



Case Study

Cloud Design & Migration

Dated: Feb-12-2024

| | |
|--|-----------|
| Executive Summary | 3 |
| Problem Statement | 3 |
| Our Solution | 6 |
| High Level Architecture | 6 |
| Key components of New Architecture | 7 |
| CloudFlare: Firewall / DDOS Protection | 7 |
| AWS Elastic Load Balancer (ELB) | 7 |
| AWS Auto-scaling Group (ASG) | 7 |
| AWS Elastic Compute (EC2) | 7 |
| AWS Relational Database (RDS) | 7 |
| AWS Elastic File System (EFS) | 7 |
| Redis | 7 |
| Ci/CD Pipelines | 8 |
| Improvements Observed | 9 |
| Memory Utilization | 9 |
| CPU Utilization | 9 |
| Database CPU Utilization | 10 |
| Database RDS Connections | 10 |
| Conclusion | 12 |
| Need Server / Cloud Optimization or Management | 12 |

Executive Summary

The customer, a well-known E-Commerce B2B information portal, was facing performance problems for their web portal. Additionally, their current architecture was based on the dedicated servers with fixed resources.

When they approached us for a solution, we used our proven mechanism of:

1. Getting the requirements. We call this step 'Customer's goal list'
2. Review of the current system. This step is called 'Establish the baseline'.
3. Design and implement the solution. Another name for this step is 'Implement, test, close'

We created an auto-scaling solution on the cloud that cost half of what the client was paying before. It also allowed them to handle much more traffic optimally. Following table shows a quick comparison of different metrics between before and after:

| Metric | Before Implementation | After Implementation |
|-------------------------------|-----------------------|----------------------|
| Cost per month | \$4,800 + Tax | \$2,370 + Tax |
| Load time (static page) | 7 seconds | 2.5 seconds |
| Load time (DB-heavy page) | 23 seconds | 4 seconds |
| Max visitors concurrent users | 50 | 400 |

Problem Statement

Our customer is a B2B Ecommerce information portal specializing in robotics and automated systems for assorted verticals. Company was established in 2020 and is headquartered in the United States of America.

Their information portal was based on Drupal and deployed on an architecture composed of dedicated servers. However, during the periods of high traffic, the portal would be sluggish and frequently would result in 50x errors (meaning the requests were not processed completely).

These machines were also expensive to run and adding additional machines would result in considerable OPEX (operating expenses).

Key points can be summarized as follows:

1. Drupal is a database heavy technology and the existing portal based on an older version of Drupal
2. Static content was not served using a CDN

3. Caching was not employed
4. Architecture based on the dedicated machines were not capable of auto-scaling*
5. The OS of the Dedicated machines was old (at least by one major release version). Older version of PHP and libraries were also found
6. Static page load time was around 7 seconds
7. Heaviest web application page had a loading time of 23 seconds
8. Cost per month was over \$4,800 + tax
9. Static CSS + JS files were not optimized (over 20 includes)
10. Images were not optimized and sprites were not used (development related optimizations).
11. CI/CD Deployment process was not optimized. Deployment to production would result in a 20-30 minute outage.
12. MySQL Memory consumption is 2% and remains constant (looks like DB optimization is missing). It means DB would chew up CPU by running queries instead of serving content out of its cache.
13. Production web server and database server were coupled on the same machine. This would result in RAM starvation as load increases
14. Log rotation was off and this resulted in ridiculously large logs.

```

/var/www/...com/logs$ ls -l
total 27332484
-rw-r--r-- 1 root root 4099132 Oct 12 15:32 a.log
-rw-r--r-- 1 root root 143412413 Nov 19 16:36 proxy_access_r_log
-rw-r--r-- 1 www-data www-data 27321666129 Nov 19 16:38 proxy_access_ssl_log
-rw-r--r-- 1 www-data www-data 98814279 Nov 19 15:45 proxy_error_log
-rw-r--r-- 1 root root 9546379 Nov 19 16:06 proxy_error_r_log
-rw-r--r-- 1 root root 410864448 Oct 12 15:18 o.log
/var/www/.../logs$
    
```

15. PHP configuration for the FPM component was not optimized. It was taking too much of the CPU power.

```

top - 16:30:32 up 52 days, 20:25, 2 users, load average: 0.57, 0.60, 0.76
Tasks: 262 total, 2 running, 161 sleeping, 0 stopped, 0 zombie
%Cpu(s): 12.4 us, 0.1 sy, 0.0 ni, 87.4 id, 0.1 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 32939804 total, 4945720 free, 2843976 used, 25150108 buff/cache
KiB Swap: 999420 total, 996860 free, 2560 used, 29554856 avail Mem

  PID USER      PR  NI   VIRT   RES   SHR  S  %CPU  %MEM    TIME+  COMMAND
 22743 www-data  20   0 958888 741736 84056 R 100.0  2.3   6:41.31 php-fpm
 17756          20   0  44232   4164  3376 R   0.3  0.0   0:00.03 top
 22861 gitlab-+  20   0 2358480 658224 14828 S   0.3  2.0   5100:47 mysqld
 25923 root      20   0 3046180 86824 54284 S   0.3  0.3   1:54.00 dockerd
    1 root      20   0  78036   9092  6544 S   0.0  0.0   1:13.97 systemd
    2 root      20   0     0     0     0 S   0.0  0.0   0:01.44 kthreadd
    4 root      0 -20     0     0     0 I   0.0  0.0   0:00.00 kworker/0:0H
    6 root      0 -20     0     0     0 I   0.0  0.0   0:00.00 mm_percpu_wq
    7 root      20   0     0     0     0 S   0.0  0.0   0:26.21 ksoftirqd/0
    
```

16. Old builds were not removed and hence took up a lot of space.

```
@vps298:/var/www/██████████.com/builds$ ls -l
total 20
drwxrwxr-x 19 gitlab-runner gitlab-runner 4096 Oct 25 22:18 395172837
drwxrwxr-x 19 gitlab-runner gitlab-runner 4096 Oct 27 22:11 396903067
drwxrwxr-x 19 gitlab-runner gitlab-runner 4096 Oct 28 22:12 397828072
drwxrwxr-x 19 gitlab-runner gitlab-runner 4096 Nov 11 22:12 407345414
drwxrwxr-x 19 gitlab-runner gitlab-runner 4096 Nov 18 22:11 412082852
@vps298:/var/www/██████████/builds$ █
```

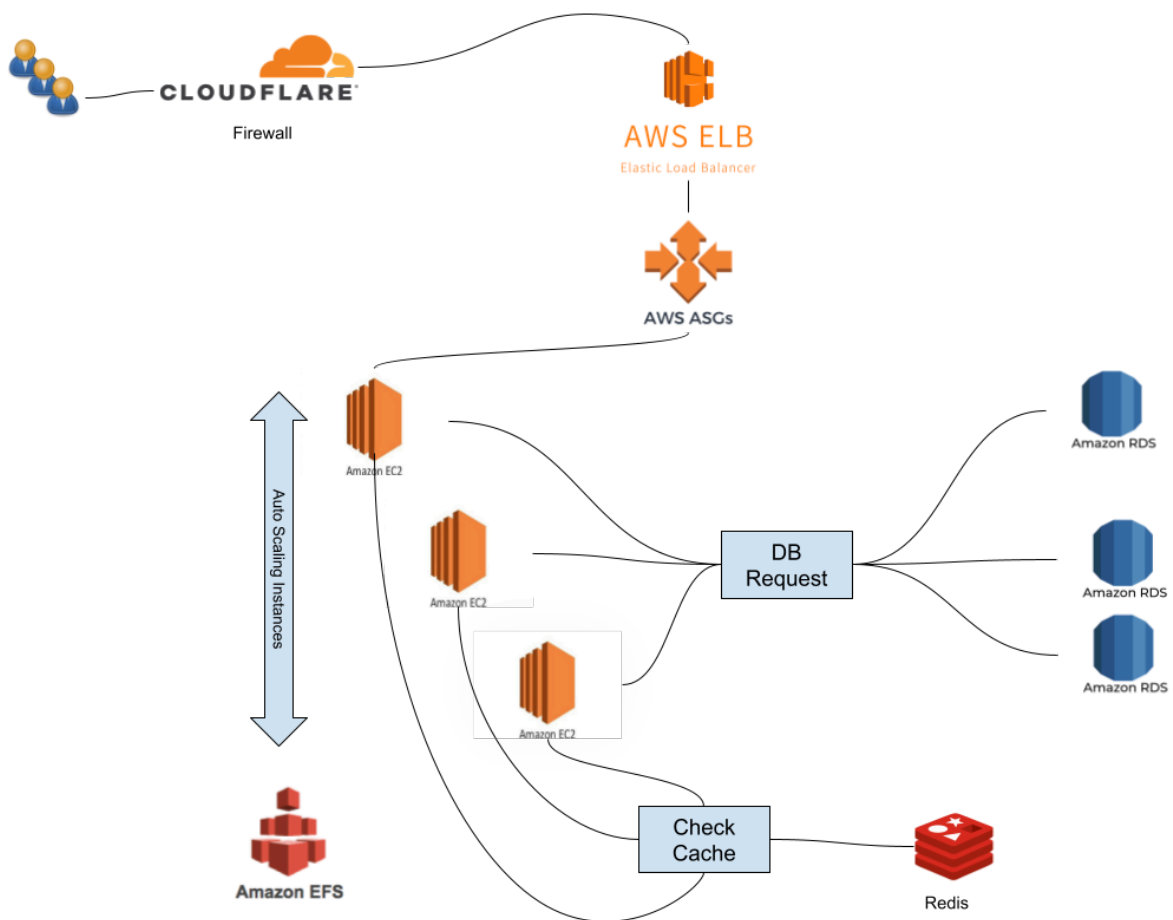
Our Solution

In order to address issues listed above, we devised a comprehensive plan for the creation of new architecture and migration of the existing website, email and media. During the design process, we considered following key items:

1. Robustness
2. Scalability
3. Security
4. Dynamic scaling
5. Cost effectiveness
6. Disaster Recovery
7. Using managed services
8. Comprehensive alerting mechanism
9. Agile Project management approach

High Level Architecture

The following diagram shows a high level architecture for our solution.



Key components of New Architecture

CloudFlare: Firewall / DDOS Protection

We opted to use Cloudflare as a Firewall / DDOS protection. Additionally, we used Cloudflare also to manage our DNS settings. Key factor involving this selection was the cost of Cloudflare as opposed to AWS WAF. SSL certificate is also provided by Cloudflare.

AWS Elastic Load Balancer (ELB)

This component accepts the incoming request from Cloudflare and forwards it to the auto-scaling group. This is a managed component. It hides our running VM from the outside world. This allows us to provide segregation from the outside world making it further secure.

AWS Auto-scaling Group (ASG)

This component manages the number of VMs (called EC2 in AWS lingo) depending on the heuristics defined. As traffic increases, ASG will spool up more instances to meet the additional demand. Similarly as load decreases additional instances will shut down automatically. At a minimum two instances will always be running.

AWS Elastic Compute (EC2)

This component is the VM machine that runs Software code. This component will host a copy of the Drupal application. A minimum of 2 such components will always be running. We decided to go with Ubuntu 20.04 OS as it had most compatibility with the Drupal version

AWS Relational Database (RDS)

This component is the managed Database Service. We used RDS running MySQL. We used Master-Slave architecture. All write operations would use Master and Read read operations will be served by the Slaves. This helps in increasing the concurrent number of users.

AWS Elastic File System (EFS)

This component is used for providing a uniform disk space for reading and writing media files. This is a managed service and offers automatic scalability and easy backup configuration.

Redis

This component is used for caching frequently used information. Key purpose is to reduce load on Database. Redis is an open source tool mostly used for this purpose. We used AWS's Elastic Cache with Redis engine for this purpose.

CI/CD Pipelines

We used AWS Pipelines that pulls updated code from Github repository and creates the deployment package and pushes the code to EC2 instances.

Improvements Observed

After migration to the new architecture, we observed various metrics under various load conditions and noticed that there was barely any load.


Memory Utilization

Following graph shows memory consumption of the running instance(s).



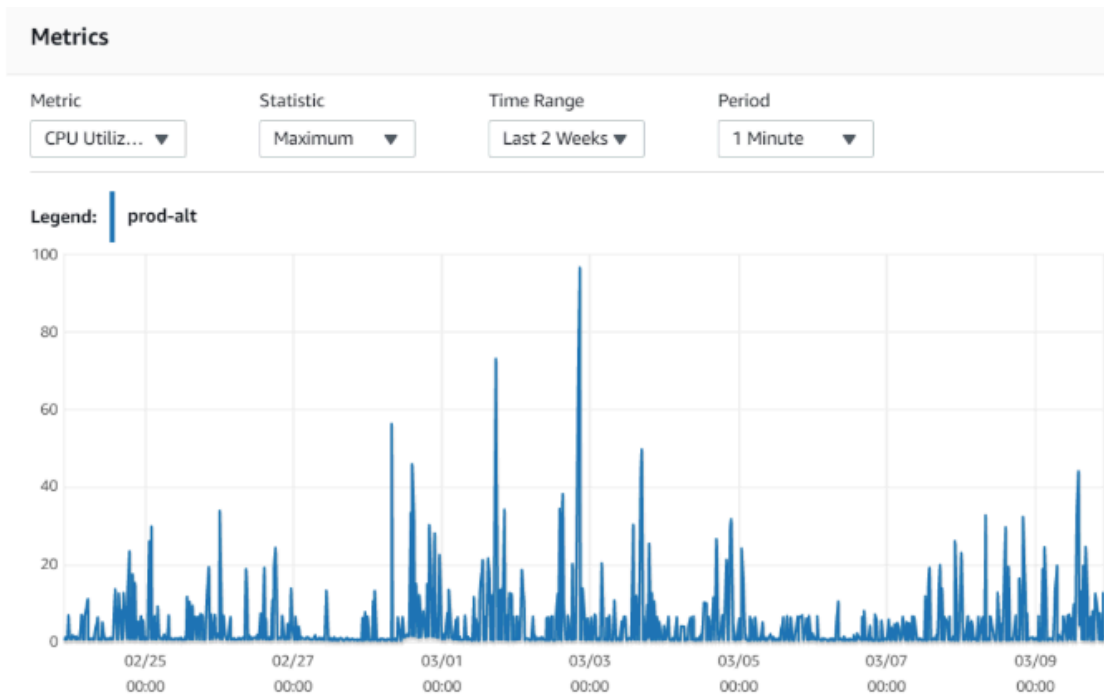
As shown in the above graph, memory consumption barely goes over 50%

CPU Utilization



The above graph shows the CPU utilization over the period of 10 days. It is evident, mostly utilization remains less than 50%.

Database CPU Utilization



Above graph shows the CPU utilization of the RDS database instance. Average utilization remains less than 40%

Database RDS Connections



Above graph shows the number of RDS connections. Low number of connections shows the effectiveness of cache, as most of the requests were served by cache.

Conclusion

This case study illustrates how the FAMRO-LLC team approached the client's project. We thoroughly reviewed, discussed and experienced the problem(s) faced by the client. We interviewed their engineers, content writers and top management and compiled their feedback.

We used collected information to create a solution that addresses customer's current pain areas. After implementation we thoroughly tested the solution beyond the current load predictions to ensure that the solution is future ready.

Our solution helped our client to save up to **50.1%** in infrastructure costs per month. Loading time for static page went down by **64.3%**. Loading time for database heavy page went down by **82.6%**

Need Server / Cloud Optimization or Management

If you are experiencing similar issues or looking for a team to manage your infrastructure - please don't hesitate to contact us.

You can:

Phone: +971-505-208-240

Email: farhan@famro-llc.com

LinkedIn: <https://www.linkedin.com/company/famro-llc/about>