A Literature Survey on Cloud AutoScaling Mechanisms

[1]Gopesh Banker, [2]Gayatri Jain
[1]Student, [2]Assistant Professor
Department of Computer Engineering, LJIET, Ahmedabad

_____

***Abstract -* In recent years, cloud computing has evolved as a much efficient, proactive and widely accepted technology. Mostly virtualization techniques are used to serve the requests from the clients. However, major issue arising for all cloud users is to handle the dynamic number of requests from time to time and using efficiently all the existing resources. This problem is overcome by the auto-scaling mechanism of Virtual Machines. Basic auto-scaling mechanisms are naïve as they only check for individual VM healths. Here I come up with a survey of variety of auto-scaling mechanisms and comparing them.**

***Keywords -* Cloud computing, cloud resource scaling, dynamic scaling, virtual machine utilization**

_____

I. INTRODUCTION

Cloud Computing is the most widely used platform now a days. It provides its users, to dynamically allocate and deallocate resources at run time. This property of automatically scaling a resource is termed as auto-scaling in the cloud infrastructure. Most cloud service providers charge their users on a pay per hour model for each of the resource they consume. Hence efficient utilization of existing resources is very much essential for cloud users. These resource include the Virtual Machine (VM), the network bandwidth or the storage capacity allocated to the virtual machine. Hypervisors are being used to create these virtual machines. Here in this paper, we only consider on auto-scaling mechanisms of Virtual Machines. Ideally, the load on the VM decides whether to scale up or scale down a resource. If the CPU utilization reaches a particular upper threshold level, the cloud infrastructure scales up the VM resource i.e. a new VM is added to the cluster of VMs and the load balancer starts sending new request on this newly created VM and vice-versa. But, just considering a threshold value for VM utilization is a naïve solution and cannot be used in real world.

Another factor arising is the VM boot up time and the VM shutdown time. Once you start a VM, it doesn't start from the next minute itself, it takes some time for the booting process. Same is the case for shutting down a VM. Moreover, it also depends on the configuration of the resource. Cost is also an issue in the cloud environment because any VM created is charge on an hourly cycle and if the scaling mechanism is not proper, the cost of the user would be highly impacted. In this paper, I hereby provide a detailed literature survey of such auto-scaling mechanisms. The paper focuses on various mechanisms for scaling VM resource in a cloud infrastructure.

II. RELATED WORK

2.1 Cloud Auto-Scaling with deadline and Budget Constraints[1]

As a computing platform, clouds own distinct characteristics compared to utility computing and grid computing. Ming Mao, Jie Li and Marty Humphrey[1] have identified some of them which can largely affect the way people use cloud platforms, especially in cloud scaling activities. They describes various metrics to be accounted for while scaling a VM resource:

Unlimited resource limited budget

Ideally, a cloud infrastructure offers its users unlimited power of computing tasks and storage capacity. This enables users to scale to an extremely large size, but are not free, they are charged on an hourly basis for each of the resource being utilized. Hence, a cloud auto-scaling mechanism should explicitly consider user budget constraints while acquiring a resource.[1]

Non – Ignorable VM instance acquisition time

Even if you can scale cloud resources anytime, it doesn't mean that cloud scales it very fast. Based on their research, they estimate the instance acquisition time to be of 10 minutes approximately before the resource is ready to be used.[1] Same is the case for VM shutting down time which takes about 2-3 minutes in Windows Azure[6]. This implies that users have to consider two issues in cloud dynamic scaling activities. First, count in the computing power of pending instances. If an instance is in pending status, it means it is going to be ready soon. Ignoring pending instances may result in booting more instances than necessary, therefore waste money. Second, count how long the pending instance has been acquired and how long further it needs to be ready to use. If the startup time delay can be well observed and predicted, application admin can acquire machines in advance and prepare early for workload surges.[1]

Full Hour Billing Model

The hourly billing model saves money when the users shut down the machines. But, VMs resources are always billed on an hourly basis. A 1 minute usage is charged for a whole hour and a 60 minute usage is also charged for the same whole hour and if the resource was started and shutdown two times in an hour, then it is charged for two VM resources. Hence, a better policy is to only shutdown the VM resource when it approaches its full hour operation, once an instance is started.[1]

Multiple Instance Types

Instead of offering single instance type, cloud providers now offer various instance types for its users. They can use different types of instances based on their performance and usage. For example, EC2 instances are grouped into three families, standard, high-CPU and high-memory. Standard instances are suitable for all general purpose applications. High-CPU instances are well suited for computing intensive application, like image processing. High-memory instances are more suitable for I/O intensive application, like database systems and memory caching applications[1]. One important thing is that instances are charged differently and not necessarily proportional to its computing power. For example, in EC2, c1.medium costs twice as much as m1.small. But it offers 5 times more compute power than m1.small. Thus for computing heavy jobs it is cheaper to use c1.medium instead of the least expensive m1.small. Therefore, users need to choose instance type wisely. Choosing cost-effective instance types can both improve performance and save cost[1].

From all these parameters taken into account, Ming Mao, Jie Li and Marty Humphrey proposed a solution to select a VM instance type that is suitable for the current job processing i.e. either they are High CPU intensive, High I/O intensive or High Memory intensive. They implemented this architecture on the Windows Azure platform where they ran a real scientific application (MODIS)[1].

Based on the policy to shutdown the VM instance once it reaches its one hour billing cycle, they find out the average processing time to be as follows.

Table 1 Average processing time.[1]

	Mix Avg 30 jobs/hour STD 5 jobs/hour	Computing Intensive Avg 30 jobs/hour STD 5 jobs/hour	I/O Intensive Avg 30 jobs/hour STD 5 jobs/hour
General 0.085$/hour Delay 600s	Average 300s STD 50s	Average 300s STD 50s	Average 300s STD 50s
High-CPU 0.17$/hour Delay 720s	Average 210s STD 25s	Average 75s STD 15s	Average 300s STD 50s
High-IO 0.17$/hour Delay 720s	Average 210s STD 25s	Average 300s STD 50s	Average 75s STD 15s

The instance acquisition for VMs is shown in the below figure 1.

Figure 1 : Instance acquisition Time[1]

Integrated and Autonomic Cloud Resource Scaling[2]
Typical Cloud resource scaling solutions possess following limitations:

1. Thresholding mechanisms used to trigger scaling decision process are not Cloud-ready. That is, they are not suitable for huge scale and dynamic Cloud environment. They are naïve or rudimentary causing spurious scaling of resources[2].
2. In a Cloud, resources from three different domains, compute, storage and network, are acquired/released (scaled up or down) on demand (typically in virtualized form, including application or software resources). But in current solutions resources from these three domains are scaled in isolation rather than in combination resulting in spurious auto-scaling[2].
3. The performance metrics that trigger scaling are considered in isolation. For example, compute resources are scaled based on CPU/compute resource metric only, but not in correlation with, for example, metrics from the network domain[2].
4. While existing solutions may scale compute or storage resources automatically when specified condition or policy is satisfied, network resources are scaled statically.[2]
5. Network resources are not scaled in relation to compute and storage domains and vice versa. For example, when virtual compute resources are scaled, virtual load-balancer (LB) or firewalls (FW) may also need auto-scaling.[2]

IACRS Tenant Interface Input
The interfaces and relevant parameters provided in IACRSTI are as follows (we call the input IACRS Policy or IACRSPol) :

- Resource R to be monitored and auto-scaled (typically this will be an existing resource URI).
- Performance metric to be monitored (PMT).
- Reference to a resource (end-point 2: EP2), if the metric is binary (such as response time, delay or jitter from one resource to another).
- Group ID of the group R belongs to and optionally an indication whether R is a parent of the group.
- Threshold Policy (TP) for R and PMT:
  - ThrU = #%, ThrbU = #%, ThrL = #%, ThroL = #%, Duration = # sec (see Section B below for definition).
- Auto-scaling policies (SP) for a resource and metric:
  - Resource acquisition (SP-RA) policy: <increment in #%> (INCR), <max resources in #> (MAX), where increment is the number of resources to be acquired, which is calculated as <original # of resources in a group> * INCR rounded to the next higher number and MAX is the maximum number of resources that can be acquired when no more resources will be acquired irrespective of load on all the resources in a group.
  - Resource release (SP-RR) policy: <decrement in #%> (DECR), <min resources in #> (MIN), where decrement is the number of resources to be released, which is calculated as < # of resources in a group> * DECR rounded to the next higher number and MIN is the minimum number of resources that should be retained when no more resources will be released irrespective of load on all the resources in a group. If SP-RR is not specified, then SP-RA will be used with INCR replaced with DECR and MIN defaulting to 1.

The scaling algorithm developed is as below

Metric value trending up with persistence
- When the PMT_Value crosses above the ThrU for the first time the duration clock starts to tick and we set a variable Tick_Up_Start to 1.
- If PMT_Value remains above the ThrU or oscillates around it, but remains above the ThrbU for the specified duration (such as 5 minutes), then stop (reset) clock tick and set Tick_Up_End.
- If PMT_Value persisted for duration since crossing ThrU and remains above ThrbU sincelast time Scale value was 00 or 01 (this latter condition is to prevent repeated pattern such as the one shown in Figure 2), then set the variable Scale to 11 (which may potentially trigger scaling):
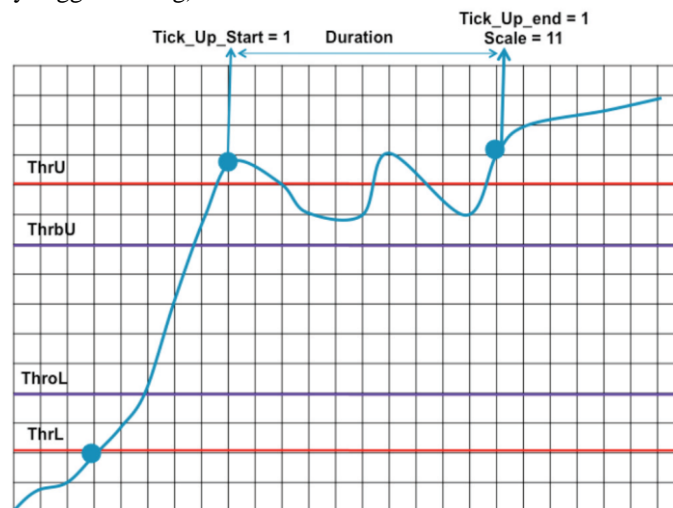

Figure 2 : Metric value trending up with persistence[2]

Metric Value Trending Down from above
- If metric value remains below the ThrbU for the specified duration, then variable Scale will be set to 01 (see below group/correlation heuristics on how Scale value 01 is utilized). The lock for this duration starts to tick when the PMT_Value crosses below the ThrbU for the first time, since it was above ThrU.
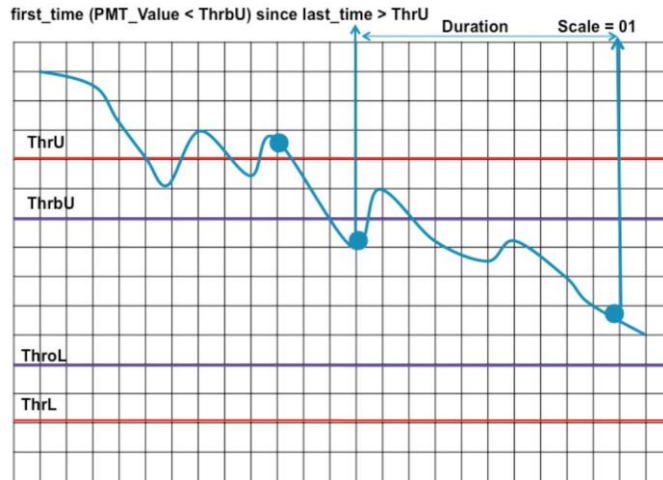

Figure 3 : Metric Value Trending Down from above[2]

Metric value trending down with persistence

The heuristics algorithm is similar as PMT_value trending up with persistence except that the algorithm is applied around ThrL and ThroL and Scale value is set to 00 when PMT_Value crosses below ThrL for the first time since Scale value was 11 or 01 and remain around ThroL for the specified duration. A PMT_Value behavior is shown in Figure 4.
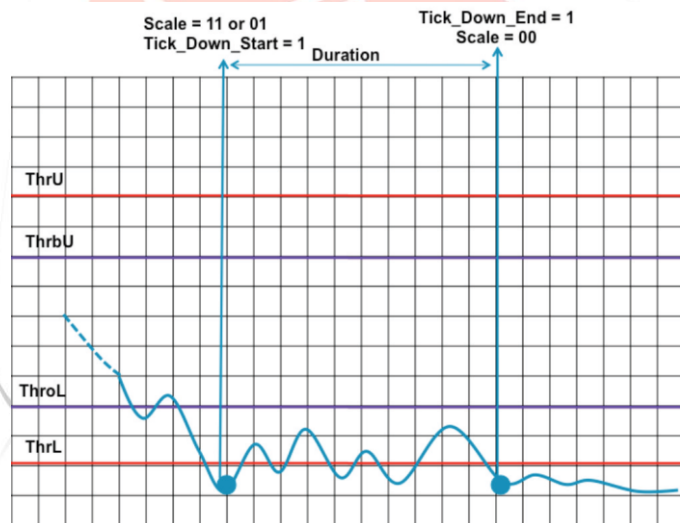

Figure 4 : Metric value trending down with persistence[2]

Non-scaling Metric Value behavior
- The PMT_Value for non-scaling behavior graph is shown in Figure 5, in which the Scale value will retain its last value.
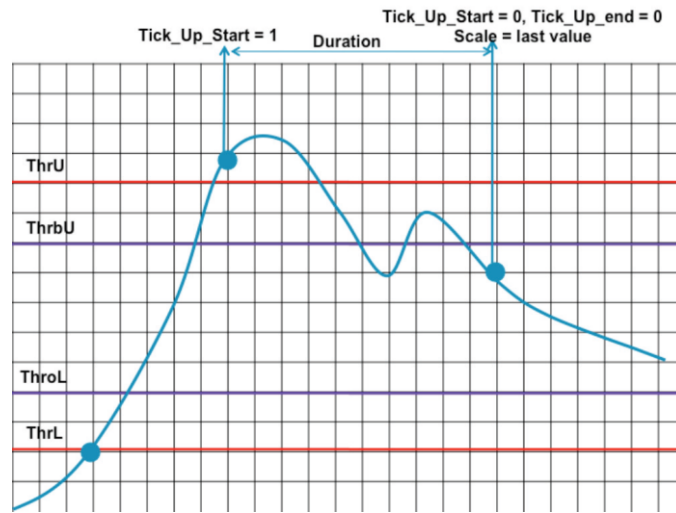
Figure 5 : Non-scaling Metric Value behavior[2]

***An Auto-Scaling Mechanism for Virtual Resources to Support Mobile, Pervasive, Real-Time Healthcare Applications in Cloud Computing*[3].**

In this article, Yong Woon Ahn and Albert M. K. Cheng, Jinsuk Baek, Minho Jo, Hsiao-Hwa Chen[3], propose a novel auto-scaling mechanism in order to dynamically adjust the number of VMs to handle deadline-critical real-time data, which varies in size over time. In consequence, the resizing of the virtual resources for processing the given data is achieved on an on-demand basis. The key mechanism is to predict the volume of future data. Although it is not necessarily trivial to predict the exact moment at which a large volume of data will be delivered from sensor-based medical devices, most objects monitored by sensor-based devices typically show symptoms before transitioning to an abnormal state.[3]

Basically, it is the task of a VM manager to scale VMs up or down. Here, Real Time Application Servers (RTA servers) are configured in the VMs and they are responsible for the scaling of VMs.
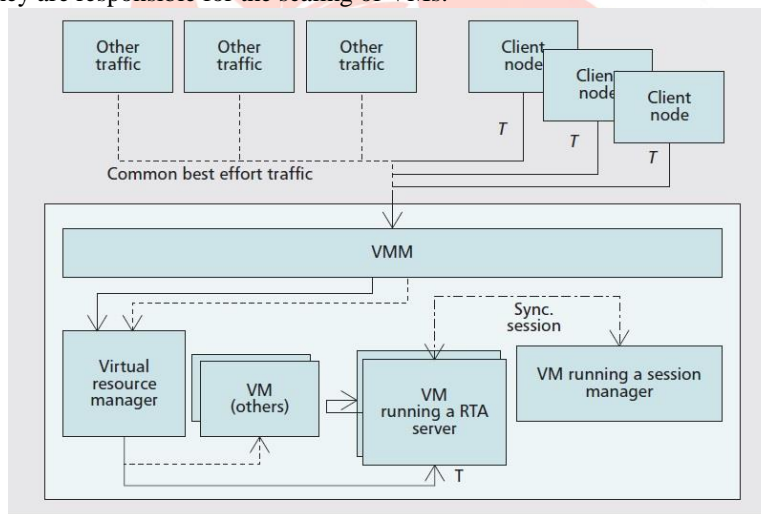

Figure 6 : Physical machine architecture with virtual RTA servers[3]

At any particular time, the VCPU load is represented by a state N. If, the VCPU load increases or decreases, the state transition is made in forward or backward direction. We can represent this transition in an angular format ranging between -90° and +90°, where -90° represents 0% load and +90° represents 100% load which is a degree representation D[i]. In order to provide more accurate predictions, the number of transitions should vary depending on the value ofD[i]. Let us assume that there are M sections between –90° and 90°. If the value of D[i] is equal to or greater than w90°/M/2, but smaller than (w+1)90°/M/2, and D[i] is bigger than 0°, the analyzer moves the state w transitions forward. If the value of D[i] is equal to or smaller than w(–90°)/M/2, but still smaller than (w+1)(–90°)/M/2, and D[i] is smaller than 0°, the analyzer moves the state w transitions backward. Figure 7 shows an example, when there are four different sections, including F1, F2, B1, and B2.[3]

Figure 7 : State diagram for transitions of the proposed auto-scaling mechanism[3]

If the root RTA server reaches State N, it determines whether it needs more computing resources or not by checking the states of other child RTA servers. If it is needed, it launches a new child RTA server, where the overloaded pending real-time tasks will be assigned. On the other hand, if D is smaller than 0°, it now makes w backward transitions. When it reaches at State 1, it terminates one of its child VMs and intercepts the workload of the terminating child RTA.[3]

*Efficient Auto-scaling in the Cloud using Predictive Models for Workload Forecasting***[4]**

Here, Nilabja Roy, Abhishek Dubey and Aniruddha Gokhale[4] have proposed an autoscaling algorithm, which is dependent on the workload prediction. The algorthim uses the receding horizon control and iterates over the number of look ahead steps and calculates the cumulative costs. For every future time step, it computes the cost of selecting each possible resource allocation by calculating the response time for that particular machine configuration. Once the response time is calculated, it is used to calculate the cost of the allocation which is a combination of how far the estimated response time is from the SLA bounds, cost of leasing additional machines and also a cost of re-configuration. The cost of reconfiguration is computed based on the number of machines that need to updated. Obviously reconfiguration will incur some costs and thus the algorithm will try to reduce the amount of reconfiguration. Each of these cost components will have weights attached to them which may be varied depending on the type of application and its requirements. Applications are required to specify which factors are more important to them, and our auto-scaling algorithm will honor these specifications in making the decisions.[4]

III. CONCLUSIONS

This paper summarizes different scaling mechanisms in a cloud environment. It gives the basic idea of understanding the working of cloud platform where cloud service providers provide the IaaS architecture to its users. Based on the user requirement, the scaling mechanism for VMs can be utilized. Current mechanisms provided by the cloud providers like AWS[5] cannot be utilized for real time usage. So these strategies can help come out of such dilemma scenarios. We also aim at proposing a new mechanism that takes into account the three parameters viz : CPU utilization, VM response time and Future load predictions.

REFERENCES

[1] M. Mao, J. Li, and M. Humphrey, "Cloud Auto-Scaling with Deadline and Budget Constraints," Proc. 2010 IEEE/ACM Int'l. Conf. Grid Computing, Oct. 2010, pp. 41–48.

[2] M. Hasan, E. Magana, A. Clemm, L. Tucker, and S. Gudreddi, "Integrated and autonomic cloud resource scaling, " in iEEE Network Operations and Management Symposium (NOMS 2012), Apr. 2012, pp. 1327-1334.

[3] Y. W. Ahn, A. M. Cheng, M. Jo, and H.-H. Chen, "An Auto-Scaling Mechanism for Virtual Resources to Support Mobile, Pervasive, Real-Time Healthcare Applications in Cloud Computing," IEEE Netw., 2013.

[4] N. Roy, A. Dubey, and A. Gokhale, "Efficient autoscaling in the cloud using predictive models for workload forecasting," in Cloud Computing (CLOUD), 2011 IEEE International Conference on. IEEE, 2011, pp. 500-507.

[5] http://aws.amazon.com/autoscaling/

[6] http://azure.microsoft.com/

[7] http://en.wikipedia.org/wiki/Cloud_computing/

[8] http://aws.amazon.com/ec2/