WINTECHCON 2018

September 28, 2018

No reboot Operating System Update using OpenStack Technologies

Sanchita Sinha

(IBM India Pvt LTD)

Problem Statement

- 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.
- Few prior solutions available to avoid system reboot while applying update were
  - concurrent update-enabled fixes
  - Kernel patching
- But, above mentioned solutions had several limitations
  - no scalability,
  - no capability of maintenance level update
  - not persistent on reboot
  - deployment of only limited Kernel related fixes on a running VM
  - could not be generalized to all Kernel or Kernel Extensions fixes

Problem Statement

- 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 without a system reboot and without causing any application downtime.
- Live update is scalable, persistent on reboot, not limited to only certain kernel/kernel extensions, has the capability of maintenance level update.
- To Perform live update operation, one major challenge is the admin effort to setup the environment.
  - 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.
  - The VM which needs to undergo live update, should have all its disks multipath enabled.

Problem Statement

- Hence Admin needs to set multipath for all disk in each VM.
- 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.
- Live update is required to have provided additional disks for surrogate bootup and mirroring operations.
- All these adds on cost in form of maintenance of additional storage, human resources, additional time, building skills etc.

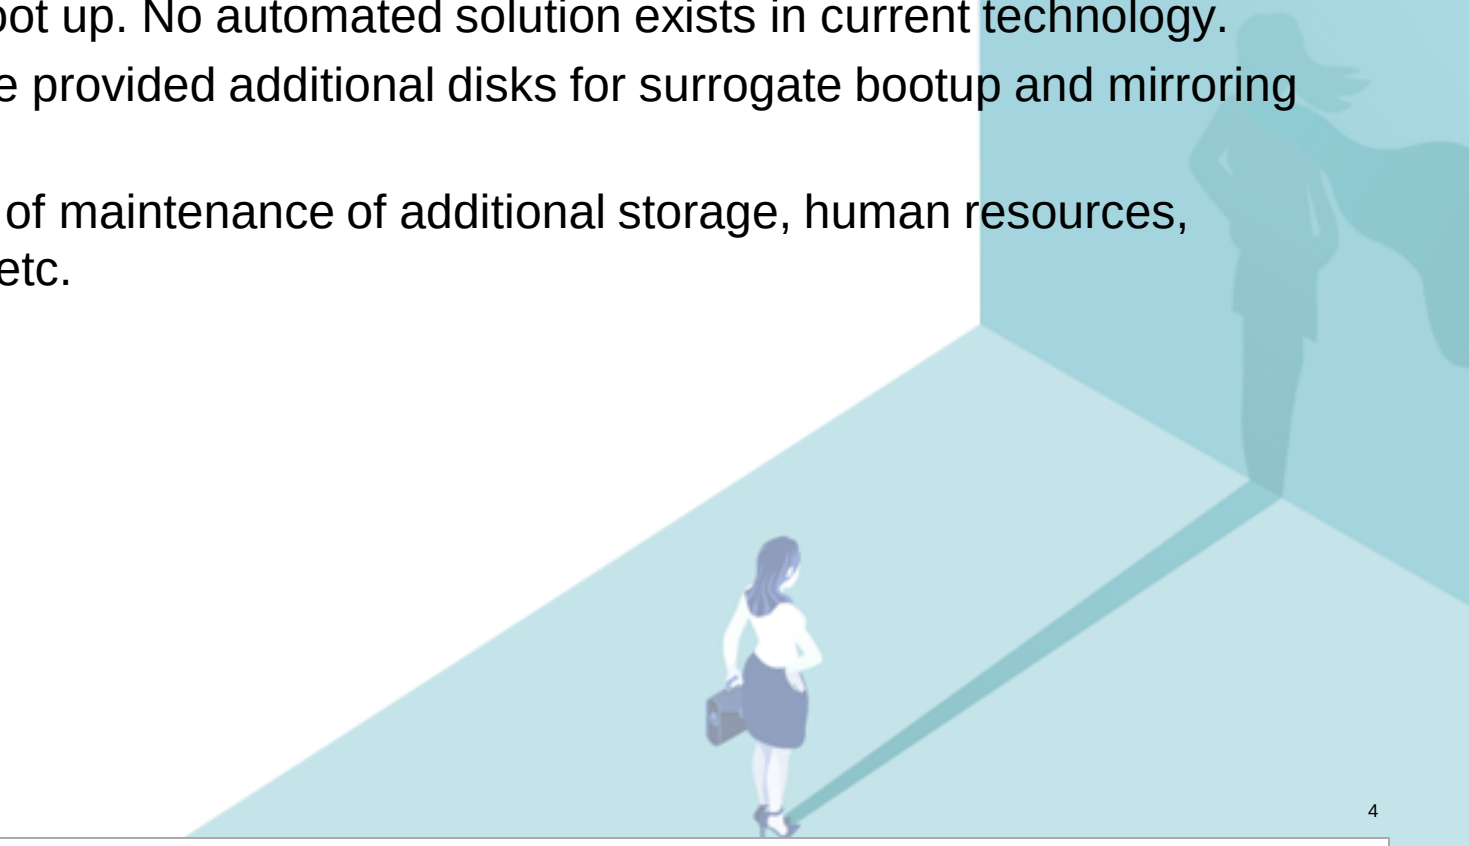

Your solution approach

- 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.
- 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.

Your solution approach

- 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.
- 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.
- 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.
- 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.

Your solution approach

- 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.

o add text)

Your solution approach

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

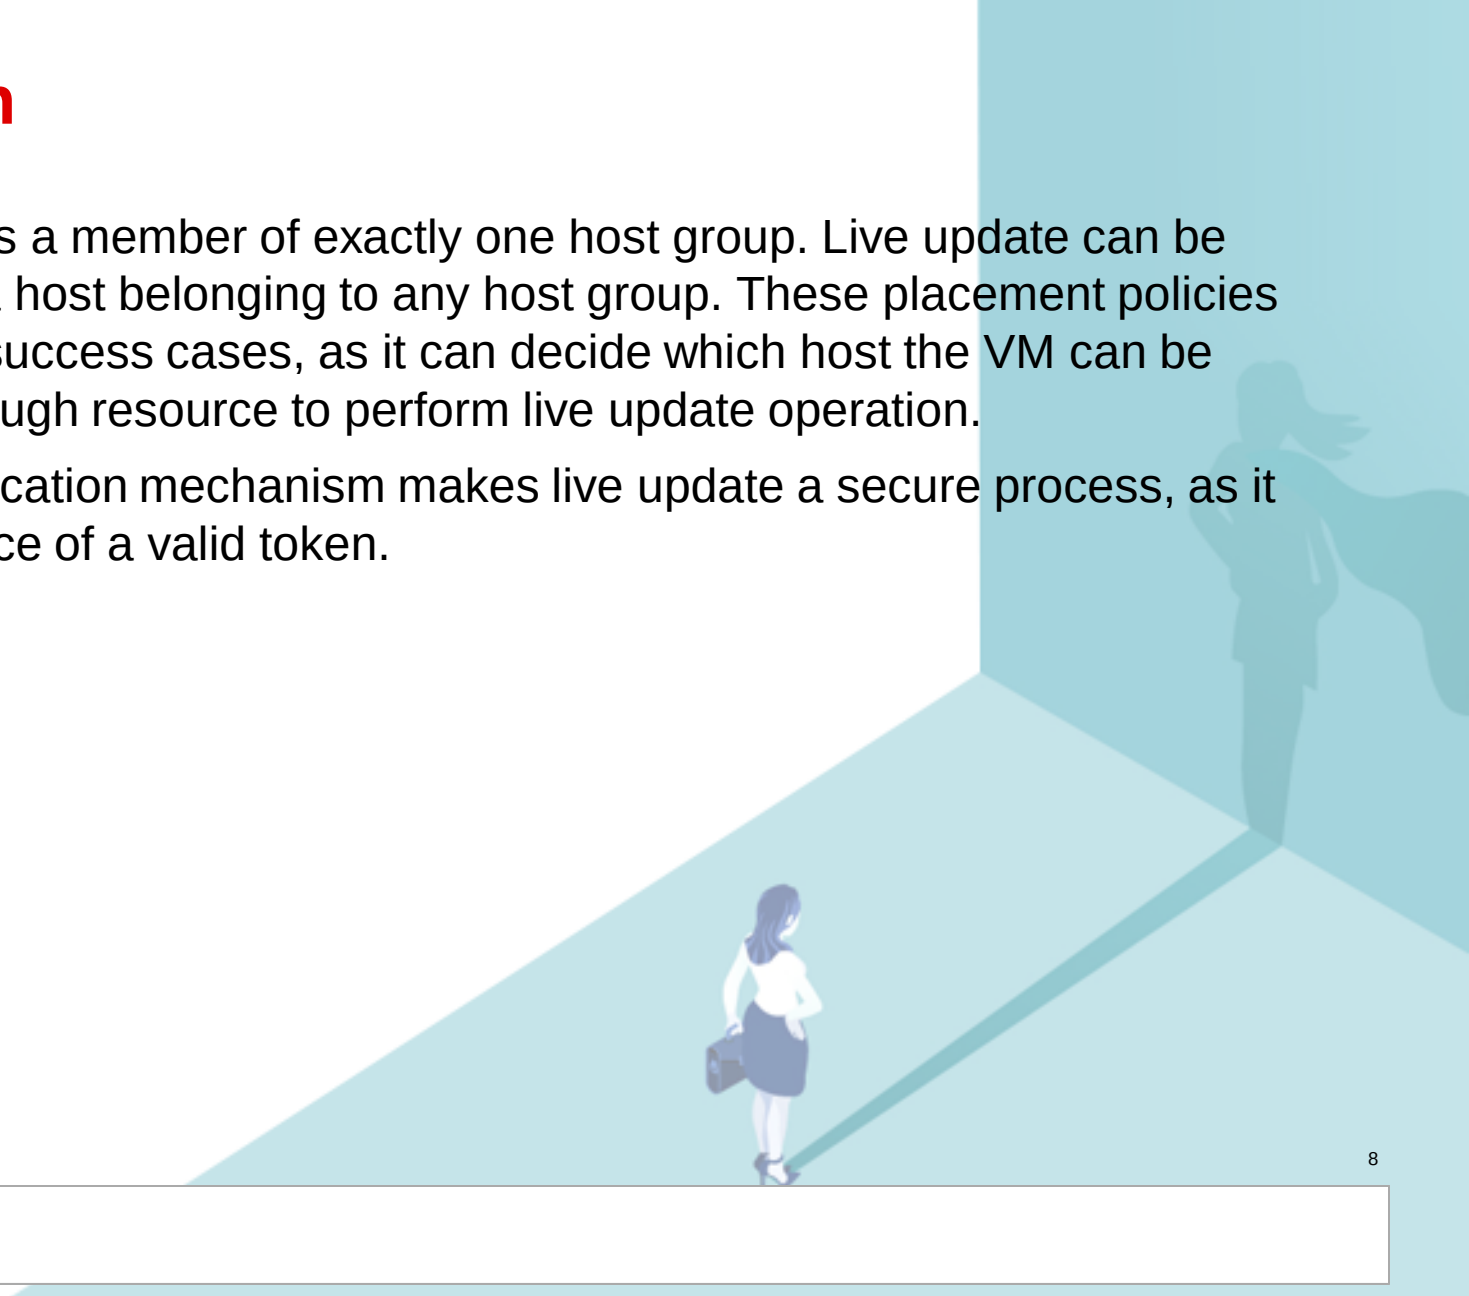

Your solution approach

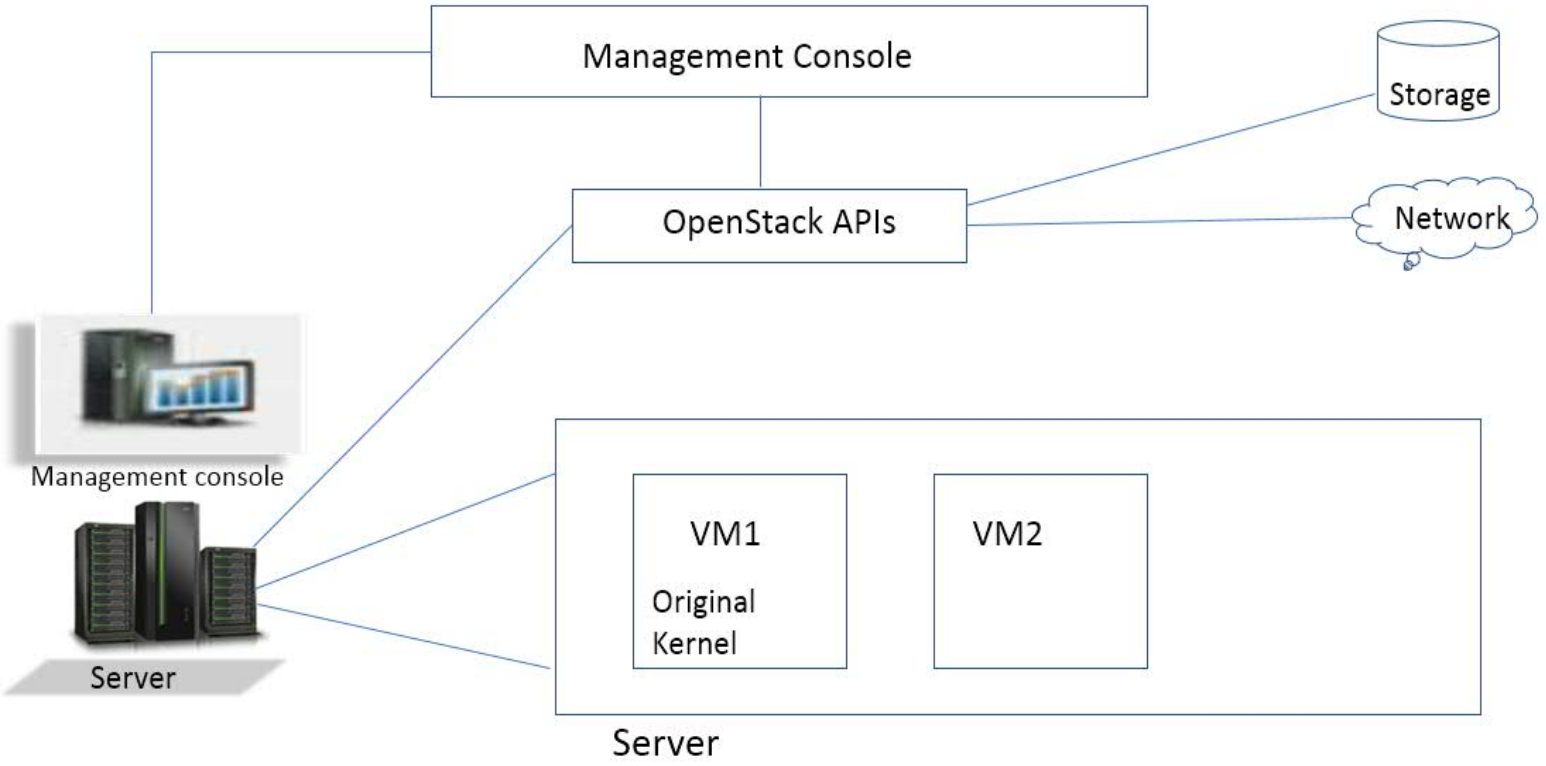

Figure 1: OpenStack managed Servers

Your solution approach

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.

Your solution approach

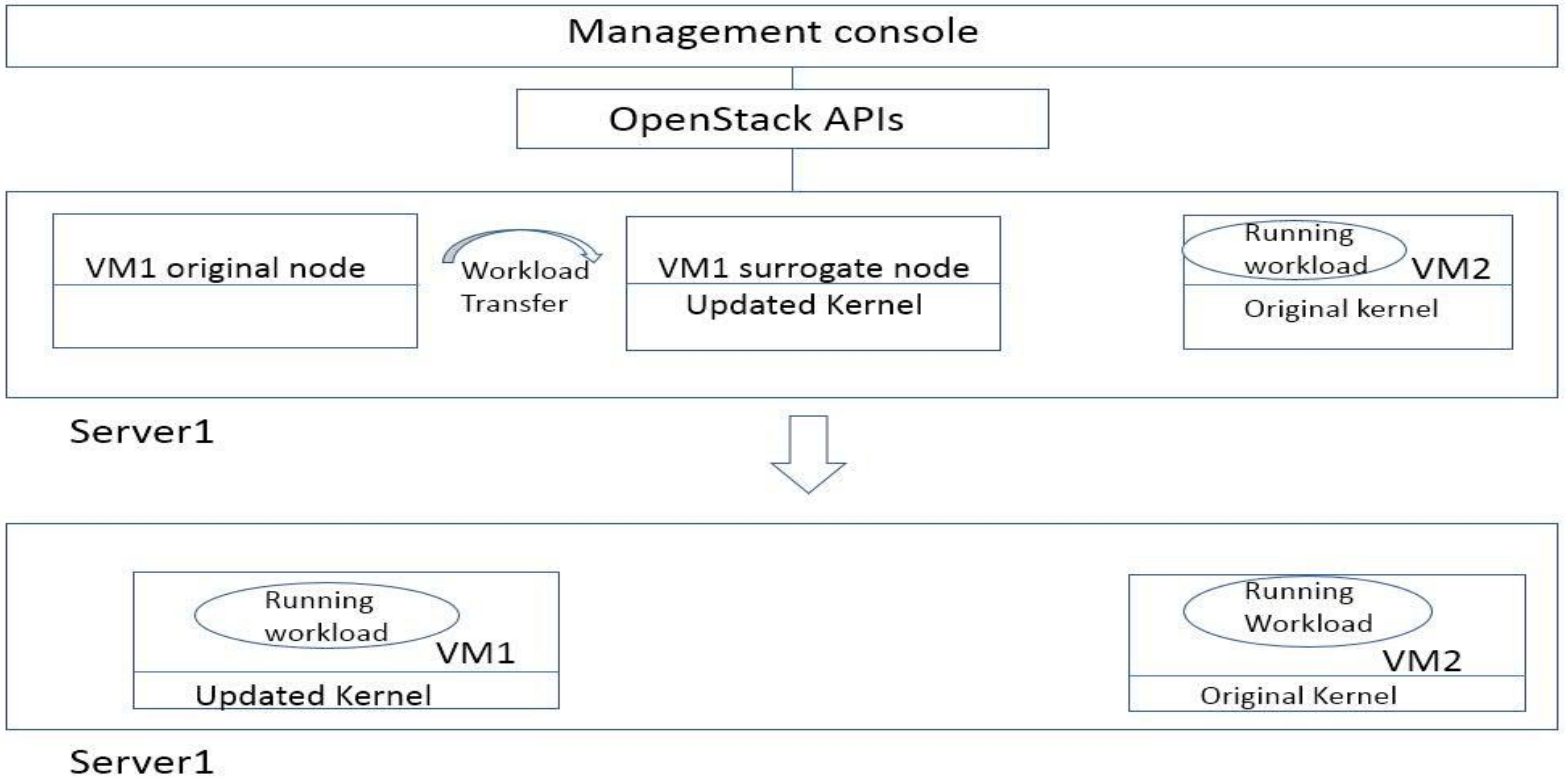

Figure 2: Live update Using OpenStack Technologies

Your solution approach

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.

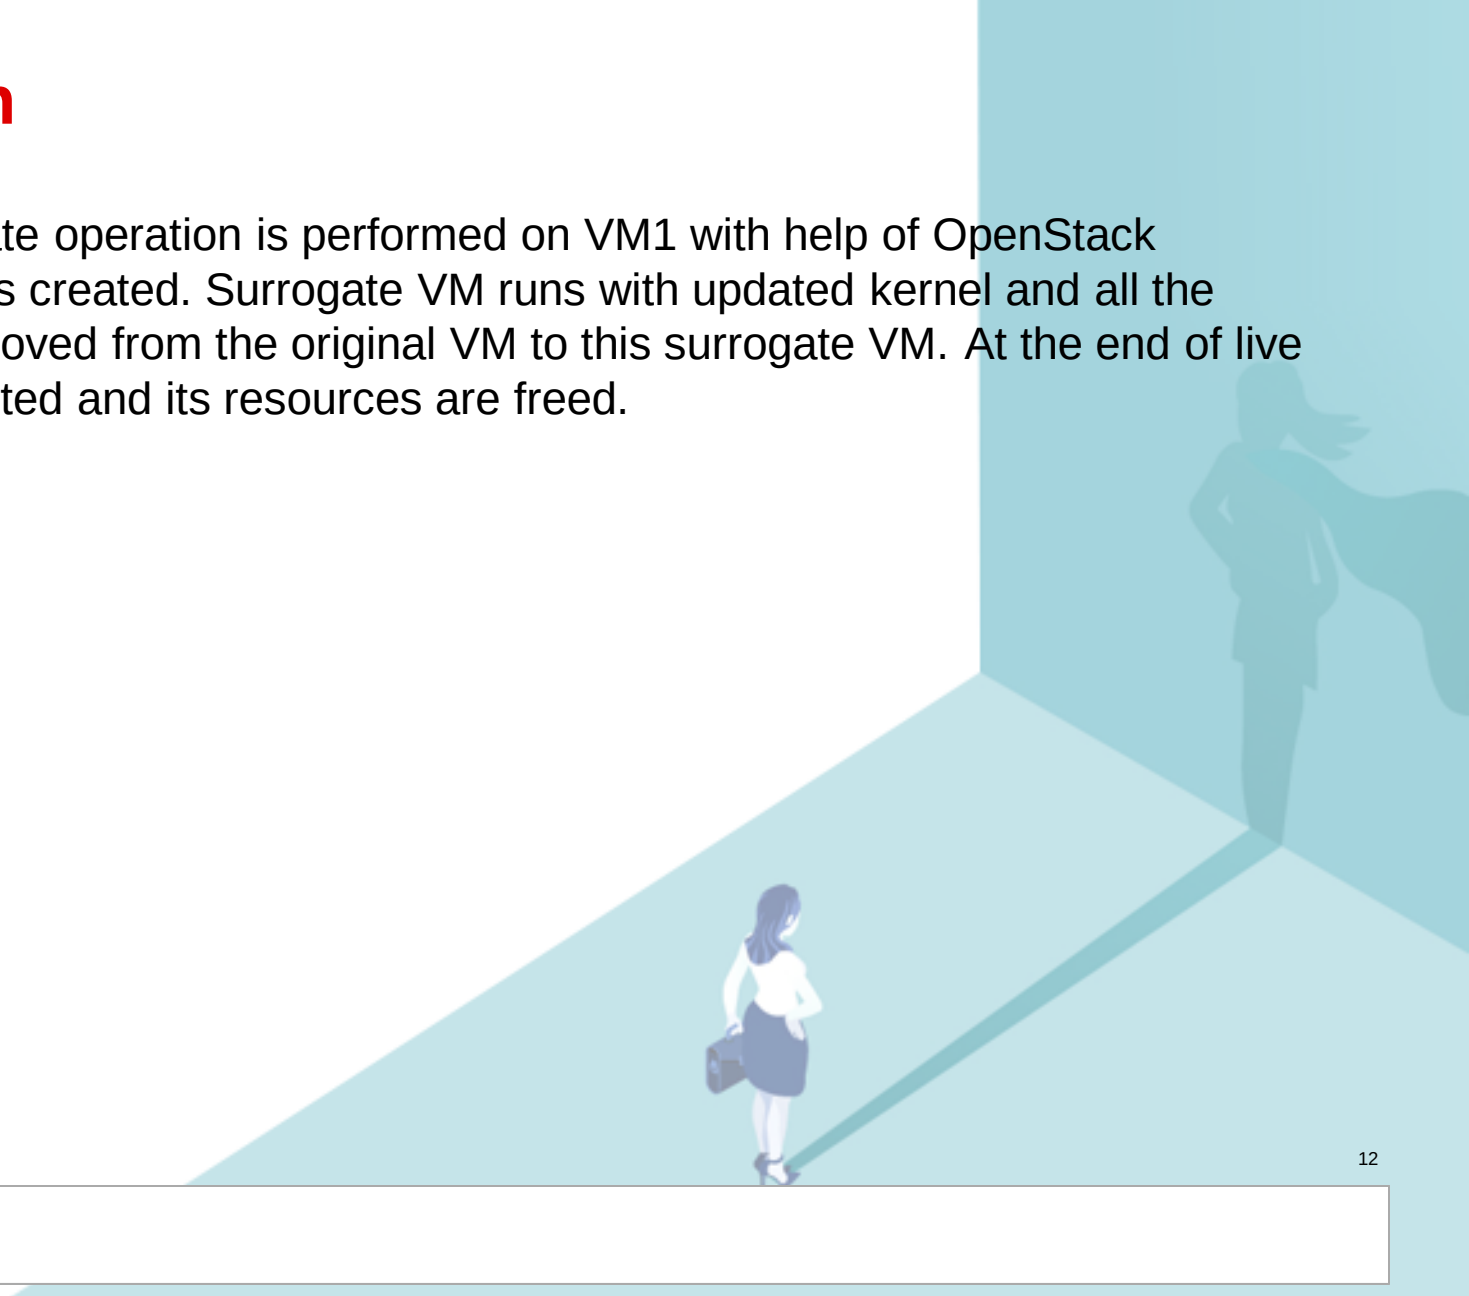

Experimental Results

- No reboot Operating system update
- No application downtime during system update
- Update available for kernel/kernel extension/ service packs
- Reduces manual effort
- No additional storage cost
- Ease of management of storage and fibre channel fabric
- No manual multipath setup
- All additional disk cleaned up post operation completion
- Easy deployment of surrogate LPAR
- Simultaneous update possible for several VM together
- Complete Roll-back in case of update failure
- Authentication based secure process

Experimental Results

- Roll back mechanism is available to withdraw the only the fix without changing the actual build label
- Mechanism is available to go back to original build level in case of service packs
- All supported applications running on VM are unaffected post update
- Running process and their states are completely moved post update
- No closure of application required
- All updates are persistent on system reboot

Conclusions

This presentation 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.

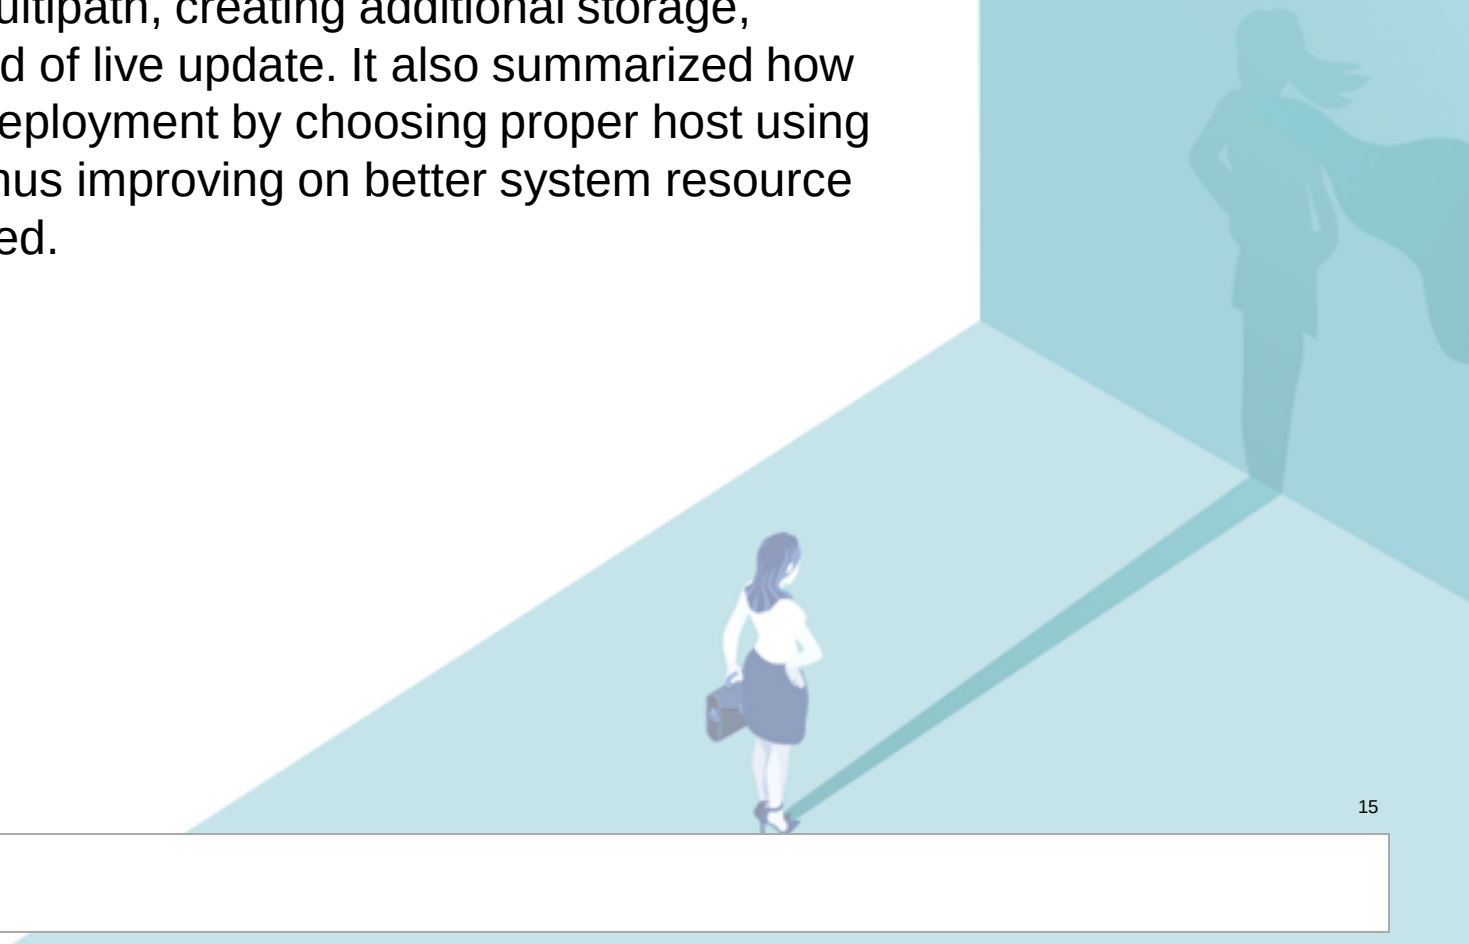

References

- [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/