6 VNF High Availability

6.1 Service Availability

In the context of NFV, Service Availability refers to the End-to-End (E2E) Service Availability which includes all the elements in the end-to-end service (VNFs and infrastructure components) with the exception of the customer terminal such as handsets, computers, modems, etc. The service availability requirements for NFV should be the same as those for legacy systems (for the same service).

Service Availability = total service available time / (total service available time + total service recovery time)

The service recovery time among others depends on the number of redundant resources provisioned and/or instantiated that can be used for restoring the service.

In the E2E relation a Network Service is available only of all the necessary Network Functions are available and interconnected appropriately to collaborate according to the NF chain.

General Service Availability Requirements

- We need to be able to define the E2E (V)NF chain based on which the E2E availability requirements can be decomposed into requirements applicable to individual VNFs and their interconnections
- The interconnection of the VNFs should be logical and be maintained by the NFVI with guaranteed characteristics, e.g. in case of failure the connection should be restored within the acceptable tolerance time
- These characteristics should be maintained in VM migration, failovers and switchover, scale in/out, etc. scenarios
- It should be possible to prioritize the different network services and their VNFs. These priorities should be used when pre-emption policies are applied due to resource shortage for example.
- VIM should support policies to prioritize a certain VNF.
- VIM should be able to provide classified virtual resources to VNFs in different SAL

6.1.1 Service Availability Classification Levels

The [ETSI-NFV-REL] defined three Service Availability Levels (SAL) are classified in Table 1. They are based on the relevant ITU-T recommendations and reflect the service types and the customer agreements a network operator should consider.

Table 1: Service Availability classification levels

SAL Type	Customer Type	Service/Function	Notes
Level 1	Network Operator Control Traffic Government/ Regulatory Emergency Services	• Intra-carrier engineering traffic • Emergency telecommunication service (emergency response, emergency dispatch) • Critical Network Infrastructure Functions (e.g VoLTE functions DNS Servers,etc.)	Sub-levels within Level 1 may be created by the Network Operator depending on Customer demands E.g.: • 1A - Control; • 1B - Real-time; • 1C - Data; May require 1+1 Redundancy with Instantaneous Switchover

Level 2	Enterprise and/ or large scale customers (e.g. Corporations, University) Network Operators (Tier1/2/3) service traffic	• VPN • Real-time traffic (Voice and video) • Network Infrastructure Functions supporting Level 2 services (e.g. VPN servers, Corporate Web/ Mail servers)	Sub-levels within Level 2 may be created by the Network Operator depending on Customer demands. E.g.: • 2A - VPN; • 2B - Real-time; • 2C - Data; May require 1:1 Redundancy with Fast (maybe Instantaneous) Switchover
Level 3	General Consumer Public and ISP Traffic	• Data traffic (including voice and video traffic provided by OTT) • Network Infrastructure Functions supporting Level 3 services	While this is typically considered to be "Best Effort" traffic, it is expected that Network Operators will devote sufficient resources to assure "satisfactory" levels of availability. This level of service may be pre-empted by those with higher levels of Service Availability. May require M+1 Redundancy with Fast Switchover; where M > 1 and the value of M to be determined by further study

Requirements

- It shall be possible to define different service availability levels
- It shall be possible to classify the virtual resources for the different availability class levels
- The VIM shall provide a mechanism by which VNF-specific requirements can be mapped to NFVI-specific capabilities.

More specifically, the requirements and capabilities may or may not be made up of the same KPI-like strings, but the cloud administrator must be able to configure which HA-specific VNF requirements are satisfied by which HA-specific NFVI capabilities.

6.1.2 Metrics for Service Availability

The [ETSI-NFV-REL] identifies four metrics relevant to service availability:

- Failure recovery time,
- Failure impact fraction,
- Failure frequency, and
- Call drop rate.

6.1.2.1 Failure Recovery Time

The failure recovery time is the time interval from the occurrence of an abnormal event (e.g. failure, manual interruption of service, etc.) until the recovery of the service regardless if it is a scheduled or unscheduled abnormal event. For the unscheduled case, the recovery time includes the failure detection time and the failure restoration time. More specifically restoration also allows for a service recovery by the restart of the failed provider(s) while failover implies that the service is recovered by a redundant provider taking over the service. This provider may be a standby (i.e. synchronizing the service state with the active provider) or a spare (i.e. having no state information). Accordingly failover also means switchover, that is, an orederly takeover of the service from the active provider by the standby/spare.

Requirements

- It should be irrelevant whether the abnormal event is due to a scheduled or unscheduled operation or it is caused by a fault.
- Failure detection mechanisms should be available in the NFVI and configurable so that the target recovery times can be met
- Abnormal events should be logged and communicated (i.e. notifications and alarms as appropriate)

The TL-9000 forum has specified a service interruption time of 15 seconds as outage for all traditional telecom system services. [ETSI-NFV-REL] recommends the setting of different thresholds for the different Service Availability Levels. An example setting is given in the following table 2. Note that for all Service Availability levels Real-time Services require the fastest recovery time. Data services can tolerate longer recovery times. These recovery times are applicable to the user plane. A failure in the control plane does not have to impact the user plane. The main concern should be simultaneous failures in the control and user planes as the user plane cannot typically recover without the control plane. However an HA mechanism in VNF itself can further mitigate the risk. Note also that the impact on the user plane depends on the control plane service experiencing the failure, some of them are more critical than others.

Table 2: Example service recovery times for the service availability levels

SAL	Service Recovery Time Threshold	Notes
1	5 - 6 seconds	Recommendation: Redundant resources to be made available on-site to ensure fast recovery.
2	10 - 15 seconds	Recommendation: Redundant resources to be available as a mix of on-site and off- site as appropriate. • On-site resources to be utilized for recovery of real-time services. • Off-site resources to be utilized for recovery of data services.
3	20 - 25 seconds	Recommendation: Redundant resources to be mostly available off-site. Real-time services should be recovered before data services

6.1.2.2 Failure Impact Fraction

The failure impact fraction is the maximum percentage of the capacity or user population affected by a failure compared with the total capacity or the user population supported by a service. It is directly associated with the failure impact zone which is the set of resources/elements of the system to which the fault may propagate.

Requirements

- It should be possible to define the failure impact zone for all the elements of the system
- At the detection of a failure of an element, its failure impact zone must be isolated before the associated recovery mechanism is triggered
- If the isolation of the failure impact zone is unsuccessful the isolation should be attempted at the next higher level as soon as possible to prevent fault propagation.
- It should be possible to define different levels of failure impact zones with associated isolation and alarm generation policies
- It should be possible to limit the collocation of VMs to reduce the failure impact zone as well as to provide sufficient resources

6.1.2.3 Failure Frequency

Failure frequency is the number of failures in a certain period of time.

Requirements

- There should be a probation period for each failure impact zones within which failures are correlated.
- The threshold and the probation period for the failure impact zones should be configurable
- It should be possible to define failure escalation policies for the different failure impact zones

6.1.2.4 Call Drop Rate

Call drop rate reflects service continuity as well as system reliability and stability. The metric is inside the VNF and therefore is not specified further for the NFV environment.

Requirements

- It shall be possible to specify for each service availability class the associated availability metrics and their thresholds
- It shall be possible to collect data for the defined metrics
- It shall be possible to delegate the enforcement of some thresholds to the NFVI
- Accordingly it shall be possible to request virtual resources with guaranteed characteristics, such as guaranteed latency between VMs (i.e. VNFCs), between a VM and storage, between VNFs

6.2 Service Continuity

The determining factor with respect to service continuity is the statefulness of the VNF. If the VNF is stateless, there is no state information which needs to be preserved to prevent the perception of service discontinuity in case of failure or other disruptive events. If the VNF is stateful, the NF has a service state which needs to be preserved throughout such disruptive events in order to shield the service consumer from these events and provide the perception of service continuity. A VNF may maintain this state internally or externally or a combination with or without the NFVI being aware of the purpose of the stored data.

Requirements

- 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
- It should be possible to specify whether the NFVI or the VNF/VNFM handles the service recovery and continuity

- If the VNF/VNFM handles the service recovery it should be able to receive error reports and/or detect failures in a timely manner.
- The VNF (i.e. between VNFCs) may have its own fault detection mechanism, which might be triggered prior to receiving the error report from the underlying NFVI therefore the NFVI/VIM should not attempt to preserve the state of a failing VM if not configured to do so
- The VNF/VNFM should be able to initiate the repair/reboot of resources of the VNFI (e.g. to recover from a fault persisting at the VNF level => failure impact zone escalation)
- It should be possible to disallow the live migration of VMs and when it is allowed it should be possible to specify the tolerated interruption time.
- It should be possible to restrict the simultaneous migration of VMs hosting a given VNF
- It should be possible to define under which circumstances the NFV-MANO in collaboration with the NFVI should provide error handling (e.g. VNF handles local recoveries while NFV-MANO handles geo-redundancy)
- The NFVI/VIM should provide virtual resource such as storage according to the needs of the VNF with the required guarantees (see virtual resource classification).
- The VNF shall be able to define the information to be stored on its associated virtual storage
- It should be possible to define HA requirements for the storage, its availability, accessibility, resilience options, i.e. the NFVI shall handle the failover for the storage.
- The NFVI shall handle the network/connectivity failures transparent to the VNFs
- The VNFs with different requirements should be able to coexist in the NFV Framework
- The scale in/out is triggered by the VNF (VNFM) towards the VIM (to be executed in the NFVI)
- It should be possible to define the metrics to monitor and the related thresholds that trigger the scale in/out operation
- Scale in operation should not jeopardize availability (managed by the VNF/VNFM), i.e. resources can only be removed one at a time with a period in between sufficient for the VNF to restore any required redundancy.