CLOUDFLARE®

How to Build Infrastructure to Deliver a Superior Online Experience

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

Index

Introduction

Providing a superior online experience for a global customer base is no longer optional. As demand increases for web-based services and applications, businesses must satisfy customer needs while ensuring that their websites and applications remain as secure, fast, and reliable as possible.

With this shift, enterprises face new challenges and opportunities for growth — from anticipating and meeting customers' digital needs to mounting a strong defense against web-based attacks, overcoming latency issues, preventing site outages, and maintaining network connectivity and performance.

Building a superior online experience doesn't just require a single tool or product suite, but the integration of a comprehensive security posture and performance features designed to cut down on latency and improve network reliability — as outlined in the diagram below:

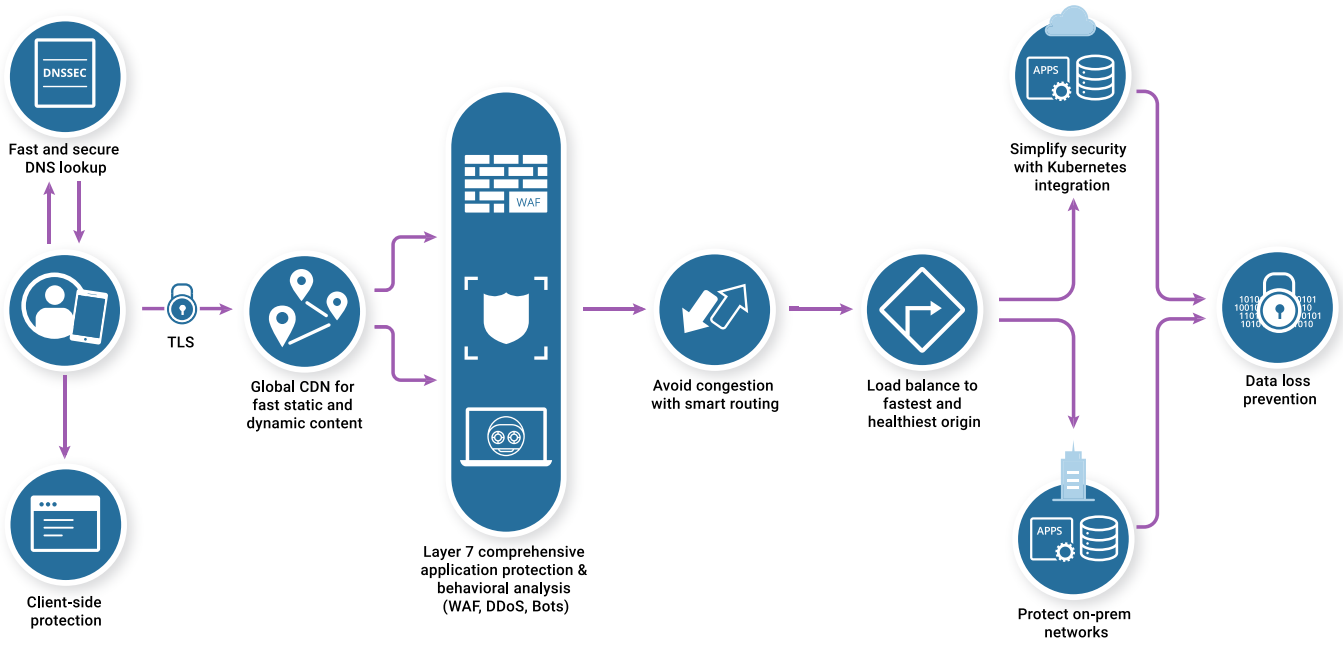

Here are five key steps modern enterprises need to take in order to meet customer needs and provide a secure and seamless user experience.

.

STEP 1

Ensure Secure, Fast, and Reliable Customer Connections

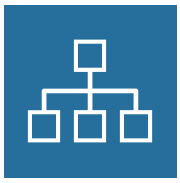

DNS

DNS is an essential component of every Internet-based business, yet it is often overlooked until something breaks. As DNS attacks become more prevalent, businesses are starting to realize that the lack of a resilient DNS creates a weak link in their overall security strategy. The millions of dollars spent on building and securing web properties are of no value if their applications are unavailable and their customers can't find them.

DNS challenges

High latency: Businesses may face web performance problems when their webpages frequently load assets from more than one domain, increasing the time required to resolve each requested domain.

In-house DNS infrastructure: Self-hosted DNS is costly to maintain, may add latency due to slower DNS resolution for a globally-distributed customer base, and is not fully protected against sophisticated DNS attacks.

Small-network DNS providers: When selecting a DNS solution, businesses often make the mistake of choosing a provider that does not have a large network and does not perform DNS resolution at all data centers. This can restrict performance and reliability, particularly for companies that need to reach customers across various regions of the globe.

What to look for in a DNS provider

Integrated security solutions: Because the DNS threat landscape is so diverse, effectively mitigating DNS attacks requires an integrated security strategy that includes DNSSEC, DDoS attack mitigation, and a DNS firewall. For large enterprises that prefer to maintain their own DNS infrastructure, a DNS firewall can be implemented in conjunction with a secondary DNS. This setup adds a security layer to the on-premise DNS infrastructure and helps ensure overall DNS redundancy.

Fast DNS resolution: For businesses considering cloud-based managed DNS providers, it is essential to select a provider that can maximize performance and availability with fast DNS resolution and geo-based or dynamic routing.

Redundancy: Businesses that choose to host their DNS records with a single provider are more vulnerable to outages since they depend on a single point of failure. In order to maximize resiliency, businesses need to not only enlist the help of multiple separate managed DNS providers, but also ensure that those providers do not share the same nameserver facilities.

Client-side security

Today, up to 70%[1] of the code that executes and renders on users' browsers comes from external, unmonitored JavaScript integrations, opening up new and expansive avenues for client-side attacks like Magecart, cross-site scripting (XSS), credit card skimming, website defacement, and worse.

Server-side security tools have limited, if any, visibility into client-side threats and no way of preventing attacks or patching these vulnerabilities. For companies with a web presence, it is vital that they deploy and maintain dedicated client-side protection in order to secure their websites against these common and rapidly-evolving threats.

Client-side attacks

Cross-site scripting (XSS): XSS attacks occur when an attacker attaches, or inserts, malicious code onto a legitimate website, often with the purpose of stealing user login credentials, accessing other sensitive information, or taking control of a user's browser.

Magecart attacks: Magecart attacks fall under the umbrella of 'data skimming,' in which attackers insert malicious code into websites and scrape confidential user data (e.g. credit card numbers, passwords, etc.) from online payment forms. This kind of attack may be more difficult for businesses to detect, since attackers can disguise malicious code within harmless code or encode stolen data so that it can be returned to the attacker undetected.

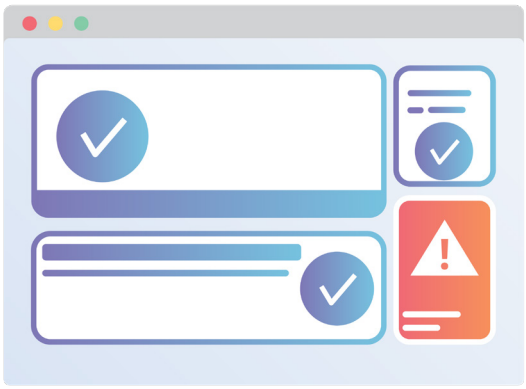

Spoofing: Spoofing, or disguising malicious communication by impersonating a trusted source, allows attackers to steal sensitive user data, reroute traffic to cause a DDoS attack, or gain unauthorized access to an organization's system or network.

[1] Bermingham, Mark. "Redefining Client-Side Security with the Tala Security Certified Module for NGINX Plus," NGINX, https://www.nginx.com/blog/redefining-client-side-security-tala-certified-module-nginx-plus/

What to look for in a client-side security solution

End-to-end protection: Rather than homing in on any one kind of client-side threat, businesses need to protect backend infrastructure as well as frontend processes.

Minimal impact on performance: While deploying and managing stringent security protocols is a top-of-mind concern for many businesses, it is crucial that those security products do not tamper with web performance, as sluggish websites can put off potential customers and result in increased bounce rates and decreased conversions.

TLS

Businesses that store and transmit sensitive data must find ways of protecting that data from leaking, misuse, and theft. TLS network protocols — also referred to as 'SSL' — help accomplish this by encrypting communications over public networks and authenticating trusted parties. These protocols are designed to uphold customer privacy and shield data from third-party monitoring and tampering.

TLS challenges

Managing SSL/TLS certificates: Although SSL/TLS certificates are designed to verify the identities of trusted parties, they may become compromised and manipulated by bad actors. Since these certificates can typically be purchased by anyone, an attacker can use the certificate to impersonate a trusted party, bypass security procedures, and gain access to confidential data.

Staying up-to-date: Each subsequent version of SSL/TLS attempts to patch known vulnerabilities and help strengthen websites' defenses against attacks. By using outdated versions of these encryption protocols — such as TLS 1.1 or 1.2 — businesses may inadvertently allow malicious parties to target existing vulnerabilities and carry out attacks. Today, only 22% of the Alexa Top 1,000 websites use the latest version of TLS.[2]

[2] Holz, Ralph, Amann, Johanna, Razaghpanah, Abbas, and Vallina-Rodriguez, Narseo. "The Era of TLS 1.3: Measuring Deployment and Use with Active and Passive Methods," The University of Sydney, https://arxiv.org/pdf/1907.12762.pdf

Ensuring compliance with regulations and standards: Many businesses are held to data privacy standards by the General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA), both of which establish guidelines to help protect customer data against theft or misuse. Businesses that neglect to properly encrypt customer data traffic may end up overlooking potential threats like malware and data exfiltration.

What to look for in a TLS solution

Ease of implementation: For businesses that need to manually configure SSL/TLS protocols, an accidental misconfiguration can prevent customers from accessing those businesses' websites. Select a provider that allows for streamlined implementation of SSL/TLS, enables easy testing and rollback (should something go wrong), and automatically keeps TLS protocols up to date in order to patch the latest known vulnerabilities.

Flexibility: As businesses' security needs evolve, they may need an SSL/TLS provider that offers multiple types of certificate configurations — from those issued by a certificate authority to self-signed certificates.

Compliance fulfillment: Comprehensive SSL/TLS inspection allows businesses to identify potential threats that may be disguised in encrypted data traffic. By quickly locating and combating these threats, they can protect customer data from tampering and theft while ensuring compliance with regulations like the GDPR and CCPA.

STEP 2

Accelerate User Experience

Global CDN

For global businesses looking for easy and effective ways to extend their reach, CDNs are a logical alternative to investing in remote hosting facilities. They allow businesses to reach a vast, distributed consumer base without having to build out an expansive global infrastructure, which can be costly and difficult to manage and maintain.

CDN challenges

Security and performance tradeoffs: While a CDN should be resilient in the face of network issues and protect sensitive data from theft and loss, integrating security features can lead to performance tradeoffs as additional latency is introduced.

Lack of real-time analytics: Real-time analytics can either be difficult to obtain or are significantly delayed, which affects a business's ability to gain insight into the performance of their web properties.

Monolithic architecture: Legacy CDNs were built with a monolithic architecture, and configuration changes often take a while to populate across their network. These kinds of delays can frustrate developers, who need to quickly implement and iterate on these policy changes.

What to look for in a CDN provider

Performance gains: When evaluating a CDN, businesses should first run tests to determine that performance in key markets is comparable. Every CDN provider offers a mix of results — some better in other regions. Following those initial tests, businesses should assess whether any degraded performance is acceptable or negligible considering the gains in other regions.

Real-time analytics: How easy is it to collect analytics and data that provides information on usage and the user experience of your web domains around the world. Also, what is the delay in sourcing that information.

Developer-friendliness: Developer and engineering teams have a seat at the table in deciding the role that a CDN plays in the construction of their digital infrastructure. Businesses should look for a provider that offers seamless API integration and doesn't require outside vendors' professional services for custom configurations.

Faster routing

Businesses with global customer bases need to ensure the performance of their web applications. There are a number of factors that contribute to latency, from cumbersome security protocols to less-than-optimal network conditions for users around the globe.

Smart routing can help alleviate these issues by selecting routes based on network congestion and reliability. It runs on top of BGP and tests available paths to determine which one provides superior performance. By avoiding unreliable network connections, it can also avoid dropped packets, which typically result in slow service and network disruption.

Routing challenges

Network congestion: Network congestion can be limited to a certain geographic area that lacks sufficient infrastructure, or it can affect an entire ISP's network.It increases as more users flock to the Internet for communications and services during high-traffic events, like national holidays or global catastrophes.

Increased network latency: Higher network latency directly impacts the speed and performance of a website, which can in turn degrade the overall user experience and lead to fewer conversions and decreased revenue.

Cost-prohibitive solutions: Private leased lines and MPLS (multiprotocol label switching) networks may be purchased from large service providers to get exclusive access routes for business communications. While there are some advantages to leasing private lines, however, they are often increasingly cost-prohibitive to implement and may impact profitability and budget resources for other business-critical services.

How to select tools for optimal network path selection

Asset performance improvements: With the right provider, a business should be able to deliver web traffic over the fastest links available, resulting in a measurable increase in asset speed and performance and improved end-user experience.

Real-time metrics and analytics: Businesses should be able to assess any potential routing issues with real-time metrics and analytics, allowing them to avoid congestion, optimize their routes, measure global performance improvements, and deliver increased uptime — regardless of user location, device, and current network conditions.

Network monitoring tools: In order to reduce latency and circumvent network congestion, businesses may implement network monitoring tools to identify potential congestion points and help prioritize traffic, ensuring that the most critical workloads reach their intended destinations and that no single application becomes a bandwidth hog.

Mobile optimization

Like those browsing websites on desktop computers, mobile users expect page load times to be under 3 seconds on their mobile devices, making mobile optimization a crucial component of any successful web-based business. If businesses can't meet user expectations with a fast and seamless mobile experience, their customers will go elsewhere.

Mobile optimization challenges

Poor image optimization: Images that are not properly sized and formatted for small mobile screens can degrade the user experience, as they distort page views and make webpages more difficult to navigate. The larger the file size of an image is, the longer it takes to download. As a result, users may see increased page load times, since many devices don't have good enough screen resolution or a large enough screen to make very high resolution images necessary.

Increased bandwidth utilization: Large file sizes consume more bandwidth, leading to higher data fees from hosting providers. Businesses that fail to implement a proper image optimization solution for their mobile properties may encounter numerous risks, from increasing bandwidth costs to loss of customer interest, conversions, and revenue.

Cumbersome internal development: Web developers are often tasked to support the optimization of assets for mobile devices. This may involve the creation of manual processes to replicate images for a variety of device types. This is a non-ideal solution, as it becomes one more system or process that needs to be managed and maintained over time. The alternative — implementing a dedicated third-party solution for image optimization — is often costly and may lead to diminishing ROI in terms of feature utilization.

What to look for in a mobile optimization solution

CDN integration: Mobile optimization should seamlessly integrate with existing CDN services, allowing businesses to take advantage of any cache benefits and reduce dependencies on that service for redundant image files. The right CDN is one that helps improve mobile performance by detecting user device requirements, "virtualizing" images for smaller file sizes, and caching images as close to mobile users as possible.

In-house vs. third-party maintenance: When considering mobile optimization solutions, carefully consider what kind of long-term maintenance and updates are required, compared with solutions that can be built and managed in-house. Businesses should also consider whether these optimization solutions can be extended to support third-party storage locations outside of their domains.

STEP 3

Strengthen Your Security Posture for Your Infrastructure

Web application firewall

Even for businesses that actively strive to safeguard their infrastructure and data, it can be extremely difficult to operationally implement security efforts — especially in a world where every security gap creates an opportunity for attack.

With a WAF in place, businesses can protect against zero-day attacks and shield their applications against common threats like cross-site request forgery (CSRF), cross-site scripting (XSS), and SQL injection attacks. A WAF also enables businesses to maintain granular control over their security policies by setting rules that can protect vulnerabilities in their applications and mount a defense against emerging threats.

WAF challenges

Resource-intensive onboarding and management: Up-to-date infrastructure patching is a key component of developing good cyber hygiene and protecting business-critical applications. However, even the largest security teams are usually unable to patch their entire infrastructure due to the sheer number of patch fixes and significant velocity of patch releases by numerous vendors. While WAF solutions help alleviate this problem, they can often be time- and resource-intensive to onboard and manage, as many of today's WAFs require a number of highly-skilled security professionals to handle this process.

Lack of flexibility: Appliance-based hardware WAFs represent an extremely outdated security solution in today's threat landscape. Since applications and data mostly reside in a hybrid environment — both within on-premise infrastructure and in the cloud — a cloud-based WAF has become a vital component of any business's layered defense strategy. Unlike hardware WAF appliances, cloud-based WAFs can provide protection against vulnerability-based attacks, no matter where the application and infrastructure are hosted.

Agility: In a world where protecting assets and data is a continuous race between security teams and bad actors, agility is key. Hardware-based WAFs fail at providing an agile mechanism to create rules and propagate them quickly across all infrastructure. The time between a vulnerability being announced to deploying a patch to protect against the attack attempting to exploit that vulnerability is critical.

What to look for in a WAF solution

Ease of use: When it comes to selecting, onboarding, and managing a WAF, usability is crucial. Onboarding a WAF should not take weeks or months, and managing it should not require an army of professionals. Additionally, businesses may want to consider a WAF provider that offers seamless API integration.

Real-time threat intelligence: One of the key shortcomings of a hardware-based WAF is that it lacks real-time context for threats and attacks. Businesses may be able to integrate threat intelligence feeds with a hardware-based WAF, but this will only give them a reactive solution — not a proactive one. With an attack landscape that is rapidly evolving, real-time threat intelligence context is critical for businesses to stay up-to-date on the latest threats. WAFs should include real-time context for global and diverse threats, paying attention not only to the size of the threat intelligence data set, but the diversity of the data.

Comprehensive coverage: While attackers often try to exploit common vulnerabilities, including the OWASP Top 10, they are increasingly interested in critical and zero-day vulnerabilities. A comprehensive WAF solution should include managed rulesets that are regularly updated to automatically thwart attacks attempting to exploit zero-day and other critical vulnerabilities.

Bot mitigation

Malicious bots can wreak havoc on web-based businesses, not only compromising sensitive data and disrupting the overall customer experience, but directly impacting a business's operating costs as well. And, as bot attacks are becoming increasingly sophisticated, it can be harder to distinguish genuine user activity from automated bot activity — exposing businesses to greater risks than ever before. Sites may become compromised when targeted by malicious bot activity, which can overwhelm web servers, skew analytics, prevent users from accessing webpages, steal user data, distribute spam, damage brand integrity, and impact customer retention and revenue.

Not all bots are harmful, however. By implementing a bot management solution, businesses can distinguish between useful and harmful bot activity and prevent malicious behavior from impacting user experience.

Bot challenges

High infrastructure costs: Any web traffic represents a tangible cost to a business, since they need to host the content, deploy the servers, and pay for the storage and compute. Unfortunately, those costs rise when web properties are targeted by malicious bot activity. While good bots are essential to a business for SEO, customer support, and other useful tasks, bad bots can cause excessive bandwidth charges by scraping content and disrupting services.

Poor user experience: Customers acutely feel the effects of bad bots on a business. They may get locked out of their accounts, charged with fraudulent transactions, or simply be unable to access a business's website at all. Servers that have been overloaded by bot activity will be unable to deliver fast page loads to legitimate users, leading to increased abandonment rates in virtual shopping carts, higher bounce rates, decreased conversions, and overall loss of customer engagement, retention, and revenue.

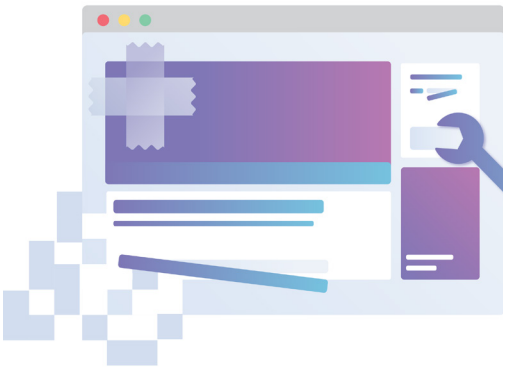

Skewed analytics: In addition to degrading user experience and causing infrastructure costs to skyrocket, malicious bots may skew analytics and paint a false picture of a business's web performance. Bad bot traffic tends to be low-quality and can negatively affect a business's aggregate analytics data (for instance, by artificially inflating page views), preventing a business from gaining valuable insights into their traffic patterns and performance metrics.

What to look for in a bot mitigation solution

Accurate detection: Before businesses can combat bad bots, they need to be able to accurately identify bot activity on their web properties. Some of the most accurate detection methods pair threat intelligence with behavioral analysis, fingerprinting, and machine learning, which can help businesses detect malicious activity without interfering with real users' activity or disrupting the user experience.

Seamless integration: Even a comprehensive bot mitigation strategy is rendered useless if it requires lengthy and involved configuration. Bot mitigation solutions should be easily and quickly integrated with any technology stack, security strategy (including DDoS attack prevention), and CDN, so that customers receive the benefits of upgraded attack protection without seeing a measurable impact on user experience.

Diverse mitigation methods: As malicious bot actors grow more complex and sophisticated from year to year, businesses need to adapt their mitigation strategies to ensure that bot activity can be detected and thwarted as soon as possible. Since no one tactic is capable of preventing every kind of bad bot behavior, it is essential to implement a wide range of detection and mitigation methods, including one or more of the following: blocking all bot traffic, whitelisting good bots, challenging suspected bots with CAPTCHAs, maintaining daily logs of all site traffic, implementing additional authentication for all users, and redirecting bots to alternative content.

DDoS attack mitigation

By consuming all available bandwidth between targeted devices and the Internet, DDoS attacks not only cause significant service disruptions, but have a tangible and negative impact on business as customers are unable to access a business's resources.

To protect web servers, a reverse proxy helps prevent attackers from being able to identify and target your servers' IP addresses. For more complex Layer 7 DDoS attacks, a web application firewall (WAF) can act as a reverse proxy to shield targeted servers from certain types of malicious traffic.

Some companies build or deploy their own reverse proxies, but this requires intensive software and engineering resources, as well as a significant investment in physical hardware. An easier and more cost-effective way to reap the benefits of a reverse proxy is to use a CDN that offers global server load balancing, allowing businesses to mitigate DDoS attacks closer to the source without impacting performance.

Of course, it's not enough to just protect web servers from DDoS attacks. Enterprises often have on-premise network infrastructure hosted in public or private data centers that need protection from these threats as well.

Legacy approaches to DDoS mitigation

Scrubbing: Scrubbing requires rerouting network traffic to centralized scrubbing servers in designated geographic locations in an attempt to filter or 'scrub' out malicious traffic from non-malicious traffic. Re-routing all traffic to a geographically distant scrubbing center can add considerable latency, which is often unacceptable for most applications.

On-premise hardware boxes: Another DDoS mitigation technique uses on-premise hardware boxes to scan traffic and filter out malicious requests. The scanning hardware also introduces network latency and inhibits performance due to the bottleneck nature of re-routing network traffic through the boxes to complete the scanning process. On-premise anti-DDoS appliances often have a bandwidth limit by default, which is based on the combination of the organization's network capacity and the box's hardware capacity.

What to look for in a DDoS mitigation provider

Mitigation capacity and time-to-mitigation: Businesses should assess their existing capacity for mitigating DDoS attacks without impacting site functionality. The traditional approach to absorbing the spikes in traffic generated by DDoS attacks has been to build out costly on-premise server farms that are easily overwhelmed by volumetric attacks. A more effective approach is to deploy a cloud-based mitigation solution that offers unlimited capacity to protect against DDoS attacks and can provision services at the network edge.

Always-on vs. on-demand protection: With on-demand mitigation services, traffic must be rerouted to a cloud mitigation service every time a potential DDoS attack is detected, and users only pay for DDoS mitigation when it is needed. Stopping a DDoS attack can take longer because traffic spikes must reach certain thresholds before analysis begins and someone manually turns on the mitigation service.

By comparison, always-on mitigation continuously routes and filters all site traffic, so only clean traffic reaches the user's servers at all times. While more expensive than on-demand services, always-on mitigation provides automatic, uninterrupted protection and leads to faster response times. For businesses facing a constant barrage of attacks, always-on mitigation may be a more cost-effective choice than on-demand protection.

Integrated security and performance: DDoS attacks cause sluggishness and outages that not only degrade performance, but also damage a business's ability to achieve sustainable growth. Businesses need to consider integrated security and performance solutions that provide a stalwart defense against DDoS attacks without negatively impacting site performance or consumer experience.

STEP 4

Ensure High Availability of Your Applications by Building a Resilient Infrastructure

Load balancing

Maximizing server resources and efficiency can be a delicate balancing act. Servers that become overloaded or are too geographically distant from end users can have a detrimental effect on business, as increased latency and server failure can result in lost revenue, broken customer trust, and brand degradation.

Cloud-based load balancers distribute requests across multiple servers in order to handle spikes in traffic. The load balancing decision takes place at the network edge, closer to the users — allowing businesses to boost response time and effectively optimize their infrastructure while minimizing the risk of server failure. Even if a single server fails, the load balancer can redirect and redistribute traffic among the remaining servers, ensuring that customers never experience significant latency or see a site outage. The load balancer also allows for active health checks, which allows businesses to identify underperforming servers and take preemptive measures before a breakdown actually occurs.

Load balancing challenges

Application downtime and latency: Even brief delays can have a noticeable effect on engagement and conversion rates. Latency becomes noticeable to users around 30 milliseconds, and even delays as short as 100 to 400 milliseconds can negatively affect consumer behavior.[3] Businesses that fail to implement a functional load balancing solution may see conversion rates and revenue decrease when their web properties inevitably run into performance issues.

High costs and limited flexibility: Hardware load balancers are expensive, complex, and inherently unable to scale as a business grows. Enterprises must purchase hardware devices upfront by estimating the amount of traffic their web properties will receive, which means that they are often either required to pay for unused capacity or forced to suffer through suboptimal load times and web performance until they can increase the number of devices in use.

Complex traffic management: For businesses looking to maximize their web performance and reliability, getting visibility into traffic patterns, managing regional traffic, and monitoring the health of origin servers is essential. Without that visibility, businesses may not be able to detect when malfunctioning load balancing solutions inadvertently route traffic to a server that is experiencing problems, causing significant delays and downtime for users.

[3] Brutlag, Jake. “Speed Matters,” Google AI Blog, https://ai.googleblog.com/2009/06/speed-matters.html

What to look for in a load balancer

Vendor-agnostic solution: A load balancer that offers multi-cloud and hybrid cloud support can help businesses circumvent vendor lock-in and complex configuration. Instead of supplanting existing load balancing services, a standalone cloud-based solution can integrate with cloud vendors' native load balancers or traditional hardware appliances to maximize flexibility and minimize misconfigurations. Essentially, it should enable businesses to dynamically distribute traffic across their infrastructure — whether their origin servers are hosted on-premise or in multi-cloud or hybrid cloud environments.

Active health checks and detailed analytics: When selecting a load balancing solution, visibility is key — not only so businesses can ensure the health of their servers and applications, but so that they can get ahead of potential latency and downtime, too. With detailed analytics on traffic patterns and origin health, businesses should be able to identify underperforming servers and optimize their infrastructure for high availability and uptime.

CDN integration: In an ideal configuration, a load balancing solution works in tandem with a CDN to minimize latency and bandwidth consumption. By caching static content at the network edge, a CDN is able to deliver content from the closest server to the end user, both enhancing overall web performance and reducing the total number of requests sent to the origin server.

STEP 5

Detect Anomalous Behavior and Secure Web Properties at the Edge

Data Loss Prevention (DLP)

With the advent of cloud computing, data breaches have become one of the most significant threats to modern businesses. Often the result of targeted attacks, internal system glitches, and simple human error, these breaches can expose sensitive customer data, violate data privacy regulations, and cost millions of dollars in fines and lost revenue.

In order to stay ahead of potential leaks and the loss of sensitive data, businesses may implement a number of cybersecurity strategies and products. Data loss prevention (DLP) solutions can also assist businesses in complying with the GDPR and CCPA, among other regulations designed to enforce data privacy policies and protect user data from unauthorized use.

DLP challenges

Complex deployment: Legacy DLP solutions, while robust in nature, can be complex and time-consuming to set up. Businesses need to tailor DLP policies to fit specific groups of users and business use cases, which can be a painstaking process that requires extensive external support and administration. And, as DLP rules prevent users from transmitting data beyond strictly defined boundaries, they can hamper employee productivity and collaboration by inadvertently blocking access to necessary data.

Lack of comprehensive data protection: DLP policies need to cover both regulated data (sensitive data that should remain classified) and unregulated data (all publicly known information, which may include some confidential data). However, many legacy DLP products do not account for the protection of unregulated IP data, leaving businesses vulnerable to substantial loss if a breach occurs.

Limited visibility: In order to ensure total protection of confidential data, prevent insider threats, and remain compliant with local privacy regulations, businesses need visibility into the ways their data is being accessed and transmitted. Focusing too much on existing or expected threats may impede a business's ability to stay ahead of unexpected malicious behavior.

What to look for in a DLP solution

Streamlined deployment and management: Hardware-based DLP systems are not only complex to set up and clunky to manage, but require constant updates against ever-evolving threats. Cloud-based solutions can help reduce deployment costs while remaining flexible against new data threats and giving businesses greater visibility into their data usage and management.

Flexibility: Businesses need to adopt a DLP solution that is flexible enough to accommodate various user groups and use cases, while remaining easy to implement, manage, and maintain. Rather than depending on hardware-based, legacy DLP solutions, enterprises should consider cloud-based alternatives that offer greater flexibility and control over the policies and protection methods put in place to defend confidential company and customer data.

Protection vs. prevention: Although legacy DLP systems primarily focus on data loss prevention, they cannot safeguard confidential company data against every external and internal threat. Next-generation solutions allow businesses to not only mitigate these threats, but recover from data breaches in a faster and more efficient manner.

Edge programmability

Edge computing enables businesses to shift application development to the network edge, bringing computation as close to end-users as possible and minimizing latency, server resources, and bandwidth usage. With a serverless architecture, businesses can offload infrastructure configuration and management to a third party, freeing up their developers to focus their efforts on building and deploying applications, enabling custom configuration for existing applications, and helping optimize application development and security.

Edge programmability challenges

Latency and cold starts: Because serverless computing runs functions on an as-needed basis, it can take up to several seconds to start up a function. These 'cold starts' may introduce unwanted latency and require businesses to take proactive measures — like finding workarounds to minimize the duration and frequency of cold starts — in order to ensure that users don't see a measurable decline in web performance.

Lack of global scale: For web-based businesses with a widely-distributed user base, applications must be deployed on a global scale. Moving applications to the network edge helps businesses reach users more efficiently while also improving application performance and delivery.

Inefficient resource application: Building applications and creating custom programming are involved processes that require dedicated internal resources, adequate server capacity with a central cloud provider, and time for testing. When done poorly, this can increase the cost of prototyping and delay businesses from getting their applications to market quickly.

What to look for in an edge computing solution

Simplified scalability: Since edge computing is designed to minimize latency for end users by delivering assets from the network edge, it is crucial that businesses select a provider with a global network. This not only enables businesses to expand their reach and provide a faster user experience, but simplifies the deployment process by running code at the edge to reach customers in all locations, rather than on a region-by-region basis.

Improved developer experience: Developing new applications and augmenting existing ones should be a streamlined process, allowing developers to deploy code quickly and easily — without needing to depend on technical operations teams to do so.

Reduced infrastructure costs: With serverless architecture in place, businesses no longer need to pay for unused server space or idle CPU time. Instead, they can shift more request handling to the network and only pay for the resources they need, allowing them to significantly reduce infrastructure costs.

How Cloudflare helps businesses deliver a superior online experience

Creating a superior online experience requires the right security and performance strategy — one that not only enables enterprises to accelerate content delivery, but ensures network reliability and protects their web properties from site outages, data theft, network vulnerabilities, and other critical attacks.

Built on a network that spans 200+ cities in over 90 countries around the world, Cloudflare provides a scalable, integrated global cloud platform that helps businesses deliver security, performance, and reliability for their on-premise, cloud, and SaaS applications. To learn how you can protect and secure your online business with Cloudflare, visit Cloudflare.com.

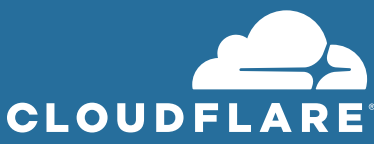

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com

REV: 200408