LOADBALANCER

WHITE PAPER

The top three networking trends for 2023

and the inescapable implications for load balancing and Application Delivery Controllers (ADCs)

> *Current trends in cloud computing, the Internet of Things, and digital immunity are amplifying many of the most widely experienced networking challenges. Application Delivery Controllers (ADCs) won’t solve all these problems, but they will help to relieve many of the pressures likely to be faced by network managers, IT directors and CTOs in 2023.*”

Executive Summary

Every year, new technology trends emerge that increase the complexity of network management and create new considerations for network design. 2023 will be no exception. Current trends in cloud computing, the Internet of Things, and digital immunity are amplifying many of the most widely experienced networking challenges.

Three trends, in particular, are likely to exacerbate network management in 2023:

1. **Cloud migrations are proceeding more cautiously than anticipated,** with many organizations now operating a combination of public cloud, private cloud and on-premise data centers. These hybrid networks are complicated to manage and will demand more flexible approaches to reduce risks and costs.

2. **Resilience is becoming the top priority in network design.** With the inevitability of security incidents, bugs, product failures and human error, there is growing recognition that networks need to be designed with digital immunity in mind to ensure greater application resilience and business continuity.

3. **The Internet of Things is expanding rapidly,** and more and more organizations are planning to harness it over the next five years. For network managers, this increases the pressure to deliver networks that offer greater scalability and 100% service availability.

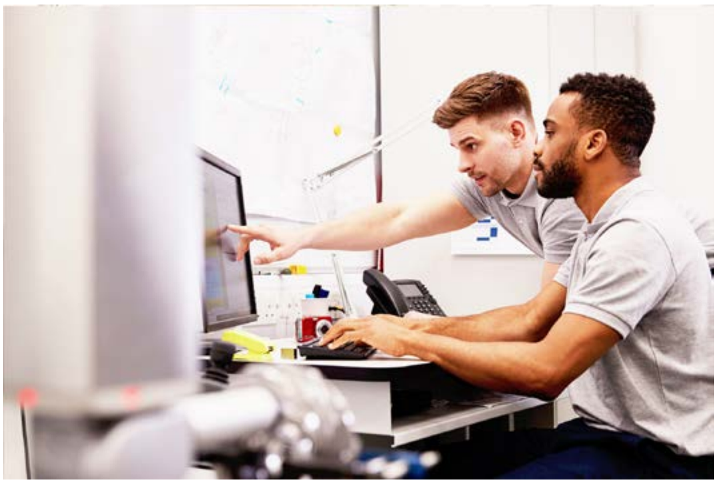

These three networking trends have inescapable implications for the use of load balancing technology within corporate networks, regardless of whether these networks are on-premise, in the cloud or in hybrid environments. Application Delivery Controllers (ADCs), which can be physical devices or virtualized appliances, distribute traffic across multiple servers in datacenters and the cloud, preventing a failure in a single server from causing application downtime, or a degradation in application performance for users.

Load balancers are often already in use across almost all corporate networks, but organizations can use them in a more agile, strategic and cost effective way to help them alleviate some of the challenges that arise from current technology trends. Application Delivery Controllers (ADCs) won't solve all of these problems, but they can help to relieve many of the pressures likely to be faced by network managers, IT directors and CTOs in 2023.

1. Cloud migrations are proceeding more cautiously

Cloud-first as a strategy promised many things, not least lower costs, greater flexibility, and simplified network management. In reality, however, the cloud has sometimes failed to live up to the hype. Consequently, in 2023, organizations will increasingly be adopting a more cautious approach to cloud migrations and, instead, combining cloud-based and on-premise infrastructure in hybrid networks, that are inevitably more complicated to manage.

The networking trend

The growing trend of operating hybrid networks is being driven, in particular, by the realization that cloud strategies are not always as cost-effective as anticipated. The 2022 State of the Cloud Report, published by Flexera, reveals that organizations are facing public cloud bills that are, on average, 13% higher than budgeted, and these organizations expect their cloud spending to increase by 29% in the twelve months ahead. [1] To avoid this 'bill shock' in 2023, many companies are now delaying planned cloud migrations and instead retaining their on-premise network infrastructure for longer than anticipated.

Other organizations are even moving some applications from the cloud back to on-premise equipment. One survey of 350 enterprises, conducted in 2019, found that 74% of companies had moved at least one application back to their on-premise networks from the cloud, for reasons including poor performance in the cloud, costs and security concerns. [2] Although cloud platforms have matured over the last three years, there are still many well-publicized cloud migration failures, and Gartner stated very recently that three out of four companies do not have a 'fit for purpose' cloud strategy. [3] For most organizations, therefore, continuing to maintain on-premise networks is as much a necessity as a safety net.

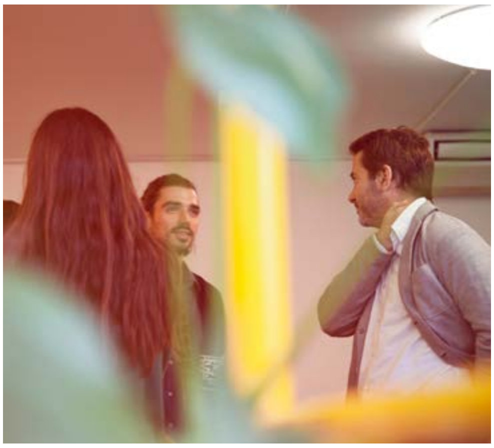

Those organizations that are migrating applications and data to the cloud are understandably taking measures to reduce risk. Many are wary of entrusting all their data and applications to a single cloud provider, so they are opting for multiple public and private clouds instead. An independent research report published by Virtana in May 2022 revealed that 82% of organizations plan to grow their multi-cloud environments. [4] Multi-cloud strategies are prudent because although IT outages in the large cloud platforms are rare, they do happen. Also, when organizations become 'locked in' to a single cloud provider, they are more exposed to cost increases. However, the more cloud platforms an organization has, the more complex network management becomes.

[1] Trends in Cloud Computing: 2022 State of the Cloud Report. Flexera (report), 2022 [cited 2022 Dec 6]. Available from: https://www.flexera.com/blog/cloud/cloud-computing-trends-2022-state-of-the-cloud-report/

[2] The Bi-Directional Cloud Highway: User Attitudes about Securing Hybrid- and Multi-Cloud Environments. IHS Markit (whitepaper), 2019 [cited 2022 Dec 7]. Available from: https://www.fortinet.com/content/dam/fortinet/assets/analyst-reports/ar-2019-ihsm-fortinet-wp-q2.pdf

[3] Cloud Strategy Cookbook. Gartner (infographic), 2022 [cited 2022, Dec 7]. Available from: https://www.gartner.com/en/doc/748733-infographic-cloud-strategy-cookbook

[4] The State of Multi-cloud Management 2022. Virtana (research report), 2022 [cited 2022 Dec 6]. Available from: https://www.virtana.com/the-state-of-multi-cloud-management-2022/

The implications for load balancing

The current trend for hybrid networks enables organizations to adopt a more cautious approach to migrating to the cloud and reduces the associated risks. However, it can significantly heighten the complexity of managing core networking devices and technologies, including load balancers. An organization with two public cloud services, a private cloud and a legacy data center could, for example, have four or more different load balancing solutions to manage and maintain.

However, a more strategic use of Application Delivery Controllers (ADCs) can:

Simplify network management

Network complexity can be eased by using a single load balancer technology for all on-premise and cloud deployments. An organization with its own datacenter and two cloud providers would not have to train its IT team in the use of three different load balancing technologies, making it easier for staff to manage the entire network environment. It would also be easier to recruit new staff and get them up-to-speed as there would be a common load balancing framework across local and cloud-based systems.

> *Network complexity can be eased by using a single load balancing technology for all on-premise and cloud deployments.*”

Improve control in phased cloud migrations

Recognizing that not all cloud migrations go to plan, a single load balancing platform for on-premise and cloud will also give network managers the ability to implement a phased and carefully controlled migration. Using the global server load balancer (GSLB) feature in the load balancers, they can start by sending 10% of live traffic to AWS, for example, and then monitor how things are working. If no problems are encountered, the network manager can then gradually increase the proportion of traffic going to AWS until 100% of traffic is handled in the cloud. If any issues are experienced, having a single load balancing software provider will allow all traffic to be instantly routed back to the on-premise datacenter while glitches are ironed out, maintaining performance for users.

Reduce risk from cloud lock-in

With a single load balancing platform, network managers will find it easier to move applications to the cloud, and then back again to on-premise networks or over to alternative cloud networks, whenever needed. This gives them the agility they need to reduce risks e.g. if cloud costs rise, cloud performance is inadequate, or business needs change.

2. Resilience is becoming the top priority in network design

Focus in recent years has been on IT security and cyber security measures that help to prevent an attack at the network or application layer. Increasingly though, organizations are realizing that they need to place equal emphasis on IT resilience measures to lessen the impact of potential security incidents, as well as bugs, equipment failures, and human errors, all of which could lead to downtime. The premise is that as we cannot anticipate and prevent all risks, we must take measures to mitigate them.

The networking trend

Network resilience has been thrown into the spotlight by research group Gartner, which has put digital immunity at the very top of its list of the top 10 strategic technology trends for 2023. [5] The firm encourages organizations to build robust digital immune systems, which it defines as "resilient systems that mitigate operational and security risks," to reduce unplanned and unwelcome downtime.

It is difficult to accurately calculate the level of application downtime currently experienced by businesses globally. This is primarily because most application failures are not reported and are only noticed by the communities of users that are trying to access them at the time. However, in a survey of 3,000 business and IT leaders, published in 2022, Veam found that as many as 40% of servers (which equates to two out of five) experienced outages in a 12 month period. [6]

One thing is certain: downtime can lead to substantial financial losses, an erosion in customer loyalty, and user frustration. When Amazon experienced downtime for 59 minutes in 2021, the estimated financial cost was 34 million. [7] According to Gartner, however, the primary motivation that will drive organizations to improve their digital immunity in 2023 is to ensure that the customer experience is not "compromised by defects, system failures, or anomalies, such as software bugs or security issues". Gartner anticipates that if organizations invest in creating a digital immune system, they could decrease downtime by up to 80% and thereby deliver a significant improvement in the customer experience. [8]

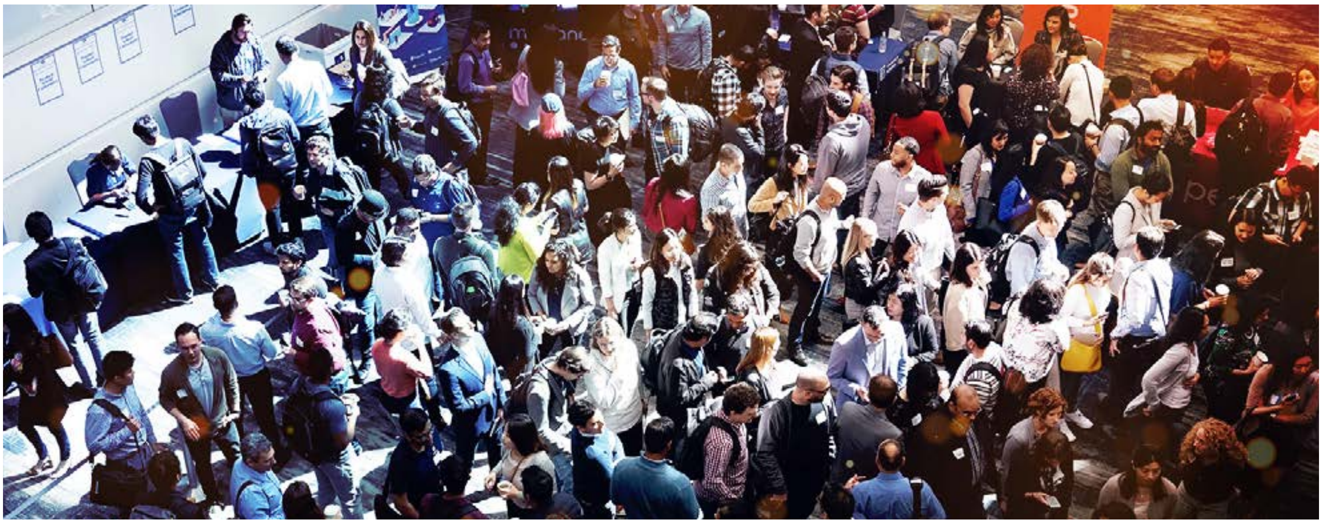

[5] Gartner Top 10 Strategic Technology Trends for 2023. Gartner (blog), 2022 [cited 2022 Nov 21]. Available from: https://www.gartner.com/en/articles/gartner-top-10-strategic-technology-trends-for-2023

[6] Real-World Statistics on Downtime and Data Loss in 2022. Veam (report), 2022 [cited 2022 Dec 7]. Available from: https://www.veeam.com/blog/data-loss-2022.html#:~:text=In%202021%2C%20organizations%20reported%20a,high%2Dpriority%20and%20normal%20applications

[7] Amazon 'missed out on $34m in sales during internet outage'. Independent (article), June 2021 [cited 2022 Dec 7]. Available from: https://www.independent.co.uk/news/business/amazon-down-internet-outage-sales-b1861737.html

[8] What Is a Digital Immune System and Why Does It Matter? Gartner (Internet), 2022 [cited 2022 Nov 21]. Available from: https://www.gartner.com/en/articles/what-is-a-digital-immune-system-and-why-does-it-matter

The implications for load balancing

Any organization that is planning to improve its network resilience—whether by reinforcing existing business continuity systems or as part of a new digital immune system strategy—will undoubtedly need to consider the use of load balancing. The primary role of an Application Delivery Controller (ADC) is, after all, to improve resilience by enabling traffic to failover to an alternative server in the event of an incident or outage, preventing downtime and avoiding disruption to the end user.

Used effectively, load balancers can:

Form the foundation of a digital immune system

Load balancers are a fundamental tool for achieving 'auto remediation', one of the six prerequisites for a strong digital immune system, according to Gartner. [viii] Used within on-premise, cloud and hybrid networks, they automatically detect and respond to issues, diverting user traffic away from failed servers to operational servers without any intervention from IT staff. A resilient network design will also use load balancers to direct traffic between datacenters and between private and public cloud platforms and on-premise infrastructure to provide uninterrupted service for users, should a failure event occur.

Improve the 'observability' of network performance

In adopting a digital immune system strategy, organizations will want to make full use of Application Delivery Controllers (ADCs) to gain visibility into what is happening in their networks. This 'observability' — another of Gartner's prerequisites for a strong digital immune system [viii] — enables organizations to be more proactive in monitoring the health of the load balancers and the services that are being balanced by them. They can see what is happening across the network, not just at the application level. Potential problems caused by emerging bottlenecks or unexpected increases in traffic can then be detected before they cause server overloads and outages, and steps can be taken to prevent any degradation in application performance.

> “*In adopting a digital immune system strategy, organizations will want to make full use of ADCs to gain visibility into what is happening in their networks.*”

Provide an additional layer of security

Many load balancers offer a web application firewall (WAF), either as a standard feature (included in the standard license fee) or as an optional extra (for an additional cost). Complementing other network security systems, WAF solutions are designed to protect the application layer and can be set up with a specific set of rules for individual applications, depending on the perceived risks and likely attack vectors for these systems. Deploying a WAF can, therefore, provide a useful additional layer of security, strengthening the resilience of on-premise and hybrid networks.

Make security updates easier to manage

Delaying critical software updates and neglecting to install patches is one of the biggest risks to security, so network managers should be constantly on the look-out for better ways to manage software updates for all their IT products and systems—and this includes load balancers. When organizations have multiple load balancers across many sites, making manual updates can mean months of effort and service interruptions, and by the time they are finally completed, they may need to be done again.

Recognizing this, some load balancing vendors are now introducing new, centralized ways of managing software updates to their load balancing suite using clustered pairs of load balancers that are non-interrupting, meaning critical updates can take place without any subsequent downtime. Centralized load balancing management portals make it much easier for network managers to manage security and feature updates for all their load balancers simultaneously, from a single window.

3. Use of the Internet of Things is expanding rapidly

Following COVID-19—and perhaps as a result of the pressures faced during COVID-19—more and more organizations are adopting the Internet of Things (IoT). In parallel, the number of connected devices in existing IoT networks is increasing steadily, creating challenges for network scalability. Many network managers will need to build new, more flexible infrastructure to achieve the high service availability that users expect, while keeping costs in check.

The networking trend

According to Statista forecasts, there will be two billion more Internet-connected devices in 2023 than in 2022, and ten billion more within five years. That makes 21.3 billion Internet-connected devices worldwide by 2027. [9] Global management consulting firm McKinsey previously estimated that 127 new devices connect to the Internet every second, [10] a figure that has undoubtedly increased in recent years and is likely to grow further in 2023.

The foremost implication of the growth in IoT is the resulting explosion in data volumes, which network managers need to be able to accommodate. The latest figures from Statista indicate that the total data volume generated by connected IoT devices worldwide will reach 79.4 zettabytes (ZBs) by 2027. [11] This staggering volume of data is in addition to the growing repositories of 'big data' that organizations already store and new data streams being used for artificial intelligence and machine learning.

Besides the volume of IoT data, the vital importance of IoT data must also be considered. Increasingly, IoT is being used for mission-critical applications that deliver crucial, real-time information to users or the general public. For these applications, high service availability is absolutely imperative, as is the need to prevent data loss. In the medical sector, where the market for IoT is expected to grow from USD 89.07 billion in 2021 to USD 446.52 billion by 2028, [12] IoT data can literally be life-saving. Remotely monitored blood pressure or diabetes readings from patients, for example, need to be received and acted upon quickly, without fear of delays or data loss that could prevent patients from receiving the care and medication they need.

[9] IoT connected devices worldwide 2019-2030. Statista (Internet) [cited 2022 Nov 21]. Available from: https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/

[10] What's new with the Internet of Things? McKinsey (Internet) [cited 2022 Dec 7]. Available from: https://www.mckinsey.com/industries/semiconductors/our-insights/whats-new-with-the-internet-of-things

[11] Data volume of internet of things (IoT) connections worldwide in 2019 and 2025. Statista (Internet) [cited 2022 Dec 7]. Available from: https://www.statista.com/statistics/1017863/worldwide-iot-connected-devices-data-size/

[12] Internet of Things [IoT] in healthcare Market Size & Trends, 2021-2028. Fortune Business Insights (Internet) [cited 2022 Dec 8]. Available from: https://www.fortunebusinessinsights.com/internet-of-things-iot-in-healthcare-market-102188

The implications for load balancing

Few organizations will attempt to capture and manage IoT data on their on-premise networks. Most will build their IoT systems from the ground up in the cloud, but if flexibility and scalability are not part of the initial IoT system design, organizations could face escalating costs, inadequate service availability, or data loss. Application Delivery Controllers (ADCs) play a central role in IoT networks and can be used in a variety of ways to ensure IoT systems achieve the anticipated benefits.

In IoT environments, load balancers can:

> *ADCs play a central role in IoT networks and can be used in a variety of ways to ensure IoT systems achieve the anticipated benefits.*”

Prevent data loss between IoT devices and the cloud

Many organizations, including hospitals in particular, need to be able to ensure that data collected from connected devices on their premises is quickly and reliably transferred to cloud-based systems. In typical usage, load balancers work as a reverse proxy, channeling traffic from outside a corporate network to web servers and application servers on the inside (whether these servers are held in a public or private cloud, or on-premise data center). However, in the case of IoT, load balancers can also work as a forward proxy and be used to make sure that data collected by internet-connected devices on-site is sent reliably to cloud-based applications and storage systems without data loss.

Provide greater control over IoT cloud costs

Accurately estimating and then managing IoT cloud spend will be one of the biggest challenges faced by companies adopting or expanding their use of IoT in 2023. Part of the reason for this cost control issue is the way that cloud-native load balancers are billed. Organizations that use AWS load balancers in the AWS cloud, for example, are charged not only on metered uptime but also on the number of new and active connections established, as well as bytes and rules processed. As data volumes increase, costs can soar. By using a platform-agnostic Application Delivery Controller (ADC) organizations have greater flexibility to switch their IoT traffic easily from one public cloud platform to another public cloud, data center, or private cloud to access more competitive rates.

Optimize the usage of an IoT network with dynamic load balancing

Load balancers do not choose how they balance egress (outbound) traffic when multiple ISPs are in the mix — they simply route a response back over their default gateway, and this can cause problems with BCP38 and ingress filtering. These issues can be overcome with Policy Based Routing (PBR), which will ensure that responses are sent through the correct gateway based on the interface and the source IP address that is responding to the request, thus overcoming the problems of asymmetric routing. In this instance, the load balancer balances inbound traffic across the configured backends whilst listening on IP addresses from multiple ISPs and sending return traffic through PBR across the correct return path. If GSLB (Global Server Load Balancing) is being utilized to serve IP addresses, then the load balancer can stop returning server A's IP addresses if there is an issue with ISP-A thus allowing traffic to fall over safely to ISP-B, maintaining IoT service availability and de-risking IoT architecture.

Want to know more?

If you would like more information about how to optimize the use of load balancers and address the challenges that you may face in 2023, get in touch.

LOADBALANCER

About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.

Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

Follow us: @loadbalancer.org