

Zeta Interactive AWS Migration One Byte at a Time



Executive Summary

[Zeta Interactive](#) is a Big Data and analytics company that helps leading brands acquire, grow, and retain customers. Founded by David A. Steinberg and John Sculley (former CEO of Apple Computer and Pepsi-Cola) in 2007, Zeta uses Big Data, advanced analytics, and machine learning to power marketing programs for hundreds of middle market and Fortune 500 brands. Zeta Interactive has rapidly grown into a recognized leader in the emerging [Big Data](#) and Customer Lifecycle Marketing sectors. With over 600 employees worldwide, the company is headquartered in New York City, with Centers of Excellence in Silicon Valley, London, and Hyderabad, India, and additional offices in Boston, MA; Boulder, CO; Detroit, MI; Scottsdale, AZ; Middlesbrough, Bristol, UK, and Chennai, India. Zeta was named by Forbes as one of America's 50 Most Promising Private Companies in 2014.

The Challenge

Originally built as a digital marketing service, Zeta pivoted four years ago to focus on providing customers with a software platform to manage and analyze relationships and deploy marketing campaigns. It reported 2014 revenue of more than \$100 million and projects that will more than double this year. In mid-2015, Zeta raised \$125 million in funding to accelerate acquisitions and fuel innovation.

With acquisitions on the horizon and an aggressive growth trajectory, this software-as-a-service (SaaS) provider needed assurance that their infrastructure can scale to meet their customer demands. This would have needed significant new CapEx spend on hardware to increase capacity of their infrastructure co-located at an Equinix datacenter. Zeta decided that it was time for [AWS migration](#) to take advantage of the scalability, high availability, and OpEx financial model. Zeta determined that a phased migration was the best approach to meet their needs.

"We have had a fantastic partnership with ClearScale for some time now. The ClearScale team of experts has risen to every single challenge and are highly knowledgeable, exceptionally dedicated, and responsive."

Bharat Goyal, Senior Vice President, Engineering. Zeta Interactive

The ClearScale Solution

Zeta Interactive partnered with ClearScale to design an appropriate architecture for their [AWS deployment](#) and support them through the migration process. ClearScale worked closely with Zeta to understand their current processes and technology. ClearScale designed the network architecture and developed automation scripts so that Zeta could integrate with [AWS services](#) such as Auto Scaling for failover. Puppet was also implemented for user management tasks. ClearScale and Zeta have been working hand in hand to migrate components in stages to AWS.

Network Architecture

Availability, redundancy, and security are major requirements for Zeta's SaaS product offering. ClearScale designed a network architecture to serve as the foundation for meeting these requirements. The production environment is deployed in a single Virtual Private Cloud (VPC) distributed across three Availability Zones (AZs). This mitigates risk related to data center outages. A second Production VPC is deployed in another AWS Region as a standby site. This enables Zeta to failover services to the secondary AWS Region in the event that the primary AWS Region becomes unavailable, as would be required for Disaster Recovery (DR).

Three subnets are provisioned within each AZ: DMZ (Public) subnet, Main (Private) subnet, and Database (Private) subnet. Each subnet has specific security configuration and are managed independently. The Development environment is deployed under a separate AWS account for security and to prevent any cross contamination into Production.

Production and Development VPCs are connected to Equinix (Zeta's main datacenter). Zeta uses AWS Direct Connect to provide a dedicated 10G connection between the primary Production VPC in AWS and the Equinix datacenter.

Automation

In the past Zeta utilized a great deal of manual process for deploying services. Cloud automation is required to integrate with AWS services and provide scalability and failover capabilities. ClearScale reviewed Zeta's manual processes and developed a series of automation scripts to build a bootstrap process for deploying services. This enabled Zeta to take advantage of AWS Auto Scaling services and integrate with existing Zeta infrastructure.

Components of the bootstrap process:

- Configure hostnames of the Elastic Compute (EC2) instances
- DNS registration with Route53 and integration with on-prem DNS solution
- Setup initial users on the EC2 instances
- Attach Elastic Block Storage (EBS) volumes to EC2 instances

The bootstrap process is a key component that enables Auto Scaling failover capabilities. For example, if an EC2 instance is down in the first AZ, Auto Scaling will initiate the bootstrap process to create an EC2 instance in the second AZ. To preserve the data from the original instance, a snapshot of the EBS volume is taken, moved to the second AZ, and mounted to the new instance. Service is returned to normal operations.

Puppet is a configuration management application that can automate the provisioning, configuration, and ongoing management of servers and software. Zeta's Puppet implementation is primarily used for centralized user management. Users, access groups, and SSH keys are defined for each environment using YAML files. AWS metadata tags are assigned to every EC2 instance deployed. This identifies the environment membership (e.g., Production, Staging, Development). Puppet agents are automatically installed and configured as part of the bootstrap process for deploying EC2 instances. The puppet agents use the metadata tag assigned to the instance for creating user resources on the server. Subsequent puppet runs continue to manage user creation and deletion as the YAML files are updated on the source.

A Phased Methodology for Migration

Zeta implemented a phased approach for their migration. ClearScale designed and implemented a robust network architecture connected to Zeta's datacenter in Equinix. ClearScale supported the migration by developing the automation frameworks for integration and deployment. ClearScale also built Proof of Concepts (PoCs) for components such as PostgreSQL RDS and MySQL RDS, so that Zeta could perform extensive testing prior to migration. Zeta migrated individual system components in phases and consulted ClearScale for best practices. Each service is built, tested, and cut over independently, providing a methodical approach for migration. This approach allowed Zeta to build confidence in their AWS deployment and bring their operations teams up to speed through hands-on experience, each step of the way.

The Benefits

Zeta Interactive now has the capability to expand their infrastructure in response to the aggressive growth schedule of the business and deliver highly available services to their SaaS customers. Zeta has achieved a higher level of operational efficiency and service provisioning through cloud automation. Leveraging AWS to implement a multi-AZ, multi-Region infrastructure design provides Zeta and their customers' business continuity in the event of disaster. Zeta can focus their efforts on delivering new products and features that drive their customers' growth and success.