

Zendesk Scalability and Performance

At Zendesk, our goal is to help bring companies and their customers closer together—with software that’s easy to use, easy to customize, and easy to scale. To help make customer conversations easy and more productive, it starts with ensuring Zendesk runs smoothly while providing customers with a service that scales to meet their ever-changing needs.

As our customers grow and their needs change, the technology that they depend on everyday needs to scale to meet demand for performance and reliability. Zendesk is built upon an enterprise-level operations and technology architecture that exceeds industry standards and future-proofs your business.

A Redundant and Distributed Environment Built For Global Operations

Zendesk uses a fully redundant, distributed, and automated environment consisting of geographically separate data centers running multiple, self-sustaining instances of the Zendesk application. To ensure high performance for our customers around the world, Zendesk runs on four data centers, with a fifth one coming online soon.

Each data center increases capacity and reduces latency for customers. Multiple transcontinental data centers reside in the United States and Europe. Zendesk offers customers the option to select their preferred data center location.

High Flexibility Yields High Performance With A Pod Architecture Design

Zendesk is built on top of a flexible and efficient data center network. Each data center consists of one or more pods. This pod architecture presents maximum flexibility with the benefits of faster provisioning, user customization, and improved scalability. A new Zendesk customer is assigned to a specific pod in the data centers. Each pod has all the resources necessary to run Zendesk independently of the other pods. The resources include application servers, web servers, and database servers, as well as resources to handle other work such as email processing and reporting.

Each pod contains several database clusters—each consisting of a master database server and two database servers that continuously mirror the master. If the master unexpectedly encounters a problem, one of the spares immediately takes its place. This automated system of backup servers significantly reduces the possibility of downtime.

A Sharded Database For Maximum Scalability

The building blocks of each database cluster are shards, which are smaller, logical databases. A sharded architecture allows Zendesk to rebalance load across physical database clusters and pods—enabling a higher level of performance and scalability. Today, Zendesk has hundreds of shards distributed across database clusters worldwide. Each shard supports a certain number of customers. The number varies depending on the data volumes generated by the customers. If a customer demands more capacity, Zendesk can fully dedicate a single shard or a database cluster to the customer.

ZENDESK OPERATIONS ON A GIVEN DAY



SCALE
3.2 million tickets

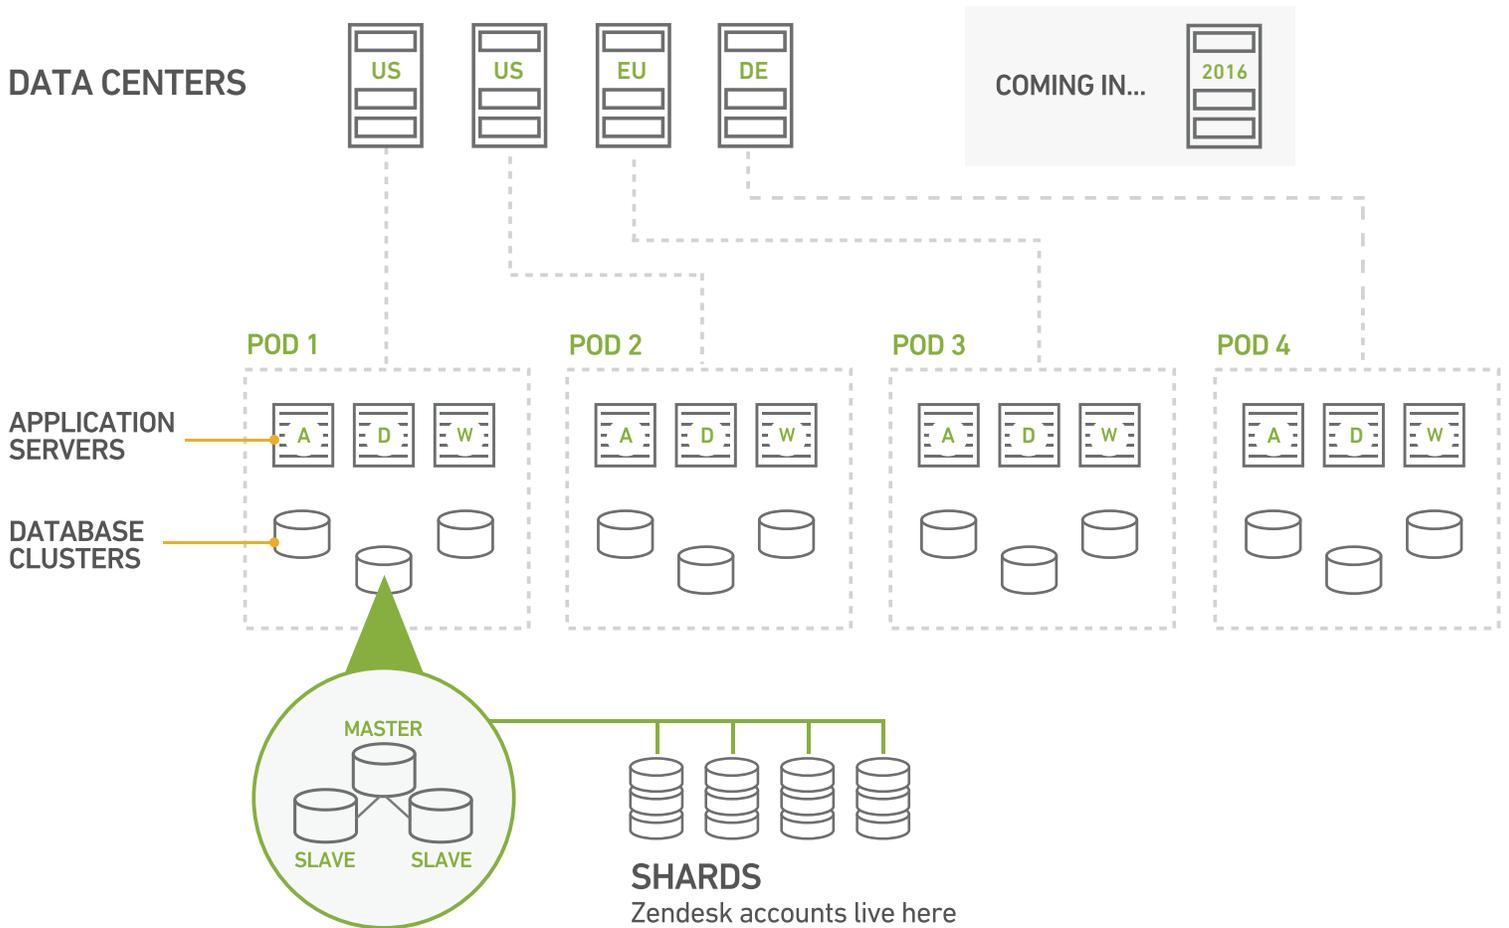


PERFORMANCE
865 million API calls



EFFICIENCY
Average request response time:
158 millisecond or 1/6 of a second

Data Center Technical Architecture



ABOUT ZENDESK

Zendesk is a cloud-based customer service platform. It is designed to be easy to use, easy to customize, and easy to scale. Frustrated with the state of enterprise customer service software in 2007, three Danish entrepreneurs sought out to create beautifully simple software that could change how any company interacted with its customers. Today more than 65,000 companies use Zendesk. Learn more at www.zendesk.com.