Guide

# Edge Cloud Platform Buyer's Guide

Your companion to evaluating edge cloud platforms. Use this guide to ensure you ask the right questions and investigate the capabilities that are important to you.

fastly

# Introduction

When we talk about a “modern CDN” we mean one that can do much more than people often think CDNs are capable of. In the 1990s, CDNs had a much simpler set of responsibilities – moving static objects like an image or HTML page more quickly across the internet to an end user. Since then, the internet has evolved (significantly!), creating an opportunity for CDNs to take on a larger and more complex role with new, critical functionality. Now, it’s about more than accelerating websites and applications; great CDNs also protect them at the edge, and provide advanced tools to make teams more productive. Part of this evolution in CDNs has been a shift from serving as a single-purpose tool to an edge cloud platform – a one-stop-shop for caching, delivery, security, storage, and development at the edge.

Today, many vendors call their offerings a “platform”, but not all platforms are created equally. In this Buyer’s Guide we’ve provided you with everything you need to evaluate an edge cloud platform and to ensure you ask the right questions when investigating the capabilities that are important to you.

# Better platforms lead to better business outcomes

Many organizations have tolerated subpar CDN and WAF performance for years, avoiding a migration to superior platforms due to fear of the risk and cost associated with the migration process. But evolving expectations, better understanding of the cost of security failures, and concerns about limitations to innovation imposed by vendors, are collectively driving organizations to be more open to considering a switch. Key drivers for embarking on a platform migration include a desire for improved content delivery (with enhanced site and application performance + lower latency), increased visibility, superior, more agile security solutions and bot management. The best edge cloud platforms can address all of these requirements and more, offering solutions that don’t lock you in and that open new possibilities.

If you’re considering moving more of what you do to the edge, or moving it to a better edge platform, it is important to approach your vendor and platform evaluation holistically – that is, to consider not just the product offerings and specs (WAF, CDN, DDoS etc.) across vendors, but also to investigate the integrity and capabilities of the underlying network and platform these services run on. The capabilities of WAFs, CDNs and edge computing are constantly evolving, so it’s critical to ensure a platform is built to support you today and can adapt to meet your changing needs in the future.

# Evaluating an edge cloud platform

Vetting an edge cloud platform should entail a comprehensive look at the platform itself along with an examination of the network and infrastructure that support it. While there are certainly any number of approaches to implementing or enhancing an edge platform strategy, we take pride in the way we’ve built the Fastly platform and believe that in order for a platform to truly support an organization’s business goals, it must satisfy three core competencies.

*You can begin your platform evaluation by analyzing target platforms through the lens of the following three characteristics that we deem essential to any modern best-in-class edge platform.*

# 1. A superior underlying network

## What makes a network "better"?

Not all edge platforms are created equal; comparison of their architecture, performance, and the future potential of the networks they're built on can help to differentiate them.

When assessing the underlying network, its core capabilities should include low latency, high reliability, extreme resilience, and the capacity to effortlessly scale. Remember, everything is built on top of this network, so if you aren't starting with the absolute fastest, lowest latency, and highest flexibility network, then anything you do on the network is operating from a deficit.

A network's capabilities aren't always easy to gauge as there's no single metric for performance, resilience, or scalability. However, here are some capabilities you can look for as indicators of a "better" network.

### Network evaluation checklist:

- ☐ The network delivers ultra-low latency performance
- ☐ The network supports modern protocols like HTTP/3 and QUIC
- ☐ The network is fully software defined, and not dependent on hardware switches or routers
- ☐ The network receives regular upgrades and performance improvements
- ☐ The network has advanced capabilities for instant traffic path optimization and routing as well as auto-healing
- ☐ The network can respond in real time to performance degradation or connectivity issues
- ☐ The network offers infrastructure-agnostic routing
- ☐ There is a single network for the entire platform and all of its products

## How a better network makes a difference

**An ultra-low latency and higher performance network**

*When you select a platform that operates on a software-defined network, it ensures higher efficiency in several ways.*

Often, hardware switches and routers become bottlenecks in a network's architecture since their performance is hard-capped and cannot be improved through software updates. This introduces a form of technical debt into the network architecture – you have a physical component that is less capable of adapting without significant costs and associated resources needed to perform replacements and upgrades. Similarly, the switches and routers' capabilities are limited and can't be infinitely customized to support architecture changes and improvements. Finally, in the case of a failure, hardware switches and routers make it difficult to route around them.

When network switching and other capabilities instead operate on a **software-defined network**, this enables continuous performance improvements with software updates, more granular and custom routing capabilities, and improved overall resilience. A network running on software can intelligently respond to issues and failures and can route around a failure or interruption without reliance on a specific box in a server rack.

A software-defined network can also be infrastructure-agnostic with routing decisions; network logic can control which infrastructure (cloud, hybrid, multi-cloud) items are routed to based upon real-time information like context, traffic, demand, or network conditions. When a network operates with this depth of 'intelligence', it can truly support you at the highest speeds and with the highest possible performance.

A higher performance network will also support modern protocols like HTTP/3 and QUIC, which reduce latency and delays and contribute to a better digital experience for end users. Being fully software defined also means the network is more future-ready; software updates can enable support for new protocols and features that are yet to be defined without having to replace and upgrade hardware components to achieve compatibility.

**A more reliable, resilient network, with auto-healing**

When you select a network that is truly software-defined, it can support greater automation that instantly responds to common performance degradation or connectivity issues with internet transit providers. This capability allows the network to side-step slow areas of the internet and always find global 'fast lanes', resulting in consistent speed, resiliency, and reliability. Fastly's Autopilot makes this possible - our network routes around "internet weather" and other problems thanks to our resilient architecture.

**A unified network that scales instantly**

A platform that runs on a single network has additional advantages. From a performance perspective, it eliminates the latency tax of services living on disparate networks. Take the example of DDoS protection; when a network is unified, protection occurs on the same network that the content is served from, removing the need to send traffic to different networks for scrubbing and a subsequent return. In terms of egress, this means never being charged to move data from storage on the network to another service like your CDN.

A unified network also makes it simpler to scale instantly. If you remove the bottlenecks of hardware components, and when everything is running on one network, it's easier to ensure that all customers experience the highest level of performance for all of their traffic – static, dynamic and media traffic, as well as TLS and DDoS protected traffic. For example, on Fastly's network, every POP location runs every service we offer, so whatever our customers use us for will operate seamlessly from anywhere, and scale instantly within the Fastly global network. Maybe one of your pages or sites goes viral, or you're under attack and performing security checks against a massive bot attack, or you're serving a huge piece of content that is extremely popular – no matter what needs to scale, and no matter where in the world it needs to scale, there's network flexibility and capacity ready to handle it.

# 2. A powerful, easy-to-use, developer-friendly platform

## How do you evaluate a platform?

It is not enough for a platform to be built on top of a thoughtfully-designed architecture, or to provide a lot of new capabilities. The platform must also be easy to use. DevOps, SecOps, application developers, and everyone else should experience a sense of relief and measurable productivity improvements.

The keys to those productivity and efficiency gains are integrations, CI/CD compatibility, and robust logging and observability that empower teams to do more. You are looking for a general flexibility and usability in the platform that benefits DevOps, SecOps, and application developers. Beyond that, you're looking for ways to help those teams work together to unblock and empower each other in ways that were previously impossible.

As in the previous section, it's not simple to measure 'ease-of-use' for an edge cloud platform - there's no single metric to look at without a deeper analysis of a platform's capabilities. But here is a list of capabilities that are great indicators of a platform's power, ease-of-use, and the amount of control it provides to you.

### Platform evaluation checklist

- ☐ The platform is highly configurable, letting you make changes and updates to your CDN service yourself without waiting on professional services
- ☐ The platform is highly programmable, providing a flexible serverless compute environment for App Developers to innovate on
- ☐ The platform is API-first, allowing for API-based control and configuration across its set of capabilities and products
- ☐ The platform works with the CI/CD workflows and tooling that developers already use; Amazon S3, Google Cloud Storage, Drupal, and more
- ☐ The platform supports DevSecOps practices with single-click integrations with the most common dev and ops alerting engines (chat-ops, project management, etc) + integrates with data collection, infrastructure monitoring, and deployment automation tools
- ☐ Logging is real-time and complete; it is not delivered with lag, or limited to a sample or subset of your total data
- ☐ Logging can be streamed to other endpoints with pre-built integrations for commonly used tools

## How a better platform makes a difference

**API-first, CI/CD compatibility and more configurability and control**

A platform that offers configurability and programmability is one that gives you (and your teams) the access and control to change what you want, when you want, with real-time visibility. You and your team should have more control of your platform, and the ability to make changes or rollbacks immediately, without delay.

When a platform is built in such a way that it allows the services running on it to be fully configurable and customizable, it is game changing for DevOps productivity. Teams can deploy custom code and configs on the platform and effortlessly roll back changes if necessary. An API-first approach to building also allows customers to integrate easily with tools they already use like Terraform and other commonly used elements of a CI/CD workflow.

**Flexible serverless compute environment to innovate on**

A platform should also provide App Developers with a flexible serverless compute environment to innovate on. The environment should be highly scalable, allowing developers to build and deploy logic at the edge, thereby solving the most latency-sensitive use cases closer to end users.

**Real-time logging, observability, and insights**

When a platform runs on a fully software-defined network, this removes the 'black box' limitations that often add serious limitations to logging and reporting. Many platforms offer either incomplete data reporting or limited access to sample sets of data. Worse still, networks with inferior architecture often cause delays in data delivery.

Immediate access to data is critical for decision making - delays of even a minute or two can mean the difference between catching something early or discovering it too late. With timely insights into traffic and how users are interacting with services, Dev teams can better control the end-user experience. This is thanks to the ability to take an agile real-time approach to troubleshooting and performance evaluations.

A platform should provide you with pre-built integrations with the tools you already use, and the capability to stream logs anywhere you like. It's your data, and you should have easy, instant access to all of it in the ways that work for you.

**Empowering for DevOps and DevSecOps**

Legacy WAFs and bot solutions frequently force security teams outside of their toolchain environments when identifying and remediating issues. Rule testing and validation efforts often require teams to switch between systems and tools, hindering fast, effective response times in the face of new threats. These cumbersome processes inhibit a streamlined DevSecOps program.

A platform should integrate seamlessly with existing tooling and workflows. Engineers and SecOps teams need a platform that folds into the way they already work; think alerting engines, ChatOps, project management tools, and incident tracking systems just for starters. On the security side, legacy WAFs usually force teams outside of their toolchain environments to identify, diagnose, and resolve issues, which slows down time-to-resolution on issues. To ensure a strong security posture, teams need real-time incident alerting that works on the tools they already use to guarantee fast and prioritized incident remediation.

The platform you select should have a clear track record of solving these common DevSecOps pain points, improving the performance and productivity of your teams, and reducing alert fatigue and burnout. You can read more about how Fastly has worked to successfully improve the DevSecOps workflows.

# 3. Platform superiority translates to higher performance and better products

## How do you evaluate the product suite?

Ultimately, a powerful underlying network and a developer-friendly platform don't mean anything if those benefits don't translate into differentiators and benefits across the platform's product suite. The right platform will enable the best in performance, reliability, and flexibility, allowing you to drive highly-personalized experiences, instantly, from the edge.

You should see measurably improved performance with faster websites and applications, as well as more efficient performance with lower egress charges. You should have more flexible deployment options for your WAF, and measurably faster time-to-protection on threats through access to instant insights. You should see overall improvements to accuracy and protection, and more effective bot protection. You should have faster and more secure edge computing and storage solutions that outperform the competition.

Here are a few indicators to look for within the product suite of an edge cloud platform that serve as indicators about the quality of the platform and network upon which they're built.

(Note – we also have product-specific Buyer's Guides for CDN, WAF, and edge computing available if you want to focus on comparing at the product level.)

### Product suite evaluation checklist:

- ☐ Proof of industry-leading performance
- ☐ Very high percentage of WAF customers are deployed in blocking mode
- ☐ Security product deployment is proven to be easy, fast, and flexible across environments
- ☐ WAF and bot rule building is simple and extremely fast, and doesn't need much tuning
- ☐ Available edge computing that is secure by default and high performance
- ☐ Extremely high performance data storage at the edge
- ☐ Transparent business continuity planning (BCP) to ensure product resilience and availability

## How a better platform enables a superior product suite

**Industry-leading performance proven through faster and higher performance products**

The platform your products operate on should enable the fastest and highest performance possible across your product portfolio, helping you deliver these benefits to your end users. Fastly is transparent about our low-latency performance in content delivery and truly instant content purging, as well as the cold start times for our edge computing, and our 30x faster KV Store for storage at the edge. Look for simple and undeniable metrics like these for product performance that can only be achieved when the underlying platform and network are not creating bottlenecks.

For security products like a WAF, you can also look for proof of easy deployments and metrics around the percentage of customers who have it deployed in blocking rather than logging mode. When a vendor has a high percentage of customers who run in blocking mode, this can be understood as a clear vote of confidence in the accuracy and ease-of-use of a WAF. It shows that the team believes it can protect without causing a lot of false positives.

WAFs can also be judged based on how easy it is to create and deploy new rules while maintaining accuracy. You should look into how much tuning and maintenance time can be eliminated for your security team. Customer references that confirm deployment was simple across environments should help provide you with even further confidence.

Edge computing capabilities should be similarly high performance, but they should also be secure by default. Look for strong and granular isolation and sandboxing. Fastly Compute has per-request sandboxing, protecting more by default for anything it runs, showing that there doesn't need to be a trade-off between performance and security. For reference you can read about how Fastly Compute's KV Store can deliver anywhere from 10x to 30x faster performance than the competition with a secure by default solution.

**Continuous investment into reliability and resilience**

You should have confidence in the resilience, reliability and availability provided by the network itself, and the products you use on top of it. Look for more than just assurances – you should see a strong track record of reliability, but also transparency about how a company responds and how they ensure that things are always becoming more reliable. Business Continuity Planning (BCP) is a standard approach to ensure a constant improvement and iteration on product resilience and availability. It means an organization isn't just making improvements – they're actively searching for potential points of failure and fixing them before they become problems. For reference, here is one example where you can find Fastly's approach to resilience and BCP efforts.

# What's next: Compare at the product level

Now that you understand what makes a powerful platform and what questions you can ask to evaluate your platform options, let's dig into the product-level differentiators that you should investigate. We said it above, but we'll reiterate here as well; the core elements we discussed that make up a best-in-class platform are important by themselves and also critical to enabling the product benefits we'll discuss below.

Across CDN, WAF, and edge computing services there are clear differences, and many of these can be traced back to the platform differentiators we talked about above. The Buyer's Guides linked below for CDNs, WAF, and edge computing provide checklists to help you with your product-level evaluations.

# Related resources

## CDN Buyer's Guide

Your Definitive CDN Buyer's Guide: Get help navigating the content delivery landscape with this comprehensive guide.

**Download the guide now →**

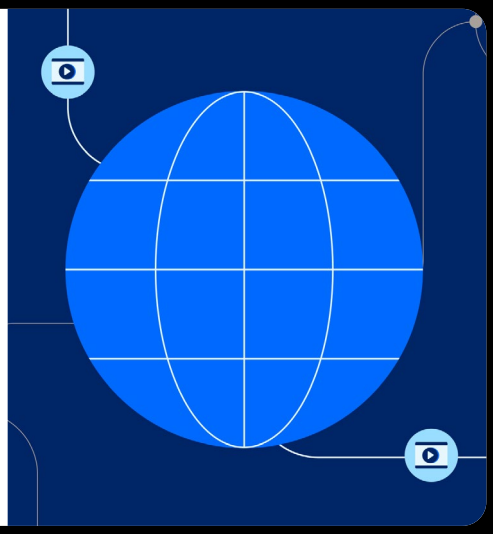

## Edge Computing Buyer's Guide

Learn how to do more on the edge. Your guide to choosing an edge solution that's right for you today, and can grow with you in the future.

**Get the guide now →**

## Essential WAF Buyer's Guide

The guidance you need to select a modern-day WAF. Everything you need to protect your apps and APIs today, and into the future.

**Get the WAF guide now →**

## How Fastly delivers an ROI that sounds impossible

See Fastly's real-world impact with Forrester's Total Economic Impact™ (TEI) study. See how Fastly showed a 189% ROI over 3 years*. Get the report

**Get the report →**

# Bonus content

The Cloud Native AppSec Playbook

A guide for engineering, operations, and security teams providing the "how" and "why" of cloud native application

Four things every security director should know about GraphQL

Educate your security teams on GraphQL and keeping their environments secure.

Simplified Image Optimization

Your companion for a faster, more engaging web experience.

Guide to the Modern CDN

Learn about the difference a modern CDN can make.

5 Reasons Companies Need a Modern CDN

Get five strategies for evaluating your current content content delivery strategy.

Appsec Guide for Multi-Layer Security

Discover 8 tactics for a unified approach to AppSec.

The Modern Application Development Playbook

Learn how computing at the edge allows orgs to maintain innovation speeds while satisfying DevOps needs.

fastly