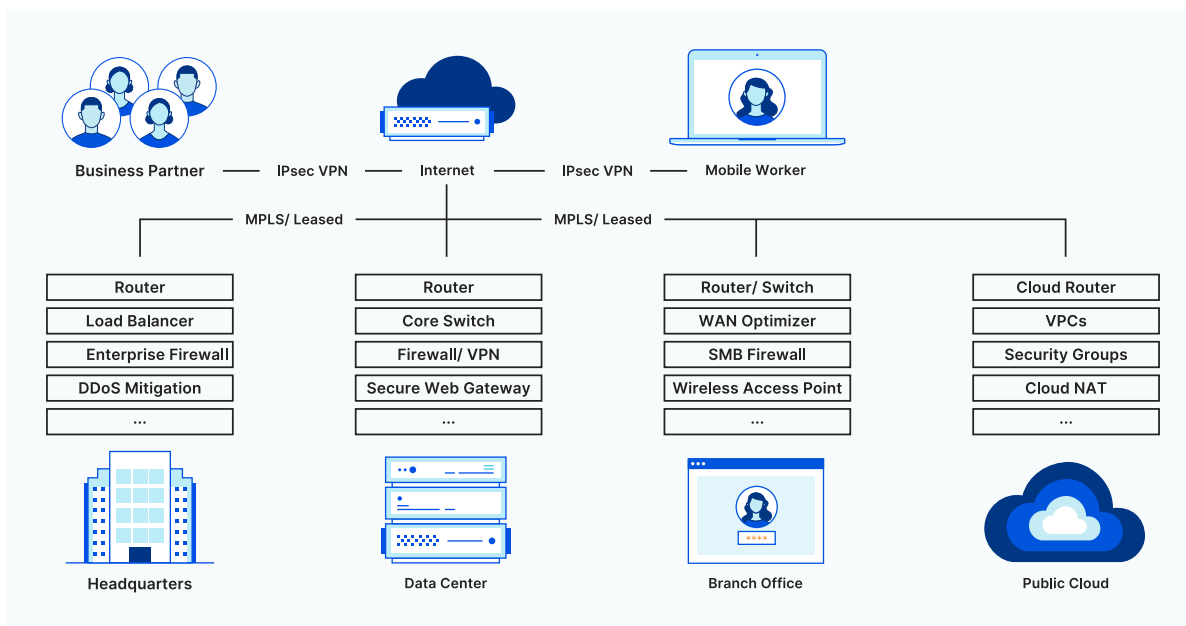

The death of network hardware appliances — and what it means for your cloud migration

During cloud migration, many networking functions remain on-premise, creating capacity limitations, high total cost of ownership, support challenges, and security gaps. This paper outlines those challenges, quantifies their consequences, and proposes cloud-based solutions for improving the speed, affordability, and security of hybrid cloud infrastructure.

Cloud migration has proven to be an effective strategy for reducing infrastructure costs, improving the availability of data and applications, and increasing operational agility. However, this migration rarely happens in one fell swoop. Many large organizations find themselves with a complex, heterogenous mixture of cloud and on-premise infrastructure:



Such hybrid infrastructure is not necessarily a bad thing, but it does bring complications. Specifically, it creates situations where various networking functions—such as DDoS mitigation, load balancing, firewall, and VPN—remain on-premise.

These network hardware appliances just aren't up to the task of securing and accelerating critical infrastructure in a cloud-focused world. They've always been a hassle—an expensive, often unruly mess of equipment strung together with spider webs of cables. But add the cloud to the picture, and security gaps quickly emerge—along with performance penalties and additional support challenges.

This paper outlines those risks of maintaining network hardware in a world moving to cloud-based services and offers strategies for avoiding them.

Network hardware appliances span a variety of specific functions, and are used somewhat differently from organization to organization. Common examples include:

Security	Performance and Reliability
DDoS mitigation	Load balancing
Firewall	Traffic acceleration/ WAN Optimization
Virtual private network	

When this hardware is deployed on-premise, the resulting architecture generally suffers from four categories of risks: **capacity limitations**, **high total cost of ownership**, **support challenges**, and **security gaps**. The first two categories have always posed problems to some degree. The other two are exacerbated by cloud migration.

Capacity limitations

It should be no surprise that by their very nature, network hardware appliances can become overburdened during unexpected traffic surges — whether that traffic is legitimate or not. But several recent trends mean reaching those limits is a more common concern.

Consider DDoS mitigation. The largest DDoS attack in history is claimed to have reached a maximum volume of 2.3 Tbps in February of 2020 ([source](#)) — before that, other attacks in the past two years have reached 1.7 and 1.3 Tbps. All of these attacks would overburden many times over the most advanced DDoS mitigation hardware boxes on the market, which typically provide a fraction of the capacity required to mitigate such attacks.

The largest DDoS attack in history is claimed to have reached a maximum volume of 2.3 Tbps in February of 2020

Not all organizations will attract attacks of such scale — but not all organizations can or do implement the most advanced DDoS mitigation hardware, either. A Cloudflare study found that roughly 2.4% of network layer DDoS attacks in Q2 2020 had maximum volumes of over 100 Gbps ([source](#)), which would overburden many purportedly high-capacity hardware-based mitigation solutions.

Furthermore, attack volume does not take into account legitimate traffic that might reach your data center at the same time. Should a smaller attack arrive during a high-traffic period — such as the Black Friday shopping weekend, when ecommerce daily pageviews double overnight, on average ([source](#)) — the resulting traffic surge still might be enough to push security hardware past its breaking point.

DDoS mitigation is just one example of on-premise hardware's capacity limitations. Other examples include:

- **Load balancers:** Individual on-premise load balancers can easily be overburdened by sudden spikes in legitimate traffic. When this happens, it can take a long time to provision and install additional hardware. The alternative is maintaining enough capacity for the worst-case scenario, but this approach requires the organization to continually run a lot of hardware at a high cost.
- **Virtual Private Networks (VPNs):** VPN usage has become much harder to predict in advance. A May 2020 survey of US-based employers found that 53% of full-time employees are working from home as a result of the Covid-19 pandemic — a 7x increase over 2019 ([source](#)) — with 22% expecting to remain at home after the pandemic. Should these levels of remote access exceed an on-premise VPN's capacity, employees will suffer from login difficulties, latency, and ultimately decreased productivity.

When faced with these problems, one response is to buy more, newer, higher-capacity hardware. But such an approach introduces a host of other problems.

Costs of ownership

Like capacity limitations, it should come as no surprise that data center hardware is expensive. For example, the hardware required to attain approximately 100 Gbps of DDoS mitigation capacity might cost between \$400,000 and \$500,000 up front.

What's more, these costs are just one part of a hardware appliances' total cost of ownership. Consider the following expenses:

- **Team costs:** Purchasing, operating, and maintaining hardware to defend against threats at every layer of the OSI model — and to provide the level of performance and reliability expected from modern websites and Internet applications — requires team members who are experts in every one of those networking functions. Building a team with this breadth and depth of expertise is an expensive proposition. Skills gaps also make it hard to accomplish: a 2020 ISACA survey found that 62% of companies say their IT team is understaffed ([source](#)).
- **Maintenance costs:** A 2019 Forrester report defines data center hardware as 'aging' after three years ([source](#)), yet warranties for those entire periods often require extra expenditure. The alternative is unexpected — and thus unbudgeted — repairs from the original manufacturer or a third party. Hardware malfunctions can also cause data center downtime, which has an average opportunity cost of over \$8,800 per minute ([source](#)).
- **Replacement costs:** Replacing a hardware appliance every three years requires organizations not only to repay their initial investment, but to dedicate resources to shipping and installing new hardware. Delaying these replacements often results in more frequent malfunctions — and thus additional maintenance costs.

Contrast this model with cloud-delivered networking services. They are possible to operate with a nimbler team, do not impose maintenance and shipping costs, and do not force organizations to choose between costly upgrades and an increase in malfunctions.

Hardware malfunctions can cause data center downtime, which has an average opportunity cost of over \$8,800 per minute

Support challenges

Supporting network hardware appliances is not just an expensive proposition, but also a logistical challenge. Hardware requires frequent patching in order to keep up with the latest vulnerabilities and attack tactics — a process that often relies on manual implementation, and is thus susceptible to human error.

The more hardware appliances an organization uses, the higher the chances it will eventually neglect a patch due to inattention or concerns about affecting vital systems. In June of 2020, enough companies failed to promptly install a VPN appliance patch that US Cyber Command, a division of the Department of Defense, had to send out a warning that the vulnerability could be exploited by hostile nation-states ([source](#)). In fact, patching hardware can be so complex that an entire category of software exists to help companies keep up to date ([source](#)).

In June of 2020, enough companies failed to promptly install a VPN appliance patch that US Cyber Command, a division of the Department of Defense, had to send out a warning that the vulnerability could be exploited by hostile nation-states

And the consequences of just one missed patch can be significant. Not only will the hardware remain vulnerable, but once a patch is released, the corresponding vulnerability becomes a higher-profile target for opportunistic attackers. Contrast this situation with cloud-based security services, in which fixing vulnerabilities and installing updates happens automatically by default, and can take as little as thirty seconds to propagate depending on the cloud provider's network speed.

Other maintenance challenges with hardware include:

- **Troubleshooting:** In a hardware-only scenario, troubleshooting often forces IT teams to go through the arduous process of unplugging load balancers, firewalls, and other on-premise appliances one at a time to discover where the problem lies.

This process is further complicated by the use of cloud services. Hardware-reliant organizations often manage access to those services through the centralized data center and all of its individual appliances. When employees are unable to access a particular service, IT teams have an extra place to check in order to diagnose the issues. This extra effort adds up quickly in large organizations — a 2019 survey found that companies with 1000+ employees subscribe to over 200 SaaS applications on average ([source](#)).

- **Physical maintenance:** When a hardware appliance does break, IT teams must physically unplug it, order a replacement, test the replacement, and reinstall it — another arduous process.

Security gaps

Even if an organization had the resources required to continually provision and maintain the latest, highest-capacity on-premise hardware, the resulting infrastructure would still suffer from critical security deficiencies — especially in a world trending towards the cloud.

Consider employee access management. While VPN hardware can establish encrypted tunnels between remote employee devices and applications hosted in an internal data center, it cannot monitor and secure user activity after establishing this tunnel.

Should the employee's device become compromised by malware, or should a phishing attack compromise their VPN credentials, an attacker might be able to use that VPN access to access a wide variety of sensitive information. Both phishing and malware continue to pose serious risks — in April 2020, Google reported that they blocked 18 million COVID-19-related malware and phishing emails every day ([source](#)). (Cloud-based VPNs also suffer from this problem, but the fact remains that on-premise hardware alone cannot provide truly secure remote access to internal applications.)

Should the employee's device become compromised by malware, or should a phishing attack compromise their VPN credentials, an attacker might be able to use that VPN access to access a wide variety of sensitive information.

Cloud services and SaaS applications further complicate security for hardware-centric infrastructure. In a hybrid cloud model, for example, an organization operates a mixture of on-premise and cloud infrastructure. The organization cannot simply send security hardware to a cloud provider. If it wishes to continue using on-premise hardware for its own data center, different parts of its infrastructure will be protected in different ways, giving security teams less visibility into and control over incoming attacks.

Cloud-based services can overcome both of these challenges by unifying data centers and cloud services under a single software-defined layer. A detailed explanation of this approach is beyond the scope of this paper — to learn more, explore the following articles:

- [What is a Zero Trust network?](#)
- [What is Secure Access Service Edge?](#)

Cloud-based security and performance services: Advantages and challenges

Delivering network services through the cloud avoids many of the problems associated with hardware: capacity limitations, costs, support challenges, and security gaps.

- **Capacity:** Due to the cloud's distributed nature and software-defined nature, organizations can provision additional capacity easily as their business scales .
- **Cost:** The add-on costs of hardware are either nonexistent or easier to plan for in advance. What's more, cloud services are typically classified as operating expenditures, not capital expenditures, which offers tax and accounting benefits for many businesses.
- **Support:** Logistical and resource needs are handled by the service provider. In addition, there is no chance of missing a patch, since updates occur automatically.
- **Security:** Software-defined networking services can unify different infrastructure under a single protective layer.

However, cloud networking services present their own risks if not deployed thoughtfully:

Risk	Description
Latency	Some cloud-based network functions rely on specialized cloud-based data centers — e.g. scrubbing centers for DDoS mitigation. Backhauling traffic to those data centers can add significant latency depending on its location relative to the destination server. This problem compounds when an organization uses different providers for different networking functions. When traffic must hop from provider to provider, latency can be measured in hundreds of milliseconds.
Support	When an organization uses different providers for different functions, troubleshooting remains an issue. It can be hard to tell which provider is the cause of congestion or outages.
Costs	When an organization uses different providers for different functions, the time (and thus the money) required to manage them can still be high.

To avoid these problems, consider the following strategies:

- **Look for providers that work with both cloud and on-premise infrastructure.** This capability allows IT and security teams to set consistent controls and monitor global traffic from a single place.
- **Look for cloud providers offering multiple networking functions that work together.** This often reduces the number of network hops traffic must make, resulting in reduced latency—and thus faster application performance—for end users. Also, when you troubleshoot network problems, you have one vendor to call, not many. Finally, bundling multiple functions together often results in lower costs.
- **Look for cloud providers that can perform multiple networking functions from every location in their network.** Providers that expand their service portfolios by acquisition do not always integrate those new services fully, which means certain functions can only be delivered through certain data centers. Consider providers who offer these functions across the entirety of their network to avoid the same problems listed above.
- **Look for cloud providers with a broad global presence.** This capability supports the previous one, ensuring end users are always close to the network no matter where they are. It also creates a large network surface with which to absorb DDoS traffic and conduct other networking functions that require a large capacity.

How Cloudflare can help

Cloudflare has built a global cloud platform that delivers a broad range of services — making organizations more secure, enhancing the performance of their applications, and eliminating the cost and complexity of managing individual network hardware. This platform serves as a scalable, easy-to-use, unified control plane to deliver security, performance, and reliability across on-premise, hybrid, cloud, and software-as-a-service (SaaS) applications.

Crucially, every data center in Cloudflare's 200+ city global networking can deliver every one of these services, reducing the latency that can complicate cloud implementations. To learn more, visit www.cloudflare.com.