

How to Overcome the Darkside of Moving IT Apps to the Cloud

The Emergence of the Darkside of Public Cloud Networks

The last several years have seen tremendous growth in public cloud networks. And for good reason. There is the real possibility of decreasing time to market for cloud-based business applications. This has caused many enterprises to make the leap to the cloud.

However, once the rosy glasses disappear, a dark side emerges. This often takes the form of performance problems, cost over runs, and increased downtime which makes the return on investment (ROI) thin or non-existent. In fact, some [companies](#) are moving all of their operations back to on-premises.

So, how do you overcome issues with the cloud? Is the answer to move everything back to on-premises? Maybe. Another option might be to split operations between the cloud and on-premises. For instance, according to the [State of Hybrid Cloud](#) survey report published by Virtana, 72% of enterprises that moved their operations to the cloud have decided to move at least one application back to an on-premises location. This is because an estimated 41% of the applications should never have been moved to the cloud in the first place.

[An IDC survey](#) found a similar result where, “71% of respondents expect to move all or some of their workloads currently running in public clouds back to private IT environments over the next two years.”

Let's take a quick look at some of the issues affecting the public cloud and then discuss how you can overcome them to achieve your perfect balance.

Cloud Networking – What's the Problem?

The fundamental issue with cloud networking is that too many businesses are leaping to the cloud without creating a solid plan of what should be moved and why. Basically, the business teams are getting ahead of the technology teams. Once this happens, the whole project can be jeopardized or hampered with triple (or more) costs as applications are moved to the cloud, moved back to on-premises while the technology teams regroup, and then specific applications are moved to the cloud for a second time.

So, how do we overcome this? The first thing to is to understand what potential problems you may face as you move to the cloud. For instance, the white paper [Blinded by the Cloud – The Real Cost of Moving Applications to the Cloud](#) details several of these issues. The core list is summarized below:

- Security concerns
- Performance issues
- Single vendor lock in
- Complexity due to use of a multi-cloud architecture
- The actual cost may be much higher than you were quoted

The next step is to address each core issue. For instance, are you concerned about single vendor lock in? If so, what is your solution? Are you planning a multi-cloud solution – basically using a combination of two or three public cloud vendors (AWS, Google Cloud, Azure, etc.)? If so, you need to address the complexity of replicating and synchronizing data across multiple clouds. You should keep in mind that a multi-cloud approach will result in double, triple, or more, costs.

Another set of questions involve, “How do you plan to address the security for your instance?” While the vendor has responsibility in this area, so do you? Therefore, before you make the leap, you need to address all of these concerns (and more) in your plan.

If you need more help, you can refer back to the [Blinded by the Cloud – The Real Cost of Moving Applications to the Cloud](#) paper as it not only summarizes the five problems but offers solutions as well.

How to Make the Cloud Visible

This next step is to make sure you plan network visibility into your solution. Whether you have a pure cloud, pure physical on-premises, or mixture (hybrid) of both network types, you need to be able to observe issues and problems as they occur. For instance, the single most important activity an enterprise can do to protect itself from cybersecurity threats is to implement a visibility architecture. This is because if you can't see the problem, how do you intend to fix it? When your network is breached is not the time you want to find out that you have blind spots across your network.

Shockingly, Enterprise Management Associates (EMA) found in a 2022 research report ([Network Visibility Architecture for the Hybrid, Multi-Cloud Enterprise](#)) that 66% of the companies surveyed failed in their attempts to implement a visibility architecture. This means that those enterprises were blind to potential problems and probably had encountered quite a few issues.

Your analysis tools then must not only see packets inside the network, they must also recognize patterns of activity — good, bad, and unusual. But on-premises, cloud, hybrid, and multi-cloud networks are making visibility more difficult than ever. This is why you need a visibility architecture.

So, what is a visibility architecture?

Three Layers to a Visibility Architecture

A visibility architecture consists of three fundamental sections:

- A data access layer
- A data control plane layer
- And a monitoring tool layer

Each layer consists of hardware and/or software dedicated to a specific set of functions. This is illustrated in Figure 1.




Figure 1: Three Layers of a Visibility Architecture

The first layer is the data access layer. This section splits off (or creates) copies of packets and then forwards them to the control plane layer. Within the access layer you will find taps, virtual taps (software taps for the cloud), SPAN ports, bypass switches, and aggregation taps.

Here are some important facts to consider about the access layer:

- Taps are dedicated devices that can't be hacked (because there is no IP address) and fail to wire (if power to the tap is lost, traffic continues to pass into the network)
- Taps can often be managed, making administration easier than SPANs
- Taps don't drop packets, making them the only viable solution to ensure your tools see 'all the packets'
- SPANs are often considered "free", but nothing in life is really free. SPANs must be programmed and reprogrammed as the network changes. And growing complexity increases the likelihood of not seeing all the packets

The control plane layer optimizes packets received from the access layer, and then forwards them to the tools within the monitor layer. Network packet brokers (NPB) are the main component of the control layer. The vendor and model of the packet broker chosen will have a significant impact on how well you can optimize your network data. For instance, packet brokers that support advanced features often yield the highest ROI.

Advanced packet broker features include the ability to:

- Aggregate packets from multiple sources
- Filter packets by OSI Layer and send packets to tools base on OSI Layer filtering
- Load balance packets sent to the analysis tools at the monitoring layer
- Regenerate traffic
- Remove duplicate packets
- Strip off unnecessary header information
- Perform SSL/TLS decryption
- Generate metadata using NetFlow

The monitoring layer is where threat analysis, network and application performance management, and network troubleshooting occur. In this layer you will see a vast array of specialized security analysis tools, like Security Information Monitoring (SIM), Security Event Monitoring (SEM), or a combination of both as in Security Information & Event Monitoring (SIEM).

Highly specialized in functionality, tools at the monitoring layer typically require skilled personnel to successfully operate. But even then, they're only as good as the data they receive from the network.

Active Monitoring Delivers Unique Insights

In addition to the basic visibility architecture, you should strongly consider investing in an active monitoring solution. An example would be the Keysight Hawkeye product. These types of solutions let you test current and future network and application situations so that you have visibility into network and application performance. As an example, the [Virtana 2021 State of Hybrid Cloud report](#) found that 36% of survey respondents encountered provisioning issues as they migrated to their public cloud instances. Active monitoring can help prevent this situation.

Active monitoring delivers four important use cases for cloud and hybrid cloud networks:

- Immediately test for application and performance complaints
- Create network and application performance baselines for on prem and cloud
- Periodically test cloud performance as applications are moved to the cloud
- Validate your cloud provider's SLA performance for an independent assessment

To get the performance answers you need, you will want to create a proactive cloud monitoring solution. Basically, this is a monitoring solution that uses software agents and probes that you can place across your cloud and physical infrastructure. Data can be collected either passively or actively by probes within the network to give you an instant status of what is, and what is not, happening on the network. With this solution, you can use visibility technology to actively test your solution before migration, during migration, and after migration.

During the migration process, proactive performance monitoring of both your on-premises and cloud environments will be useful. Test the performance yourself to characterize how it is actually working at all phases. With the right tool, this testing can be fairly painless. An alternative is to copy and export cloud data back to your on-premises performance monitoring tools (assuming that you are operating a hybrid cloud environment) for analysis there. Many organizations that just blindly port services and applications to the cloud encounter cloud computing issues quickly, particularly performance issues.

Business intelligence applications are one example of a problem area. After porting the service from your completely controllable on-premises environment to a public cloud instance, you may find that it runs slower (after you receive multiple customer complaints). The “lift and shift” concept failed. The result is often an increase in more CPU, RAM, and interconnect bandwidth. This creates an unplanned and perpetual cost increase.

To get a true indication of network performance, network tools (e.g. NPM and APM) need to have a good, i.e. large, amount of traffic coming in to them. This often makes you dependent upon peak busy hours and so forth. With proactive monitoring solutions, you can place probes anywhere in your network to test with whenever you want to. The synthetic traffic provides you the network and/or application loading of a “busy hour” and the flexibility to perform evaluations during the network maintenance window.

Your DevOps teams can pretest how an application will perform on the network under load before your users do to create faster and better network upgrade rollouts. For instance, you can simulate lots of application traffic, a combination of traffic types, and/or lots of different protocol usage that could stress test the network without having to wait until the busy hour.

Once the migration starts, you can measure the ambient latency, throughput, and performance problems on a per-hop basis within the network to see how it is performing. This lets you analyze both your on-premises solution as well as your cloud solution and can be especially important if you have a hybrid solution right now but are in the (often multi-year) process of transitioning from the physical to the virtual (cloud) world. A proactive testing and monitoring approach gives you the confidence that your new application rollouts will be successful in either network.

A fourth common problem is that you often cannot natively tell how your applications are truly performing or even how your cloud instance is performing. Is it meeting or exceeding the service level agreement (SLA) that was put in place? Your cloud vendor will probably tell you that it is, but you have no independent data for a “check and balance” strategy on what they are delivering.

A proactive monitoring solution can provide SLA and customer experience information for a wide range of applications including voice, video, web services, and critical enterprise applications. The information gathered can then be used to inform management about which goals are being met. If goals are not being met, you can use the impartial data you have collected and contact your vendor to have them either fix any observed network problems or give you a discount if they are failing to meet agreed upon SLAs.

Proactive monitoring also allows you to perform SLA validation during business hours, since it is not service disrupting. This allows you validate the cloud vendor’s SLA performance at will.

Don't Be Afraid of Reversion to On-Premises Solutions

As has been mentioned earlier, once you have made the leap to a public cloud solution, you may find out that you’ve made a mistake. If this is so, it’s okay. You simply need to make a reassessment as to how your business goals and network architecture fit together. It’s better to get the architecture right as soon as possible, so that you can capitalize on the benefits and not spend time “fighting fires.” Since the cost economics are not working out as advertised, it doesn’t make sense to waste money just because of the cloud hype.

Many companies are performing this activity. A Forbes article ([Why Is Cloud Migration Reversing From Public To On-Premises Private Clouds?](#)) mentions that “many companies that aggressively migrated their work from on-premises clouds are looking to move work back to on-premises and private clouds.”

As an alternative, maybe a hybrid cloud architecture might be the best idea? [Forrester Research](#) and other analysts are suggesting this hybrid approach. Physical taps can be combined with cloud taps to capture data across the whole hybrid architecture. Physical and virtual packet broker solutions can then be used to filter and optimize that data before it is sent to physical and/or virtual security and monitoring tools for analysis. Active monitoring solutions can be thrown in to give you more network insight. This blended approach gives you visibility across your whole integrated physical and virtual environments.

Conclusion

As cloud deployments continue to shift and settle, it may be that we will see a higher proportion of hybrid cloud solutions. While the scenario might sound a little scary to some, this should actually be a good thing. Someone once said, “You need the right tool for the right job.” A hybrid IT solution could, depending upon your business particulars, enable you to leverage the best of both (cloud and on-premises) worlds.

Visibility in your network is not optional. Failing to build a scalable visibility architecture to provide visibility for today and tomorrow could ultimately result in limited visibility, leaving you susceptible to costly security incidents and performance problems.

Reach out to Keysight Technologies and they can show you how to optimize your public cloud, physical on-premises, and hybrid cloud solutions.

Learn more at: www.getnetworkvisibility.com.