location, either before it's up and running or before you start monitoring in earnest. This baseline will be very useful when problems occur and you need to start pinpointing the cause.

Metrics

Decide which metrics make sense for your business and end users. There are a couple of old standby metrics that still make a lot of sense for modern, cloud-run businesses:

- **Jitter** gets noticed quickly by users. It's the percentage of packets with delay variation between source and destination, and users experience it as choppy video or audio.
- **Latency** should include round-trip time to get the real picture of how packets are traveling from source to destination. The internet is asymmetric, and latency measurements should take that into account.
- **Packet loss**, the percentage of network packets lost between source and destination, isn't the same as data loss. Users will notice even 1% of packet loss, since it results in slow applications. It can also seriously affect the quality of VoIP and video streaming apps.
- **Capacity** gauges the actual application delivery path, including WiFi, and reflects the maximum transit rate between source and destination. It's important to understand both utilized and available capacity and how they affect end-user experience.
- **Quality of service** guarantees are tied to routing priority for traffic over specific network ports or protocols. If QoS is enforced along the entire network path, it can ensure that higher-priority workloads get to their destination first. So critical applications will still work well for users, even if email or music streaming is slow.

In this cloud-driven world, monitoring these criteria is ever more important. And when you're supporting remote locations, you should continuously measure jitter, latency, packet loss, capacity and quality of service to get a full picture of how networks and applications are performing at any given location. You may consider other metrics depending on your industry, or prioritize some of these. For example, a remote call center location using VoIP will have jitter and packet loss as high-priority metrics, and IT may set up reporting and alerting on those accordingly.

It's helpful to understand how these metrics work together, too. For example, network congestion isn't just about bandwidth constraints. It could also have to do with the overall capacity of that particular network. Capacity is the highest achievable bandwidth you can get on the most congested hop along the network connection between client and app. So adding more bandwidth won't necessarily fix network congestion problems.

SLAs. Now that SaaS, cloud and internet providers are all part of your infrastructure, those relationships and agreements are essential. Do your research before signing the SLA, if you can. You may inherit an SLA for a provider or application, too, in which case, make the most of it. SLAs go beyond simple availability and may include uptime guarantees, disaster recovery plans and other promises, along with penalties if the provider fails to meet goals. Know which metrics you're tracking so you can align those goals with the SLA.

SLA management can become a weighty task if you've got agreements with various network providers, at various remote locations, for multiple public clouds and SaaS providers. Try to streamline where you can—for example, if you deploy SD-WAN, take that SLA into account more than the one associated with the WAN or public internet connection.

Prepare for Growth and Change, at Any Location

As distributed and cloud-based as businesses already are, there's still more room for growth and change. Cloud continues to grow, as does SaaS and the use of WiFi. Ideally, you'll use what you've learned in managing and monitoring one remote office or branch to replicate that success at the next new store, hospital or bank.

In addition, continue to improve reporting that can spur further action and improvement. For example, once you can monitor every ISP and see its performance over time, you can decide which ones you will use most or for the most important workloads—or which ones you might stop using altogether. The more you know your infrastructure, the faster you can solve problems. It'll be easier to pinpoint a hardware problem if you're confident in the performance of the network it's attached to.

AppNeta does all of this and more—in fact, remote location monitoring is one of our specialties. It is possible for remote offices and external sites to be under your control, and our monitoring tools can help.