

# Not All Cloud Services Are Built Alike



# Content

|    |                                 |
|----|---------------------------------|
| 2  | Introduction                    |
| 4  | Protection Throughout the Stack |
| 10 | Conclusion                      |

# Introduction

Nothing is more critical than trust when you've decided to adopt cloud services supporting mission-critical systems and capabilities. You need to trust that the cloud vendor is available, scalable, resilient, reliable, and as committed to security and compliance as you are.

This is especially true when it comes to workforce identity and access management (IAM). Your workforce IAM infrastructure essentially serves as the central point through which your team monitors and secures access to all your applications.

The stakes are high, and hosting core workforce IAM infrastructure on a cloud platform can sometimes feel like a leap of faith, especially for organizations that are accustomed to managing their own proprietary identity infrastructure. If on-premises software or servers experience downtime or an outage, IT admins can take immediate action to solve the problem internally—a scenario not always possible when relying on external third parties.

Furthermore, organizations in highly regulated industries are often cautious when considering cloud infrastructure, as they may be required to report to regulators any situation where the security and privacy of data is compromised, or when critical information systems are rendered inaccessible. In such cases, it's only natural for organizations to want to maintain full control over their IT functions.

However, the benefits of outsourcing IT infrastructure to cloud partners have been proven repeatedly. It doesn't just significantly reduce the total cost of operations and ownership, it also provides superior reliability and security compared to legacy, on-prem platforms.

Therefore, the question you should be asking isn't whether you should work with a cloud vendor to supply your workforce IAM capabilities; the question is which vendor you should work with. Who can you trust to deliver constant uptime, with scalability and flexibility, and the security and protection your users deserve?

Okta's Workforce Identity Cloud was designed to foster trust. We've built the software, operationalized the processes, and hired the talent it takes to offer customers a suite of best-in-class solutions that adhere to the highest standards. Okta is:

- Built for scale: the service scales up and down seamlessly according to your needs;
- Always available: the service is designed for zero downtime. No maintenance windows required;
- Secure: the service is more heavily audited than anything that can be built and maintained in-house; and
- Constantly innovating: the service enables new capabilities and rapid innovation, and insulates your organization from the constantly changing IT landscape.

This technical whitepaper offers a comprehensive overview of the software and operational architecture that enables Okta Workforce Identity Cloud to run a scalable, highly available, on-demand IAM service with 99.99% uptime.

# Protection Throughout the Stack

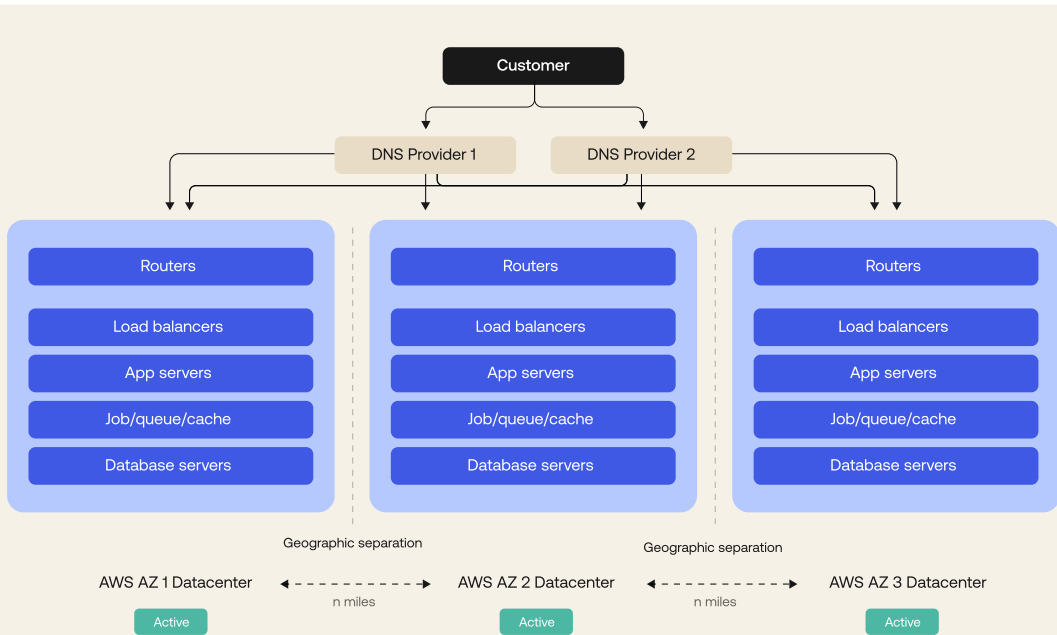
A highly available service begins with incorporating redundancy at every step.

## Resilient architecture

One of the most critical aspects of Okta’s WIC architecture is that it is completely multi-tenant. With this design, customers share the same underlying environment. This creates an economy of scale, allowing Okta to make the infrastructure extremely robust in terms of redundancy, monitoring, and processes.

Each Okta environment is called a cell. Okta cells contain all of the infrastructure required to operate an instance of WIC. More information about cells can be found in the An Insider Look: How Okta Builds and Runs Scalable Infrastructure whitepaper. Because each cell can operate independently, they form the basis of our availability strategy by helping to fully isolate each cell during outages.

Within each cell, we run a fully optimized system consisting of an Edge, containing routers, proxy load balancers and firewall services, an application tier where our software runs, and a set of database services. Everything is hosted in Amazon Web Services (AWS) or Google Cloud Platform (GCP) across multiple availability zones in an active-active-active configuration and geographically separated regions. The service is designed for high scale, high throughput, and high availability.



### Stateless

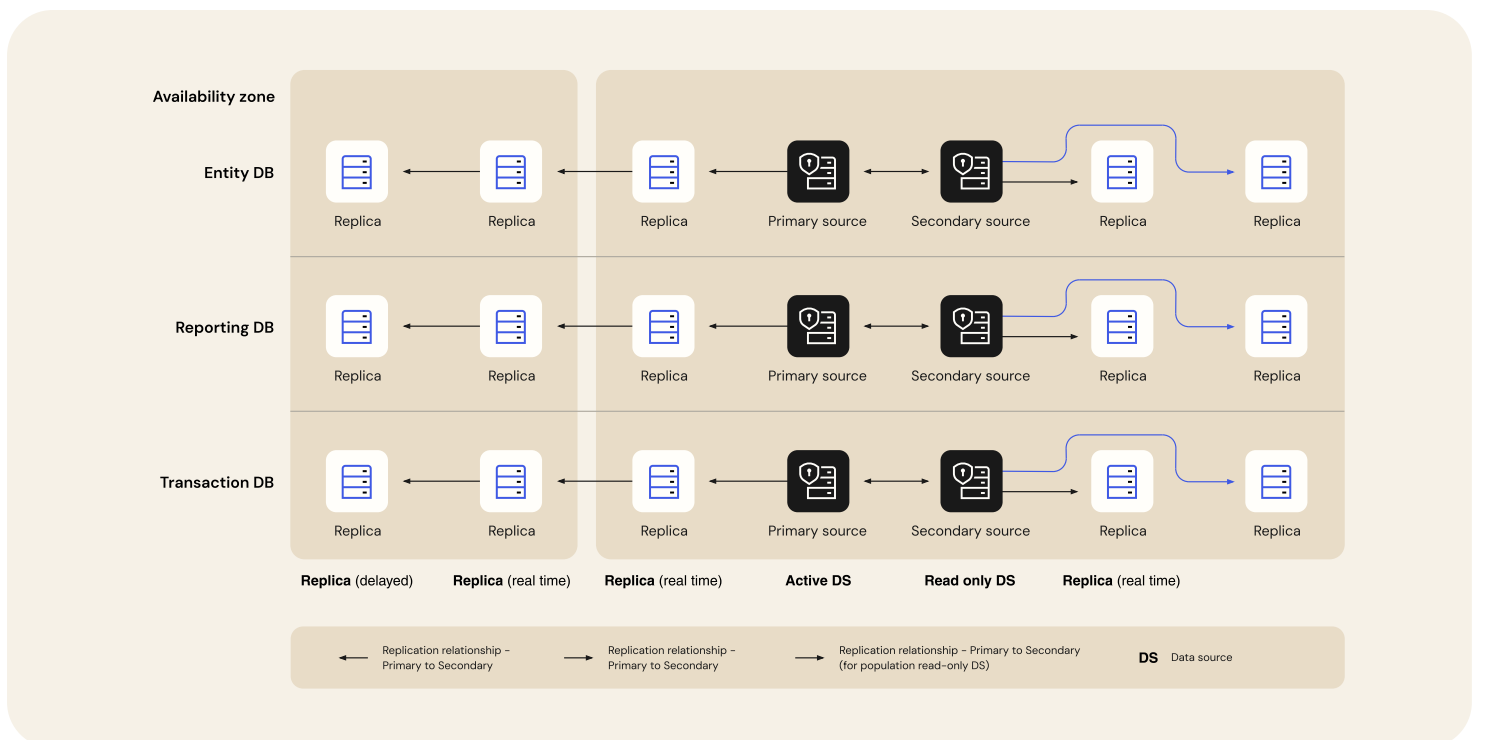
Except for databases, all components are functionally stateless. As a result, above the database tier, any server in the stack can handle any request since each request is independent of one another

That means that all of the components of our system can be scaled up at will simply by spinning up new instances in AWS or GCP. This allows Okta to automatically adjust to changes in demand and route around failed services to other active systems.

Requests are queued and processed by available application servers so that application server failures are transparent to users. If an application server fails, the authentication is automatically forwarded to a functioning server and retried.

### Functionally optimized databases

We run a highly available configuration with read replicas so there is no single point of failure. If one primary database goes down, another is promoted. Replicas are live across multiple availability zones in two distinct regions for both AWS and GCP. Based on the current scale and resiliency requirements, the number of read replicas can be seamlessly increased as needed.



## **No single points of failure**

### **Subprocessor failover**

We're not perfect, but we're working towards it. In 2016, Okta suffered an outage due to a vendor who fell victim to one of the largest distributed denial-of-service attacks in history. Okta's customer support team scrambled to help our customers restore service, and afterwards, a full root cause analysis (RCA) was performed to identify any other single points of failure that may exist in the service. As a result, Okta identified and implemented backup vendors for our sub processors that help us serve our customers' critical identities. From SMS and voice authentication to round-the-clock customer support, Okta is ready to react. As part of the vendor risk-assessment program, Okta also assesses sub processors for disaster recovery readiness.

### **Employee failover**

Recovery of the service also requires contributions from Okta team members who are able to support it. Okta's disaster recovery and business continuity plans therefore include provisions to ensure that, in the case of a regional disaster, employees are geographically distributed. Okta's Customer Support team also operates on a "follow the sun" model, ensuring seamless, round-the-clock assistance is available to customers. Each team also maintains multiple methods of communication, so in the event of concurrent loss of production and corporate services, engineering teams are able to communicate and focus on restoring the service.

### **Data backup**

Replication across multiple availability zones and regions ensures that Okta is able to function as a highly available service, but it does not provide adequate protection in the event of a catastrophic failure of AWS or GCP infrastructure or data corruption. Therefore, Okta also takes frequent database snapshots and stores them in a read-only, protected account. These encrypted, protected backups serve as a failsafe that enables Okta to restore the service across fresh infrastructure, if needed. Restoration is supported at a cell level and not tenant level.

## **No single points of failure**

### **Subprocessor failover**

We're not perfect, but we're working towards it. In 2016, Okta suffered an outage due to a vendor who fell victim to one of the largest distributed denial-of-service attacks in history. Okta's customer support team scrambled to help our customers restore service, and afterwards, a full root cause analysis (RCA) was performed to identify any other single points of failure that may exist in the service. As a result, Okta identified and implemented backup vendors for our sub processors that help us serve our customers' critical identities. From SMS and voice authentication to round-the-clock customer support, Okta is ready to react. As part of the vendor risk-assessment program, Okta also assesses sub processors for disaster recovery readiness.

### **Employee failover**

Recovery of the service also requires contributions from Okta team members who are able to support it. Okta's disaster recovery and business continuity plans therefore include provisions to ensure that, in the case of a regional disaster, employees are geographically distributed. Okta's Customer Support team also operates on a "follow the sun" model, ensuring seamless, round-the-clock assistance is available to customers. Each team also maintains multiple methods of communication, so in the event of concurrent loss of production and corporate services, engineering teams are able to communicate and focus on restoring the service.

### **Data backup**

Replication across multiple availability zones and regions ensures that Okta is able to function as a highly available service, but it does not provide adequate protection in the event of a catastrophic failure of AWS or GCP infrastructure or data corruption. Therefore, Okta also takes frequent database snapshots and stores them in a read-only, protected account. These encrypted, protected backups serve as a failsafe that enables Okta to restore the service across fresh infrastructure, if needed. Restoration is supported at a cell level and not tenant level.



## **Identity that's always on**

### **No maintenance windows**

Traditionally, the ability to provide a highly available service is in direct conflict with delivering continuous innovation to that service. At Okta, we knew we couldn't make that compromise. By combining a canary-based deployment architecture with an automated testing and deployment process, we can deliver continuous innovation with no planned downtime or maintenance windows required.

Any development committed to our mainline code kicks off with a new build and a series of unit, functional, and security tests. Once the fully automated process completes and ensures that the code has passed our rigorous testing, the software is ready to be deployed.

The operations team then deploys the fully tested code to our customer-facing environments. Product updates are first published into the sandbox environment to allow customers to perform their own acceptance testing. After one week, the changes are published to production. Throughout this rolling deployment process, we ensure upgraded app servers only talk to upgraded versions of the data backup, avoiding the risk of data corruption.

### **Proactive monitoring**

Disasters are never caused by one single event; they typically come about through a chain of events that must all happen together. Identifying and stopping this chain early is key to a successful high-availability program, and that's why Okta invests heavily in proactive monitoring of the service.

The first layer of defense against the unexpected is robust instrumentation and monitoring for all components of the system. We split this into two categories: external monitors and internal monitors.

For external monitoring, Okta uses multiple services with globally distributed test agents to constantly monitor our application. This provides us with a constant feed of real data on how our service is operating. Because we are fully multi-tenant, the data is real and applicable to all of our customers.

We also use internal monitors to show us when things are having problems, but more importantly, we use them to show us what systems are having problems. They are more sensitive than the external monitors, so they frequently warn of problems before they affect site performance or availability. Okta uses internal monitors on all subsystems and software components for maximum visibility.

### **Rapid disaster recovery**

As the market leader in IAM, Okta is used to control access to all types of mission-critical data for our customers, and this data must be available all the time. Okta's architecture is therefore designed to be highly resilient and always available. But even the best systems can fall to unforeseen events. In these cases, Okta provides a one-hour recovery time objective (RTO) for authentications and read-only actions and a 24-hour RTO for full restoration of the service, which exceeds the availability of many other cloud services. This way, as we deal with an incident, you don't have to worry about data loss. For customers that require a faster recovery time, Okta offers Enhanced Disaster Recovery that reduces the RTO to five minutes for authentications and read-only actions. You can learn more about Enhanced DR [here](#).

Backing up your data is only half of the battle, and the middle of a disaster is not the right time to test whether you can successfully restore your service. That's why Okta performs disaster recovery testing at least quarterly and continuous testing for [Enhanced Disaster Recovery](#) customers. These tests ensure that when the team is required to respond to an emergency, they will have honed the required skills and the data will be ready.

Disaster recovery is a core control of compliance programs like PCI, SOC2, and ISO 27001, each of which require evidence of regular testing for in-scope systems. To help our customers demonstrate their ability to meet these requirements, Okta publishes attestations for each of our tests through a self-service documentation portal. These attestations include when the test was performed, its duration, the result, and any actions Okta is required to take to ensure your data is ready when you need it.

In addition, Okta publishes statistics on [status.okta.com](https://status.okta.com), which includes not only the current status, but the history of any outages and performance issues that affect the service, whether or not they were caused by Okta. We also provide detailed information about the issue and what Okta is doing to prevent it from happening again.

## Conclusion

At Okta, we understand that when your identity system goes down, business stops. That's why we've designed our system to be resilient to the known and the unknown, and we stand behind that statement with industry-leading service level agreements and transparency. Across the entire system lifecycle, from design to deployment, maintenance and disposition, Okta is built to be always on.

### **About Okta**

Okta is the World's Identity Company. As the leading independent Identity partner, we free everyone to safely use any technology—anywhere, on any device or app. The most trusted brands trust Okta to enable secure access, authentication, and automation. With flexibility and neutrality at the core of our Okta Workforce Identity and Customer Identity Clouds, business leaders and developers can focus on innovation and accelerate digital transformation, thanks to customizable solutions and more than 7,000 pre-built integrations. We're building a world where Identity belongs to you. Learn more at [okta.com](https://okta.com).