No reboot Operating System update using OpenStack technologies

Sanchita Sinha
IBM Systems
IBM India Pvt Ltd
Bangalore, India
sancsinh@in.ibm.com

***Abstract*—** At Several occasions there might be need for Operating System to upgrade itself by installing bug fixes (Kernel, Kernel Extensions) or by moving to a newer Service Pack. This affects business critical workloads as the system needs a reboot for the newer version to take effect. This issue can be solved by using the live update operation which allows the application of bug fixes and Services Packs without requiring a system restart.

This paper will provide an insight on different challenges and requirements needed for live update operation and how one can leverage the OpenStack Technologies to overcome those.

To Perform live update operation, one major challenge is the admin effort to setup the environment. This paper will provide a detailed insight into how OpenStack technologies provide solutions to resolve the challenges, greatly reduce manual effort and provide cost effective solutions.

Keywords—Virtual Machine, Operating system, live update, OpenStack, no reboot

I. Introduction

Live update operation is the process of application of bug fixes or patches in form of new kernel/ kernel extensions or even upgrade to a newer Operation System Version from the current version; all without causing any application downtime.

Prior to live update operation there were few solutions available which too did not require system reboot while applying update. For example, concurrent update-enabled fixes too did not require reboot. But there were many limitations. They were not scalable. There was no capability of maintenance level update. Kernel patching also was not an efficient solution in terms of scalability, persistence and maintenance update. It allowed the deployment of limited Kernel related fixes on a running VM. These patches were limited to some fixes and could not be generalized to all Kernel or Kernel Extensions fixes. Hence a comprehensive solution was required which is scalable and persistent on reboot.

The live update operation does not have the same limitations as in the case of kernel patching. Kernel fixes can be cumulative and across different Kernel components. Kernel Extensions that do not require a checkpointed state can also be updated using the live update operation. Using live update, it is possible to update to a newer Service Pack without workload downtime to suit 24x7 work schedules with Minimized Software Incidents.

The VM which needs to undergo upgrade, is called the original VM. Normally, after applying updates to the original VM a reboot is often required. To avoid this, as part of the live update operation, a new surrogate VM is created. The surrogate VM boots with the updated kernel and all the processes and their resources are moved from the original VM to the surrogate VM.

The updates that are installed by using the live update operation are always committed. Therefore, the changes cannot be rejected later. If there is any requirement to return to the previous operating system level, back up should be taken to an alternate disk before the live update operation starts. To move back to the previous level, system can be booted from the alternate disk.

In next sections, problem description section provides details about the challenges faced for adoption of live update in terms of setting up the environment like storage requirements, setting up the multipath. Results and Contribution section provides details into how OpenStack technologies provide automated and cost-effective solutions to address all challenges. Images section describes the live update process using OpenStack technologies, followed by conclusion of this paper and references.

II. Problem Description

A. In Production environment several VMs might need to be installed with the patch fix. Or VMs might need to be upgraded to new Operating System Version for supporting more features.
   - With existing technologies, both requires significant manual effort, as this needs to be done one by one for each VM.

B. The VM which needs to undergo live update, should have all its disks multipath enabled.
   - Admin need to set multipath for all disk in each VM.

C. When live update deploys the surrogate partition, it also requires giving access to the required storage for the surrogate to boot up.
   - No automated solution exists in current technology.

D. Live update is required to have provided additional disks for surrogate bootup and mirroring operations.

III. Results and Contribution

Highly scalable modern cloud management and deployment of critical enterprise workloads can be achieved using OpenStack technologies. OpenStack technologies provide automated solutions for all problems described in above section. It also provides cost effective solution for the adoption of live update operation by greatly reducing manual effort and discarding all additional storage requirements.

A. To resolve the problem of updating several VMs, one might perform live update only on one VM. This VM will be running with new kernel post live update. The Operating System image can be captured from the VM running with new kernel and this image can directly be used for future deployments. To achieve this, all volumes that belong to its boot set are included in the image generated by the capture. User can also choose which data volumes to include in the image generated by the capture. User can use cloud-init to enable the virtual machines for capture. cloud-init is a technology that take user input and configure the operating system and software on deployed virtual machines. cloud-init is widely used in OpenStack. This process can be used for deploying VM's with upgraded Operating System image as well.

B. In OpenStack based live update no manual setup of multipath is required. Since OpenStack can manage the storage system and the Fibre Channel fabric, it can deploy a partition on any managed system and give it access to the required storage. Hence all requests made by live update is served by it.

C. When surrogate boots up, live update requires to give it access to storage system. Without the help of OpenStack managed storage system, the additional disk needs to be created manually prior to performing live update. With the help of OpenStack technologies, this disk creation is automatically performed and at the end of live update deleted as well. No wastage of storage resources.

D. Live update is required to have provided additional disk for mirroring operations, this is again additional storage requirement. With the help of OpenStack technologies, this disk creation is automatically performed and at the end of live update deleted as well.

E. Since Live update always consumes 2X system resources, it is important which server the VM is deployed. OpenStack provides useful VM deployment technologies and perform better load balancing.
Using OpenStack host aggregates technology, a host group can be created. User can create as many host groups as they want and assign hosts (i.e. servers) to a group. This cloud partitioning gives better operational control depending on business requirement (e.g. multitenancy). User can now create different cloud environments for VM deployments in the form of host groups. Each host group uses a "placement policy" that defines how a host is selected within the host group.
   - **Stripping**: The host with the lowest number of VMs is selected.
   - **Packing**: The host with the greatest number of VMs is selected.
   - **CPU Allocation Balance**: The host that would end up with the smallest percentage of its processing units allocated is selected.
   - **Memory Allocation Balance**: The host that would end up with the smallest percentage of its memory allocated is selected.
   - **CPU Utilization Balanced**: The host with the lowest CPU utilization is selected.
   - **Memory Utilization Balanced**: The host with the lowest memory utilization is selected.

   A server is a host which is always a member of exactly one host group. Live update can be performed on a VM residing on a host belonging to any host group. These placement policies help leading to more number of success cases, as it can decide which host the VM can be deployed and user can have enough resource to perform live update operation.

F. OpenStack token based authentication mechanism makes live update a secure process, as it can only be performed in presence of a valid token.

IV. FIGURES

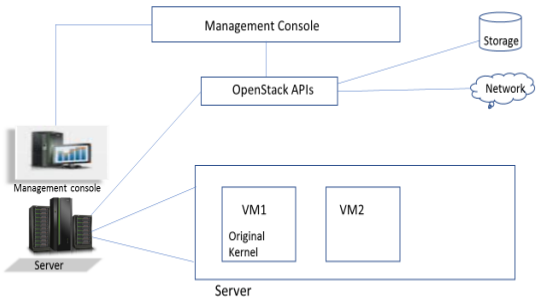

Figure 1: OpenStack managed Servers

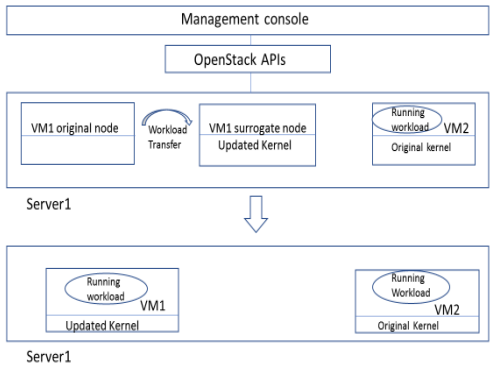

Figure 2: Live update Using OpenStack Technologies

Figure 1 Describes how Management console can leverage OpenStack Technologies and manage Servers, Storage, Network all together. Here VM1 is the node which needs to undergo live update and it belongs to host named Server1.

Figure 2 describes how the live update operation is performed on VM1 with help of OpenStack Technologies. A new surrogate VM is created. Surrogate VM runs with updated kernel and all the processes and their resources are moved from the original VM to this surrogate VM. At the end of live update operation, original VM is deleted and its resources are freed.

V. CONCLUSION

This paper has provided an insight on doing live update of Operating system without system reboot with the help of OpenStack Technologies and how OpenStack resolves several challenges faced by live update by providing automatic solution for setting up multipath, creating additional storage, removal of additional storage at end of live update. It also summarized how OpenStack helps in effective VM deployment by choosing proper host using host aggregate technologies and thus improving on better system resource availability for live update to succeed.

VI. FUTURE ENHANCEMENT SCOPE

Existing live update technology always consumes 2X system resources which can sometime arrive as an obstacle to its adoption.

Live update requires the availability of 2X CPU and memory resources on the Server which hosts the original VM. On packed systems there might not be enough free CPU and/or memory to allow the creation of the surrogate VM, or not enough to run several live updates concurrently. In case the partition is using SRIOV (Single root I/O virtualization) adapters, there is possibility that the Server does not have enough capacity to accommodate both original and surrogate VMs.

OpenStack provides automated technologies to maintain storage and network. This provides scope for VMs to migrate across Hosts. Using this VM migration technology, resource shortage issue can be solved as future enhancement.

Another Possibility for future enhancement lies for the rollback mechanism. In case, failure is observed during live update due to setup issues like network connectivity loss to management console; alternate management console can be used for all the cleanup and re-try of live update.

VII. REFERENCES

[1] https://en.wikipedia.org/wiki/OpenStack
[2] https://en.wikipedia.org/wiki/Kernel_patch
[3] https://en.wikipedia.org/wiki/Multipath_I/O
[4] https://developer.openstack.org/api-guide/quick-start/