ETSI GR NFV-EVE 020 V5.1.1 (2022-10)

GROUP REPORT

Network Functions Virtualisation (NFV) Release 5; Evolution and Ecosystem; Report on NFV support for Network Function Connectivity eXtensions

Disclaimer

Reference

DGR/NFV-EVE020

Keywords

MANO, network, NFV, virtualisation

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - APE 7112B
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° w061004871

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The declarations pertaining to these essential IPRs, if any, are publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI Directives including the ETSI IPR Policy, no investigation regarding the essentiality of IPRs, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are trademarks of ETSI registered for the benefit of its Members. **3GPP™** and **LTE™** are trademarks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners. **oneM2M™** logo is a trademark of ETSI registered for the benefit of its Members and of the oneM2M Partners. **GSM**® and the GSM logo are trademarks registered and owned by the GSM Association.

Foreword

This Group Report (GR) has been produced by ETSI Industry Specification Group (ISG) Network Functions Virtualisation (NFV).

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

1 Scope

The present document studies the integration of Dyncast in the NFV architectural framework. Dyncast despatches data flows according to network status and computing resources. The impact on NFV such as necessary extensions of NFV-MANO components is analysed.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ETSI GS NFV-IFA 011 (V3.3.1): "Network Functions Virtualisation (NFV) Release 3; Management and Orchestration; VNF Descriptor and Packaging Specification".

[i.2] ETSI GR NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[i.3] ETSI GS NFV-TST 008 (V3.3.1): "Network Functions Virtualisation (NFV) Release 3; Testing; NFVI Compute and Network Metrics Specification".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the terms given in ETSI GR NFV 003 [i.2] apply.

3.2 Symbols

Void

3.3 Abbreviations

For the purposes of the present document, the abbreviations given in ETSI GR NFV 003 [i.2] and the following apply:

NOTE: An abbreviation defined in the present document takes precedence over the definition of the same abbreviation, if any, in ETSI GR NFV 003 [i.2].

NFCX Network Function Connectivity eXtensions

4 Overview

The present document presents Dyncast in order to improve resource optimization, which dynamically dispatches packets based on computing resources and network status.

Traditional networks provide unicast, multicast, and broadcast, which are widely used. Network Function Connectivity eXtensions (NFCX) introduces the extension of network connectivity capabilities to benefit distributed computing systems. The NFCX can be regarded as a new communication model by combining computing and network resources to efficiently convey the multi-point-to-multi-point traffic model. The typical examples include the Dyncast, Convergecast and Collectivecast.

Dyncast assumes there are multiple equivalent VNF instances which provide the same service. These VNF instances are deployed on different NFVI-PoPs or the same NFVI-PoP. Meanwhile, these instances are interconnected and collaborate with each other to achieve a holistic objective. When a service request arrives, the first packet of this data flow is delivered to the proper VNF instance according to computing resource consumption and network status. The subsequent packets from the same flow follow the first one.

Clause 5 gives details of use cases, where clause 5.1 introduces the use case that VNF instances are deployed on different NFVI-PoPs and clause 5.2 describes the use case that VNF instances are on the same NFVI-PoP. Clause 6 presents analysis of use cases. Clauses 7 and 8 show recommendations and conclusions of the present document respectively.

5 Use Cases

5.1 Inter-Site Multiple VNFs/VNFCs Connectivity via Dyncast

5.1.1 Motivation

This use case describes that Dyncast performs inter-site dynamic network connectivity between multiple equivalent VNFs. It is assumed that the multiple VNFs are deployed in different NFV-PoPs.

Figure 5.1.1-1 illustrates a high-level architecture for inter-site VNFs connectivity via Dyncast.

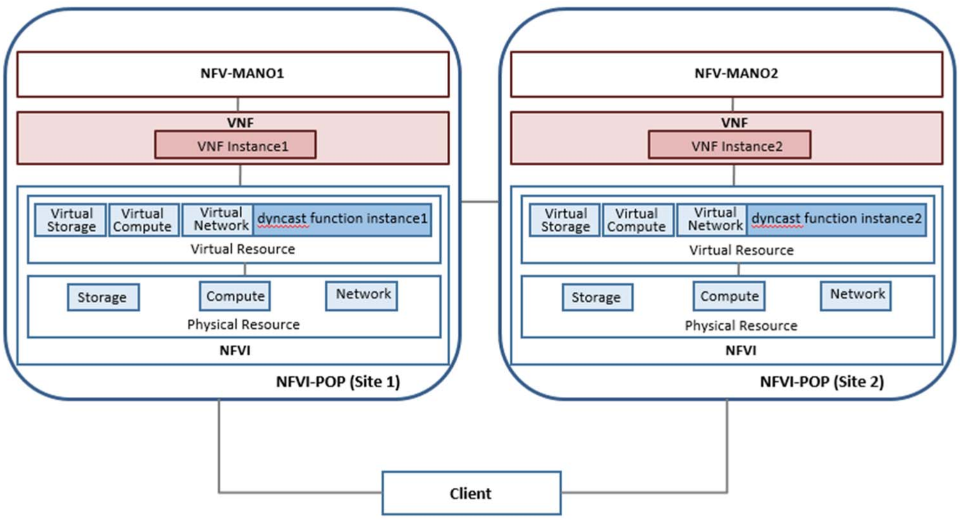

Figure 5.1.1-1: Inter-site VNFs connectivity via Dyncast

In this case, multiple equivalent VNF instances deployed in different NFV-PoPs are abstracted and virtualised into one group to provide a unified virtual interface for users. Dyncast performs dynamic connectivity, optimal selection and dispatching between different VNF instances which form a pool of virtual network functions.

In this case, using Dyncast is close to introducing a load balancer implemented at the NFVI level. Dyncast directs traffic between different NFVI-PoPs to support cross-site VNFs collaboration, which means that a proper site can be selected through Dyncast during session establishment, and/or a session of one site can be shifted to other sites through Dyncast during the session duration phase in case the original site fails to work properly, e.g. network congestion.

Therefore, multiple equivalent VNF/VNFC instances distributed across multiple sites and interconnected via Dyncast can achieve better scalability, utilization efficiency of system resources and service performance through collaboration.

5.1.2 Detailed User Story

5.1.2.1 Summary

In this case, multiple equivalent VNF instances are deployed in different NFVI-PoPs and considered as one group to provide one kind of service for clients. Dyncast performs dynamic connectivity between clients and different VNF instances according to computing resources and network status.

5.1.2.2 Actor(s)

Table 5.1.2.2-1 describes the use case actors and roles.

Table 5.1.2.2-1: Use case #1, actors and roles

#	Actor	Description
1	NFV-MANO1	NFV-MANO1 managing the NFVI (site1) hosting all resources involved
2	NFV-MANO2	NFV-MANO2 managing the NFVI (site2) hosting all resources involved
3	VNF instance1	VNF instance1 is the first possible target
4	VNF instance2	VNF instance2 is the second possible target
5	Dyncast function instance1	Dyncast function instance1 implementing Dyncast functionality at site 1. See note
6	Dyncast function instance2	Dyncast function instance2 implementing Dyncast functionality at site 2
7	Client	The user who sends a request of network function, and a client can be a network function connected to an NFV environment
NOTE: The Dyncast function instance1 performs dynamic and anycast network connectivity, which dynamically direct flows coming from client to a suitable instance of the VNF.		

5.1.2.3 Pre-Conditions

Table 5.1.2.3-1 describes the pre-conditions.

Table 5.1.2.3-1: Use case #1, pre-conditions

#	Pre-condition	Description
1	VNF instance1 and VNF instance2 implementing the same function are instantiated at site1 and site2 respectively.	
2	Dyncast function instance1 and Dyncast function instance2 are deployed and running at site1 and site2 respectively.	
3	NFV-MANO1 and NFV-MANO2 are both running.	
4	NFV-MANO1 notify Dyncast function instance1 about the network function information when VNF instance1 is instantiated. See note.	
5	NFV-MANO2 notify Dyncast function instance2 about the network function information when VNF instance2 is instantiated.	
NOTE: The network function information can be function name, the locator indicating where the instance is running (e.g. unicast IP address), the required monitoring parameters applicable to the VNF, etc. The function name can be used as the global unique identifier of a function.		

5.1.2.4 Description

Table 5.1.2.4-1 describes the basic flow for function addressing and dispatching. The Dyncast function instance perceives all the packets of a flow coming from the client. When Dyncast function instance receives the first packet, it selects which target of VNF instances to perform the request base on the information on the network status and the monitoring parameters collected by the NFV-MANO, and then continuously steers the traffic toward the target VNF instance for the subsequent packets, following the decision for the first packet.

Table 5.1.2.4-1: Base flow for function addressing and dispatching

#	Flow	Description
1	NFV-MANO1 -> Dyncast function instance1	NFV-MANO1 provides the collected monitoring parameters information of the VNF instance1 to Dyncast function instance1. See note 1.
2	NFV-MANO2 -> Dyncast function instance2	NFV-MANO2 provides the collected monitoring parameters information of the VNF instance2 to Dyncast function instance2. See note 2.
3	Dyncast function instance1 -> Dyncast function instance2	Dyncast function instance1 and Dyncast function instance2 exchange the function information and the associated monitoring parameters information of the VNF instances deployed and running in NFVI-POPs. See note 3.
4	Client -> Dyncast function instance1	Dyncast function instances can receive and identify service request from the client by an IP anycast address for a service. When the client sends a request to a network function, a Dyncast function instance receives the first packet of the flow and selects the target between the VNF instance1 and VNF instance2, according to the monitoring parameters information and network status information. See note 4. In this case, the Dyncast function instance1 receives the first packet and selects VNF instance2 as the target, because VNF instance2 can provide better service according to the information of the monitoring parameters and network status. See note 5.
5	Dyncast function instance1 -> Dyncast function instance2	The Dyncast function instance1 transfers this first packet of the flow to Dyncast function instance2.
6	Dyncast function instance2 -> VNF instance2	The Dyncast function instance2 forwards this first packet of the flow to the selected VNF instance2.
7	Client -> Dyncast function instance1	The Dyncast function instance1 receives the following subsequent packets of the flow coming from the client.
8	Dyncast function instance1 -> Dyncast function instance2	In accordance with the decision for the first packet, the Dyncast function instance1 constantly steers the traffic toward the same target of VNF instance2 for the subsequent packets. Therefore, the Dyncast function instance1 transfers the following subsequent packets of the flow to Dyncast function instance2.
9	Dyncast function instance2 -> VNF instance2	The Dyncast function instance2 forwards the subsequent packets of the flow to the selected VNF instance2.
NOTE 1: The monitoring parameters information is related to the usage of NFVI resources by VNF. e.g. CPU/GPU/memory utilization, etc. NOTE 2: Step 1 and step 2 can be performed concurrently. NOTE 3: Different mechanisms can be used in updating the monitoring parameters information, e.g. extended BGP, extended IGP or controller based mechanism. NOTE 4: The network status is related to the network performance, e.g. latency, packet loss rate, throughput, etc. NOTE 5: In principle, the client can be logically connected to any Dyncast function instance. In this case, the Dyncast function instance1 is selected by underlay as the receiver of the client's request through its anycast address.		

5.1.2.5 Post-Conditions

Table 5.1.2.5-1 describes the post-conditions.

Table 5.1.2.5-1: Use case #1, post-conditions

#	Post-condition	Description
1	VNF instance2 processes and serves the request.	

5.1.3 Analysis

As shown in the flow in clause 5.1.2.4, Dyncast addressing and dispatching demands taking into account monitoring parameters information.

Table 5.1.3-1 lists the Dyncast monitoring parameters information, including the usage of compute and memory resources which can be derived from the performance metrics defined in ETSI GS NFV-TST 008 [i.3], and the quality of the networks between Dyncast nodes which can be measured by Dyncast monitoring modules.

Table 5.1.3-1: Dyncast monitoring parameters information

Parameters type	Parameters name	Description
compute	cpu_utilization	The percentage of CPU utilization quantifies the ratio of time of processor usage to the time in the measurement interval, which can be derived from Processor Utilization metric (see clause 6.6 in ETSI GS NFV-TST 008 [i.3]).
memory	memory_utilization	The percentage of memory utilization in the measurement interval that can be derived from Memory Used (see clause 8.6 in ETSI GS NFV-TST 008 [i.3]), and the amount of memory allocated.
network	mean latency	The average round-trip time in the measurement interval between Dyncast nodes, round-trip time (RTT) is the time it takes for a packet to go from the sending endpoint to the receiving endpoint and back.
network	peak latency	The peak round-trip time in the measurement interval between Dyncast nodes, round-trip time (RTT) is the time it takes for a packet to go from the sending endpoint to the receiving endpoint and back.
network	mean throughput	The average available bandwidth in the measurement interval between Dyncast nodes (bytes per second).
network	peak throughput	The peak available bandwidth in the measurement interval between Dyncast nodes (bytes per second).
network	mean packet loss rate	The average rate of packets transmission loss in the measurement interval between Dyncast nodes.
network	peak packet loss rate	The peak rate of packets transmission loss in the measurement interval between Dyncast nodes.

The information of virtualised compute and memory resources performance can be provided explicitly to Dyncast by NFVI, and the networks resources status can be measured by Dyncast. The metrics currently specified in ETSI GS NFV-TST 008 [i.3] can support Dyncast monitoring. Therefore this use case can be implemented without additional changes to ETSI GS NFV-TST 008 [i.3].

5.2 Intra-Site Multiple VNFs/VNFCs Connectivity via Dyncast

5.2.1 Motivation

This use case describes that Dyncast performs intra-site dynamic network connectivity between multiple equivalent VNF instances. It is assumed that the multiple VNF instances implementing the same function are deployed in one site.

Figure 5.2.1-1 illustrates a high-level architecture for intra-site VNFs connectivity via Dyncast.

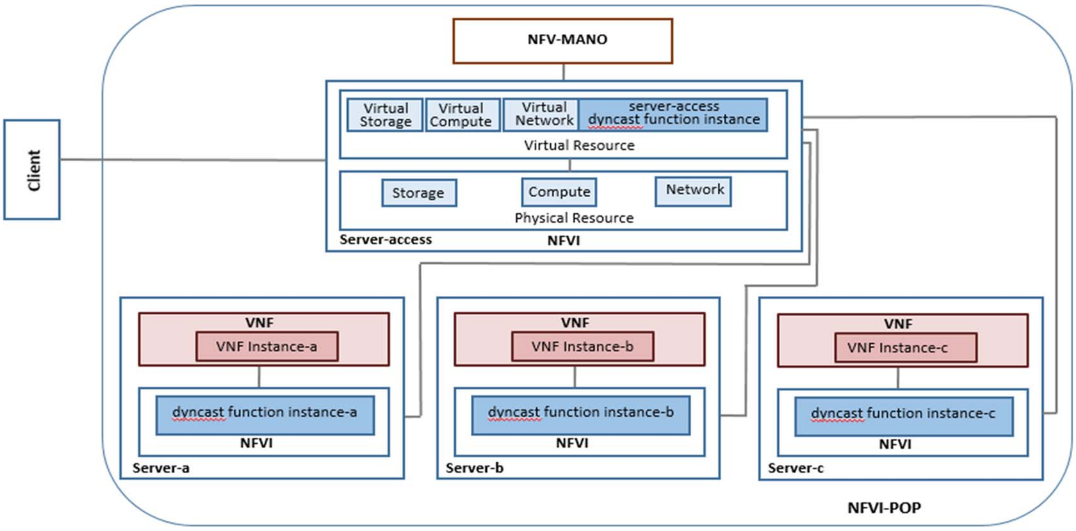

Figure 5.2.1-1: Intra-site VNFs connectivity via Dyncast

In this case, multiple equivalent VNF/VNFC instances deployed in a NFVI-PoP are abstracted and virtualised into one group to provide a unified virtual interface for users, and using a server-access Dyncast function instance is close to introducing a load balancer implemented at the NFVI level, which performs dynamic connectivity, optimal selection and dispatching between the different VNF/VNFC instances in a specific site to support intra-site collaboration. It means that a VNF/VNFC instance can be selected through a server-access Dyncast function instance during session establishment.

Therefore, multiple equivalent VNF/VNFC instances distributed within a single site and interconnected via Dyncast will collaborate and achieve better scalability, utilization efficiency of system resources and service performance.

In this case, the server-access Dyncast function provides selection and dispatching in a NFVI-level balanced way, compared with legacy application-level based load balancing, this provides higher-performance load balancing since Dyncast selects the target network function instance and directly forwards the data at the networking layer rather than getting them processed throughout the full protocol stack.

5.2.2 Detailed User Story

5.2.2.1 Summary

In this use case, it is assumed that multiple equivalent VNF instances implementing the same network service are instantiated on one NFVI-PoP. The use case illustrates how a request can be dispatched from the client to the most appropriate instance, using a Dyncast dynamic routing functionality.

5.2.2.2 Actor(s)

Table 5.2.2.2-1 describes the use case actors and roles.

Table 5.2.2.2-1: Use case #2, actors and roles

#	Actor	Description
1	NFV-MANO	NFV-MANO managing the NFVI hosting all resources involved
2	VNF instance-a	VNF instance-a is the first possible target
3	VNF instance-b	VNF instance-b is the second possible target
4	VNF instance-c	VNF instance-c is the third possible target
5	server-access Dyncast function instance	server-access Dyncast function instance implementing Dyncast functionality on server-access. See note 1
6	Dyncast function instance-a	Dyncast function instance-a implementing Dyncast functionality on server-a. See note 2
7	Dyncast function instance-b	Dyncast function instance-b implementing Dyncast functionality on server-b
8	Dyncast function instance-c	Dyncast function instance-c implementing Dyncast functionality on server-c
9	Client	The user who sends a request of network function, and a client can be a network function connected to an NFV environment
NOTE 1: The server-access Dyncast function instance performs dynamic and anycast network connectivity, which dynamically directs flows coming from client to a suitable instance of the VNF. The server-access is a specified access server in a site to receive all flows that need to be taken over by Dyncast functionality, which can be specified and configured by NFV-MANO. NOTE 2: The Dyncast function instance-a only receives the flow that need to be processed by co-locate VNFs, the same applies to Dyncast function instance-b and Dyncast function instance-c.		

5.2.2.3 Pre-Conditions

Table 5.2.2.3-1 describes the pre-conditions.

Table 5.2.2.3-1: Use case #2, pre-conditions

#	Pre-condition	Description
1	VNF instance-a, VNF instance-b and VNF instance-c implementing the same function are instantiated on server-a, server-b and server-c respectively, which means that VNF instance-a, VNF instance-b and VNF instance-c are equivalent VNFs.	
2	The server-access Dyncast function instance is deployed and running on server-access.	
3	The Dyncast function instance-a, Dyncast function instance-b and Dyncast function instance-c are deployed and running on server-a, server-b and server-c respectively.	
4	NFV-MANO is deployed and running.	
5	NFV-MANO notify server-access Dyncast function instance about the network function information when VNF instances are instantiated. See note.	
NOTE: The network function information can be function name, the locator indicating where the instance is running (e.g. unicast IP address), the required monitoring parameters applicable to the VNF, etc. The function name can be used as the global unique identifier of a function. NFV-MANO provide information only to server-access Dyncast function instance, but it does not need to provide information to the VNF instance-a, VNF instance-b and VNF instance-c.		

5.2.2.4 Description

Table 5.2.2.4-1 describes the basic flow for function addressing and dispatching inside a site. The server-access Dyncast function instance deployed on server-access perceives all the packets of a flow coming from the client. When server-access Dyncast function instance receives the first packet, it selects which target of VNF instances to perform the request base on the information on the network status and the monitoring parameters collected by the NFV-MANO, and then continuously steers the traffic toward the target VNF instance for the subsequent packets, following the decision for the first packet.

Table 5.2.2.4-1: Base flow for function addressing and dispatching

#	Flow	Description
1	NFV-MANO -> server-access Dyncast function instance	NFV-MANO continuously provides the collected monitoring parameters information of the VNF instances to server-access Dyncast function instance. See note 1.
2	Client -> server-access Dyncast function instance	The server-access Dyncast function instance can receive and identify service request from the client through an IP anycast address for a service. See note 2. When the client sends a new request to a network function, the server-access Dyncast function instance receives the first packet of the flow and selects the target between the VNF instance-a, VNF instance-b and VNF instance-c, according to the monitoring parameters information and network status information. See note 3. The server-access Dyncast function instance selects VNF instance-b as the target, because VNF instance-b can provide better service according to the information of the monitoring parameters and network status.
3	server-access Dyncast function instance -> Dyncast function instance-b	The server-access Dyncast function instance transfers this first packet of the flow to Dyncast function instance-b.
4	Dyncast function instance-b -> VNF instance-b	The Dyncast function instance-b forwards this first packet of the flow to the selected VNF instance-b.
5	Client -> server-access Dyncast function instance	The server-access Dyncast function instance receives the following subsequent packets of the flow coming from the client.
6	server-access Dyncast function instance -> Dyncast function instance-b	In accordance with the decision for the first packet, the server-access Dyncast function instance constantly steers the traffic toward the same target of VNF instance-b for the subsequent packets. Therefore, the server-access Dyncast function instance transfers the following subsequent packets of the flow to Dyncast function instance-b.
7	Dyncast function instance-b -> VNF instance-b	The Dyncast function instance-b forwards the subsequent packets of the flow to the selected VNF instance-b.
NOTE 1: The monitoring parameters information is related to the usage of NFVI resources by VNF. e.g. CPU/GPU/memory utilization, etc. NOTE 2: Among Dyncast function instances in one site, only the server-access Dyncast function instance deployed on server-access receives all the traffic flow of the network function request from the client, and performs the dynamic selection among VNF instance-a, VNF instance-b and VNF instance-c. It is recommended that the server-access Dyncast function instance receives request based on an IP anycast address rather than IP unicast address, because there are possibly equal VNFs deployed on other POPs (e.g. user case1), anycast enables any server-access Dyncast instance deployed on different POPs to receive request of client. NOTE 3: The network status is related to the network performance, e.g. latency, packet loss rate, throughput, etc.		

5.2.2.5 Post-Conditions

Table 5.2.2.5-1 describes the post-conditions.

Table 5.2.2.5-1: Use case #2, post-conditions

#	Post-condition	Description
1	VNF instance-b processes and serves the request.	

5.2.3 Analysis

As shown in clause 5.2.2.4, Dyncast addressing and dispatching demand Dyncast monitoring parameters which are listed in Table 5.1.3-1. These parameters include the usage of computing and memory resources and the quality of network service between Dyncast nodes, which can be derived from the performance metrics defined in ETSI GS NFV-TST 008 [i.3].

6 Analysis

Use Case 1 shows that the Dyncast connects multiple equivalent VNF instances which are deployed at different NFV-PoPs. Users could regard these VNF instances as a group, which shares the same IP anycast address. When users send the first packet, the Dyncast selects the proper VNF instance for users by network status and computing resources. The subsequent packets follow the first one. NFV-MANO and NFVI provide parameters about network status and compute resources, which are specified in ETSI GS NFV-IFA 011 [i.1] and ETSI GS NFV-TST 008 [i.3]. Therefore, this use case can be implemented without additional changes to ETSI GS NFV-IFA 011 [i.1] and ETSI GS NFV-TST 008 [i.3].

Use Case 2 presents that the Dyncast connects multiple equivalent VNF instances which are deployed at the same NFV-PoP. VNF instances connected by the Dyncast can be considered as a group, which have the common IP anycast address. When users send the first packet, the Dyncast selects the proper VNF instance for users by network status and computing resources. The subsequent packets follow the first one. NFV-MANO and NFVI provide parameters about network status and computing resources, which are specified in ETSI GS NFV-IFA 011 [i.1] and ETSI GS NFV-TST 008 [i.3]. Therefore, this use case can be implemented without additional changes to ETSI GS NFV-IFA 011 [i.1] and ETSI GS NFV-TST 008 [i.3].

The Dyncast provides dynamic network connectivity for equivalent VNF instances by considering network status and computing resources. Dyncast has no effect on the NFV components and can be implemented without any change to existing NFV-MANO and NFVI specifications.

7 Recommendations

Dyncast can be implemented in the NFVI without any change to existing NFV-MANO and NFVI specifications. There is no recommendation.

8 Conclusions

Dyncast assumes there are multiple equivalent VNF instances to provide a single service and VNF instances are interconnected and collaborate with each other to balance the service load. The Dyncast dispatches packet flows based on computing resource consumption and network status, which can select the optimal VNF instance for the service request. Information about computing resources and network status used in the Dyncast can be obtained from the NFVI. Metrics about computing resources and network status include cpu_utilization, memory_utilization, latency, throughput and packet loss rate, which have been specified in ETSI GS NFV-TST 008 [i.3]. Therefore, Dyncast can be implemented as part of the NFVI.

History

Document history		
V5.1.1	October 2022	Publication