

3 Virtualization Facilities (Host OS, Hypervisor)

3.1 Requirements on Host OS and Hypervisor and Storage Requirements: - The hypervisor should support distributed HA mechanism - Hypervisor should detect the failure of the VM. Failure of the VM should be reported to

the VIM within 1s

- The hypervisor should report (and if possible log) its failure and recovery action. and the destination to whom they are reported should be configurable.
- The hypervisor should support VM migration
- The hypervisor should provide isolation for VMs, so that VMs running on the same hardware do not impact each other.
- The host OS should provide sufficient process isolation so that VMs running on the same hardware do not impact each other.
- The hypervisor should record the VM information regularly and provide logs of VM actions for future diagnoses.
- The NFVI should maintain the number of VMs provided to the VNF in the face of failures. I.e. the failed VM instances should be replaced by new VM instances

3.2 Requirements on Middlewares Requirements: - It should be possible to detect and automatically recover from hypervisor failures

without the involvement of the VIM

- Failure of the hypervisor should be reported to the VIM within 1s
- Notifications about the state of the (distributed) storage backends shall be send to the VIM (in-synch/healthy, re-balancing/re-building, degraded).
- Process of VIM runing on the compute node should be monitored, and failure of it should be notified to the VIM within 1s
- Fault detection and reporting capability. There should be middlewares supporting in-band reporting of HW failure to VIM.
- Storage data path traffic shall be redundant and fail over within 1 second on link failures.
- Large deployments using distributed software-based storage shall separate storage and compute nodes (non-hyperconverged deployment).
- Distributed software-based storage services shall be deployed redundantly.
- Data shall be stored redundantly in distributed storage backends.
- Upon failures of storage services, automatic repair mechanisms (re-build/re-balance of data) shall be triggered automatically.
- The storage backend shall support geo-redundancy.