



ENGINEERING SCALE AND PERFORMANCE FOR CUSTOMER IDENTITIES



WHITE PAPER

TABLE OF CONTENTS

03

04

04

04

05

05

05

06



INTRODUCTION

More than ever, customers demand secure and seamless experiences. They want consistent interactions on all devices, user-friendly interfaces and data security, but they also need instant access to the content, products and services of brands they interact with. And when we say instant, we really mean it. One study found that people will abandon a website that takes more than 3 seconds to load¹. These levels of service are even more difficult to maintain in environments with tens (or even hundreds) of millions of customers and thousands of logins per minute.

Enterprises must be able to handle massive volumes of customers while securely providing social login, single sign-on (SSO), privacy and preference management and other analyst-recommended customer identity and access management (CIAM) capabilities. CIAM solutions are able to achieve high performance and availability while enabling engaging experiences for millions of customers. There are a number of new business drivers that are making scalability for customers more urgent:

1. **Customer experience expectations are rising.** Customer experience leaders are raising the bar, in large part by providing high-performance interactions with digital properties.
2. **Mobile interactions are more frequent.** Years ago, customers may have interacted with a brand's digital properties once a day from their computers. But today's customers are mobile and interactions are far more frequent. Push notifications, geolocation and other back-end services can use resources even when a customer isn't using an app.
3. **Customer profile data is more complex.** Customer profiles no longer consist only of basic attributes like name, email address and phone number. They now contain unstructured geolocation and browser data, device IDs, sensitive personally identifiable information (PII) and much, much more.

All this considered, scale is more important than ever for customer-facing enterprises. Let's explore the key things that enterprises must keep in mind when meeting scale and performance requirements for customers.


¹ Sean Work, "How Loading Time Affects Your Bottom Line," Kissmetrics Blog, accessed April 17, 2017, <https://blog.kissmetrics.com/loading-time>



ENGINEER FOR PEAK USAGE

Peak usage scenarios may seem few and far between, but they're both the most costly and the most likely times for systems to lag and fail. Imagine a retailer lagging or having an outage on Black Friday, or maybe a tax service provider going down on April 15th. An outage for only five minutes during peak usage times like these could be as costly and damaging to a brand as an outage for hours on any other day.

There are other, less predictable scenarios that may result in peak usage. For example, consider a marketing campaign or other event going viral, or the Fed raising interest rates for financial institutions. In these scenarios, there's no planning runway. Outages during those peak usage scenarios can be just as costly as outages during the ones that are more predictable, but ramping up on a moment's notice isn't always possible.



BALANCE HORIZONTAL AND VERTICAL SCALE

Having the ability to scale both vertically and horizontally is important. Some organizations may be tasked with working in virtualized environments and require horizontal scale, while others may be on higher-performance hardware and have more room to scale vertically. Sometimes, however, a balance is required between the two to get the most reliability from scaling out (horizontally), or to get the most performance from scaling up (vertically). The most important factor is that your CIAM solution can support both approaches.

LIMIT RESOURCE CONSUMPTION

Whether a surge in requests from an app is caused by a denial of service (DoS) attack, peak usage or a bug or oversight, that one app can cause identity services to lag across all of your other applications. It's important to enable centralized policies that can prioritize and preserve the performance of mission-critical applications. Without this capability, a partner or otherwise less-critical application could potentially put the performance of your higher-value applications at risk. Imagine if your new beta application with 10,000 users sends a surge of requests and causes lags or outages in your e-commerce application with 10 million users.



PROVIDE INSTANT ACCESS TO CUSTOMER DATA

Remember a few years ago when solid state drives (SSDs) were the next big thing that would enable your computer to boot up in five seconds vs 30 or more seconds? Storing customer data in memory (versus on disk) provides the same type of speed enhancements as applications access your customer data. The requirement there, especially when storing tens or hundreds of millions of customer identities, is that you have to have a way to efficiently compress and store customer data with sophisticated compression algorithms. Once achieved, all customer identities can be stored in memory, and accessing millions of customer attributes simultaneously is nearly instantaneous. These efficiencies provide customers with blazing fast experiences that drive engagement.

ENABLE HIGH-PERFORMANCE ACTIVE-ACTIVE DATACENTER DEPLOYMENTS

Many businesses have a need to replicate customer identity repositories across geographies. It's widely accepted that multi-master, as opposed to master-slave architectures, is the best approach. Unfortunately, it can still be difficult to meet this requirement while achieving the performance standards, responsiveness and availability necessary to keep customers engaged. When replicating over a wide area network (WAN), you want to be limited only by the latency in the connection, not by additional processing or transmission of unnecessarily large datasets.

Few systems do WAN replication well. They may use a completely alternate mechanism to achieve the WAN replication than what's used on servers on the same local area network (LAN), which can lead to delays of minutes or longer. Or, they may have replication protocols that are excessively chatty, which wastes WAN bandwidth and reduces the maximum modification rate that can be maintained without developing cross-site backlogs. For these reasons, it's imperative to test your replication over WAN for the desired efficiency before you commit to a CIAM solution.

UTILIZE ENTRY BALANCING FOR YOUR IDENTITY DATA

Sharding, or entry balancing, is an important element of achieving scale and performance in very large customer use cases. Entry balancing means putting some identities, or entries, in one directory and others in another. It's important to be able to entry balance based on attributes in identity profiles. You might split identities down the middle and put those with last names beginning with A to M in one replicated repository and the rest in another, or you might choose the repository where identities are stored based on customer geography or another profile attribute. In either case, it's important to expose a unified, virtual view of the entire dataset to applications.

In addition to scale and performance benefits, the ability to store profiles in different geographies, based on a customer's citizenship for example, can help to meet data residency requirements set by privacy regulations such as General Data Protection Regulation (GDPR).



CONCLUSION

Customers are demanding. Even small lags in their experience can cause them to move to a competitor. Meeting the availability and performance service-level agreements (SLAs) required by customer-facing enterprises takes a specific set of capabilities and best practices that include sophisticated data compression, replication efficiency, horizontal scalability and more. It's important to keep these requirements in mind as you plan for customer scale and performance. Some CIAM solutions can meet customer SLAs while others can't. Finding the right CIAM solution can help you meet your scale and performance goals, enhance customer experiences, drive engagement and ultimately grow revenue for your business.