Article title: Serverless Architecture For Bulk Email Management

Authors: Priyatham Sai Chand Bazaru[1]

Affiliations: Mahatma Gandhi Institute of Technology[1]

Orcid ids: 0000-0002-7764-626X[1]

Contact e-mail: priyathamsaichand@gmail.com

License information: This work has been published open access under Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. Conditions, terms of use and publishing policy can be found at <https://www.scienceopen.com/>.

Preprint statement: This article is a preprint and has not been peer-reviewed, under consideration and submitted to ScienceOpen Preprints for open peer review.

DOI: 10.14293/S2199-1006.1.SOR-.PPTDBKS.v2

Preprint first posted online: 21 December 2021

Keywords: Cloud Computing, Serverless, Email, Distributed Systems, Computer Science

Serverless Architecture for Bulk Email Management

Bazaru Priyatham Sai Chand
Dept. of Computer Science
Mahatma Gandhi Institute of Technology
Hyderabad, Telangana
Email: bpriyathamsaichand_cse180508@mgit.ac.in

Abstract—Sending emails in large quantities can be tedious considering free services do not cover bulk email and paid services can be costly and are not easy to customize. Traditional email client used for basic emailing services fail to be useful in larger volumes of emails to target people or spread information to consented individuals. This paper proposes a serverless architecture to tackle such problems by using one such offering from the Amazon Web Services(AWS) which can be easily replaced by a software architects choice of service. The constraints help to make an architecture using components that can fit most of the needs of a serverless backend and extend it to scenarios such mobile notifications, One Time Password (OTP) systems or other means of communication to minimize single point of failure and also decrease the dependency on physical servers for such operations offering a comparable solution within the cloud. The architecture proposed is tested to find the time taken to send the emails of various quantities and see how it affects the cost. The architecture was successful able to send multiple emails in a quick and single invocation and has demonstrated a higher level of scalability compared to conventional methods.

I. INTRODUCTION

Email has become one of the most important means of communications. Sending emails through conventional software like Gmail is limited to a mere thousands which is not sufficient to fit into the business needs and managing the insights of sent emails is crucial to companies to further enhance their approach.

Serverless computing is any computing platform that hides server usage from developers and runs code on-demand automatically scaled and billed only for the time the code is running[1]. This reduces the time and knowledge required to deploy pieces of software while offering significant reduction in pricing compared to hosting the software on dedicated hardware. Various cloud platforms such as AWS, Microsoft Azure and Google Cloud Platform have adapted to this way of thinking making it easily accessible. A robust serverless architecture for sending bulk emails of the magnitude of thousands to millions using scalable model which is implemented using AWS lambda[2] and Simple Email Service(SES)[3] for the SMTP protocol. Challenges such as able to trace the bounce emails as well as the statistics for the emails sent. The prototype is built using the React framework along with underlying services from AWS. These tightly integrated services offer the best way to simulate the architecture and help in sending customized for every recipient based on email templates along with storing bounce emails.

The effectiveness of the scalability offered by the architecture is analyzed using X-ray traces to simulate and test the

bottlenecks and how to overcome them is discussed in the paper. With serverless technology, the cloud provider abstracts away the server management, provisioning servers with fine granularity, on demand, and with a pay-per-use model[4].

This paper proposes an architecture for sending and management of the emails through serverless application using the AWS services. The paper outlines the accuracy and the reliability of the design and future aspects for this architecture.

II. OVERVIEW OF BULK EMAIL MANAGEMENT SYSTEM

The bulk email management system is used to send emails of magnitude in thousands and more and broadcast the message to multiple email addresses with a single operation. All the required fields such as the senders' email addresses, subject, the html body and to addresses are selected and the sending is initiated with a button. Upon successful sending to the API endpoint the emails are processed and sent through the AWS simple email services to all the users and the emails that were bounced are returned after the sending operation. It makes use of a Javascript based framework called ReactJS to input the fields and send it to the API hosted on the AWS service where an array of services are used to scale and perform the sending and retrieving of bounces.

III. ARCHITECTURE DESIGN

The architecture demonstrated in the Fig 1 with the main functionalities include sending the emails and tracking the bounces. Initially, all the data sent to the lambda using the API gateway which handles the request and response for the whole operation. Upon reaching the lambda, JavaScript Object Notation (JSON) objects from the string parameters preprocessed, the html part of the email along with the subject made into SES templates accessed by other lambdas and the emails batched no more than 50 emails per batch and placed in a Simple queue Service[5] queue. The ordering of the messages is important hence, we use a First In First Out queue, which triggers another lambda that sends emails to all the 50 recipients using a templated email function by accessing the template that we created earlier. We set a flag to mark the end of the batches deleting the template by the end of the execution.

The emails use a configuration set that triggers events such as bounce, complaints or seen to other services for logging and other analysis. The configuration set sends all the bounce receipts to SNS which triggers a lambda function that stores the receipts useful information as a document in a NoSQL

Fig. 1. Architecture for bulk email management

database such as dyanamoDB once all the emails sent and all bounce receipts stored in the database. Another lambda triggers using an API retrieves all the records and sends the bounced emails as response to API gateway which forwards it the application shown to the user.

The emails that are bounced are identified by the configuration set and are sent to the SNS with email address and the type of bounce which invokes a lambda trigger that stores the Simple Notification Service(SNS)[6]notifications after filtering the message into the dynamo DB.

All the bounced emails are automatically updated to a email supression list which saves the names of bounced emails and accepts them but doesn't send the emails after it is declared bounce once. Thereby, saving the bounce rate of the service.

Additionally, All the other statistics are sent to Cloudwatch metrics which can be used to analysis the statistics of sending emails at different time frames and track the bounce rate and can be useful in finding out the effectiveness of the campaign by the number of people who opened the email.

The architecture keeping in mind the constraints on SES service. The advantages of this is two folds. It helps to work within the limits of the service. Moreover, these constraints can be fine tuned to deal with the limits of SMTP server if an independent server or other services are used.

IV. CONSTRAINTS

The architecture is based a few constraints from the SES service and other bottlenecks that affect the sending of emails. The SES through which the emails are initially sent through sandbox to test and later needs to be upgraded by the AWS authority to permit the user to request increase in the number of emails to be sent. For the purpose of this architecture 50,000 emails have been requested which was granted by the authority.

- 14 emails/second
- 50,000 emails per day
- request data 5kb

There are two performance monitoring services available to AWS application developers: CloudWatch and X-ray.

V. ANALYSIS

Basic message transfer from Author to Recipients is accomplished by using an asynchronous store-and-forward communication infrastructure in a sequence of independent transmissions through some number of mail transfer agents[7]. Two of the most common ways to analyze the architecture was to consider the time required to send all the emails and the cost associated with it. Both of the metrics are analysed in the following section.

A. Performance

There are two performance monitoring services available to AWS application developers: CloudWatch and X-ray.[8] The application is analyzed to see the time required to send the emails and the time taken to run the functions between the request and response to the API endpoint. This can be seen through the cloudwatch[9] logs and X-ray to find out the issues with latency over various regions and finding the best regions and how to optimize the functions. In this case, we used X-Ray as a monitoring and display service that automatically samples the entry and exit of function instances, called segments, using unique trace identifiers (**trace_id**) [8][10].

Here we take the case of sending multiple emails and the time taken is sampled over 30 data points to represent in Figure 2. The time is calculated by considering the cloudwatch logs that can be used to determine the time taken to execute the last function of lambda as show in Figure 1 to send the email and the time leading to that function from the start of the API gateway request log that is stored in the logs. This experimented several times to cancel out the variations of cold start that is a common phenomenon in serverless functions. All the calculations are done to avoid any overlap between different invocations and only a single request is sent within the same region.

Fig. 2. Time taken to execute and return the response to the number of emails that are sent.

The number of emails sent is directly proportional to the time taken. The time taken for 10,00 emails averaging at around 30 minutes is sufficient and satisfactory considering that the time taken is derived from the last email that is sent in the list of emails.

B. Cost

The cost of the services can also be a motivating factor to implement such an architecture. In our use case The following architecture can be tested with the AWS free tier and later expanded to the needs of the users. With \$0.01 dollars per 1,000 emails but one needs to be wary about the associated services like the SES, SNS and the cloudwatch metrics to analyze the emails add up to cost but since they are calculated

based on the GB of data it needs to pre-process the rates will not exceed the cost of emails.

Here we take the example, provided by the SES platform to estimate the cost required to implement the architecture.

You use Amazon SES to send about 250,000 emails per month. You receive 1,000 emails per month. You don't use dedicated IP addresses. Every message you send and receive is 32KB in size which results in a total of \$25.98 per month [11] which is significantly less than competitors such as SendGrid or MailChimp who offer their own SMTP server or schedule emails to fit into the constraints of other providers to carry out email campaigns.

Amazon provides sample pricing calculations, but as your workload varies, so will the billing.[12] The SQS requests can exceed the free tier if not monitored carefully and add up to additional costs for the next 1000 requests or more based on the usage of the system.

VI. CONCLUSION

To conclude, the architecture has been effective in sending bulk emails of various volumes with runtime that can be considered as sufficient or optimal and cost-effective. The primary goal of the architecture to be able to manage the entire system without a server being provisioned directly i.e. serverless is achieved and is promising to further develop and enhance its capabilities. The statistics displayed can be used for reach through the customers and the key idea of the architecture is to make use of serverless functions and services to manage emails with an optional fronted interface which was the case in this paper. Thus, making it a fast application to deliver emails to the users' inboxes in a short period of time. Without provisioning any IP addresses and acquiring the compute power needed to send the emails beforehand thereby making it more effective and possesses a higher level of abstraction than the conventional method.

REFERENCES

- [1] P. Castro, V. Ishakian, V. Muthusamy, and A. Slominski, "The rise of serverless computing," *Communications of the ACM*, vol. 62, no. 12, pp. 44–54, 2019.
- [2] Aws lambda. [Online]. Available: <https://aws.amazon.com/lambda/>
- [3] Amazon simple email service. [Online]. Available: <https://aws.amazon.com/ses/>
- [4] S. Eismann, J. Scheuner, E. Van Eyk, M. Schwinger, J. Grohmann, N. Herbst, C. L. Abad, and A. Iosup, "Serverless applications: Why, when, and how?" *IEEE Software*, vol. 38, no. 1, pp. 32–39, 2020.
- [5] Amazon simple queue service. [Online]. Available: <https://aws.amazon.com/sqs/>
- [6] Amazon simple notification service. [Online]. Available: <https://aws.amazon.com/sns/>
- [7] D. Crockier, "Internet Mail Architecture," RFC 5598, Jul. 2009. [Online]. Available: <https://rfc-editor.org/rfc/rfc5598.txt>
- [8] W.-T. Lin, C. Krintz, R. Wolski, M. Zhang, X. Cai, T. Li, and W. Xu, "Tracking causal order in aws lambda applications," in *2018 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2018, pp. 50–60.
- [9] Amazon cloudwatch. [Online]. Available: <https://aws.amazon.com/cloudwatch/>
- [10] Amazon x-ray. [Online]. Available: <https://aws.amazon.com/xray/>
- [11] Amazon ses pricing. [Online]. Available: <https://aws.amazon.com/ses/pricing/>
- [12] A. Eivy and J. Weinman, "Be wary of the economics of" serverless" cloud computing," *IEEE Cloud Computing*, vol. 4, no. 2, pp. 6–12, 2017.