



Think Global, Compute Local

5 edge computing use cases

ebook



TABLE OF CONTENTS

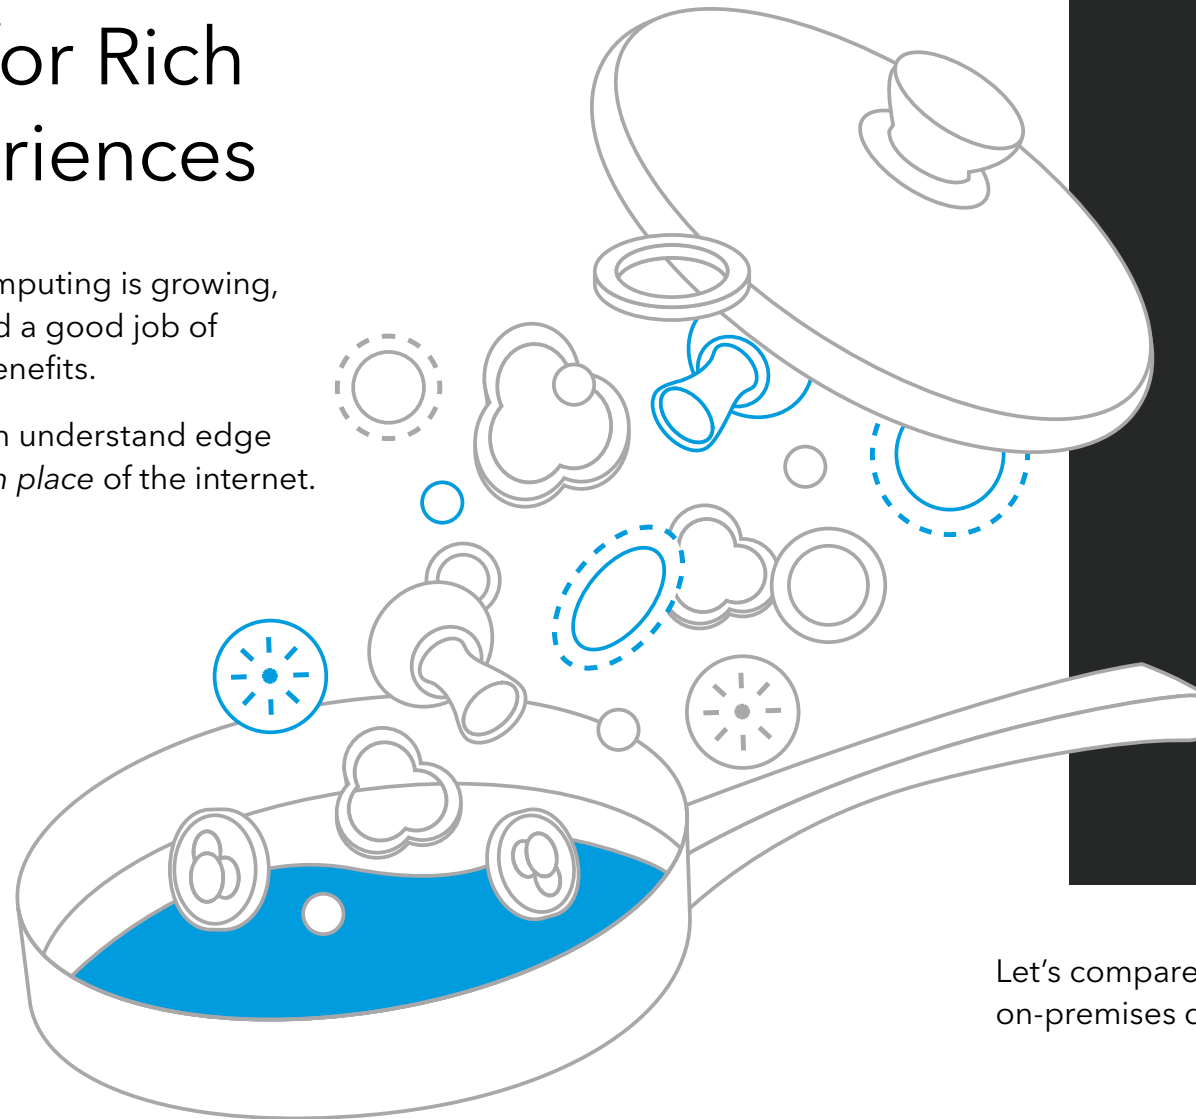
Introduction	03
Chapter 1: Geolocation	07
Chapter 2: A/B testing	08
Chapter 3: Dynamic content	10
Chapter 4: Third-party services	11
Chapter 5: Privacy compliance	14

Edge Computing

A Recipe for Rich User Experiences

The buzz about edge computing is growing, and it's time someone did a good job of explaining its business benefits.

To clear things up, we can understand edge computing as the *mise en place* of the internet.



EVERYTHING READY AND WITHIN ARM'S REACH


In cooking, the French term *mise en place* means "everything in its place."

It refers to the efficient prep and setup of ingredients before cooking. Without *mise en place*, harried cooks would waste time looking for garlic while the onions are frying. They'd still be cutting vegetables when the oil is reaching its burning point. The wait time for entrées would be doubled or tripled. But with a proper *mise en place* setup, everything you need is ready and within arm's reach.

Edge computing operates on a similar principle. If you need something quickly (i.e., milliseconds), it must be within arm's reach of every user. The average website juggles countless microservices to ensure great user experiences, strong security, incremental learning and optimization, and more. Each of these operations adds delays that hamper the user experience.

Let's compare how the internet works when totally centralized (in the cloud or with on-premises data centers) versus when edge computing is incorporated.

The shape of the internet of **yesteryear**



When a website visitor opens a page, they are making a content request. That request is sent to the application server. The origin server fulfills the request by sending content to its destination – the visitor’s device.

It takes time for the request to be received and fulfilled. The distance between the origin server and the visitor’s device adds latency, i.e., slower load times. An origin server is slower to deliver content to a user who is far away than it is when delivering content to a user who is nearby.

The shape of the internet **today**



Since the rise of streaming video and music, content delivery networks (CDNs) have taken much of the load off of origin servers. A CDN is a group of servers located across the world in strategic locations to deliver content. By storing content on these intermediate servers, businesses improve page load times and can handle large surges in traffic.

One of the key functionalities of a CDN is **offload**.

WHAT IS OFFLOAD?

Offload moves a portion of resource-intensive content or requests to run on a separate platform. In the case of a CDN, that platform is closer to the user, and this reduces latency they might experience otherwise, while also offering the ability for core infrastructure to scale further due to lighter demand.

When a CDN like Akamai’s gets a content request, the **cached** content is not far from the end user. This results in:

1. Faster load times: Because we’re closer to the user, we know the network condition and will map to an optimal server.
2. Better infrastructure management: Expect better load balancing thanks to the reduced overhead of managing high traffic demands.

“Growing demands placed on web infrastructure have given rise to the next generation of CDN solutions ...”



WHAT IS A CACHE?

In a traditional CDN, the cache stores data so future requests for that data can be served faster. The data stored in a cache could be from an earlier request, bypassing the need to retrieve it from the origin server again.

Consider Black Friday. You need the infrastructure to support a traffic influx. With a CDN, you can host and cache part of your content in the cloud, providing you with the breathing room to respond to converting users. The same holds true for videos, large images, and any large static content.

There are variables other than content size that affect latency, such as intermediate routers and switches, congested peering points, and inefficient routes. More advanced high-value features, such as A/B testing and personalization, are not static and traditionally come with a performance penalty.

These growing demands placed on web infrastructure have given rise to the next generation of CDN solutions intended to manage both static content *and* application logic. This is edge compute.

The shape of the internet to come

A host of UX-critical functions are lightweight enough to move from the origin to the edge. Moving these functions to the edge keeps the logic closer to the user, smoothing and speeding every interaction.

Edge servers exist close to the end user to avoid round trips for every tiny function. If servers nearby can perform microservice functions, then you can reduce latency, improve the balance between performance and personalization, and avoid the bottleneck of relying on origin servers.

WHAT IS THE EDGE?

The edge *is* your CDN, but it's a CDN that you can run code on.



In this guide, we'll demonstrate the top use cases for how edge computing can save businesses money while enhancing the user experience.

These use cases are:

1. Geolocation
2. A/B testing
3. Dynamic content
4. Third-party services
5. Privacy compliance

Geolocation

Location-based personalization is much more than selling raincoats when and where it's raining. Geolocation can help you to increase user engagement, reduce bounce rate, and improve conversions by delivering location-specific content like localized site versions, nearby stores, or special products and offers.

By moving a portion of your website functionality to the edge, you can speed up performance, determine location while remaining compliant, and deliver a stellar experience.

Meet your users at the edge

Besides benefiting the user experience, geolocation on the edge can help businesses remain compliant in heavily regulated industries. For example, online retailers navigate a complex landscape of promotions and discounts that vary from region to region.

Edge compute geolocation benefits

- Dynamically personalize by providing different experiences based on a user's geolocation (e.g., the U.S. version versus the UK version)
- Reduce the friction of conversion by pre-populating online forms with information relevant to users (e.g., a list of phone area codes, ZIP codes, city, state, etc.) based on the IP geolocation of the user
- Automatically maintain compliance with regional policies

An international bank only has customers in a limited set of countries. Their long list of logic was hard-coded into their cloud infrastructure.

They now manage the code themselves in a nimbler **edge environment**.

A/B testing

A/B testing involves several moving parts: in-the-moment segmentation, traffic routing, tracking, data collection, and dynamic content. This lets companies continually refine the user experience and, by extension, business performance.



Latency: the saboteur variable

This type of experimental decision logic is typically implemented on the server side or in the client's browser. Broadly speaking, a developer would choose the server side to prioritize security and the client side to prioritize speed.

Whichever method you choose, these services eliminate the ability to cache – meaning either everything goes back to the origin with every request, or it fragments the cache, reducing effectiveness and efficiency – which drags down the customer experience.

The added latency associated with an experiment could be enough to negatively impact performance. Even worse, latency can skew results in favor of an inferior page element that wouldn't normally perform better.

Performance is a real-world condition and should not become the secret confounding variable while researchers and analysts assume "all else being equal." Getting A/B test results wrong isn't just a disappointing slip-up: a 1%-2% difference in conversion rate could represent a multimillion-dollar missed opportunity over the lifetime of a decision.

A wellness retailer uses **low-latency A/B testing** at the edge to segment users, assemble content, and deploy test codes to continually optimize and drive conversions.

Create real “real-world conditions” at the edge

The logic that ran server-side can now run at the edge, thus reducing round trips to the origin and gaining the ability to cache what previously could not be cached. This results in faster page performance, decreased traffic to the origin, and therefore more accurate results.

To run a separate segmentation service in parallel with other operations and prevent page-load delays, businesses should:

- Segment users
- Save decisions across sessions so users return to the same experience
- Maintain copies of objects in the cache so a user quickly gets the right content

With these functions running at the edge, all the origin infrastructure must do is create one copy of each variant. One result of unburdening the origin and speeding page performance could be more accurate results, which could represent a big difference in business performance.

BENEFITS OF RUNNING EXPERIMENTS ON THE EDGE

Experimentation is critical for learning about real customer preferences and behaviors before your competition does.

The decision about which content to serve to a user segment is written at the edge, as opposed to the origin, resulting in:

- Faster, and more consistent, end-user experiences
- Increased offload
- Reduced origin requests and resulting compute



Dynamic content

Companies want to deliver more personalized user experiences, but the functions of identifying the user and determining which content to present typically reside at the origin.

Segmenting your traffic with different experiences dilutes your cache and adds speed-inhibiting functions. This is the same for logic, personalization, and experimentation, which hinder performance if left to run at the origin.

The trade-off usually means personalized content is uncacheable. Uncacheable content causes low offload and negatively impacts performance. Edge compute can be used to detect the characteristics of inbound requests, rapidly identify unique users, and retrieve personalized content.

Making a decision about what content to serve will obviously happen *before* the content can be retrieved. If you can move decision-making closer to the user, everything else speeds up. In other words: Make the decision close to the user and serve content from the cache.

Performing this logic at the edge enables highly personalized user experiences while increasing offload and making possible a faster, more consistent experience.

EDGE COMPUTE DYNAMIC CONTENT BENEFITS

Personalization speaks directly to your customers to catch them in their most receptive moods and moments. Personalization can increase conversions, improve retention, expand social sharing, and elevate revenues. Personalized content is retrieved from the edge, as opposed to the origin, resulting in:

- Faster, and more consistent, end-user experiences
- Increased offload

Chapter 4

Third-party services

Enterprise web experiences run a resource-clogging gamut of third-party services, including analytics, tracking, social widgets, chatbots, payment providers, mile-high marketing stacks, and more.



Take control of third-party applications

The scale, reliability, and speed of these third-party applications are out of your hands because very often they each have their own infrastructure – wherever that may be – or they run locally on the user’s device.

There are downsides to microservices living in the vendor’s origin. If your visitor is in Miami and makes a content request of your San Francisco server, that is already a long trip. Add vendor-native third-party services in New York City, Boston, Chicago, Toronto, and Los Angeles, and you’ve created a hellish itinerary and a sluggish experience on the user’s end, even if some of these requests are parallel or nonblocking.

Edge computing does with third-party applications what Amazon does with shipping logistics. If Amazon relied on manufacturers to handle shipping themselves, then one-day shipping would be impossible.

By operating dozens of distributed fulfillment centers around the U.S., Amazon keeps inventory close to every possible U.S. customer. Thus, they can earn the differentiation of fastest delivery in the business.

For Akamai, edge computing moves the “inventory” (third-party service) from the “manufacturer” (vendor’s infrastructure) and distributes it among our 4,100+ “fulfillment centers” (edge servers).

Running code on a user’s device also has its own big downsides:

- Creates a processor burden for the user’s device
- Drains the device’s battery life
- Gives users permission to inspect the code, opening potential exploitation

These downsides have been accepted as trade-offs for maintaining site performance. Moving the code to the edge doesn’t require these trade-offs.

EDGE COMPUTE THIRD-PARTY SERVICE BENEFITS IN BRIEF

- Reduced dependency on third-party infrastructure because the code is edge-native
- Improved user experience
- The output of an edge function can be cached and used for subsequent requests



A scalable solution

Migrating even a few of these functions from “the main itinerary” to the edge frees up third-party-bound resources for everything else, creating a better experience. Because the code is nearby on the edge, latency is low. The code is more reliable because it can scale to meet changing traffic volumes.

HOW AKAMAI ENABLED A SMOOTH VACCINE ROLLOUT

Akamai and the Queue-it virtual waiting room service partnered for sign-ups for COVID-19 vaccinations. Moving the waiting room service to the edge reduced the risk of a website crash due to traffic spikes at a critical moment in history.

Businesses can handle traffic spikes by offloading certain functions to the edge. These might include ticket reservations to a hot concert or holiday shopping rushes. Third-party services don't need to resolve the trade-off of speed versus functionality; the two can coexist on the edge.



An online retailer saw its product launches consumed by users bypassing purchase limits.

The brand outmaneuvered the malicious actors by moving its queuing application logic from the **client to the edge**, guarding the waiting room application from exploitation.

Privacy compliance

Following legislation like GDPR and CCPA, websites now need to get consent for third-party tracking. Third-party cookies are code snippets set by a website other than the one you are currently on. Cookies stay in a user's browser so a site can "remember" information between visits.

Moving consent-tracking functionality to the edge is an important step to balance privacy compliance with running a high-performing digital experience.

When the user consents to be tracked, cookies that enable a personalized user experience are added to the session. If the user does not consent, the cookie is discarded and the user has a more generic experience.

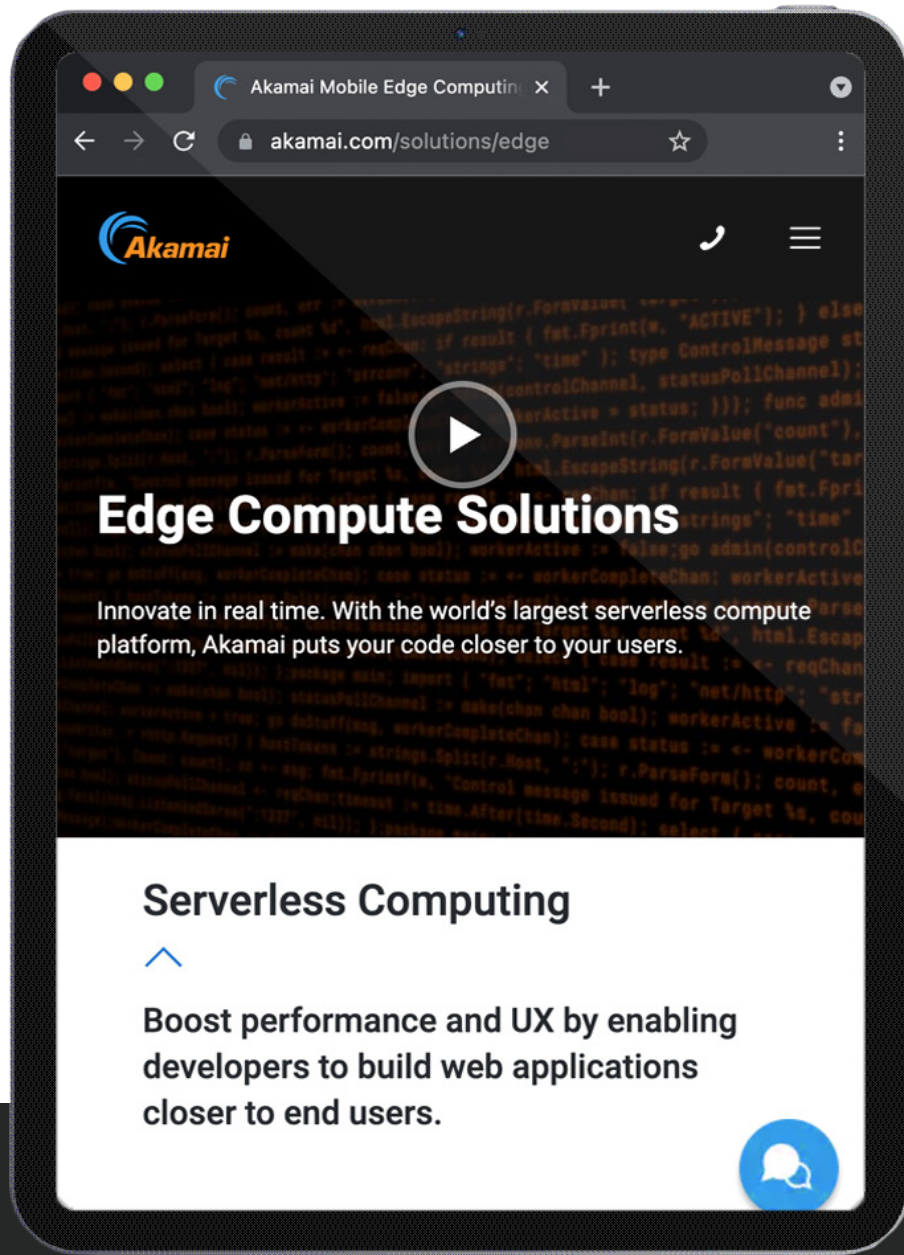
Performing these consents at the edge improves offload and enables cacheability, allowing rapid lookups. By restricting the movement and storage of personally identifiable information (PII), users can choose when, where, and for how long their data is accessible to third-party applications. This improves the user experience while helping ensure privacy compliance.



A global analytics company uses an edge microservice to determine a user's **consent history** according to regional regulatory requirements.

BENEFITS IN BRIEF

- Filter out unnecessary traffic back to origin, resulting in a reduction in origin infrastructure and costs
- Increase offload
- Remain compliant without compromising end-user performance
- Decrease vulnerability to exploitation



Write the code; deploy on Akamai

The rise of edge computing marks a significant shift in how data and business logic will be distributed and consumed in the future.

Akamai enables you to bring your application out to the edge, deploying code on the world's most distributed low-latency edge computing platform. This puts your business logic only milliseconds away from your users, where you can make decisions that can shape their digital experience faster. And because it's your code, your team can innovate upon the customer experience at the speed your audience and business demand.

We've mentioned the powerful benefits of moving certain workloads to the edge. But the evolution of edge computing is only beginning. Join Akamai as we work with the world's leading companies to build the next evolution of digital experiences.

Get closer to your users.

[Learn more](#)

Akamai powers and protects life online. The most innovative companies worldwide choose Akamai to secure and deliver their digital experiences – helping billions of people live, work, and play every day. With the world's largest and most trusted edge platform, Akamai keeps apps, code, and experiences closer to users – and threats farther away. Learn more about Akamai's security, content delivery, and edge compute products and services at www.akamai.com and blogs.akamai.com, or follow Akamai Technologies on [Twitter](#) and [LinkedIn](#). Published 11/21.

