Fast VM Migration in Edge Cloud: A VM Slicing and Caching Approach

Dongbiao He*, Teng Ma*, Jinlei Jiang, Guangwen Yang
Department of Computer Science and Technology, Tsinghua University, Beijing, China
{*hdb13@mails., mt16@mails., jjlei@, ygw@*}*tsinghua.edu.cn*
$*$ *are students*

Abstract

In recent years, there has been a significant interest in edge computing for it has the potential to greatly reduce delay and backhaul traffic load of mobile applications by moving cloud services to the edge. Edge computing is a promising solution related to IoT or AR/VR applications, but there is a conflict for high requirement in VM migration for service handoff and low bandwidth available owing to a wide area network (WAN) connection[2, 3]. For example, the disk state of the VM may also need to be migrated in addition to memory state when shared storage is not available across edge cloud sites. This requirement calls for new approaches for VM migration which could be adapted to the edge clouds.

Now, we are building a edge cloud platform which enables fast VM migration, designed for supporting cloudlet-based applications. The basic solution is to combine the technology of VM image slicing with caching. First of all, we divide the image file into several data chunks. In each edge servers, we cache the most popular set of chunks. Ideally, if there is a migration request, the chunks only contain user data would be transmitted to the remote edge server. Finally, the reconstructed image file would take over the jobs for serving the users.

Our design contains three major components for fast VM migration in service handoff between edge servers: the image data chunks manager for disk slicing, the cache coordinator for disk chunks allocation and the VM migration engine for fast data chunks transmission.

Image data chunks manager is built with the design of Liquid[1]. VM images are split into fixed size data blocks. Each data block is identified by its unique fingerprint, which is calculated during the slicing process. A specific VM image is represented by a sequence of fingerprints which refer to the data blocks inside the VM image. A cache coordinator is placed centrally (e.g., data center) to calculate the reference counting extensively to track usage of each data chunks in the edge servers. The disk data chunks distribution could be managed with different caching policies in various requirement. It could also remove the unused garbage data chunks when the edge server is running out of space to increase the utilization of the storage.

VM Disk state migration usually occupies the largest component of the entire migration time, leading to low efficiency in service handoff for edge cloud applications. The disk migration over slow WAN links can be optimized by VM disk slicing and caching with following advantages: First, data deduplication techniques can be easily applied to reduce the total size of data need to be transmitted. The VM migration engine would be responsible for calculating the minimum data needed from the source. Second, in avoid of reducing the resending ratio of the modified disk data chunks, the chunks remain unchanged during the migration would be send in advance. Namely, data chunks may be sent based on their dirty rate, with higher frequency blocks not being sent until the end. Third, by distributing different VM data chunks into several edge servers, there would be a number of potential edge servers caching the data chunks required. Therefor, a multiple sources VM disk transmission could be established to further increase the bandwidth available for the migration over multiple transmission paths.

We use following steps to live migrate each VM following the design. Step 1) Calculate the data chunks required of the destination and find the potential data chunks sources. Step 2) Establish the connections between the destination host with data chunks sources, and transfer the disk state needed. Step 3) Transfer the memory state of the VM to the destination edge servers as it continues running, and the steps remained would be the same with the previous design. The whole process should not disrupt any other active network connections between the clients and the servers involved. We plan to evaluate resource (e.g., bandwidth, storage) utilization and migration time reduction of the design. To evaluate performance, we will emulate the edge cloud environment using commodity x86 servers, for example, by limiting the bandwidth of the network.

References

[1] Zhao, X., Zhang, Y., Wu, Y., Chen, K., Jiang, J., Li, K. (2014). Liquid: A scalable deduplication file system for virtual machine images. IEEE Transactions on Parallel and Distributed Systems, 25(5), 1257-1266.

[2] Mashtizadeh, A. J., Cai, M., Tarasuk-Levin, G., Koller, R., Garfinkel, T., Setty, S. (2014, June). XvMotion: Unified Virtual Machine Migration over Long Distance. In USENIX Annual Technical Conference (pp. 97-108).

[3] Ha, K., Abe, Y., Eiszler, T., Chen, Z., Hu, W., Amos, B., Satyanarayanan, M. (2017, October). You can teach elephants to dance: agile vm handoff for edge computing. In Proceedings of the Second ACM/IEEE Symposium on Edge Computing (p. 12). ACM.