

Driving Frictionless Business enabling Global Expansion ensuring Customer Experience at Optimized Cloud Infrastructure Costs

How a publishing major achieved consistent customer delight with 40% less cloud infrastructure costs

Overview

The journey to build frictionless business starts with organizations migrating their systems and applications from on-premise to Cloud for significantly improved business agility, scalability and compelling customer experience. However, frictionless business can only be achieved if the Day after Cloud is managed efficiently. Our digital publishing client was expanding its customer base from 0.5 million to 1 million as part of their global business expansion programme. While they had migrated their IT Ops to the Cloud, they continued to face challenges in controlling costs or ensuring that their Business - Critical Applications were always on. They looked to us for help in establishing the Day after Cloud best practices to mitigate these challenges.

Challenges

In the pursuit of expanding its business internationally, we helped our client in the migration of its entire IT estate to the Cloud. This was a multi-Cloud and multi-region environment; multiple zones were set up to cater to 0.5 million users (teachers, students, support staff and sales community) through different business applications. The next critical step in this journey was the seamless management of the entire IT estate (Applications & Infrastructure) in Cloud. Our client wanted to ensure that their applications were always available to manage increasing user base globally, and creating the Best Practices for smooth Cloud Operating Model, optimizing cloud costs and usage. This was an extremely complex task as it entailed:

- Creation of independent zones across various AWS data centers (US, Ireland, Sydney, Seoul, China) to cater to countries in these geographies
- Time zone differences and therefore IT maintenance was challenging
- Compliance and adherence to business and Data Privacy Laws for each country
- Maintaining consistency of user experience across the exponentially increasing user base
- Always-on availability of applications with capability to handle increasing user loads
- Reduction of annual cost of IT infrastructure which was high despite using Cloud environment
- Resolving hundreds of EC2 and RDS instances in production and non-production environments quickly, to maintain overall efficiency of Cloud Ops



Our Approach and Solution Delivery

Our consultative approach focused on identifying the pain points keeping in mind the following objectives:

| | |
|---|---|
| Our team took the following three initiatives to achieve these objectives | 1 Server consolidation to drive efficiency in server usage and optimize cloud infrastructure costs |
| | 2 Application performance monitoring tools to improve application performance |
| | 3 Creating best practices and discipline in the operating model for optimal usage of the cloud infrastructure. |

Server Consolidation: For optimal usage of server capacity for different applications, we decided to set up multiple applications on a single EC2 instance in each geography and zone. Similarly, all the application dbs were setup under a single RDS instance per zone. This strategy helped in drastic reduction of EC2/RDS instances.

Application Performance Improvement: Application performance tool monitored application performance with detailed data analysis that helped identify inefficiencies in API implementation and DB operations. Various techniques (eg. lazy loading, caching, query optimization, batch processing, asynchronous operations) were applied to improve the code which resulted in drastic reduction in CPU consumption, memory and DB utilization and improving Application performance. Most APIs now respond with a transaction speed of 50ms.

Disciplined Approach for Driving Optimal Usage: The following measures were initiated by the Relevance Lab team to streamline and optimize the use of Cloud infrastructure to bring down costs:

- Weekly utilization and performance reports
- Evaluation of newer instances
- Instances and RDSs having predictable usage throughout the year were brought under reserved instances from the on-demand mode
- Modification of Operations strategy, for e.g., eliminating swim lanes, optimizing back-up schedule and conducting regular audit performance resulting in decreasing load balances, storage and data transfer and Cloud front distribution.

Results and Benefits

58%

EC2 and RDS instances

41%

Hosting cost reduction

65%

Reduction of IT infrastructure cost

100%


Availability of all applications

1 million customers

Experienced application performance and reduced cost of Cloud infrastructure.

Reduction in Cloud Infrastructure Costs

Our consultative approach focused on identifying the pain points keeping in mind the following objectives:



About Relevance Lab

Relevance Lab is a next-generation technology services company — **Born in the Cloud, Transformed by Data, and Re-Imagined for GenAI.** We help enterprises cut through the noise of rapid AI adoption by focusing on business outcomes, building on enterprise-grade foundations, and partnering through every step of their AI journey. With deep cloud expertise, strong product engineering DNA, and proven platforms in DevOps and Research Computing, we enable organizations to drive frictionless business and confidently evolve into truly AI-Powered Enterprises.