Okta's Architecture

Scaling Okta to Billions of Users





Index



A Paradigm Shift in Scale for Identity and Access Management	3
Our Proven Ability to Rapidly Scale	5
Preparing for Billions of Users	7
Beyond Architecture	11
Okta's Vision and Focus on Customer Success	13

A Paradigm Shift in Scale for Identity and Access Management

Billions of users? No, that's not a typo.

Identity and access management is not just about maintaining a profile for every person on the planet. It's about an identity – for each individual employee, customer or partner – for every organization in the world. How many separate organizations maintain data about your own identity? 20? 50? 100?

When Okta thinks about scale, and where we need to be in the future, we think about it in the tens of billions of users. Our vision is to connect every user, every application, every organization and every device. From the start, we made decisions for Okta's platform architecture to support this broad vision.

This approach brings a sea change from on-premises identity platforms of the past or a build-it-yourself approach. On-premise platforms are expensive, time consuming to set up, and hard to maintain. These platforms are deployed per company, and the onus for scaling them as needed is on that individual customer. Okta's architecture, on the other hand, is designed to scale With dynamically system-wide. proprietary techniques on top of today's leading cloud infrastructure technology, we have designed a platform with the potential for limitless scale.

Three pillars to Okta's secure, always on architecture



Scalability

Capability to automatically handle a growing amount of work and potential to be enlarged to accommodate that growth

Reliability

Ability to perform its intended functions and operations without experiencing failure

Security

Processes, tools and policies to prevent, detect, amd respond to threats





Uptime in 2018

Zero Planned Downtime

Zero Planned Downtime

Uptime in 2019

This is only one part of the equation. Today's users expect a secure, seamless experience while IT and development teams adapt to increasing demand. Interruptions, downtime and security incidents can severely hurt an organization's productivity. That is why Okta views scalability, reliability and security as three equally important pillars of our always on architecture.

By maximizing isolation in a multi-tenant architecture, Okta provides 99.99% uptime and zero planned downtime. In fact, we maintained 99.9999% uptime in 2019, 99.9955% uptime in 2018 and 99.9975% uptime in 2017, even as we scaled 290% in authentications per month. You can see the current status of Okta's availability at any time at <u>status.okta.com</u>.

Along the way, we have always architected Okta for greater usage than needed. At 1 million users, we were ready for 5 million. At 10 million, ready for 50 million. At 100 million, ready for 500 million. We don't plan to stop there. As we onboard customers with greater and greater scale requirements, we have the team and technology in place to get us to even greater levels.

Our Proven Ability to Rapidly Scale

The Okta Identity Cloud is built on the industry's most reliable, secure and scalable platform, period. We knew from day one that we had to be more reliable than anything we connected to, and today we're proud to have a proven track record.

In the past year, the volume of authentications on Okta nearly doubled. Our customer base grew to over 8,400, with an increasing percentage coming from large-scale employee use cases with high daily transaction volumes and greater customization requirements. In addition, we now support more large B2B and B2C use cases than ever before, with even higher user counts and more sporadic transaction patterns. Users now access Okta from all 195 countries worldwide, and thousands of them are accessing applications from countries like the Philippines, South Africa and Peru.

HITACHI MAIOR LEAGUE RASERALL 300,000 Millions **30m** of baseball fans workers retail shoppers Hitachi, a multinational Albertsons, one of the Major League Baseball largest US food and drug (MLB), which relies on technology company, which has over 300,000 retailers, which Okta to handle workers authenticating leverages the Okta considerable across hundreds of platform to serve its seasonality, including domains with Okta more than 30 million authenticating a customers across 18 significant portion of its banners every week 60 million baseball fans during Opening Day

Some customers taking advantage of Okta's scale include:

Preparing for Billions of Users

Today, Okta has hundreds of millions of users on the platform. But that just tells a fraction of the story. These users access millions of applications, representing billions of application identities managed and secured by Okta.

The number of objects in our database is not as much of a limitation as transaction volumes. In particular, authentications are the greatest load on Okta, so we closely monitor the number of authentications we can handle. Today, we authenticate millions of users per hour. Additionally, Okta's service receives billions of web requests per day across API calls, HTTP requests and content delivery network (CDN) requests. These are mission critical authentications across customers, partners, and employees and include requests such as logins to core collaboration apps, MFA triggered by adaptive policies, minting of OAuth 2.0 and OIDC access and identity tokens, provisioning newly onboarded users to downstream apps and real-time deprovisioning of access.

Our engineering team continues to successfully test the platform for massive increases on current loads. They have run controlled tests for individual customer tenants to hold 100 million users with corresponding increases in authentication volume. Even with such high loads, we are still not fully utilizing Okta's scale capabilities.

Since we are aiming for tens of billions of users and authentications, we've continued to optimize our 100% cloud architecture for extreme scale. Beyond using multiple availability zones on Amazon Web Services for redundancy and high availability, along with CDNs to further optimize scale and performance from anywhere in the world, we've built our architecture with proprietary Okta innovations that we call "cells." Each cell is a self-contained instance of the entire Okta service:



Figure 3: Okta now bundles its entire 6-zone AWS architecture into "cells," which can easily be spun up for scale, performance, global footprint and other requirements.

Each cell contains hundreds of automated components, which gives us several advantages:

Risk Mitigation

Any fault in infrastructure is contained within a cell using a redundant High Availability (HA) architecture across multiple zones, so that even if an entire data center goes down, another cell in a different geography can take ownership of affected accounts within an hour.

• Staged Deployment and Rollback

We can rollout code from one node to one cell at a time, or rollback on just one cell instead of the entire service. This decreases the surface area of potential issues that could arise from a code update.

Infrastructure Provider Independence

We have the flexibility to deploy Okta on Google Compute, Microsoft Azure and additional zones and regions of AWS.

Workload Tuning

We've broken many Okta services into different tiers so we can tune them to the various access patterns our customers have. For example, we segment out big jobs and back-end service processing, run our computationally complex hashing algorithm, protect the database from chatty apps like Microsoft Office 365, and process large volumes of API calls, interactive user requests or AD/LDAP agent requests separately.

Horizontal & Vertical Scalability

By adding a cell, we can increase capacity quickly. We can also split a cell to double capacity for tenants on the original cell. In addition, not all tenants are hosted on the same cell, so we avoid the point of diminishing returns on performance.

Geographical Isolation

We can guarantee that your data stays within relevant political borders.

Today, Okta is hosted on 10 production cells and 2 preview cells, including dedicated production and preview cells for Europe and Asia Pacific, and one cell that is HIPAA and FedRAMP compliant. We are capable and ready to roll out cells for other requirements as needed, such as:



Beyond Okta's proprietary cell architecture, we've built extreme redundancy at each layer of the technology stack. Even if a SaaS, PaaS or IaaS offering used by Okta goes down, Okta remains available for its customers. This strategy extends to redundant monitoring and alerting across our infrastructure. The totality of this approach enabled Okta to remain on and functional even when entire AWS availability zones or systems have gone offline.



Okta's Technology Stack Built with Extreme Redundancy at the IaaS, PaaS, and SaaS Layers

Beyond the Architecture

An architecture built for scale is a necessity, but equally important is the people, processes and tools we have in place to test, monitor and actively scale Okta. While we run on AWS virtual infrastructure, we do not just "trust" AWS to deliver the performance our customers need. That's why we actively monitor CPU, memory, disk space and threads on all the load balancers, web application servers, job servers, cache servers and database servers that Okta runs on. In addition, we track the ongoing performance of the application, including authentication performance, response codes, agent health, job statistics and application errors.

As this data comes in and is interpreted, Okta's systems automatically deploy more servers as needed and tear down servers that are not performing up to par – all without any disruption to the service. Our operations team looks at the graphs and metrics constantly to watch for bursts of traffic and triage them. This also informs a monthly capacity planning loop, during which we decide if we need to upsize or create a new cell. In addition, we do sophisticated long-range planning to handle growing customer needs based on years of customer experiences across a variety of use cases.

These processes depend on the deep experience of the Okta operations team to understand this data and take the correct action to maintain high performance as we scale. The team is staffed 24/7, so we always have an expert monitoring the service and ready to take action as needed.

Third-party certifications are an important validation of our ability to scale securely, and Okta was the first identity provider to achieve Level 2 CSA STAR attestation. We've also obtained ISO 27001 certification for our information security management system, and are committed to expanding upon these security certifications. You can see Okta's service certifications at https://trust.okta.com/compliance.

Okta Service Certifications



ISO 27001:2013

ISO 27018

ISO 27018:2014



SOC 2 Type II





CSA STAR

FedRAMP - Authority to Operate (ATO)

Helping you meet your compliance requirements





PCI-DSS 3.2



Sarbanes Oxley (SOX)





Secure, always on architecture

It is the combination of our strategies surrounding these three important areas, as well as our focus on the intersection of our pillars, that enables Okta to deliver the most reliable, scalable and secure identity service to our customers.



Okta's Vision and Focus on Customer Success

As we described above, Okta has seen a remarkable increase in customers and usage over the last several years. Meeting these demands is not simple. It takes the right team, the right architecture, the right processes and probably most of all, a complete focus on customer success.

That's why customer success is Okta's number one priority as a company. As we've scaled up and brought on new customers that have pushed us to new volumes, we have worked closely with them to ensure their deployments and go-live dates are smooth and as uneventful as possible. We always collaborate with our customers and partners to be ready to respond and react to those events that are just not possible to predict. And in turn, we have learned from our customers to optimize how we prioritize our processing of identity events in order to build the most secure, seamless experience for all Okta users.

Okta's architecture is markedly different from the old approach of a separately scaled infrastructure for each customer. It is far more powerful, more resilient and more scalable, but it must be managed by the right team. It's this combination of technology and people that has made Okta a leader in identity and access management for workforce and customer identity use cases.



Hector Aguilar President, Technology @ Okta

Jector quilar



Jon Todd Chief Architect @ Okta

for fodd