
NeoFS Technical Specification

Architecture and Implementation details

Neo Saint Petersburg Competence Center



Date: March 26, 2026

Revision: 4e630b61

Contents

Introduction	9
Overview	9
Background	10
Technical Requirements	10
Out of Scope	10
Future Goals	10
Architecture overview	11
Design and components	11
Epoch	13
Network Map	13
Storage Policy	14
Filters	15
Selectors	16
Replicas	16
Container Backup Factor	16
Objects	17
Large objects split	17
Object Split V1	18
Object Split V2	19
Object Deletion	22
Tombstone Object	22
Containers	23
Access Control Lists	23
Basic ACL	23
Extended ACL	27
Bearer Token	32
ACL check algorithm	33
Reputation system	35
Trust	35
Algorithm	35
Subjects and Objects of Trust in NeoFS	36
Erasure Coding	36
Overview	36
Settings	37
Object model	37

Part placement	39
Reconstruction	39
Split + EC	40
Inner Ring Nodes	41
Storage Nodes	42
Address format	42
Examples:	42
Garbage Collector	42
Invalid Objects check	43
Marked Objects removal	43
Unpaid containers	43
Object Expiration	44
Storage metrics	44
Storage quotas	44
Quotas administration	44
Hard/Soft quota differences	45
Setting	45
Protocol gateways	46
HTTP	46
S3	46
Access Box scheme	46
sFTP	47
Blockchain components	48
Role of blockchain in the storage system	48
Main chain and FS chain	48
Notary service	49
FS chain Governance	50
Alphabet contracts	50
Alphabet Inner Ring nodes	50
Alphabet contracts invocation	51
Utility token distribution	52
Changing FS chain validators	53
Changing the Inner Ring list	53
NeoFS Smart Contracts	57
alphabet contract	57

audit contract	59
balance contract	60
container contract	65
neofs contract	79
netmap contract	83
processing contract	89
proxy contract	90
reputation contract	91
Balance transfer details encoding	92
Reputation model	94
Configuration	94
Managers	94
Defining a manager for a node	95
Local Trust	95
Subject and Object of a trust	95
Calculating trust	96
Transport	96
Global Trust	97
Subject and Object of a trust	97
Calculating trust	97
Transport	98
Incentive model	99
Data storage payments	99
Basic income	99
Container creation fee	101
Inner Ring candidate fee	102
Withdraw fee	102
NeoFS API v2	103
Nodes and their identification	103
Requests and Responses	104
Signing RPC messages and data structures	104
Stable serialization	104
Signature generation format	105
Signature chaining in requests and responses	106
Message body signature	107
Meta header signature	108

Verification header signature	108
Container service signatures	109
Object service and Session signatures	110
neo.fs.v2.accounting	111
Service “AccountingService”	111
Method Balance	111
Message Decimal	112
neo.fs.v2.acl	112
Message BearerToken	112
Message BearerToken.Body	112
Message BearerToken.Body.TokenLifetime	113
Message EAACLRecord	113
Message EAACLRecord.Filter	114
Message EAACLRecord.Target	115
Message EAACLTable	115
Emun Action	116
Emun HeaderType	116
Emun MatchType	116
Emun Operation	117
Emun Role	117
neo.fs.v2.audit	118
Message DataAuditResult	118
neo.fs.v2.container	119
Service “ContainerService”	119
Method Put	119
Method Delete	120
Method Get	121
Method List	122
Method SetExtendedACL	122
Method GetExtendedACL	123
Method AnnounceUsedSpace	124
Method SetAttribute	124
Method RemoveAttribute	125
Message AnnounceUsedSpaceRequest.Body.Announcement	126
Message RemoveAttributeRequest.Body.Parameters	127
Message SetAttributeRequest.Body.Parameters	127
Message Container	128
Message Container.Attribute	128

neo.fs.v2.link	131
Message Link	131
Message Link.MeasuredObject	131
neo.fs.v2.lock	131
Message Lock	131
neo.fs.v2.netmap	132
Service “NetmapService”	132
Method LocalNodeInfo	132
Method NetworkInfo	133
Method NetmapSnapshot	133
Message Filter	133
Message Netmap	134
Message NetworkConfig	134
Message NetworkConfig.Parameter	134
Message NetworkInfo	135
Message NodeInfo	136
Message NodeInfo.Attribute	136
Message PlacementPolicy	138
Message PlacementPolicy.ECRule	139
Message PlacementPolicy.Initial	140
Message Replica	141
Message Selector	141
Emun Clause	141
Emun NodeInfo.State	142
Emun Operation	142
neo.fs.v2.object	143
Service “ObjectService”	143
Method Get	143
Method Put	144
Method Delete	145
Method Head	146
Method Search	147
Method SearchV2	148
Method GetRange	150
Method GetRangeHash	151
Method Replicate	152
Message GetResponse.Body.Init	153
Message HeaderWithSignature	153

Message PutRequest.Body.Init	153
Message Range	154
Message SearchV2Response.OIDWithMeta	154
Message Header	155
Message Header.Attribute	156
Message Header.Split	157
Message Object	158
Message SearchFilter	158
Message ShortHeader	160
Message SplitInfo	161
Emun MatchType	161
Emun ObjectType	162
neo.fs.v2.refs	162
Message Address	162
Message Checksum	163
Message ContainerID	163
Message ObjectID	164
Message OwnerID	164
Message Signature	165
Message SignatureRFC6979	165
Message SubnetID	166
Message Version	166
Emun ChecksumType	166
Emun SignatureScheme	167
neo.fs.v2.reputation	167
Service "ReputationService"	167
Method AnnounceLocalTrust	167
Method AnnounceIntermediateResult	168
Message GlobalTrust	169
Message GlobalTrust.Body	169
Message PeerID	169
Message PeerToPeerTrust	170
Message Trust	170
neo.fs.v2.session	170
Service "SessionService"	170
Method Create	170
Message ContainerSessionContext	171
Message ObjectSessionContext	171

Message ObjectSessionContext.Target	172
Message RequestMetaHeader	172
Message RequestVerificationHeader	173
Message ResponseMetaHeader	173
Message ResponseVerificationHeader	174
Message SessionContextV2	174
Message SessionToken	175
Message SessionToken.Body	175
Message SessionToken.Body.TokenLifetime	175
Message SessionTokenV2	176
Message Target	176
Message TokenLifetime	177
Message XHeader	177
Emun ContainerSessionContext.Verb	178
Emun ObjectSessionContext.Verb	178
Emun Verb	178
neo.fs.v2.status	179
Message Status	179
Message Status.Detail	180
Emun CommonFail	180
Emun Container	181
Emun Object	181
Emun Section	182
Emun Session	182
Emun Success	183
neo.fs.v2.storagegroup	183
Message StorageGroup	183
neo.fs.v2.subnet	184
Message SubnetInfo	184
neo.fs.v2.tombstone	184
Message Tombstone	184
Terms and definitions	186
Glossary	186

Introduction

Overview

NeoFS is a decentralized distributed object storage system integrated with the Neo Blockchain¹.

We store and distribute users' data across a peer-to-peer network of NeoFS Nodes. Whether a business or an individual, any Neo user may join the network and get paid for providing storage resources to others, or pay a competitive price to employ NeoFS as a storage solution.

The decentralized architecture and flexible storage policies allow users to reliably store object data in the NeoFS network. Each NeoFS Node is responsible for executing the specific storage policies selected by the user, including the geographical location, redundancy level, number of nodes, type of disk, capacity, etc. Thus, NeoFS enables a transparent data placement process which gives full control over data to the users.

Deep Neo Blockchain² integration allows NeoFS to be used by Decentralized Applications (dApps) directly from NeoVM³ on the Smart Contract⁴ code level. As a result, dApps are not limited to on-chain storage and one can manipulate large amounts of data without paying a prohibitive price.

NeoFS provides native gRPC⁵ Application Programming Interface (API) and supports popular protocol gateways such as AWS S3⁶, HTTP⁷, FUSE⁸, and sFTP⁹, which allows developers to easily integrate their existing applications without rewriting code.

Together, this set of features makes it possible to utilize a dApp's Smart Contract to manage monetary assets and obtain data access permissions on NeoFS through a regular Web Browser or a mobile application.

¹<https://neo.org>

²<https://neo.org>

³<https://docs.neo.org/docs/en-us/basic/technology/neovm.html>

⁴<https://docs.neo.org/docs/en-us/basic/technology/neocontract.html>

⁵<https://grpc.io>

⁶<https://docs.aws.amazon.com/AmazonS3/latest/API/Welcome.html>

⁷https://wikipedia.org/wiki/Hypertext_Transfer_Protocol

⁸https://wikipedia.org/wiki/Filesystem_in_Userspace

⁹https://en.wikipedia.org/wiki/SSH_File_Transfer_Protocol

Background

Technical Requirements

Out of Scope

Future Goals

Architecture overview

Design and components

NeoFS heavily relies on the Neo Blockchain and its features. This allows NeoFS nodes to focus on their primary tasks — data storage and processing, while asset management and distributed system coordination are left to Neo and a set of Smart Contracts. Under this approach, the Blockchain is mainly used as a trusted source of truth and coordination data.

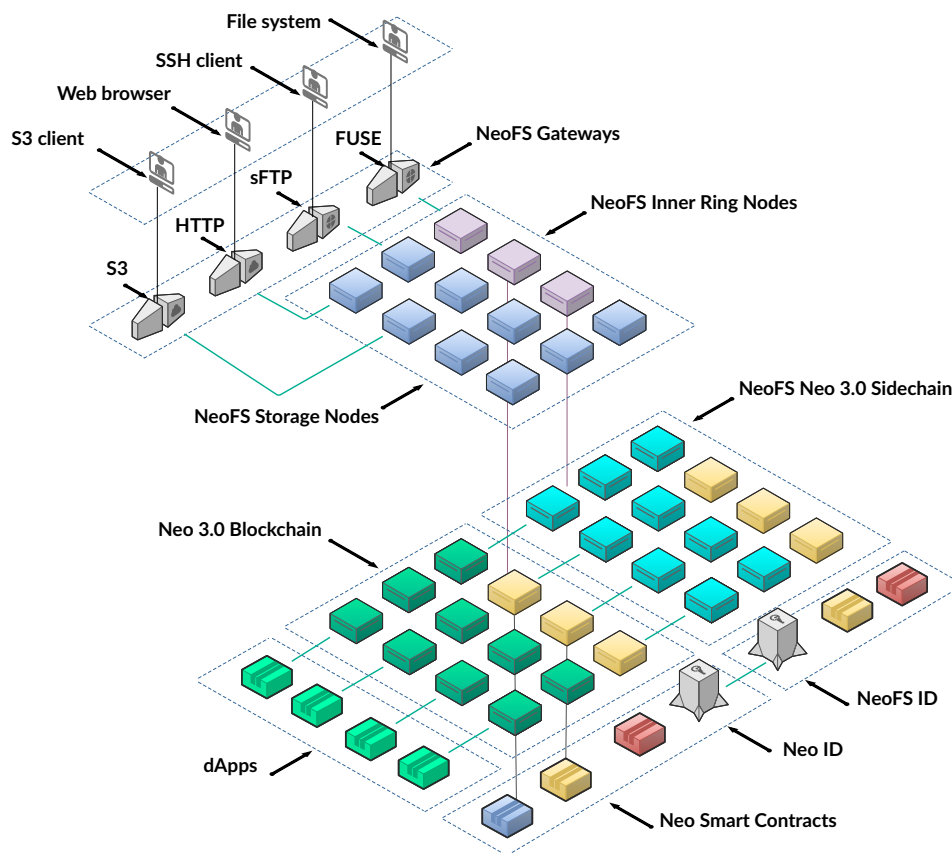


Figure 1: Architecture overview

The **N3 mainnet** hosts a NeoFS Native Contract¹⁰ concerned with user deposits and withdrawals, network settings, and other maintenance operations such as listing the keys of trusted nodes.

¹⁰<https://medium.com/neo-smart-economy/native-contracts-in-neo-3-0-e786100abf6e>

To simplify accounting operations, lessen mainnet burden, and reduce the overall network maintenance costs, NeoFS utilizes an N3-based sidechain¹¹ aka “FS chain”. FS chain runs smart contracts which control the NeoFS network structure, user settlements, balances, and other frequently changing data.

There are two types of NeoFS nodes. They are Storage nodes and Inner Ring nodes.

The first type is responsible for receiving data from a user, reliably storing it as required by the storage policy, and providing access to the data according to the applicable **Access Control Lists (ACLs)**. Such storage nodes are coordinated with smart contracts from the FS chain.

The second type does not store user data. Inner Ring nodes monitor the NeoFS network health, aggregate Storage Nodes reputation ratings, and perform data auditing, issuing penalties and bounties depending on the audit results. Inner Ring nodes listen for both mainnet and FS chain, providing a trusted and reliable way of data synchronization between the two Blockchains.

Each Storage node in the system has a set of key-value attributes describing node properties such as its geographical location, reputation rating, number of replicas, number of nodes, presence of SSD drives, etc. Inner Ring nodes generate a Network Map — a multi-graph structure which enables Storage nodes to be selected and grouped based on those attributes.

In NeoFS, a user puts files in a container. This container is similar to a folder in a file system or a bucket in AWS S3, but with a storage policy attached. The Storage Policy is defined by the user in an SQL-like language (NetmapQL), specifying how and where objects in the container have to be stored by selecting nodes based on their attributes. Storage nodes keep data in accordance with the policy, otherwise they do not get paid for their service.

All storage nodes service fees are paid in **GAS Utility Token**. After receiving GAS, a node operator may spend them to pay for their own data backups on other NeoFS nodes, or simply withdraw it for use with other services provided by the Neo Blockchain ecosystem.

An innovative feature of NeoFS is that it can be accessed directly from NeoVM on the smart contract code level. Thanks to the N3 Oracle protocol and the integration between NeoFS and Neo Blockchain, dApps are not limited to on-chain storage and can manipulate large amounts of data without paying a prohibitive price for it.

NeoFS provides native gRPC API and supports the most popular protocol gateways, allowing easy integration with other systems and applications without requiring code rewrites.

Such an architecture makes it possible to implement a dApp’s smart contract to manage digital assets and data access permissions on NeoFS and lets users access that data via regular Web Browsers or mobile applications. In the long term, we plan to add more ecosystem components that will facilitate

¹¹<https://en.wikipedia.org/wiki/Blockchain#Types>

the development of truly decentralized applications, solving almost any problems that are nowadays only possible in a centralized manner.

Epoch

For the NeoFS network to work properly, all nodes should have the same view of the network. This snapshot view must be tied to some timestamp, but in the distributed environment for NeoFS there is no reliable common source of time other than a monotonically increasing number of blocks in the Blockchain. It means we can only use discrete time model¹² and define some time period in blocks to snapshot the common view of the network. This period has to be small enough to keep the snapshot information fresh and big enough to let the information be distributed fast enough between network nodes. This regular time period is called Epoch.

During an Epoch all common information snapshots are immutable. New nodes can be registered, misbehaving nodes can be removed, some nodes can go offline, but those changes will be reflected only in the next version of the netmap issued for the next Epoch. All these changes signed by Inner Ring are propagated to the network only when the new Epoch starts.

Network Map

NeoFS Network Map, or just “netmap”, is a structured representation of all active storage nodes available in NeoFS network for the current Epoch.

Storage nodes in the netmap are identified by public key. Netmap also has additional information about each node, like network addresses and a list of attributes.

Attributes are key-value pairs with string values. Depending on the Attribute, the string value can be interpreted as a number or Hex or something else. For detailed information please see the API reference.

Here is how node information in netmap may look like:

```
key: 03e9c4847fb2f4d58161a808ff74363139c4e617cb233a2e96cc6c4c7f219dd9bf
address: /dns4/st1.storage.fs.neo.org/tcp/8080
state: ONLINE
attribute: Capacity=10000
attribute: Continent=Europe
attribute: Country=Germany
attribute: CountryCode=DE
```

¹²https://en.wikipedia.org/wiki/Discrete_time_and_continuous_time

```
attribute: Deployed=NSPCC
attribute: Location=Falkenstein
attribute: Price=0.00000042
attribute: SubDiv=Sachsen
attribute: SubDivCode=SN
attribute: UN-LOCODE=DE FKS
```

Some node attributes can be grouped (e.g. geographical ones), creating a graph representation for the network map. Strictly speaking it is a forest of rooted trees where leaves (single nodes) are shared. Every tree represents a single attribute group. For example, geographical attributes can be naturally ordered: Continent, Country, SubDiv, Location, DC (data center) or any other location identifier.

User attributes can have any name and value. They are used primarily for storage policy rules. For example, one may want to designate their own nodes with some specific attribute to be sure (with the help of storage policy) that at least one copy is always stored locally on the nodes one controls.

Storage Policy

In NeoFS, storage policy is a flexible way to specify rules for storing objects. The result of storage policy rules application to the network maps is a set of storage nodes, able to store data according to the requested policy. This result maybe called “placement”, hence sometimes you may find the term “placement policy” used to denote the same thing as the “storage policy”.

Because Storage Policy is attached to the container structure there is a compact definition for system’s internal use and some higher level language definitions for humans to use, that are translated to internal representation. For example, there is an SQL-like language to be used by humans, JSON notation to be used in software and there may be many others, like a graphical language using Blockly¹³. In our examples we will use SQL-like notation.

Storage Policy internal definition consists of four parts:

1. Filters
2. Selectors
3. Replicas
4. Container Backup Factor

The result of applying a storage policy to the netmap is a set of nodes structured by `Replicas` and used to select candidates to put objects on. The selection algorithm is deterministic, hence on dif-

¹³<https://developers.google.com/blockly>

ferent nodes or clients the same Storage Policy applied to the same version of netmap will give the identical result.

Filters

Filter is a mechanism to specify which nodes are allowed to store an object. This is done by querying node's attributes and checking if they satisfy certain condition. For example, it allows to precisely specify "Store objects in nodes from Europe, but not from Italy, which have SSD and have a good reputation".

Simple filter can compare a single node attribute with some value.

It has 3 fields:

1. `Key` — name of node attribute
2. `Value` — value to compare attribute with
3. `Op` — operation to be used for comparison

For better understanding, simple filters will be specified as `Key Op Value`.

For example:

1. `Country = Argentina` means use nodes for which `Country` attribute is equal to `Argentina`, i.e. nodes located in USA.
2. `Rating > 4.5` means use nodes with rating better than 4.5

Currently only eight operations are supported, two of which are used for creating compound filters from the other ones.

1. `EQ/NE` check if attribute is equal/not equal to the filter's value.
2. `GT/GE/LT/LE` check if numerical attribute is greater-than/greater-or-equal/less-than/less-or-equal than the filter's value.
3. `OR` checks if node satisfies at least one of the filters provided as arguments.
4. `AND` checks if node satisfies all filters provided as arguments.

Compound filter can combine simple filters to specify arbitrarily complex conditions. Consider example from the previous section: we may write `Country = Argentina AND Rating > 4.5` to filter nodes which are located in Argentina and have a good rating at the same time.

If filters are used in selectors or other filters, they should have name. Consider filter `Country = Finland OR Country = Iceland AS ColdCountry`. If we need nodes from these countries but want to vary maximum price depending on the storage type they have, we may write this filter:

`ColdCountry AND StorageType = SSD AND Price < 100`
OR
`ColdCountry AND StorageType = HDD AND Price < 10`

Selectors

Selector is a mechanism to specify which of the previously filtered nodes will be included in the container. It has 5 fields:

1. `Name` — name that can be referred to
2. `Attribute` — name of the attribute for grouping nodes. When it is set, nodes are grouped in buckets based on `Attribute` value. It can be omitted to create buckets “randomly”.
3. `Count` — number of nodes to be included in a bucket or number of buckets, depending on `Clause`.
4. `Clause` specifies how `Count` is interpreted:
 - `SAME` — choose nodes from the same bucket
 - `DISTINCT` — choose nodes from distinct buckets
5. `Filter` — name of the filter to choose nodes from. If it is omitted or is `*`, all nodes from the netmap are used.

Selector can return different set of nodes for every epoch; however, they are always the same on each storage node having the same netmap. The degree to which this set of nodes is changed also depends on how strict the filter is. For example, if we select a few nodes based on a very specific attribute, this set will always be the same. However, if all these nodes go down, data can be lost.

Replicas

Replica is an independent set of nodes where single object copy is stored. It can refer to selector (by default all nodes are considered) and can specify a number of copies to store.

Container Backup Factor

Container backup factor (CBF) controls maximum number of nodes to be included in a container’s node set. It doesn’t set strict boundaries, though. Consider placement policy which selects X nodes in 2 different countries with $CBF = 2$. In this case, we can expect container’s node set to have from X to $X * 2$ nodes in every selected country. Having less than $X * 2$ nodes is not considered as fail.

Objects

NeoFS stores all data in the form of objects, thus providing an object-based storage to the clients. These objects are placed in a flat environment (no hierarchy or directories). To access the required data, the identifying details (ID and metadata) are needed.

ObjectID is a hash that equals Headers hashes plus Payload hashes. Any object includes a system header, extended headers, and a payload. A system header is an obligatory field, while extended headers may be omitted. However, any extended header should follow a particular structure (e.g. IntegrityHeader is a must). A user can add any extended header in the form of a key-value pair, though keeping in mind that it cannot be duplicated with several values. One attribute – one value. Please note that any object initially has FileName, so that you cannot create an extended header with it as a key.

The maximum size for an object is fixed and can be changed only for the whole network in the main contract. It means that if a file is too heavy, it will be automatically divided into smaller objects. This smaller parts are put in a container and placed to a Storage Node. Later, they can be assembled to the original object. Such assembling is performed in the storage nodes upon a corresponding request for a linking object. Once your file is converted into an object (or several objects), this object cannot be changed.

One can define the format of the object in an API Specification. For more information, see API Specification¹⁴.

Large objects split

NeoFS has a limit on the maximal physically stored single object size. If there is a large object exceeding that `MaxObjectSize`, it will be split into a series of smaller objects that are logically linked together.

For each part of the original object's payload, a separate object with own `ObjectID` will be created. The large object will not be physically present in the system, but it will be reconstructed from the object parts when requested.

There are two active versions of the split objects, both are supported, but the second one is more flexible and support attributes-based ACL rules. The first one is kept for backward compatibility only. All objects participating in the split have the `Split` headers set. Depending on the place in the split hierarchy, it has different field combinations.

¹⁴<https://github.com/nspcc-dev/neofs-api/tree/master/proto-docs>

Object Split V1

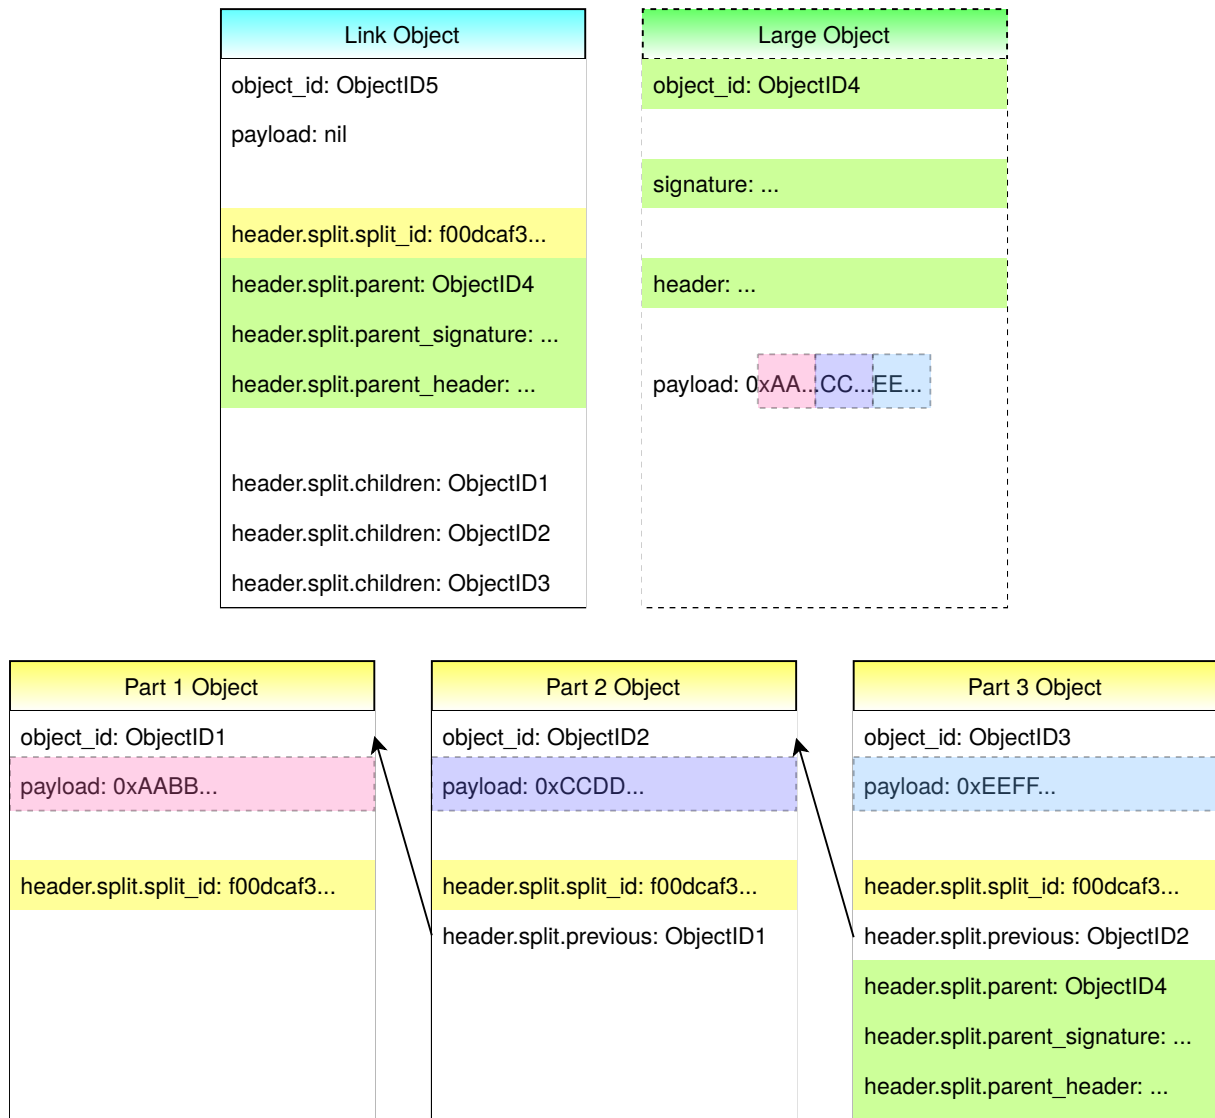


Figure 2: Large object split V1

- First part
First part object only has the split_id field set, as there is no more information known at this point
- Middle parts
Middle parts have information about the previous part in previous field in addition to the split_id

- Last part

At this point, all the information about the object under split is known. Hence, the last part contains not only the `split_id` and `previous` fields, but also the `ObjectID` of the original large object in its `parent` field, signed `ObjectID` in `parent_signature` and original object's Header in `parent_header`.

- Link object

There are special "Link objects" that have the same common `split_id`, do not have any payload, but contain original object's `ObjectID` in `parent` field, its signature in `parent_signature`, original object's Header in `parent_header` and the list of all object parts with payload in repeated `children` field. Link objects help to speed up the large object reconstruction and HEAD requests processing. If Link object is lost, the original large object still will be reconstructed from its parts, but it will require more actions from NeoFS nodes.

Object Split V2

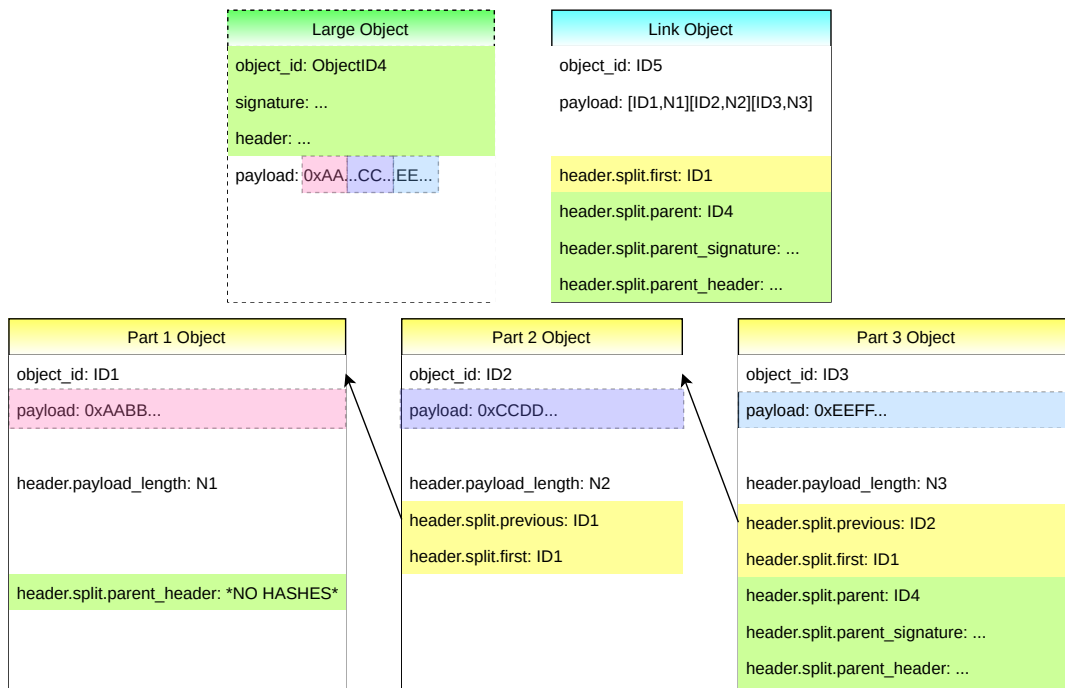


Figure 3: Large object split V2

- First part
First part object only has `parent_header` fields set. The parent header is not finished: since there is no information about payload, it cannot be measured, hashed and signed (also, no object ID can be assigned). However, it has all user information attached to the object (e.g. attributes), and plays the role of representative of the original object until all the parts are uploaded to the NeoFS.
- Middle parts
Middle parts have information about the previous part in `previous` field and the first one in the `first` field
- Last part
At this point, all the information about the object under split is known. Hence, the last part contains not only the `first` and `previous` fields, but also the `ObjectID` of the original large object in its `parent` field, signed `ObjectID` in `parent_signature` and original object's Header in `parent_header`.
- Link object
There are special "Link objects" that have the same common `first` field, but also contain original object's `ObjectID` in `parent` field, its signature in `parent_signature`, original object's Header in `parent_header` and a list of all the object parts' IDs paired with their sizes encoded in the payload. Link objects help to speed up the large object reconstruction and HEAD requests processing. If a Link object is lost, the original large object still will be reconstructed from its parts, but it will require more actions from NeoFS nodes.

All the split hierarchy objects may be physically stored on different nodes. During reconstruction, at first the link object or the last part object will be found. If it's a HEAD request, the link object or the last part object will have all the information required to return the original large object's HEAD response. For a GET request, the payload will be taken from part objects listed in the Link object's payload. As they are ordered, it will be possible to begin streaming the payload as soon as the first part object becomes available. If a Link object is lost, some additional time will be spent on reconstructing the list from `split.previous` header fields.

If the whole payload is available, a large object may be split on the client side using local tools like `neofs-cli`. In this case, the resulting object set will be signed with user's key. Such a split type can be called a "Static split".

There are attribute-based ACL rules in the NeoFS. Before the last part and the link object are created, the original object's parts should be validated according to the initial object's information about the original object.

When the large object's payload is not fully available right away, or it is too big to be split locally, the ob-

ject upload can be started in a Session with another NeoFS node and be streamed in a PUT operation, part by part. Object parts will be automatically created as soon as the payload hits the `MaxObjectSize` limit. In this case, the resulting object set will be signed with a session key signed by user's key. This split type can be called a "Dynamic split".

Object Deletion

It's hard to guarantee complete and immediate object removal in a distributed system with eventual consistency. If some nodes are offline at the time of DELETE request processing, the object may still be available there and would be replicated to other nodes.

To address this issue, NeoFS Storage nodes don't remove objects immediately, but place a removal mark in a form of a Tombstone object. In order to avoid wasting storage space on the information about data that has been already put into trash, tombstones are to be removed later by Garbage Collector.

Tombstone Object

Along with a Regular Object and a Storage Group, there is a Tombstone object type. It is like a regular object, but the payload contains a tombstone data structure. This entity is intended to synchronize object removal in a distributed system working in an unreliable environment. When one removes an Object, NeoFS Storage node actually creates a Tombstone alongside it and replicates the Tombstone among the Container nodes.

A Tombstone indicates that the given `ObjectID` does not exist any more and all requests to it must fail. Storage nodes may keep the data for a while until the Garbage Collector reaps it; though, it's up to the node's settings and implementation.

After several Epochs, we assume that deletion event has been properly spread among all Storage Nodes serving the Container, including those which haven't received it "in time" because of outage or lack of network connectivity.

The time to keep the tombstone may be varying. It is set in `__NEOFS__EXPIRATION_EPOCH` object attribute and `expiration_epoch` field in Tombstone structure. It may be set by a user directly or by an intermediate NeoFS Storage node using the default value. While the `__NEOFS__EXPIRATION_EPOCH` attribute is optional for a Regular Objects, it is obligatory for a Tombstone Object type.

Containers

In NeoFS, objects are put into containers and stored therein.

From the user's point of view, there are six verbs applicable for a Container: PUT, GET, DELETE, LIST, SetEACL and GetEACL. Also, there is an AnnounceUsedSpace operation which is intended for internal NeoFS synchronization. On an existing Container, any user is allowed to make GET, LIST and GetEACL. DELETE and SetEACL are allowed only for a Container owner, disregarding of a Container basic or extended ACL.

Any container has attributes, which are actually Key-Value pairs containing metadata. There is a certain number of attributes set automatically, but users can add attributes themselves. Note that attributes must be unique and have non-empty value. It means that it's not allowed to set - two or more attributes with the same key name (eg. Size=small, Size=big); - empty-value attributes (eg. Size= ' '). Containers with duplicated attribute names or empty values will be considered invalid.

Access Control Lists

Access control in a decentralized untrusted environment is a complicated problem. It must be verifiable by every network participant and still be open for changes to revoke unwanted access permissions or adjust to infrastructure changes.

NeoFS solves this by using **ACL** rules from the combination of sources:

- Basic ACL in the container structure,
- BearerToken ACL rules in the request,
- Extended ACL rules in the FS chain smart contract.

ACLs specifies a set of actions that a particular user or a group of users can do with objects in the container. Each request coming through a storage node gets verified against those rules and rejected if the requests's action is not allowed.

Basic ACL

Basic ACL is a part of the container structure, and it is always created simultaneously with the container. Therefore, it is never subject to any changes. It is a 32-bit integer with a bit field in the following format:

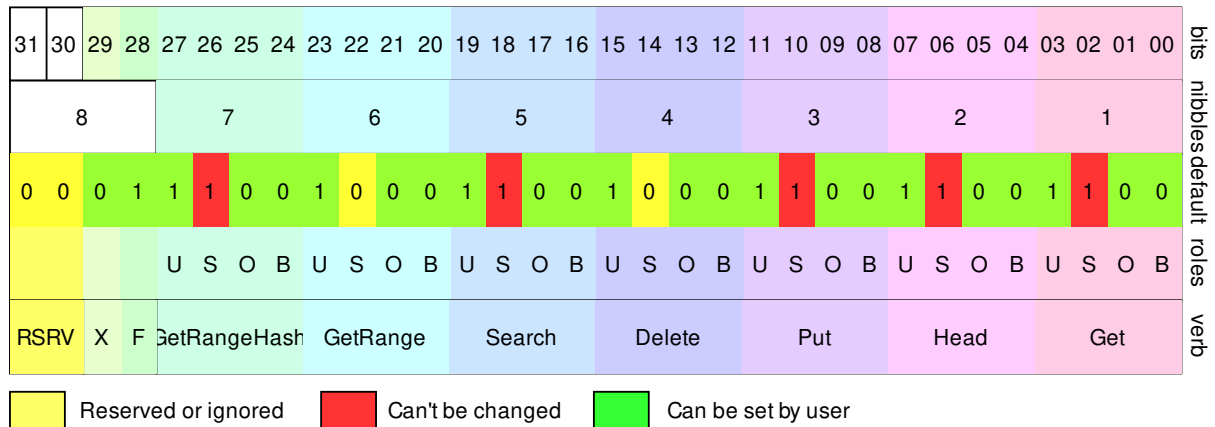


Figure 4: BasicACL bit field

Symbol Meaning Description

B	Bearer	Allows using Bear Token ACL rules to replace eACL rules
U	User	The owner of the container identified by the public key linked to the container
S	System	Inner Ring and/or container nodes in the current version of network map IR nodes can only perform Get, GetRangeHash, Head, and Search necessary for data audit. Container nodes can only do verbs required for the replication. i.e., Get, Put, Head, Search and GetRangeHash.
O	Others	Clients that do not match any of the categories above
F	Final	Flag denying Extended ACL. If set, Basic ACL check is final, Extended ACL is ignored
X	Sticky	Flag denying different owners of the request and the object If set, object in Put request must have one Owner and be signed with the same signature If not set, the object must be correct but can be of any owner. The nodes falling for SYSTEM role are exception from this rule. For them the bit is ignored.
0	Deny	Denies operation of the identified category
1	Allow	Allows operation of the identified category

Basic ACL was designed to be processed and verified really fast. It's simple enough, but covers the majority of access restriction use cases, especially when combined with a carefully tailored Storage Policy.

There are well-known Basic ACLs:

Final – with a flag denying Extended ACL:

private: 0x1C8C8CCC

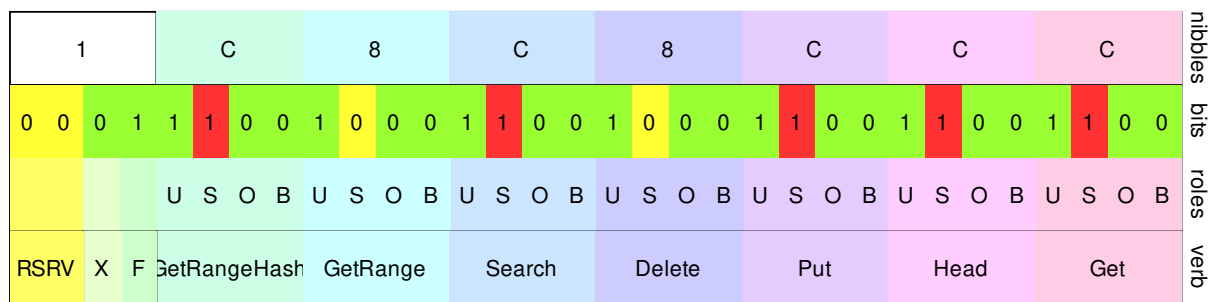


Figure 5: Basic ACL private

public-read: 0x1FBF8CFF

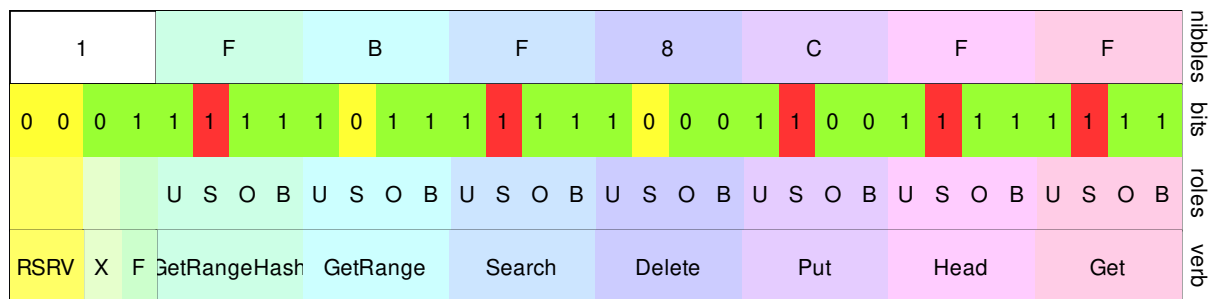


Figure 6: Basic ACL public-read

public-read-write: 0x1FBFBFFF

1	F	B	F	B	F	F	F	nibbles
0 0 0 1 1	1 1 1 1	0 1 1 1	1 1 1 1	0 1 1 1	1 1 1 1	1 1 1 1	1 1 1 1	bits
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles
RSRV X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 7: Basic ACL public-read-write

public-append: 0x1FBF9FFF

1	F	B	F	9	F	F	F	nibbles
0 0 0 1 1	1 1 1 1	0 1 1 1	1 1 1 1	0 0 1 1	1 1 1 1	1 1 1 1	1 1 1 1	bits
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles
RSRV X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 8: Basic ACL public-append

Non-final – Extended ACL can be set:

eacl-private: 0x0C8C8CCC

0	C	8	C	8	C	C	C	nibbles
0 0 0 0 1	1 0 0 1	0 0 0 1	1 0 0 1	0 0 0 1	1 0 0 1	1 0 0 1	1 0 0 1	bits
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles
RSRV X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 9: Basic ACL eacl-private

eacl-public-read: 0x0FBF8CFF

0	F	B	F	8	C	F	F	nibbles	
0 0 0 0 1	1 1 1 1	0 1 1 1	1 1 1 1	0 0 0 1	1 0 0 1	1 1 1 1	1 1 1	bits	
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles	
RSRV	X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 10: Basic ACL eacl-public-read

eacl-public-read-write: 0x0FBFBFFF

0	F	B	F	B	F	F	F	nibbles	
0 0 0 0 1	1 1 1 1	0 1 1 1	1 1 1 1	0 1 1 1	1 1 1 1	1 1 1 1	1 1 1	bits	
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles	
RSRV	X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 11: Basic ACL eacl-public-read-write

eacl-public-append: 0x0FBF9FFF

0	F	B	F	9	F	F	F	nibbles	
0 0 0 0 1	1 1 1 1	0 1 1 1	1 1 1 1	0 0 1 1	1 1 1 1	1 1 1 1	1 1 1	bits	
	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	U S O B	roles	
RSRV	X F	GetRangeHash	GetRange	Search	Delete	Put	Head	Get	verb

Figure 12: Basic ACL eacl-public-append

Extended ACL

Extended ACL is stored in the container smart contract in FS chain. This means it can be changed during container lifetime and there will be only one latest version of it in use. Only the container owner, or the bearer of a SessionToken with a Container context signed by the container owner, can change

the Extended ACL rules. Since it is stored in a form of a stable serialized protobuf structure, eACL table can be only replaced with a new version, not altered or changed in-place in any way.

Extended ACL can only specify Basic ACL rules and make them more restitutive, but it can never ease them. Extended ACL rules can never conflict with Basic ACL rules or cancel them. If something is denied at Basic ACL level, it can never be allowed again by eACL. If Basic ACL contains Allow, eACL may specify the rule to a finite list of allowed keys and Deny all others. If Basic ACL already contains Deny, eACL can do nothing. Deny in Basic ACL cannot be changed to Allow in eACL. Therefore, the records with denied GET, GETRANGE, PUT, SEARCH, HEAD for System target must be ignored. This reduces to ignoring any System target rules.

When a user creates a container with the F-bit of Basic ACL set to 0, they do not need to settle the rules immediately. For a non-existing Extended ACL request, Container contract will return a null byte array. It will be interpreted as a table with no rules.

To get the latest eACL version, a Storage Node needs to request it via RPC from the FS chain node. If an eACL can't be retrieved, the access permissions check fails.

Extended ACL rules get processed on-by-one, from the beginning of the table, based on the request operation, until matching the rule found. It means that there is no separate rule for setting denying or allowing policy. Final fallback rules must be provided by the user, if needed.

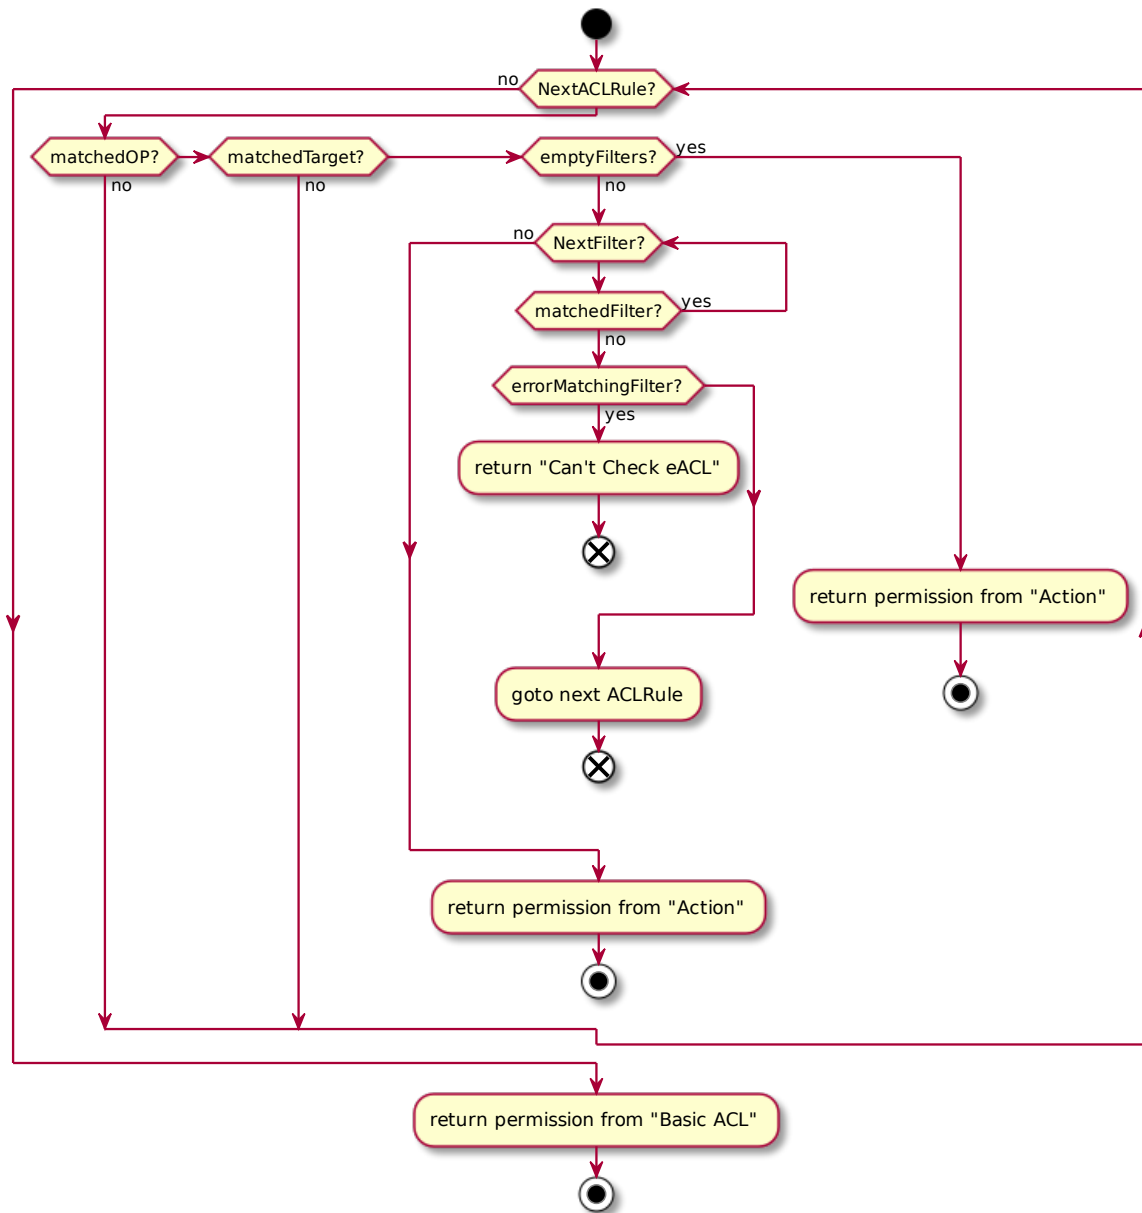


Figure 13: Extended ACL rules check

Extended ACL rules and table format may change depending on the version of NeoFS API used. Please see the corresponding API specification section for details.'

Each eACL rule record has four fields:

Field	Description
Operation	NeoFS request action verb
Action	Rule execution result action. Allows or denies access if the rule's filters match.
Filter	Filter to check particular properties of the request or the object
Target	Subject's role class or a list of public keys to match

and can be presented in different intermediate formats, like JSON, for the users' convenience.

```
{
  "records": [
    {
      "operation": "GET",
      "action": "DENY",
      "filters": [
        {
          "headerType": "OBJECT",
          "matchType": "STRING_NOT_EQUAL",
          "key": "Classification",
          "value": "Public"
        }
      ],
      "targets": [
        {
          "role": "OTHERS"
        }
      ]
    }
  ]
}
```

Note that some filters with \$Object prefix are not suitable for making denying rules on certain operations. There may be an undefined behavior on some combinations of NeoFS verbs and object attributes when eACL is set. In the table below, + means allowed to be used and - means undefined behavior, hence not allowed.

\$Object:	GET	HEAD	PUT	DELETE	SEARCH	RANGE	RANGEHASH
version	+	+	+	-	-	-	-
objectID	+	+	+	+	-	+	+

\$Object:	GET	HEAD	PUT	DELETE	SEARCH	RANGE	RANGEHASH
containerID	+	+	+	+	+	+	+
ownerID	+	+	+	-	-	-	-
creationEpoch	+	+	+	-	-	-	-
payloadLength	+	+	+	-	-	-	-
payloadHash	+	+	+	-	-	-	-
objectType	+	+	+	-	-	-	-
homomorphicHash	+	+	+	-	-	-	-
User headers	+	+	+	-	-	-	-

Let us make an example. Delete and Range operations are likely to show undefined behavior if Head has been denied for objects with particular payloadLength. They fail because they need to produce HEAD requests upon execution. If a user cannot Head, those operations cannot work properly. The full table of spawning object requests is given below.

Base/Gen	PUT	DELETE	HEAD	RANGE	GET	HASH	SEARCH
PUT	+	-	-	-	-	-	-
DELETE	-	-	+	-	-	-	+
HEAD	-	-	+	-	-	-	-
RANGE	-	-	+	+	-	-	-
GET	-	-	+	-	+	-	-
HASH	-	-	+	+	-	-	-
SEARCH	-	-	-	-	-	-	+

Also, note that user attributes cannot be used as filters in an eACL rule as it provokes an undefined behaviour. By design, when user attributes are set for a Complex Object, they are not inherited in the Part Objects and are only stored in the Link Object header. We cannot control the access for an Object of size more than `maxObjectSize`. To keep the system consistent we do not support eACL filters by user attributes for small Objects as well.

Bearer Token

BearerToken allows to use the Extended ACL rules table from the token attached to the request, instead of the Extended ACL table from the Container smart contract.

Just like JWT¹⁵, it has a limited lifetime and scope, hence can be used in the similar use cases, like providing authorization to externally authenticated party.

BearerToken can be issued only by the container owner and must be signed using the key associated with the container's OwnerID.

In the gRPC request, BearerToken is encoded in a protobuf format, but can be also presented in different intermediate formats, like JSON, for the users' convenience.

```
{
  "body": {
    "eaclTable": {
      "version": {
        "major": 2,
        "minor": 6
      },
      "containerID": {
        "value": "DIFWB4CFTayb9IAqeGwLGJdJfW6i5wWllPsF50EmazQ="
      },
      "records": [
        {
          "operation": "GET",
          "action": "ALLOW",
          "filters": [
            {
              "headerType": "OBJECT",
              "matchType": "STRING_EQUAL",
              "key": "Classification",
              "value": "Public"
            }
          ],
          "targets": [
            {
              "role": "OTHERS",
              "keys": []
            }
          ]
        }
      ],
    },
  },
}
```

¹⁵<https://jwt.io>

```
    ]
  },
  "ownerID": null,
  "lifetime": {
    "exp": "100500",
    "nbf": "1",
    "iat": "0"
  }
},
"signature": {
  "key": "AiGlnj41qh9o9uVqP9b9CArihHvXfGmljhAZNo4DceG",
  "signature": "BAwfdE1ZVL0LfREGkuXRKT2....GA="
}
}
```

BearerToken format may change depending on the version of NeoFS API used. Please see the corresponding API specification section for details.

ACL check algorithm

NeoFS tries to start with local Basic ACL checks that are fast and cheap in terms of resource consumption. This should cover the vast majority of cases. Then, if present in the request, the ACL records from BearerToken, again locally. For the rest of complex cases, the Storage Node retrieves the Extended ACL table from the Container smart contract. Thereafter, the NeoFS ACL system may slow down the request processing only in complex cases when it's inevitable.

The resulting ACL check algorithm is the following:

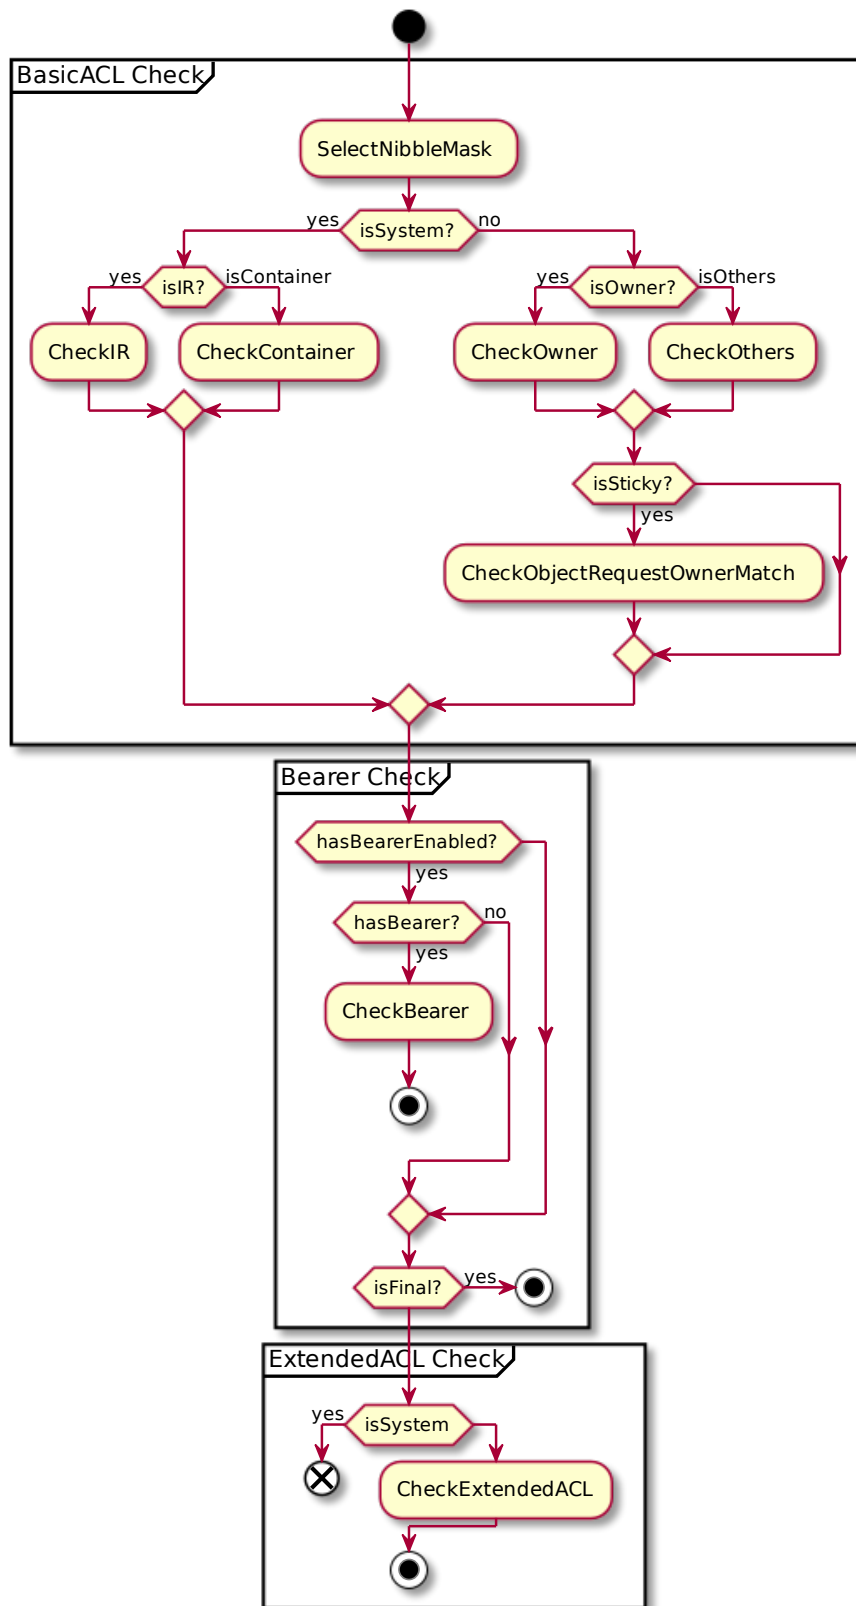


Figure 14: ACL check order

Reputation system

NeoFS reputation system is a subsystem for calculating trust in a node. It is based on a reputation model for assessing trust, which is, in turn, based on the **EigenTrust** algorithm designed for peer-to-peer reputation management. The algorithm ensures that there is a uniquely defined manager (parent node) for each network participant at each specified time (epoch). Based on the information received from its child node and other managers, it iteratively puts a complex general (Global) Trust of the entire network into the applicable child node.

The reputation system allows introducing nodes performance quality metric (trust). In a situation when most of the nodes make an honest assessment of the actions of other nodes, the system lets you calculate this metric quite accurately. This metric can be used for:

1. Filtering the list of nodes where the user intends to store information. For example, one can use only those nodes whose scores are higher than the minimum acceptable for the user.
2. Making a decision to exclude untrusted hosts from the network. For example, one can suspend a node in case it comes to the minimum level of quality in the network.

Trust

Trust in a NeoFS node is its quantitative (numerical) assessment based on the experience of interacting with that node. The higher the score is, the higher is the trust in the node and vice versa. Since the system uses a reputation-based trust model, the terms “trust” and “reputation” are considered synonymous in this document.

The Subject of trust assessment is the one who calculates trust.

The Object of trust assessment is the one whose trust is being calculated.

Reputation models are based on the nature of social media reputation. Thus, trust in a node is built up both by the estimates of the behavior of the node by another node and by the reputation of the node evaluating its behavior.

Trust in a NeoFS in a node is formed based on its interactions with the Subjects of trust assessment. Therefore, the reputation of a node changes during the NeoFS network working cycle when both the behavior of the node itself, and the behavior of other nodes change.

Algorithm

General problem statement: the Subject of assessment needs to calculate the reputation of the Object of trust assessment at a specific point of time (specific epoch).

EigenTrust is based on the notion of transitive trust: peer i will have a high opinion of those peers who have provided it with authentic information. Moreover, peer i is likely to trust the opinions of those peers, since peers who are honest about the information they provide are also likely to be honest in reporting their **Local Trust** values.

Global Trust is calculated in 3 main stages:

1. Each network member collects local statistics of network interactions with other peers, acting as the Subject of reputation assessment.
2. At the end of an epoch, each node announces its local statistics to its manager.
3. Managers exchange received information iteratively and, based on the updated data, make adjustments to the trust obtained in the previous iteration.

The algorithm uses configuration parameters that affect the result of the **Global Trust** calculation. To synchronize all network participants in terms of the values of these parameters, the nodes “read” these parameters from the Netmap contract.

Subjects and Objects of Trust in NeoFS

In NeoFS, the reputation system is used to calculate the trust in Storage Node. Thus, the Object of trust is always a Storage Node (and it is also the Subject in the local case). The Subject of Global Trust is the entire ring of Storage Nodes.

Erasur Coding

NeoFS applies erasure coding (further EC) to user data for its protection. Being an alternative to replication, EC also provides fault tolerance through data recovery. The distinctive feature is storage efficiency: EC outturn uses less storage space for the same reliability requirements in terms of number of parts that can be lost without affecting object availability.

Overview

NeoFS leverages Reed–Solomon codes to provide data EC. Original unit of data, which is user object’s payload, is divided into smaller (data) parts. Coder calculates and creates additional parity parts based on the data ones. Both data and parity parts are distributed across different storage nodes. If some data parts are lost, NeoFS uses the remaining chunks and the parity information to reconstruct the missing data and rebuild the original payload.

Settings

EC is enabled and configured in container storage policy¹⁶ (API¹⁷). Once enabled, all container objects are encoded except

- **TOMBSTONE**;
- **LOCK**;
- **LINK**.

EC is enabled by a storage policy rule with the following parameters: 1. number of data parts (**d**); 2. number of parity parts (**p**); 3. selector name to associate this rule with.

With this rule, each payload is transformed to **d+p** parts. The selector is used to select nodes for storage, similar to REP-rules.

There can be up to **4** EC rule in each container. Codes, or part payloads, are consistent for the same **d** and **p**.

Object model

For each original (parent) object, the payload is split into **d** data and **p** parity parts. Each part has the same size. Data parts contain the original payload. If its length is not divisible by **d**, the last part is aligned with zero bytes.

For each payload part, a part object is created. Following properties are inherited from the parent object:

- version;
- container ID;
- owner ID;
- creation epoch.

Then, the object is supplied with EC-specific properties:

- **__NEOFS__EC_RULE_IDX** attribute is an index of policy rule this part is created for;
- **__NEOFS__EC_PART_IDX** attribute is this part's index.

Parent header, ID and signature are set in corresponding fields of the split header.

As always, payload checksum(s), ID and signature are set last and seal the part object.

¹⁶[02-policy.md](#)

¹⁷<https://github.com/nspcc-dev/neofs-api/blob/fc8584b766d6b544337fc055dca34c29b0155d31/netmap/types.proto#L134-L180>

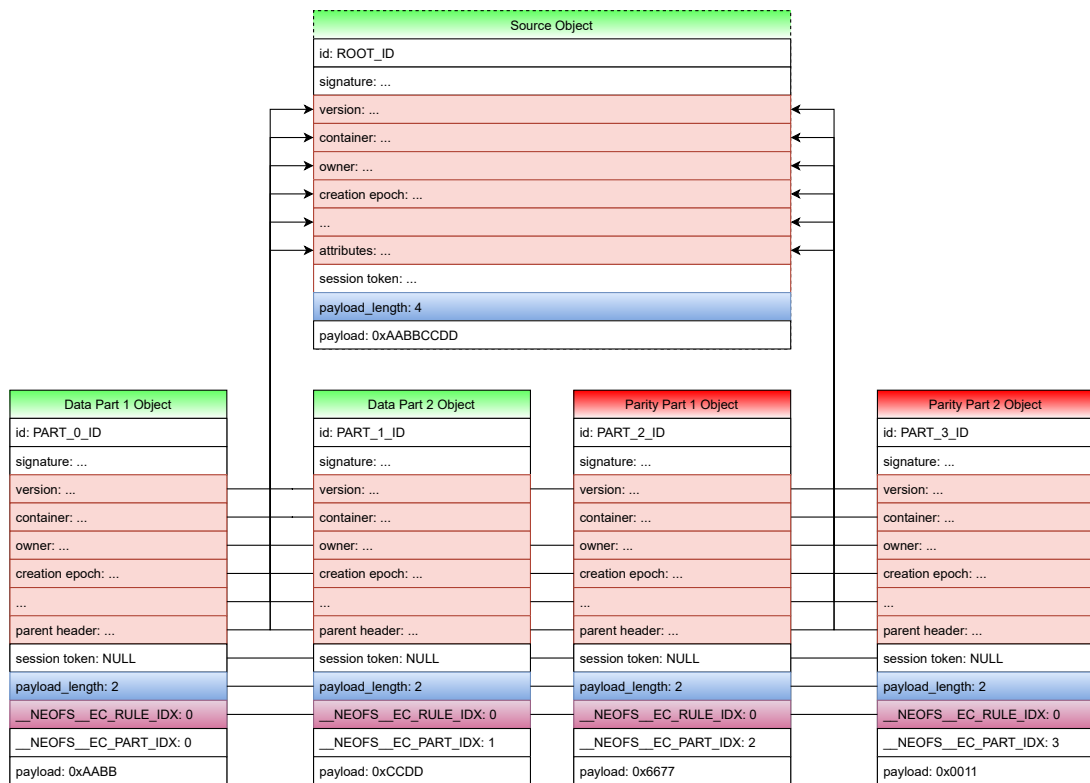


Figure 15: Object EC split

Part placement

Each part object is stored in the container in one copy. Storage nodes are selected from the network map similar to REP rules. The object for the i -th part is placed in the i -th node. If it is unavailable, the backup nodes with indexes $m * n + i$ ($n = d + p$, $m = 1, \dots, \mathbf{CBF}-1$). If all nodes for the i -th part are unavailable, nodes for the $i+1$ -th (0 for the last) part are tried, and so on.

Reconstruction

Once part objects are stored in the container, the original object remains available if at least d of any part objects are available. In other words, unavailability (including complete loss) of any p part objects does not violate availability of the original one.

The source objects are available through the regular API. This also applies to part objects. However, simple read requests require part addresses, while users request parent objects themselves.

When processing i -th policy rule (for assembling the original one in particular), in order to access j -th part object locally stored on the server, following query is used:

- address the parent object;
- add `__NEOFS__EC_RULE_IDX=i` X-header;
- add `__NEOFS__EC_PART_IDX=j` X-header.

This extension is applicable to GET, HEAD and RANGE calls.

To collect a complete object by (parent) address, you need:

1. select container nodes by container ID;
2. sort results by parent object ID;
3. received data part objects by sending GET request with parent address and `__NEOFS__EC_RULE_IDX=i` ($i = 0, \dots, d$) X-header to sorted container nodes starting from the i -th. If all data parts received, goto 5;
4. when $k \leq p$ data parts are unreachable, try to GET k parity parts the same way. As soon as at least d parts are received, original object can be restored;
5. parent header can be fetched from `parent_header`¹⁸ field of any part object;
6. parent payload is calculated from parity parts using Reed-Solomon¹⁹ decoder.

¹⁸<https://github.com/nspcc-dev/neofs-api/blob/fc8584b766d6b544337fc055dca34c29b0155d31/object/types.proto#L288>

¹⁹<https://github.com/klauspost/reedsolomon>

Split + EC

When splitting the original object by payload size, each split-part is subject to EC. As mentioned, the LINK is excluded.

Inner Ring Nodes

Storage Nodes

Address format

NeoFS uses **Multiaddress** format in a human-readable string version as a Node address in netmap. Any new node must provide correct **Multiaddress** on bootstrap stage. After bootstrap, addresses are verified by IR before new Node is a part of netmap.

Correct address *composition* and *order*:

1. Network layer(dns4, ip4 or ip6);
2. Transport layer(tcp);
3. Presentation layer(tls) - optional, may be absent.

Examples:

Correct

- /dns4/somehost/tcp/80/tls;
- /ip4/1.2.3.4/tcp/80.

Incorrect

- /tcp/80/ip4/1.2.3.4;
- /tls/ip4/1.2.3.4/tcp/80;
- /ip4/1.2.3.4/dns4/somehost/tcp/80.

Garbage Collector

Garbage Collector is a part of the Storage Engine. One GC instance runs on one shard.

GC should remove an object itself from Write Cache and Blobstor and the object's metadata from Metabase.

Upon a delete request, a Tombstone object is initialized. It contains all ObjectIDs which the deleted Object was split into. A Tombstone belongs to the Container where the deleted Object is stored. In spite of the fact that the `__NEOFS__EXPIRATION_EPOCH` attribute is assigned to the Tombstone, its value is taken from the internal field that is specific to Tombstone object type. This action makes it easier for Storage Engine data scrubbers to search through the Tombstone indices and select ones for final physical deletion. By design, expiration epoch is strictly more than the current one and never equals to it. Once created, the Tombstone is put into the Container and thus is spread therein.

Since this moment, users cannot GET the object any more. If they try to GET the object, they will get the error.

When a Tombstone is obtained, the Storage Engine produces some meta information about it, which helps it track object removal. The Graveyard entity is responsible for storing that metadata. It is a Metabase table consisting of key-value pairs. The key is the deleted Object's ID and the value depends on how much time this Tombstone has already spent in the Graveyard: - it is a Tombstone Object ID when the Object has just (i.e., on the current epoch) been inhumed; - it is a special GCMark flag evaluated to "True" if Garbage Collector has marked the inhumed Object for total deletion.

Garbage Collector's routine per epoch includes two jobs: invalid Objects check and marked Objects removal.

Invalid Objects check

We consider an Object invalid if it is expired (i.e., `__NEOFS_EXPIRATION_EPOCH` well-known attribute value is equal to the current epoch) or the Tombstone associated with this object is expired.

If an expired Object is found, GC leads it through the deletion procedure described above.

If an expired Tombstone is found, the associated Graveyard record is updated: the aforementioned GCMark toggles for the inhumed Object.

Marked Objects removal

GC searches through the Graveyard and deletes Objects which have been previously GCMark'ed for deletion. It first removes all metadata associated with this objectID and then removes the Object and the related Tombstone from Blobstor, Write Cache and filesystem.

This procedure is the same for any object type, i.e. a Storage Group removal goes through the mentioned stages as well as a Regular Object.

Unpaid containers

If network configuration contains a positive `BasicIncomeRate` (see [Incentive model](#)), container owners should pay for storage being used. Once there are not enough NeoFS tokens to pay for storage according to the defined rate, container is marked in the FS chain's Balance contract as an unpaid one. Storage nodes eventually receive notification (see Balance's `ChangePaymentStatus` event in [NeoFS Smart Contracts](#)) about it and stop serving unpaid containers. There is a safe epoch window before unpaid containers are removed from storage. Practically, it is 3 epochs now, but no expectations should be set, and unpaid data should be considered lost.

Object Expiration

In NeoFS, Objects may have an “expiration date”. When an Object expires, it is marked for deletion and isn’t available anymore. There is a well-known `__NEOFS__EXPIRATION_EPOCH` attribute which specifies the expiration date. Only a Regular Object may expire.

A Tombstone object is created upon Regular Object or Storage Group deletion. Every Tombstone object has the `__NEOFS__EXPIRATION_EPOCH` attribute as well. Thereby Storage Engine is able to filter Tombstones and select ones for total cleanup.

This attribute for a Tombstone Object is set automatically upon its creation. In case of a Regular Object, a user sets it manually.

Storage metrics

Every storage node should track taken storage by every container it serves. Statistics are sent to Container contract which, in turn, calculates total sum for every container, billing information (see [Incentive model](#)), and summary for every user. Storage nodes are free not to send any updates if no storage changes have happened since the last report.

Container contract has API (see [NeoFS Smart Contracts](#)) for fetching actual values e.g. `GetNodeReportSummary`, `GetReportByNode`, `GetBillingStatByNode`, etc. These values can be treated as a general NeoFS state metrics.

The number of reports per epoch is limited, and metrics should not be considered as up-to-date state; there is always some delay.

Storage quotas

NeoFS storage support limiting storage space used by a container and/or by a user (meaning storage sum of every container he has created or, in other words, every container that has this user as owner). Container and user quotas are independent and may overlap depending on containers state. Quotas limit the whole storage used by every storage node that serve container(s), the summary is accumulated in Container contract and is based on node’s storage reports (see [Storage metrics](#)).

Quotas administration

There is network setting (`AlphabetManagesQuotas`) that defines who is allowed to manage quotas:

- setting is missing or set to anything that cannot be converted to `true` in terms of Neo VM: only container's owner and user is allowed to set, change, cancel quotas,
- setting is set to anything that can be converted to `true` in terms of Neo VM (e.g. non-empty array with non-zero bytes): only Alphabet consensus is allowed to set, change, cancel quotas.

Hard/Soft quota differences

Quotas have two independent modes: soft and hard. Both can be set at a time, both can have any values, there is no requirement on how these values should correspond to each other.

Soft quota instructs storage nodes to warn their administrator that the limit is exceeded in logs. It does not intend to alter storage, and any storage changes are still allowed.

Hard quota instructs storage nodes to deny any object that lead to exceeding the limits. Current used storage is taken from the summary reports in the Container contract, *not* from the local statistics.

Setting

Quotas are stored and maintained inside Container contract. The contract provides API for fetching and setting values (see Container's methods in [NeoFS Smart Contracts](#) for more details):

- `SetSoftContainerQuota/SetHardContainerQuota` and `ContainerQuota`,
- `SetSoftUserQuota/SetHardUserQuota` and `UserQuota`.

To cancel previously set quota, non-positive value should be set instead.

Protocol gateways

HTTP

S3

NeoFS S3 gateway²⁰ provides API compatible with Amazon S3 cloud storage service.

Access Box scheme

S3 gateway has to authenticate user requests regarding AWS spec²¹. So we have the following scheme:

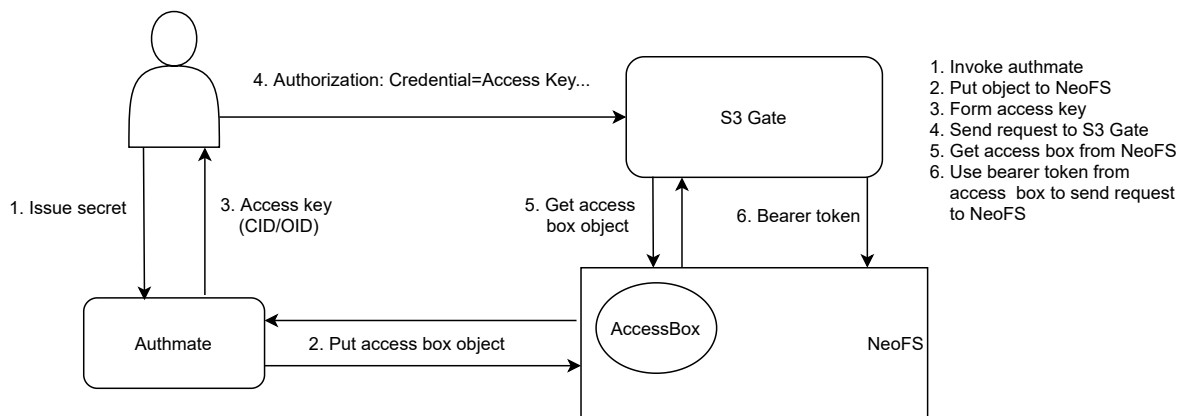


Figure 16: Access box scheme

1. A user uses the `neofs-s3-authmate`²² (Authmate) tool to get credentials (`access_key_id` and `secret_access_key`).
2. Authmate forms “AccessBox” object (see the next section) and puts it into NeoFS.
3. Authmate output contains credentials (`access_key_id` and `secret_access_key`) that can be used with AWS CLI, for example.
4. The user sends request to NeoFS S3 Gateway using standard AWS tool.
5. S3 Gateway gets “AccessBox” from NeoFS by `access_key_id` and fetches Bearer Token from it.
6. S3 Gateway uses fetched Bearer token to send a request to NeoFS on behalf of the user.

²⁰<https://github.com/nspcc-dev/neofs-s3-gw>

²¹<https://docs.aws.amazon.com/AmazonS3/latest/API/sig-v4-authenticating-requests.html>

²²<https://github.com/nspcc-dev/neofs-s3-gw/blob/master/docs/authmate.md>

Form Access Box object Actually, AccessBox is a regular object in NeoFS but properly formed. It contains an encrypted Bearer token (BT), Session tokens (STs) and a `secret_access_key`.

The credentials are formed by the following steps:

1. User provides Authmate with BT, STs and `s3gw_public_key_k`(the public keys of the gates that will be able to handle credentials) that will be used when the user sends requests via S3 Gateway.
2. Authmate:
 1. Generates a `secret_access_key` (it's 32 random bytes) and P256 (secp256r1) key pair (`authmate_private_key`, `authmate_public_key`).
 2. Forms a tokens protobuf struct that contains BT, STs, `secret_access_key`
 3. For each `s3gw_public_key_k` derives a symmetric key `symmetric_key_k` using ECDH²³ and encrypts tokens (`encrypted_tokens_k = encrypt(tokens.to_bytes(), secret_key_k, nonce_k)`).
 4. For each `encrypted_tokens_k` forms a struct called `gate_k` that contains `encrypted_tokens_k` and `s3gw_public_key_k`.
 5. Forms the final binary object AccessBox that contains `authmate_public_key`, `gate_1, ..., gate_k`.
 6. Puts the AccessBox object to NeoFS and saves its address CID/OID as `access_key_id`.
 7. Returns the pair (`access_key_id`, `secret_access_key`) to the user.

Handle S3 request On getting a request, S3 Gateway:

1. Fetches the `access_key_id` from the Authorization header.
2. Gets AccessBox from NeoFS by address (recall `access_key_id` is CID/OID)
3. Using `s3gw_private_key_k` and `authmate_public_key` derives `symmetric_key_k` and decrypts `encrypted_token_k` that has been found by matching `s3gw_public_key_k` in `gate_k` struct.
4. Checks the signature of the initial request using `secret_access_key` from tokens struct.
5. Uses BT and STs to perform requests to NeoFS.

sFTP

²³https://en.wikipedia.org/wiki/Elliptic-curve_Diffie%E2%80%93Hellman

Blockchain components

Role of blockchain in the storage system

NeoFS stores data off-chain, on Storage Nodes. Clients access it directly in a peer-to-peer fashion. This allows to maintain the quality of service (read/write speed, big data volumes) at the level of traditional storage systems. That said, NeoFS is a global-scale decentralized network and it needs a proven mechanism to serve as a source of truth and global state and remain at that scale. That's what we use the Neo blockchain for.

The Inner Ring nodes take care of the entire network health. They control the network map of Storage Nodes, manage user containers and access control lists, regulate financial operations and data audit. Therefore, Inner Ring should do the following:

- make decisions according to consensus;
- be fault-tolerant;
- have synchronized global state available to Storage Nodes;
- minimize the amount of p2p connections to protect from direct attacks.

All these requirements can be fulfilled by using the blockchain as a distributed database with business logic implemented in smart contracts. Inner Ring operations must be confirmed and audited, which makes it reasonable to use smart contract memory to store the global state of the storage system. It's what makes Inner Ring applications stateless, lightweight, and scalable.

Meanwhile, NeoFS contracts:

- store the state of the current epoch, network map, audit results, reputation estimations, and container-related data;
- manage the economy;
- control the governance of the NeoFS network.

Main chain and FS chain

Business logic of smart contracts is executed at contract method invocations from Inner Ring nodes. They are either multisigned transactions or regular transactions. However, these invocations are paid and require **GAS Utility Token** for computation. If data owners are charged for these computations, it will be economically impractical for them to use NeoFS.

To solve this problem, we divide smart contract operations into two types:

- financial transactions and governance,

- storage system associated operations.

The first type of operations should be executed in the main chain. These operations are quite rare and they require GAS as a payment asset. The main chain for NeoFS is **N3 mainnet**. Thus, NeoFS contract is deployed in the main chain. It allows to make a deposit, update network config, and reregister the Inner Ring candidate node. The main chain also manages **Alphabet nodes** of the Inner Ring using the **RoleManagement contract**.

The second type of operations can be executed on an additional chain that we call the FS chain. FS chain is a Neo N3 blockchain with a different network configuration. Validator nodes of FS chain are Alphabet nodes of the Inner Ring. FS chain GAS is managed by Alphabet contracts and should be used only for network maintenance and NeoFS contract execution. FS chain has **Alphabet**, **Audit**, **Balance**, **Container**, **Netmap**, **Reputation** contracts. No other contracts unrelated to NeoFS can be deployed in the FS chain.

Inner Ring nodes synchronize states of main chain and FS chain contracts by listening to the notification events and reacting to them.

Notary service

To make decisions according to the consensus, transactions must be multisigned by Alphabet nodes of the Inner Ring. Inner Ring nodes, however, do not have a p2p connection to each other. Therefore, multi-signature check is replaced with on-chain invocation accumulation. NeoFS contracts await for 5 out of 7 method invocations from Alphabet nodes of the Inner Ring. It leads to:

- increased number of transactions in the network;
- inability to calculate the exact price of a transaction, which leads to an increased cost of execution;
- the contract logic complication.

While in FS chain GAS is used only for utility purposes and invocation prices can be mostly ignored, Alphabet nodes of the Inner Ring in the mainnet use a real GAS asset to do withdrawal operation. High execution cost of such operation leads to high withdrawal commissions.

To solve this issue, NeoFS supports notary service²⁴. The notary service allows to build multisigned transactions natively in the blockchain. Thus, the number of transactions and their costs are reduced, contract source code is simplified. Special proxy (in FS chain) and processing (in the main chain) contracts can pay for invocation instead of Alphabet nodes of the Inner Ring. With these contracts, it becomes easier to monitor and control the NeoFS economy.

²⁴<https://github.com/neo-project/neo/issues/1573#issuecomment-704874472>

FS chain Governance

NeoFS uses its own chain as a database to store meta-information about the network: the network map, audit results, containers, key mappings, network settings, and several supplementary things.

FS chain operates on the same principles as the mainnet – there are no free transactions, the committee chooses the consensus nodes, etc. This structure provides a number of advantages. For example, one can use the same N3 tool stack to work with FS chain information.

To effectively work with FS chain, we need to solve the following problems:

- How can the mainnet committee control Inner Ring nodes and the consensus nodes of FS chain?
- How do Storage Nodes and Inner Ring nodes get FS chain **GAS Utility Token** to send transactions?

NeoFS Governance model solves these problems with seven “Alphabet” FS chain contracts and the first seven Inner Ring nodes bound to those contracts, acting as the FS chain committee.

Alphabet contracts

Alphabet contracts are seven smart contracts deployed in FS chain. They are named after the first seven Glagolitic²⁵ script letters: Az(А), Buky(Б), Vedi(В), Glagoli(Г), Dobro(Д), Yest(Е), Zhivete(Ж). These contracts hold 100,000,000 FS chain **NEO Token** on their accounts (approximately 14,285,000 for each). By storing **NEO Token** on the contract accounts, we protect it from unauthorized use by malicious FS chain nodes. Contracts do not transfer NEO and use it to vote for FS chain **Validator nodes** and to emit **GAS Utility Token**.

Alphabet Inner Ring nodes

Alphabet Inner Ring nodes are the first seven nodes in the Inner Ring list that are logically bound with one-to-one relation to the Alphabet contracts. They are the voting nodes, tasked with making the decisions in the NeoFS network. All other Inner Ring nodes take care of Data Audit, Storage Node attribute verification, and other technical tasks.

Being an Alphabet node implies running FS chain Consensus Node using the same key pair as the NeoFS Inner Ring node instance. Hence, an Alphabet node candidate must:

- Setup a NeoFS Inner Ring node instance
- Setup an FS chain full node using same key pair
- Register the same key in mainnet NeoFS Inner Ring candidates list
- Register the same key in FS chain committee candidates list

²⁵https://en.wikipedia.org/wiki/Glagolitic_script

Alphabet contracts invocation

Contracts cannot distribute the utility token or vote by themselves. To perform these operations, Inner Ring nodes invoke alphabet contract methods. An Alphabet Inner Ring node can invoke its corresponding contract only. One node invokes one contract.

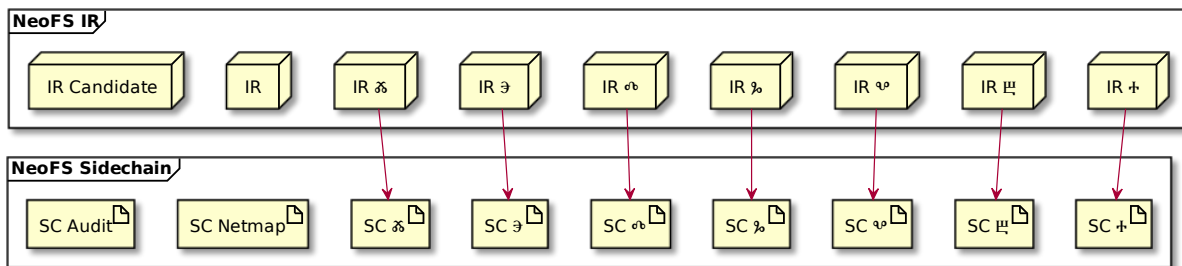


Figure 17: Inner Ring to Alphabet SC relation

Alphabetic contracts have hardcoded indexes. Contracts authenticate the method invoker by using the list of Inner Ring node keys from the Netmap Smart Contract. This scheme helps to limit malicious Alphabet Inner Ring node actions and makes network more resilient to Inner Ring nodes losses.

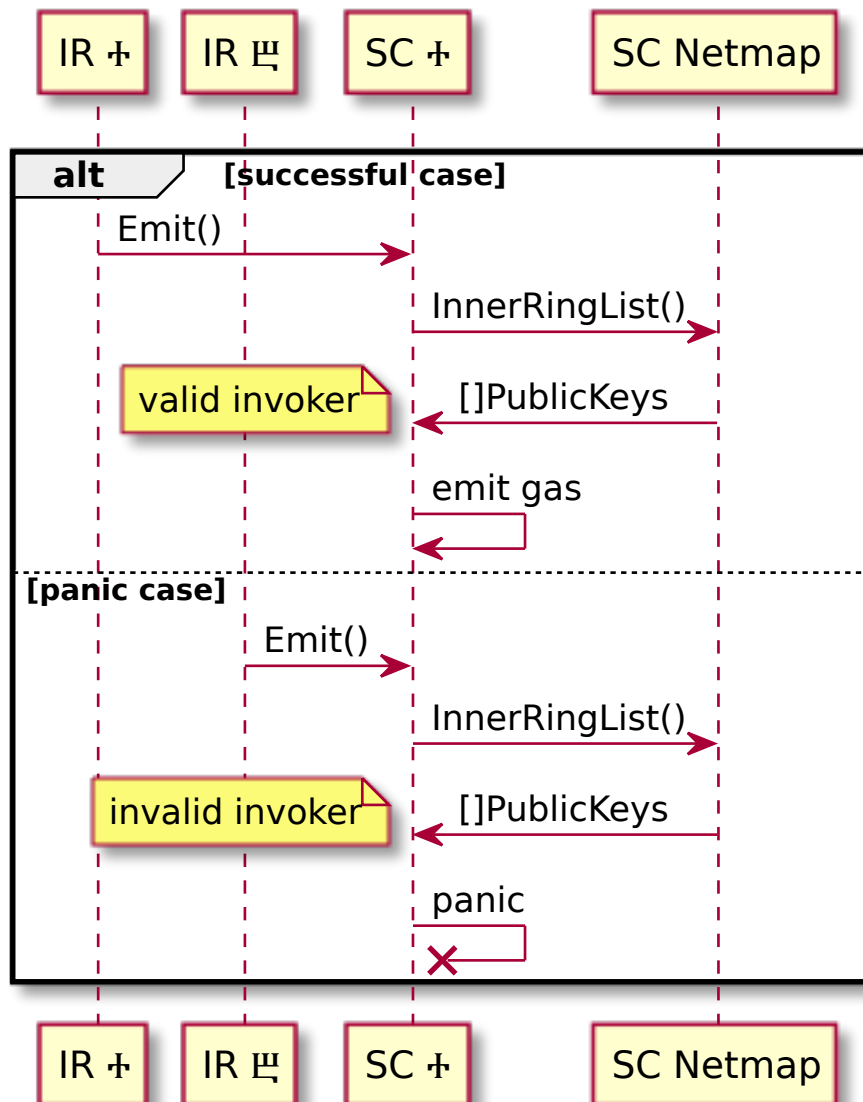


Figure 18: Alphabet SC invocation by Inner Ring nodes

Utility token distribution

Inner Ring nodes invoke `Emit()` method in corresponding alphabetical contracts. This method transfers all its **NEO Token** to its account, thereby producing utility token emission. Within the same invocation context, the contract transfers a share of the available **GAS Utility Token** to all Inner Ring wallets. Each contract will keep the $\frac{1}{8}$ 'th part on its balance as an emergency reserve.

$$InnerRingNodeEmission = G \cdot \frac{7}{8} \cdot \frac{1}{N}$$

G - contract's **GAS Utility Token** amount

N - length of Inner Ring list

After receiving **GAS Utility Token**, the nodes of the Inner Ring can periodically transfer a share to all registered Storage nodes and use the received utility token for FS chain operations: change epochs, register new containers, save data audit results, etc.

Storage nodes have a limited supply of **GAS Utility Token** to prevent malicious actions and DoS attacks. Depending on Storage Node activity and reputation records, it will receive a different utility token amount, normally enough to perform all required operations. FS chain GAS and mainnet GAS are different tokens, hence Storage Nodes don't spend rewards for internal operations and can't convert FS chain utility token into mainnet **GAS Utility Token** or vice versa.

Changing FS chain validators

Beforehand, Alphabet Inner Ring node candidates register validator keys in the list of candidates for the FS chain committee. When FS chain Netmap smart contract sends a notification regarding Inner Ring node list updates, Alphabet Inner Ring nodes invoke the `Vote ([] keys)` method of all Alphabet smart contracts in order to gather signatures and then make them vote for the FS chain Committee. Each Alphabet contract votes for the keys proposed by sending `VotesPerKey` votes for each key. Normally, there is just one key per node, hence N equals 1.

$$VotesPerKey = \frac{A}{N}$$

A - contract's NEO amount

N - length of proposed keys list

Changing the Inner Ring list

Inner Ring nodes follow a self-regulation process, allowing them to vote to substitute dead or malfunctioning nodes with new ones from the candidate list. Only Alphabet nodes prepare new Inner Ring node lists and vote for it, but all nodes listed there must confirm their participation via the same voting mechanism.

The voting procedure uses FS chain `Voting` smart contract, but the list of candidates is taken from mainnet NeoFS contract. When Inner Ring nodes agree on the updated list, it's submitted to the mainnet NeoFS smart contract and then mirrored back to Netmap smart contract on FS chain.

By using the `Emit ()` and `Vote ()` methods of Alphabet smart contracts, Inner Ring nodes take full control of FS chain. They control validator keys and utility token distribution. Thus, if the mainnet

committee will control list of Inner Ring nodes, then it will control FS chain as well.

The mainnet Committee can set the list of priority candidates for Alphabet Inner Ring nodes in the mainnet `DesignationContract`. Nodes from that list will be voted for becoming Alphabet Inner Ring nodes and substitute current Alphabet nodes, if they confirm the following requirements:

- Node's key is registered as a candidate in mainnet NeoFS smart contract
- Node's key is registered as an FS chain committee candidate
- Node is not listed as inactive in FS chain `Netmap` contract

The `DesignationContract` list may contain any number of valid candidates, and the voting process will make sure as many of them as possible are in the first seven active inner Ring nodes. If there is not enough appropriate candidates, the rest will be taken from the regular candidates list. If there are too many, only the first seven suitable nodes will be used.

The voting algorithm is the same for each Inner Ring node and starts in the following cases:

- New Epoch
- Notification from mainnet `DesignationContract` on Inner Ring nodes list change
- Notification from FS chain `Netmap` contract on inactive Inner Ring nodes list change

All Inner Ring nodes listen for notifications from the Voting contract. If they see themselves in the new Inner Ring nodes list, they confirm their participation by sending the same list in the `Prepare()` method. Only newly added nodes need to confirm their participation with a transaction. If the node is already in the active Inner Ring list, it doesn't need to send a confirmation.

When there are enough Alphabet signatures and all required candidate signatures have been sent with the `Prepare()` method, the last invocation will update the list and finish voting round.

Active Alphabet Inner Ring nodes will be waiting for the round to end and locally test invoke the `EndRound()` method. When the voting round timeout occurs and the round has not finished successfully through agreement on a new list, one of the Alphabet nodes will invoke `EndRound()` and settle the round's results.

If by the end of the voting round some newly added nodes haven't confirmed their participation, they are added to the `Netmap` smart contract's inactive list. This will trigger a new voting round without those inactive nodes.

If there are not enough candidates, Inner Ring nodes will accept the best list they can gather.

When the new list is agreed, the `Voting` smart contract sends a notification. All Alphabet nodes react with invocation of `UpdateInnerRing()` on the mainnet NeoFS smart contract. When the majority of Alphabet nodes send the update and mainnet list is updated, it will be mirrored by Alphabet Inner Ring nodes in the FS chain `Netmap` smart contract.

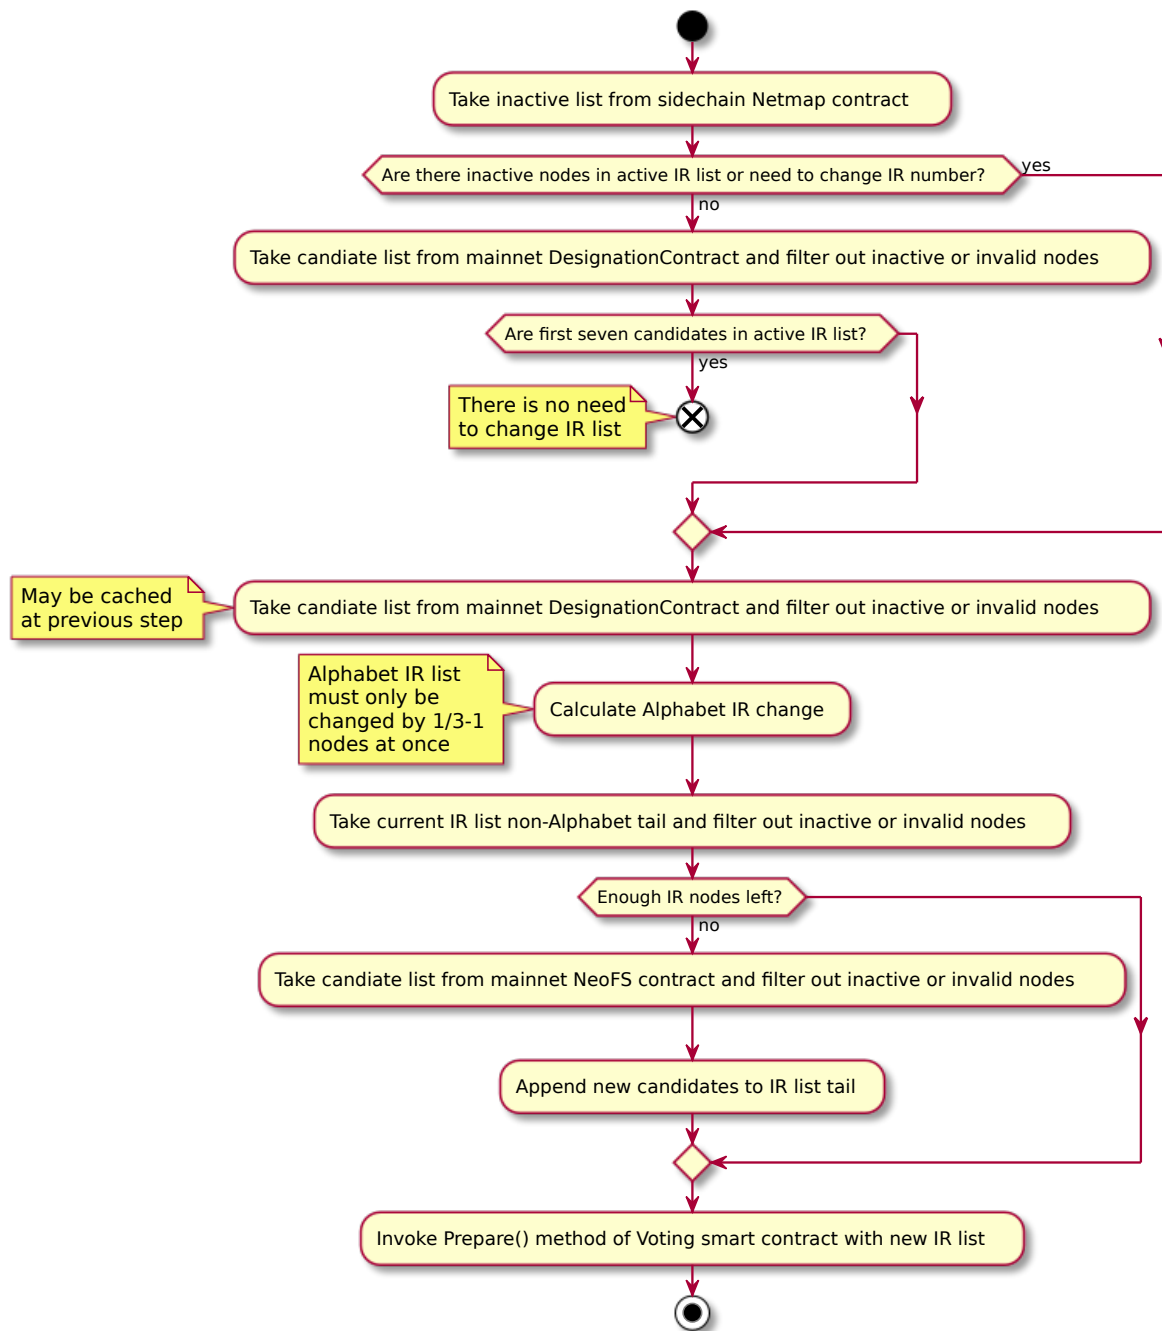


Figure 19: Inner Ring Alphabet node voting algorithm

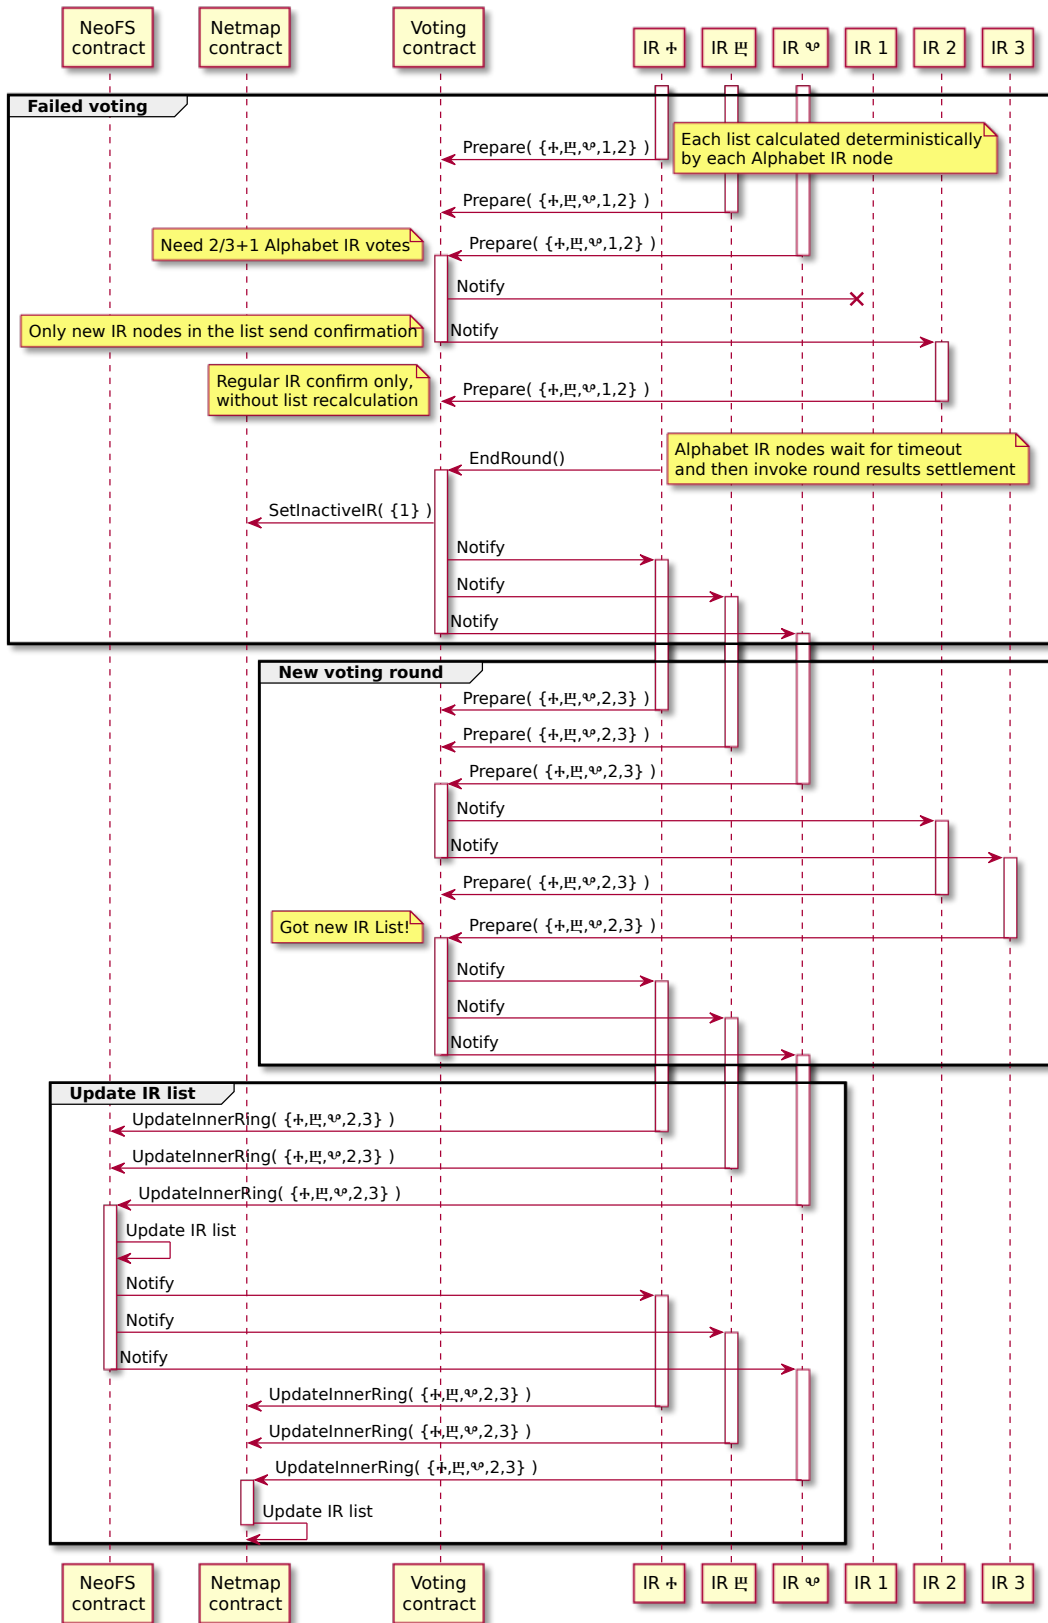


Figure 20: Inner Ring list update in mainnet and FS chain

NeoFS Smart Contracts

The API described below corresponds to neofs-contract v0.26.0²⁶.

alphabet contract

Package alphabet implements Alphabet contract which is deployed to FS chain.

Alphabet contract is designed to support GAS production and vote for new validators in FS chain. NEO token is required to produce GAS and vote for a new committee. It can be distributed among alphabet nodes of the Inner Ring. However, some of them may be malicious, and some NEO can be lost. It will destabilize the economics of FS chain. To avoid it, all 100,000,000 NEO are distributed among all alphabet contracts.

To identify alphabet contracts, they are named, names are set at contract deploy. Alphabet nodes of the Inner Ring communicate with one of the alphabetical contracts to emit GAS. To vote for a new list of side chain committee, alphabet nodes of the Inner Ring create multisignature transactions for each alphabet contract.

Contract notifications Alphabet contract does not produce notifications to process.

Contract methods

Emit

func Emit()

Emit method produces FS chain GAS and distributes it among Inner Ring nodes and proxy contract. It can be invoked only by an Alphabet node of the Inner Ring.

To produce GAS, an alphabet contract transfers all available NEO from the contract account to itself. 50% of the GAS in the contract account are transferred to proxy contract. 43.75% of the GAS are equally distributed among all Inner Ring nodes. Remaining 6.25% of the GAS stay in the contract.

Gas

func Gas() **int**

Gas returns the amount of FS chain GAS stored in the contract account.

²⁶<https://github.com/nspcc-dev/neofs-contract/releases/tag/v0.26.0>

Name**func** Name() *string*

Name returns the name of the contract set at deployment stage.

Neo**func** Neo() *int*

Neo returns the amount of FS chain NEO stored in the contract account.

OnNEP17Payment**func** OnNEP17Payment(from *interop.Hash160*, amount *int*, data any)

OnNEP17Payment is a callback for NEP-17 compatible native GAS and NEO contracts.

Update**func** Update(nefFile, manifest []*byte*, data any)

Update method updates contract source code and manifest. It can be invoked only by committee.

Verify**func** Verify() *bool*

Verify checks whether carrier transaction contains either $(2/3N + 1)$ or $(N/2 + 1)$ valid multi-signature of the NeoFS Alphabet.

Version**func** Version() *int*

Version returns the version of the contract.

Vote**func** Vote(epoch *int*, candidates []*interop.PublicKey*)

Vote method votes for the FS chain committee. It requires multisignature from Alphabet nodes of the Inner Ring.

This method is used when governance changes the list of Alphabet nodes of the Inner Ring. Alphabet nodes share keys with FS chain validators, therefore it is required to change them as well. To do that, NEO holders (which are alphabet contracts) should vote for a new committee.

audit contract

Package audit implements Audit contract which is deployed to FS chain.

Inner Ring nodes perform audit of the registered containers during every epoch. If a container contains StorageGroup objects, an Inner Ring node initializes a series of audit checks. Based on the results of these checks, the Inner Ring node creates a DataAuditResult structure for the container. The content of this structure makes it possible to determine which storage nodes have been examined and see the status of these checks. Regarding this information, the container owner is charged for data storage.

Audit contract is used as a reliable and verifiable storage for all DataAuditResult structures. At the end of data audit routine, Inner Ring nodes send a stable marshaled version of the DataAuditResult structure to the contract. When Alphabet nodes of the Inner Ring perform settlement operations, they make a list and get these AuditResultStructures from the audit contract.

Contract notifications Audit contract does not produce notifications to process.

Contract methods

Get

```
func Get(id []byte) []byte
```

Get method returns a stable marshaled DataAuditResult structure.

The ID of the DataAuditResult can be obtained from listing methods.

List

```
func List() [][]byte
```

List method returns a list of all available DataAuditResult IDs from the contract storage.

ListByCID

```
func ListByCID(epoch int, cid []byte) [][]byte
```

ListByCID method returns a list of DataAuditResult IDs generated during the specified epoch for the specified container.

ListByEpoch

```
func ListByEpoch(epoch int) [][]byte
```

ListByEpoch method returns a list of DataAuditResult IDs generated during the specified epoch.

ListByNode

```
func ListByNode(epoch int, cid [][]byte, key interop.PublicKey) [][]byte
```

ListByNode method returns a list of DataAuditResult IDs generated in the specified epoch for the specified container by the specified Inner Ring node.

Put

```
func Put(rawAuditResult [][]byte)
```

Put method stores a stable marshalled 'DataAuditResult' structure. It can be invoked only by Inner Ring nodes.

Inner Ring nodes perform audit of containers and produce 'DataAuditResult' structures. They are stored in audit contract and used for settlements in later epochs.

Update

```
func Update(nefFile, manifest [][]byte, data any)
```

Update method updates contract source code and manifest. It can be invoked only by committee.

Version

```
func Version() int
```

Version returns the version of the contract.

balance contract

Package balance implements Balance contract which is deployed to FS chain.

Balance contract stores all NeoFS account balances. It is a NEP-17 compatible contract, so it can be tracked and controlled by N3 compatible network monitors and wallet software.

This contract is used to store all micro transactions in FS chain, such as basic income settlements or container fee payments. It is inefficient to make such small payment transactions in main chain.

To process small transfers, balance contract has higher (12) decimal precision than native GAS contract.

NeoFS balances are synchronized with main chain operations. Deposit produces minting of NEOFS tokens in Balance contract. Withdraw locks some NEOFS tokens in a special lock account. When NeoFS contract transfers GAS assets back to the user, the lock account is destroyed with burn operation.

Contract notifications Transfer notification. This is a NEP-17 standard notification.

Transfer:

- name: from
type: Hash160
- name: to
type: Hash160
- name: amount
type: Integer

TransferX notification. This is an enhanced transfer notification with details.

TransferX:

- name: from
type: Hash160
- name: to
type: Hash160
- name: amount
type: Integer
- name: details
type: ByteArray

Lock notification. This notification is produced when a lock account is created. It contains information about main chain transaction that has produced the asset lock, the address of the lock account and the NeoFS epoch number until which the lock account is valid. Alphabet nodes of the Inner Ring catch notification and initialize Cheque method invocation of NeoFS contract.

Lock:

- name: txID
type: ByteArray
- name: from

- type: Hash160
- name: to
 - type: Hash160
- name: amount
 - type: Integer
- name: until
 - type: Integer

Payment notification. This notification is produced when container has been paid by its owner according to storage nodes' reports.

Payment

- name: UserID
 - type: Hash160
- name: ContainerID
 - type: Hash256
- name: Epoch
 - type: Integer
- name: Amount
 - type: Integer

ChangePaymentStatus notification. This notification is produced when container's payment status is changed. It is produced in both cases: when a paid earlier container is marked unpaid, and when an unpaid container is paid successfully.

ChangePaymentStatus

- name: ContainerID
 - type: Hash256
- name: Epoch
 - type: Integer
- name: Unpaid
 - type: Boolean

Contract methods

BalanceOf

```
func BalanceOf(account interop.Hash160) int
```

BalanceOf is a NEP-17 standard method that returns NeoFS balance of the specified account.

Burn

func Burn(from interop.Hash160, amount int, txDetails []byte)

Burn is a method that transfers assets from a user account to an empty account. It can be invoked only by Alphabet nodes of the Inner Ring.

It produces Transfer and TransferX notifications.

Burn method is invoked by Alphabet nodes of the Inner Ring when they process Cheque notification from NeoFS contract. It means that locked assets have been transferred to the user in main chain, therefore the lock account should be destroyed. Before that, Alphabet nodes should synchronize precision of main chain GAS contract and Balance contract. Burn decreases total supply of NEP-17 compatible NeoFS token.

Decimals

func Decimals() int

Decimals is a NEP-17 standard method that returns precision of NeoFS balances.

GetUnpaidContainerEpoch

func GetUnpaidContainerEpoch(cid interop.Hash256) int

GetUnpaidContainerEpoch returns epoch from which container is considered unpaid. Returns -1 if container is paid successfully.

IterateUnpaid

func IterateUnpaid() iterator.Iterator

IterateUnpaid is like [GetUnpaidContainerEpoch](#) but for every unpaid container. Iteration is through key-value pair, where key is container ID, value is epoch from which container is considered unpaid.

Lock

func Lock(txDetails []byte, from, to interop.Hash160, amount, until int)

Lock is a method that transfers assets from a user account to the lock account related to the user. It can be invoked only by Alphabet nodes of the Inner Ring.

It produces Lock, Transfer and TransferX notifications.

Lock method is invoked by Alphabet nodes of the Inner Ring when they process Withdraw notification from NeoFS contract. This should transfer assets to a new lock account that won't be used for anything beside Unlock and Burn.

Mint

func Mint(to interop.Hash160, amount int, txDetails []byte)

Mint is a method that transfers assets to a user account from an empty account. It can be invoked only by Alphabet nodes of the Inner Ring.

It produces Transfer and TransferX notifications.

Mint method is invoked by Alphabet nodes of the Inner Ring when they process Deposit notification from NeoFS contract. Before that, Alphabet nodes should synchronize precision of main chain GAS contract and Balance contract. Mint increases total supply of NEP-17 compatible NeoFS token.

NewEpoch

func NewEpoch(epochNum int)

NewEpoch is a method that checks timeout on lock accounts and returns assets if lock is not available anymore. It can be invoked only by NewEpoch method of Netmap contract.

It produces Transfer and TransferX notifications.

SettleContainerPayment

func SettleContainerPayment(cid interop.Hash256) bool

SettleContainerPayment distributes storage payments from container's owner account to storage nodes that serve container's objects. Transaction must be witnessed by the actual Alphabet multi-signature. Produces 'ChangePaymentStatus' notification on any payment status changes. If payment is successful, 'Payment' notification is thrown. If payment cannot be fully made, container is registered as an unpaid one, see [GetUnpaidContainerEpoch](#).

Symbol

func Symbol() string

Symbol is a NEP-17 standard method that returns NEOFS token symbol.

TotalSupply

func TotalSupply() int

TotalSupply is a NEP-17 standard method that returns total amount of main chain GAS in NeoFS network.

Transfer

func Transfer(from, to interop.Hash160, amount **int**, data any) **bool**

Transfer is a NEP-17 standard method that transfers NeoFS balance from one account to another. It can be invoked only by the account owner.

It produces Transfer and TransferX notifications. TransferX notification will have empty details field.

TransferX

func TransferX(from, to interop.Hash160, amount **int**, details []**byte**) **bool**

TransferX is a method for NeoFS balance to be transferred from one account to another. It can be invoked by the account owner or by Alphabet nodes.

It produces Transfer and TransferX notifications.

TransferX method expands Transfer method by having extra details argument. TransferX method also allows to transfer assets by Alphabet nodes of the Inner Ring with multisignature.

Update

func Update(nefFile, manifest []**byte**, data any)

Update method updates contract source code and manifest. It can be invoked only by committee.

Version

func Version() **int**

Version returns the version of the contract.

container contract

Package container implements Container contract which is deployed to FS chain.

Container contract stores and manages containers, extended ACLs and container size estimations. Contract does not perform sanity or signature checks of containers or extended ACLs, it is done by Alphabet nodes of the Inner Ring. Alphabet nodes approve it by invoking the same Put or SetEACL methods with the same arguments.

Contract notifications containerPut notification. This notification is produced when a user wants to create a new container. Alphabet nodes of the Inner Ring catch the notification and validate container data, signature and token if present.

containerPut:

- name: container
type: ByteArray
- name: signature
type: Signature
- name: publicKey
type: PublicKey
- name: token
type: ByteArray

containerDelete notification. This notification is produced when a container owner wants to delete a container. Alphabet nodes of the Inner Ring catch the notification and validate container ownership, signature and token if present.

containerDelete:

- name: containerID
type: ByteArray
- name: signature
type: Signature
- name: token
type: ByteArray

nodesUpdate notification. This notification is produced when a container roster is changed. Triggered only by the Alphabet at the beginning of epoch.

name: NodesUpdate

- name: ContainerID
type: hash256

setEACL notification. This notification is produced when a container owner wants to update an extended ACL of a container. Alphabet nodes of the Inner Ring catch the notification and validate container ownership, signature and token if present.

setEACL:

- name: eACL
type: ByteArray
- name: signature
type: Signature
- name: publicKey
type: PublicKey
- name: token
type: ByteArray

ContainerQuotaSet notification. This notification is produced when container's owner sets (updates) size limitation for storage used for all objects in this container.

ContainerQuotaSet

- name: ContainerID
type: Hash256
- name: LimitValue
type: Integer
- name: Hard
type: Boolean

UserQuotaSet notification. This notification is produced when container's owner sets (updates) size limitation for storage used for all objects in `_all_` containers he owns.

UserQuotaSet

- name: UserID
type: ByteArray # 25 byte N3 address
- name: LimitValue
type: Integer
- name: Hard
type: Boolean

AttributeChanged notification. Produced on successful `setAttribute` or `removeAttribute` invocations when mentioned attribute is affected.

AttributeChanged

- name: containerID
type: Hash256
- name: attribute
type: String # Name of attribute

Contract methods

AddNextEpochNodes

```
func AddNextEpochNodes(cID interop.Hash256, placementVector uint8,  
    ↪ publicKeys []interop.PublicKey)
```

AddNextEpochNodes accumulates passed nodes as container members for the next epoch to be committed using **CommitContainerListUpdate**. Nodes must be grouped by selector index from placement policy (SELECT clauses). Results of the call operation can be received via **Nodes**. This method must be called only when a container list is changed, otherwise nothing should be done. Call must be signed by the Alphabet nodes.

AddStructs

```
func AddStructs() bool
```

AddStructs makes and saves structures for up to 10 existing containers. Returns true if not all stored containers have been structurized. Recall continues the process. Returning false means all containers have been processed.

AddStructs requires Alphabet witness.

AddStructs throws NEP-11 'Transfer' event representing token mint for each handled container.

Alias

```
func Alias(cid []byte) string
```

Alias method returns a string with an alias of the container if it's set (Null otherwise).

If the container doesn't exist, it panics with `NotFoundError`.

BalanceOf

```
func BalanceOf(owner interop.Hash160) int
```

BalanceOf returns number of containers owner by given account.

BalanceOf implements NEP-11 method.

CommitContainerListUpdate

```
func CommitContainerListUpdate(cID interop.Hash256, replicas []uint8)
```

CommitContainerListUpdate commits container list changes made by `AddNextEpochNodes` calls in advance. Replicas must correspond to ordered placement policy (REP clauses). If no `AddNextEpochNodes` have been made, it clears container list. Makes “ContainerUpdate” notification with container ID after successful list change. Call must be signed by the Alphabet nodes.

ContainersOf

func ContainersOf(owner []byte) iterator.Iterator

ContainersOf iterates over all container IDs owned by the specified owner. If owner is nil, it iterates over all containers.

Deprecated: use `TokensOf` for non-empty and `Tokens` for empty owner correspondingly.

Count

func Count() int

Count method returns the number of registered containers.

Deprecated: use `TotalSupply` instead.

Create

func Create(cnr []byte, invocScript, verifScript, sessionToken []byte, name, zone string, metaOnChain bool)

Create saves container descriptor serialized according to the NeoFS API binary protocol. Created containers are content-addressed: they may be accessed by SHA-256 checksum of their data. On success, Create throws ‘Created’ notification event.

Created containers are disposable: if they are deleted, they cannot be created again. Create throws [cst.ErrorDeleted] exception on recreation attempts.

Domain name is optional. If specified, it is used to register ‘name.zone’ domain for given container. Domain zone is optional: it defaults to the 6th contract deployment parameter which itself defaults to ‘container’.

Meta-on-chain boolean flag specifies whether meta information about objects from this container can be collected for it.

The operation is paid. Container owner pays per-container fee (global chain configuration) to each committee member. If domain registration is requested, additional alias fee (also a configuration) is added to each payment.

Create must have chain's committee multi-signature witness. Invocation script, verification script and session token parameters are owner credentials. They are transmitted in notary transactions carrying original users' requests. IR verifies requests and approves them via multi-signature. Once creation is approved, container is persisted and becomes accessible. Credentials are disposable and do not persist in the chain.

Deprecated: use `CreateV2` with instead.

CreateV2

```
func CreateV2(cnr Info, invocScript, verifScript, sessionToken []byte)
↳ interop.Hash256
```

CreateV2 creates container with given info in the contract. Created containers are content-addressed: they may be accessed by SHA-256 checksum of their data. On success, CreateV2 throws 'Created' notification event.

Created containers are disposable: if they are deleted, they cannot be created again. CreateV2 throws [cst.ErrorDeleted] exception on recreation attempts.

If '___NEOFS__NAME' and '___NEOFS__ZONE' attributes are set, they are used to register 'name.zone' domain for given container. Domain zone is optional: it defaults to the 6th contract deployment parameter which itself defaults to 'container'.

If '___NEOFS__METAINFO_CONSISTENCY' attribute is set, meta information about objects from this container can be collected for it.

If cnr contains '___NEOFS__LOCK_UNTIL' attribute, its value must be a valid Unix Timestamp later the current one. On success, referenced container becomes locked for removal until specified time.

The operation is paid. Container owner pays per-container fee (global chain configuration) to each committee member. If domain registration is requested, additional alias fee (also a configuration) is added to each payment.

CreateV2 must have chain's committee multi-signature witness. Invocation script, verification script and session token parameters are owner credentials. They are transmitted in notary transactions carrying original users' requests. IR verifies requests and approves them via multi-signature. Once creation is approved, container is persisted and becomes accessible. Credentials are disposable and do not persist in the chain.

Decimals

```
func Decimals() int
```

Decimals returns static zero meaning containers are Non-divisible NFTs.

Decimals implements NEP-11 method.

Delete

func Delete(containerID []byte, signature interop.Signature, token []byte)

Delete method removes a container from the contract storage if it has been invoked by Alphabet nodes of the Inner Ring. Otherwise, it produces containerDelete notification.

Signature is a RFC6979 signature of the container ID. Token is optional and should be a stable marshaled SessionToken structure from API.

If the container doesn't exist, it panics with NotFoundError.

Deprecated: use **Remove** instead.

GetContainerData

func GetContainerData(id []byte) []byte

GetContainerData returns binary of the container it was created with by ID.

If the container is missing, GetContainerData throws [cst.NotFoundError] exception.

Deprecated: use **GetInfo** instead.

GetEACLData

func GetEACLData(id []byte) []byte

GetEACLData returns binary of container eACL it was put with by the container ID.

If the container is missing, GetEACLData throws [cst.NotFoundError] exception.

GetTakenSpaceByUser

func GetTakenSpaceByUser(user []byte) int

GetTakenSpaceByUser returns total load space in all containers user owns. If user have no containers, it returns 0.

IterateAllReportSummaries

func IterateAllReportSummaries() iterator.Iterator

IterateAllReportSummaries method returns iterator over all total container sizes that have been registered for the specified epoch. Items returned from this iterator are key-value pairs with keys having container ID as a prefix and values being [NodeReportSummary](#) structures.

IterateBillingStats

func IterateBillingStats(cid interop.Hash256) iterator.Iterator

IterateBillingStats method returns iterator over container's billing statistics made based on [NodeReport](#).

IterateReports

func IterateReports(cid interop.Hash256) iterator.Iterator

IterateReports method returns iterator over nodes' reports that were claimed for specified epoch and container.

Nodes

func Nodes(cID interop.Hash256, placementVector uint8) iterator.Iterator

Nodes returns iterator over members of the container. The list is handled by the Alphabet nodes and must be updated via [AddNextEpochNodes](#) and [CommitContainerListUpdate](#) calls.

OnNEP11Payment

func OnNEP11Payment(a interop.Hash160, b int, c []byte, d any)

OnNEP11Payment is needed for registration with contract as the owner to work.

Owner

func Owner(containerID []byte) []byte

Owner method returns a 25 byte Owner ID of the container.

If the container doesn't exist, it panics with [NotFoundError](#).

Deprecated: use [OwnerOf](#) instead.

OwnerOf

```
func OwnerOf(tokenID []byte) interop.Hash160
```

OwnerOf returns owner of the container identified by tokenID.

If referenced container does not exist, OwnerOf throws [cst.NotFoundError] exception. If the container has already been removed, OwnerOf throws [cst.ErrorDeleted] exception.

OwnerOf implements NEP-11 method.

Properties

```
func Properties(tokenID []byte) map[string]any
```

Properties returns properties of referenced container. The properties are 'name' and all KV attributes.

The 'name' property is set to 'Name' attribute value if exists. In this case, 'Name' attribute itself is not included. Otherwise, if 'Name' attribute is missing, 'name' is set to Base58-encoded tokenID. Note that container 'name' attribute is always overlapped if any.

If referenced container does not exist, OwnerOf throws [cst.NotFoundError] exception. If the container has already been removed, OwnerOf throws [cst.ErrorDeleted] exception.

Properties implements optional NEP-11 method.

Put

```
func Put(container []byte, signature interop.Signature, publicKey  
↪ interop.PublicKey, token []byte)
```

Put method creates a new container if it has been invoked by Alphabet nodes of the Inner Ring. Otherwise, it produces containerPut notification.

Container should be a stable marshaled Container structure from API. Signature is a RFC6979 signature of the Container. PublicKey contains the public key of the signer. Token is optional and should be a stable marshaled SessionToken structure from API.

Deprecated: use [CreateV2](#) instead.

PutEACL

```
func PutEACL(eACL []byte, invocScript, verifScript, sessionToken []byte)
```

PutEACL puts given eACL serialized according to the NeoFS API binary protocol for the container it is referenced to. Operation must be allowed in the container's basic ACL. If container does not exist,

PutEACL throws [cst.NotFoundError] exception. On success, PutEACL throws 'EACLChanged' notification event.

See [CreateV2](#) for details.

PutMeta

```
func PutMeta(container []byte, signature interop.Signature, publicKey  
↪ interop.PublicKey, token []byte, name, zone string, metaOnChain bool)
```

PutMeta is the same as [Put](#) and [PutNamed](#) (and exposed as put from the contract via overload), but allows named containers and container's meta-information be handled and notified using the chain. If name and zone are non-empty strings, it behaves the same as [PutNamed](#); empty strings make a regular [Put](#) call.

Deprecated: use [CreateV2](#) instead.

PutNamed

```
func PutNamed(container []byte, signature interop.Signature, publicKey  
↪ interop.PublicKey, token []byte, name, zone string)
```

PutNamed is similar to put but also sets a TXT record in nns contract. Note that zone must exist.

Deprecated: use [CreateV2](#) instead.

PutNamedOverloaded

```
func PutNamedOverloaded(container []byte, signature interop.Signature,  
↪ publicKey interop.PublicKey, token []byte, name, zone string)
```

PutNamedOverloaded is the same as [Put](#) (and exposed as put from the contract via overload), but allows named container creation via NNS contract.

Deprecated: use [CreateV2](#) instead.

PutReport

```
func PutReport(cid interop.Hash256, sizeBytes, objsNumber int, pubKey  
↪ interop.PublicKey)
```

PutReport method saves container's state report in contract memory. It must be invoked only by Storage nodes that belong to reported container. This method checks witness based on the provided public key of the Storage node. sizeBytes is a total storage that is used by storage node for storing all marshaled objects that belong to the specified container. objsNumber is a total number of container's object storage node have.

If the container doesn't exist, it panics with [cst.NotFoundError].

Remove

func Remove(id []byte, invocScript, verifScript, sessionToken []byte)

Remove removes all data for the referenced container. Remove is no-op if container does not exist. On success, Remove throws 'Removed' notification event.

If the container has '__NEOFS__LOCK_UNTIL' attribute with timestamp that has not passed yet, Remove throws exception containing [cst.ErrorLocked].

See [CreateV2](#) for details.

RemoveAttribute

func RemoveAttribute(cID interop.Hash256, name string, validUntil int,
↪ invocScript, verifScript, sessionToken []byte)

RemoveAttribute removes container attribute. Not all container attributes can be removed with RemoveAttribute. The supported list of attributes:

If name is '__NEOFS__LOCK_UNTIL', current time must be later than the currently set one if any.

The validUntil must be Unix Timestamp which has not yet pass.

RemoveAttribute must have Alphabet witness. Invocation and verification scripts must authenticate either container owner or session subject if any. If session token is set, it must be issued by the container owner and have at least one context with referenced container and 'REMOVEATTRIBUTE' verb. The sessionToken parameter must be encoded according to NeoFS API binary protocol.

If container is missing, RemoveAttribute throws [cst.NotFoundError] exception.

ReplicasNumbers

func ReplicasNumbers(cID interop.Hash256) iterator.Iterator

ReplicasNumbers returns iterator over saved by [CommitContainerListUpdate](#) container's replicas from placement policy.

SetAttribute

func SetAttribute(cID interop.Hash256, name, value string, validUntil int,
↪ invocScript, verifScript, sessionToken []byte)

SetAttribute sets container attribute. Not all container attributes can be changed with SetAttribute. The supported list of attributes:

CORS attribute gets JSON encoded '[]CORSRule' as value.

If name is '`__NEOFS__LOCK_UNTIL`', value must be a valid Unix Timestamp later than the current and already set (if any) ones. On success, referenced container becomes locked for removal until specified time.

If the name is '`S3_TAGS`', the value must be a valid JSON map, where the key is the tag name and the value is the tag value. It is an S3 gate-specific attribute. For instance: {"my-tag": "my-value"}.

If the name is '`S3_SETTINGS`', the value is not validated by the contract, but must be valid JSON. It is an S3 gate-specific attribute. The structure of it is controlled by the gate itself.

If the name is '`S3_NOTIFICATIONS`', the value is not validated by the contract, but must be valid JSON. It is an S3 gate-specific attribute. The structure of it is controlled by the gate itself.

The validUntil must be Unix Timestamp which has not yet passed.

SetAttribute must have Alphabet witness. Invocation and verification scripts must authenticate either container owner or session subject if any. If session token is set, it must be issued by the container owner and have at least one context with referenced container and 'SETATTRIBUTE' verb. The sessionToken parameter must be encoded according to NeoFS API binary protocol.

If container is missing, SetAttribute throws [cst.NotFoundError] exception.

SetEACL

```
func SetEACL(eACL []byte, signature interop.Signature, publicKey  
↪ interop.PublicKey, token []byte)
```

SetEACL method sets a new extended ACL table related to the contract if it was invoked by Alphabet nodes of the Inner Ring. Otherwise, it produces setEACL notification.

EACL should be a stable marshaled EACLTable structure from API. Protocol version and container reference must be set in 'version' and 'container_id' fields respectively. Signature is a RFC6979 signature of the Container. PublicKey contains the public key of the signer. Token is optional and should be a stable marshaled SessionToken structure from API.

If the container doesn't exist, it panics with NotFoundError.

Deprecated: use PutEACL instead.

SetHardContainerQuota

```
func SetHardContainerQuota(cID interop.Hash256, size int)
```

`SetHardContainerQuota` sets hard size quota that limits all space used for storing objects in cID (including object replicas). Non-positive size sets no limitation. After exceeding the limit nodes will refuse any further PUTs. Call must be signed by cID's owner. Limit can be changed with a repeated call. See also `SetSoftContainerQuota`. Panics if cID is incorrect or container does not exist.

SetHardUserQuota

func `SetHardUserQuota(user []byte, size int)`

`SetHardUserQuota` sets size quota that limits all space used for storing objects in all containers that belong to user (including object replicas). Non-positive size sets no limitation. After exceeding the limit nodes will refuse any further PUTs. Call must be signed by user. Limit can be changed with a repeated call. See also `SetSoftUserQuota`. Panics if user is incorrect.

SetSoftContainerQuota

func `SetSoftContainerQuota(cID interop.Hash256, size int)`

`SetSoftContainerQuota` sets soft size quota that limits all space used for storing objects in cID (including object replicas). Non-positive size sets no limitation. After exceeding the limit nodes are instructed to warn only, without denial of service. Call must be signed by cID's owner. Limit can be changed with a repeated call. See also `SetHardContainerQuota`. Panics if cID is incorrect or container does not exist.

SetSoftUserQuota

func `SetSoftUserQuota(user []byte, size int)`

`SetSoftUserQuota` sets size quota that limits all space used for storing objects in all containers that belong to user (including object replicas). Non-positive size sets no limitation. After exceeding the limit nodes are instructed to warn only, without denial of service. Call must be signed by user. Limit can be changed with a repeated call. See also `SetHardUserQuota`. Panics if user is incorrect.

SubmitObjectPut

func `SubmitObjectPut(metaInformation []byte, sigs [][]interop.Signature)`

`SubmitObjectPut` registers successful object PUT operation and notifies about it. `metaInformation` must be signed by container nodes according to container's placement, see `VerifyPlacementSignatures`. `metaInformation` must contain information about an object placed to a container that was created using `Put (PutMeta)` with enabled meta-on-chain option.

Symbol

```
func Symbol() string
```

Symbol returns static 'FSCNTR'.

Symbol implements NEP-11 method.

Tokens

```
func Tokens() iterator.Iterator
```

Tokens returns iterator over IDs of all existing containers.

Tokens implements optional NEP-11 method.

TokensOf

```
func TokensOf(owner interop.Hash160) iterator.Iterator
```

TokensOf returns iterator over IDs of all containers owned by given account.

TokensOf implements NEP-11 method.

TotalSupply

```
func TotalSupply() int
```

TotalSupply returns total number of existing containers.

TotalSupply implements NEP-11 method.

Transfer

```
func Transfer(to interop.Hash160, tokenID []byte, data any) bool
```

Transfer makes to an owner of the container identified by tokenID.

If referenced container does not exist, Transfer throws [cst.NotFoundError] exception. If the container has already been removed, Transfer throws [cst.ErrorDeleted] exception.

Transfer implements NEP-11 method.

Update

```
func Update(nefFile, manifest []byte, data any)
```

Update method updates contract source code and manifest. It can be invoked by committee only.

VerifyPlacementSignatures

```
func VerifyPlacementSignatures(cid interop.Hash256, msg []byte, sigs  
↪ [][]interop.Signature) bool
```

VerifyPlacementSignatures verifies that message has been signed by container members according to container's placement policy: there should be at least REP number of signatures for every placement vector. sigs must be container's number of SELECTs length.

Version

```
func Version() int
```

Version returns the version of the contract.

neofs contract

Package neofs contains NeoFS contract which is deployed to main chain.

NeoFS contract is an entry point to NeoFS users. This contract stores all NeoFS related GAS, registers new Inner Ring candidates and produces notifications to control FS chain.

While main chain committee controls the list of Alphabet nodes in native RoleManagement contract, NeoFS can't change more than 1\3 keys at a time. NeoFS contract contains the actual list of Alphabet nodes in FS chain.

Network configuration is also stored in NeoFS contract. All changes in configuration are mirrored in FS chain with notifications.

Contract notifications Deposit notification. This notification is produced when user transfers native GAS to the NeoFS contract address. The same amount of NEOFS token will be minted in Balance contract in FS chain.

Deposit:

- name: from
type: Hash160
- name: amount
type: Integer
- name: receiver
type: Hash160
- name: txHash
type: Hash256

Withdraw notification. This notification is produced when a user wants to withdraw GAS from the internal NeoFS balance and has paid fee for that.

Withdraw:

- name: user
type: Hash160
- name: amount
type: Integer
- name: txHash
type: Hash256

Cheque notification. This notification is produced when NeoFS contract has successfully transferred assets back to the user after withdraw.

Cheque:

- name: id
type: ByteArray
- name: user
type: Hash160
- name: amount
type: Integer
- name: lockAccount
type: ByteArray

AlphabetUpdate notification. This notification is produced when Alphabet nodes have updated their lists in the contract. Alphabet argument is an array of ByteArray. It contains public keys of new alphabet nodes.

AlphabetUpdate:

- name: id
type: ByteArray
- name: alphabet
type: Array

SetConfig notification. This notification is produced when Alphabet nodes update NeoFS network configuration value.

SetConfig

- name: id
type: ByteArray
- name: key
type: ByteArray
- name: value
type: ByteArray

Contract methods

AlphabetAddress

func AlphabetAddress() interop.Hash160

AlphabetAddress returns 2^{3n+1} multisignature address of alphabet nodes. It is used in notary-disabled FS chain environment.

AlphabetList

func AlphabetList() []common.IRNode

AlphabetList returns an array of alphabet node keys. It is used in notary-disabled FS chain environment.

AlphabetUpdate

func AlphabetUpdate(id []byte, args []interop.PublicKey)

AlphabetUpdate updates a list of alphabet nodes with the provided list of public keys. It can be invoked only by alphabet nodes.

This method is used in notary-disabled FS chain environment. In this case, the actual alphabet list should be stored in the NeoFS contract.

Cheque

func Cheque(id []byte, user interop.Hash160, amount int, lockAcc []byte)

Cheque transfers GAS back to the user from the contract account, if assets were successfully locked in NeoFS balance contract. It can be invoked only by Alphabet nodes.

This method produces Cheque notification to burn assets in FS chain.

Config

func Config(key []byte) any

Config returns configuration value of NeoFS configuration. If the key does not exist, returns nil.

OnNEP17Payment

func OnNEP17Payment(from interop.Hash160, amount int, data any)

OnNEP17Payment is a callback for NEP-17 compatible native GAS contract. It takes no more than 9000.0 GAS. Native GAS has precision 8, and NeoFS balance contract has precision 12. Values bigger than 9000.0 can break JSON limits for integers when precision is converted.

SetConfig

func SetConfig(id, key, val []byte)

SetConfig key-value pair as a NeoFS runtime configuration value. It can be invoked only by Alphabet nodes.

Update

func Update(nefFile, manifest []byte, data any)

Update method updates contract source code and manifest. It can be invoked only by the FS chain committee.

Version

func Version() int

Version returns version of the contract.

Withdraw

func Withdraw(user interop.Hash160, amount int)

Withdraw initializes gas asset withdraw from NeoFS. It can be invoked only by the specified user.

This method produces Withdraw notification to lock assets in FS chain and transfers withdraw fee from a user account to each Alphabet node. If notary is enabled in main chain, fee is transferred to Processing contract. Fee value is specified in NeoFS network config with the key WithdrawFee.

netmap contract

Package netmap contains implementation of the Netmap contract for NeoFS systems.

Netmap contract stores and manages NeoFS network map, storage node candidates and epoch number counter. Currently it maintains two lists simultaneously, both in old BLOB-based format and new structured one. Nodes and IR are encouraged to use both during transition period, old format will eventually be discontinued.

Contract notifications AddNode notification. This notification is emitted when a storage node joins the candidate list via AddNode method (using new format).

AddNode

- name: publicKey
type: PublicKey
- name: addresses
type: Array
- name: attributes
type: Map

AddPeer notification. This notification is produced when a Storage node sends a bootstrap request by invoking AddPeer method.

AddPeer

- name: nodeInfo
type: ByteArray

UpdateStateSuccess notification. This notification is produced when a storage node or Alphabet changes node state by invoking UpdateState or DeleteNode methods. Supported states: online/offline/maintenance.

UpdateStateSuccess

- name: publicKey
type: PublicKey
- name: state
type: Integer

NewEpoch notification. This notification is produced when a new epoch is applied in the network by invoking NewEpoch method.

NewEpoch

- name: epoch
- type: Integer

NewEpochSubscription notification. This notification is produced when a new contract was subscribed to new epoch event.

NewEpochSubscription

- name: contract
- type: Hash160

Contract methods

AddNode

func AddNode(n Node2)

AddNode adds a new node into the candidate list for the next epoch. Node must have [nodestate.Online] state to be considered and the request must be signed by both the node and Alphabet. AddNode event is emitted upon success.

CleanupThreshold

func CleanupThreshold() int

CleanupThreshold returns the cleanup threshold configuration. Nodes that do not update their state for this number of epochs get kicked out of the network map. Zero means cleanups are disabled.

Config

func Config(key []byte) any

Config returns configuration value of NeoFS configuration. If key does not exist, returns nil.

DeleteNode

func DeleteNode(pkey interop.PublicKey)

DeleteNode removes a node with the given public key from candidate list. It must be signed by Alphabet nodes and doesn't require node witness. See [UpdateState](#) as well, this method emits UpdateStateSuccess upon success with state [nodestate.Offline].

Epoch

func Epoch() int

Epoch method returns the current epoch number.

GetEpochBlock

func GetEpochBlock(epoch int) int

GetEpochBlock returns block index when given epoch came. Returns 0 if the epoch is missing. Do not confuse with [GetEpochTime](#).

Use [LastEpochBlock](#) if you are interested in the current epoch.

GetEpochBlockByTime

func GetEpochBlockByTime(ts int) int

GetEpochBlockByTime returns the block index when the latest epoch that started not later than the provided block time came. Returns 0 if the time precedes any known epoch.

GetEpochTime

func GetEpochTime(epoch int) int

GetEpochTime returns block time when given epoch came. Returns 0 if the epoch is missing. Do not confuse with [GetEpochBlock](#).

Use [LastEpochTime](#) if you are interested in the current epoch.

InnerRingList

func InnerRingList() []common.IRNode

InnerRingList method returns a slice of structures that contains the public key of an Inner Ring node. It should be used in notary disabled environment only.

If notary is enabled, look to NeoFSAlphabet role in native RoleManagement contract of FS chain.

Deprecated: since non-notary settings are no longer supported, refer only to the RoleManagement contract only. The method will be removed in one of the future releases.

IsStorageNode

func IsStorageNode(key interop.PublicKey) bool

IsStorageNode allows to check for the given key presence in the current network map.

IsStorageNodeInEpoch

func IsStorageNodeInEpoch(key interop.PublicKey, epoch int) bool

IsStorageNodeInEpoch is the same as [IsStorageNode](#), but allows to do the check for previous epochs if they're still stored in the contract. If this epoch is no longer stored (or too new) it will return false.

IsStorageNodeStatus

func IsStorageNodeStatus(key interop.PublicKey, epoch int, status
↪ nodestate.Type) bool

IsStorageNodeStatus checks if the given status matches for the given node in the specified epoch.

LastEpochBlock

func LastEpochBlock() int

LastEpochBlock method returns the block number when the current epoch was applied. Do not confuse with [LastEpochTime](#).

Use [GetEpochBlock](#) for specific epoch.

LastEpochTime

func LastEpochTime() int

LastEpochTime method returns the block time when the current epoch was applied. Do not confuse with [LastEpochBlock](#).

Use [GetEpochTime](#) for specific epoch.

ListCandidates

func ListCandidates() iterator.Iterator

ListCandidates returns an iterator for a set of current candidate nodes. Iterator values are [Candidate](#) structures.

ListNodes

func ListNodes() iterator.Iterator

ListNodes provides an iterator to walk over current (as in corresponding to the current epoch network map) node set. Iterator values are [Node2](#) structures.

ListNodesEpoch

func ListNodesEpoch(epoch `int`) iterator.Iterator

ListNodesEpoch provides an iterator to walk over node set at the given epoch. It's the same as `ListNodes` (and exposed as `listNodes` from the contract via overload), but allows to query a particular epoch data if it's still stored. If this epoch is already expired (or not happened yet) returns an empty iterator.

NewEpoch

func NewEpoch(epochNum `int`)

NewEpoch method changes the epoch number up to the provided epochNum argument. It can be invoked only by Alphabet nodes. If provided epoch number is less than the current epoch number or equals it, the method throws panic.

When epoch number is updated, the contract sets storage node candidates as the current network map. The contract also invokes NewEpoch method on Balance and Container contracts.

It produces NewEpoch notification.

SetCleanupThreshold

func SetCleanupThreshold(val `int`)

SetCleanupThreshold sets cleanup threshold configuration. Negative values are not allowed. Zero disables stale node cleanups on epoch change.

SetConfig

func SetConfig(id, key, val `[]byte`)

SetConfig key-value pair as a NeoFS runtime configuration value. It can be invoked only by Alphabet nodes.

SubscribeForNewEpoch

func SubscribeForNewEpoch(contract `interop.Hash160`)

SubscribeForNewEpoch registers passed contract as a NewEpoch event subscriber. Such a contract must have a 'NewEpoch' method with a single numeric parameter. Transactions that call SubscribeForNewEpoch must be witnessed by the Alphabet. Produces 'NewEpochSubscription' notification event with a just registered recipient in a success case.

UnsubscribeFromNewEpoch

func UnsubscribeFromNewEpoch(contract interop.Hash160)

UnsubscribeFromNewEpoch removes new epoch subscription made by [SubscribeForNewEpoch](#) beforehand. Does nothing if subscription has not been made. Transactions that call UnsubscribeFromNewEpoch must be witnessed by the Alphabet. Produces ‘NewEpochUnsubscription’ notification event if any unsubscription was made.

Update

func Update(nefFile, manifest []byte, data any)

Update method updates contract source code and manifest. It can be invoked only by committee.

UpdateSnapshotCount

func UpdateSnapshotCount(count int)

UpdateSnapshotCount updates the number of the stored snapshots. If a new number is less than the old one, old snapshots are removed. Otherwise, history is extended with empty snapshots, so ‘Snapshot’ method can return invalid results for ‘diff = new-old’ epochs until ‘diff’ epochs have passed.

Count MUST NOT be negative.

UpdateState

func UpdateState(state nodestate.Type, publicKey interop.PublicKey)

UpdateState proposes a new state of candidate for the next-epoch network map. The candidate is identified by the given public key. Call transaction MUST be signed by the provided public key, i.e. by node itself. If the signature is correct, the Notary service will submit a request for signature by the NeoFS Alphabet. After collecting a sufficient number of signatures, the candidate’s state will be switched to the given one (‘UpdateStateSuccess’ notification is thrown after that).

UpdateState panics if requested candidate is missing in the current candidate set. UpdateState drops candidate from the candidate set if it is switched to [nodestate.Offline].

State MUST be from the [nodestate.Type] enum. Public key MUST be interop.PublicKeyCompressedLen bytes.

Version

func Version() int

Version returns the version of the contract.

processing contract

Package processing contains Processing contract which is deployed to main chain.

Processing contract pays for all multisignature transaction executions when notary service is enabled in the main chain. Notary service prepares multisigned transactions, however their senders should have GAS to succeed. It is inconvenient to ask Alphabet nodes to pay for these transactions: nodes can change over time, some nodes will spend GAS faster. It leads to economic instability.

Processing contract exists to solve this issue. At the Withdraw invocation of NeoFS contract, a user pays fee directly to this contract. This fee is used to pay for Cheque invocation of NeoFS contract that returns main chain GAS back to the user. The address of the Processing contract is used as the first signer in the multisignature transaction. Therefore, NeoVM executes Verify method of the contract and if invocation is verified, Processing contract pays for the execution.

Contract notifications Processing contract does not produce notifications to process.

Contract methods

OnNEP17Payment

func OnNEP17Payment(from interop.Hash160, amount int, data any)

OnNEP17Payment is a callback for NEP-17 compatible native GAS contract.

Update

func Update(nefFile, manifest []byte, data any)

Update method updates contract source code and manifest. It can be invoked only by the FS chain committee.

Verify

func Verify() bool

Verify method returns true if transaction contains valid multisignature of Alphabet nodes of the Inner Ring.

Version

```
func Version() int
```

Version returns the version of the contract.

proxy contract

Package proxy implements Proxy contract which is deployed to FS chain.

Proxy contract pays for all multisignature transaction executions when notary service is enabled in FS chain. Notary service prepares multisigned transactions, however their senders should have FS chain GAS to succeed. It is inconvenient to ask Alphabet nodes to pay for these transactions: nodes can change over time, some nodes will spend FS chain GAS faster. It leads to economic instability.

Proxy contract exists to solve this issue. While Alphabet contracts hold all FS chain NEO, proxy contract holds most of FS chain GAS. Alphabet contracts emit half of the available GAS to the proxy contract. The address of the Proxy contract is used as the first signer in a multisignature transaction. Therefore, NeoVM executes Verify method of the contract; and if invocation is verified, Proxy contract pays for the execution.

Contract notifications Proxy contract does not produce notifications to process.

Contract methods

OnNEP17Payment

```
func OnNEP17Payment(from interop.Hash160, amount int, data any)
```

OnNEP17Payment is a callback for NEP-17 compatible native GAS contract.

Update

```
func Update(nefFile, manifest []byte, data any)
```

Update method updates contract source code and manifest. It can be invoked only by committee.

Verify

```
func Verify() bool
```

Verify checks whether carrier transaction contains either $(2/3N + 1)$ or $(N/2 + 1)$ valid multi-signature of the NeoFS Alphabet. Container contract's 'SubmitObjectPut' is an exception, Alphabet signature is not required, contract's 'VerifyPlacementSignatures' is called instead.

Version

```
func Version() int
```

Version returns the version of the contract.

reputation contract

Package reputation implements Reputation contract which is deployed to FS chain.

Storage nodes collect reputation data while communicating with other nodes. This data is exchanged and the end result (global trust values) is stored in the contract as opaque data.

Contract notifications Reputation contract does not produce notifications to process.

Contract methods

Get

```
func Get(epoch int, peerID []byte) [][]byte
```

Get method returns a list of all stable marshaled GlobalTrust structures known for the given peer during the specified epoch.

GetByID

```
func GetByID(id []byte) [][]byte
```

GetByID method returns a list of all stable marshaled GlobalTrust with the specified id. Use ListByEpoch method to obtain the id.

ListByEpoch

```
func ListByEpoch(epoch int) [][]byte
```

ListByEpoch returns a list of IDs that may be used to get reputation data with GetByID method.

Put

```
func Put(epoch int, peerID []byte, value []byte)
```

Put method saves global trust data in contract storage. It can be invoked only by storage nodes with Alphabet assistance (multisignature witness).

Epoch is the epoch number when GlobalTrust structure was generated. PeerID contains public key of the storage node that is the subject of the GlobalTrust. Value contains a stable marshaled structure of GlobalTrust.

Update

```
func Update(nefFile, manifest []byte, data any)
```

Update method updates contract source code and manifest. It can be invoked only by committee.

Version

```
func Version() int
```

Version returns the version of the contract.

Balance transfer details encoding

Alphabet nodes of the Inner Ring use `balance.TransferX` method to manage balances of the NeoFS users at deposit (mint), withdraw (burn), audit settlements, etc. `TransferX` method has `details` argument and `transferX` notification contains `details` field. This field contains bytes that encode the reason of data transfer. First byte of the `details` field defines the transfer type and all other bytes provide extra details.

First Byte	Description	Extra data
0x01	Mint: deposit processed by NeoFS contract.	32-byte hash of mainchain transaction which invoked <code>neofs.Deposit</code> method.
0x02	Burn: cheque processed by NeoFS contract.	32-byte hash of mainchain transaction which invoked <code>neofs.Cheque</code> method.
0x03	Lock: withdraw processed by NeoFS contract.	32-byte hash of mainchain transaction which invoked <code>neofs.Withdraw</code> method.

First Byte	Description	Extra data
0x04	Unlock: withdraw processed by NeoFS contract but cheque didn't process before timeout, so balance returned to the account.	Up to 8 bytes of epoch number when asset lock was removed.
0x10	ContainerFee: put processed by Container contract.	32-byte Container ID
0x40	AuditSettlement: payment to Inner Ring node for processed audit.	8 bytes of epoch (LittleEndian) number when settlement happened.
0x41	BasicIncomeCollection: transfer assets from data owner accounts to the banking account.	8 bytes of epoch (LittleEndian) number when settlement happened.
0x42	BasicIncomeDistribution: transfer assets from banking account to storage node owner accounts.	8 bytes of epoch (LittleEndian) number when settlement happened.

Reputation model

NeoFS reputation system is a subsystem for calculating trust in a node. It is based on a reputation model for assessing trust, which is, in turn, based on the **EigenTrust** algorithm designed for peer-to-peer reputation management.

The EigenTrust algorithm is built on the concept of transitive trust: if peer *i* trusts any peer *j*, it will also trust the peers that *j* trusts.

The Subject of trust assessment is the one who calculates trust.

The Object of trust assessment is the one whose trust is being calculated.

Configuration

To calculate reputation values using the EigenTrust algorithm, nodes need the values of the variable parameters of the algorithm. These parameters are constant but can be changed if one needs to adapt the algorithm. For synchronization and correct operation of all nodes, the values are moved to the `Netmap` contract and read from it at the beginning of each epoch.

Parameters and their keys in the global configuration of the `Netmap` contract are as follows:

Key	Type	Description
<code>EigenTrustIterations</code>	<code>uint64</code>	A number of iterations required to calculate Global Trust.
<code>EigenTrustAlpha</code>	<code>string</code>	A parameter responsible for the level of influence of the current node reputation used when calculating its trust in other nodes.

Managers

Exactly one manager is selected for each node in each epoch. Every manager is also a child node of another manager. Thus, each network participant plays two roles: a child node for one of the managers, and a manager for one of the child nodes.

Managers are required to assess the trust of a child node considering its reputation among other nodes. Also, managers are direct (and only) participants in the iterative part of the EigenTrust algorithm; in the late stage of the algorithm, they send Global Trust in their child node to the Reputation contract.

Defining a manager for a node

To unambiguously determine the manager for the current node, the following rule is used:

1. For the current epoch, a network map (a list of active participants) is obtained from the `Netmap` of the contract.
2. The resulting array is sorted using **HRW** hashing, which takes into account the current epoch number.
3. The index of the current node is determined in a uniquely sorted array - i .
4. The $i+1$ -th node is taken as a manager of the i -th node.

Thus, for the (`Netmap`, `Epoch`) pair (one for all) each network member can uniquely define a manager for any network node. In this case, forecasting a manager for node i at an epoch n , generally speaking, is a nontrivial task because of the variability of the network map. Also, the rule allows to select a manager for a node pseudo-randomly every epoch.

Taken together, all the above makes it difficult for rogue nodes to cluster for profit.

Local Trust

Local Trust trust of one node to another, calculated using *only* statistical information of their peer-to-peer network interactions. The Subject and Object of such a trust are peer-to-peer nodes.

Subject and Object of a trust

Any node is the Subject of evaluation. The Object of trust is any node other than the Subject. In NeoFS, nodes do not evaluate their own performance because the default is “every node trusts itself”.

The evaluation criterion is successful (non error in sense of network interaction) execution of one of the RPC calls:

1. Put
2. Get
3. Search
4. GetRange
5. GetRangeHash
6. Head
7. Delete

A Subject of the Local Trust assessment stores an assessment of the experience of communication with the Objects in `LocalStorage` while interacting with them. Local reputation only makes sense

in context with some epoch value, which should be taken into account when storing and transferring values.

Calculating trust

Using all these formulations, terms and statements, we recognize that Local Trust is only applicable in the context of a certain epoch. It means that peer-to-peer interactions of nodes for an epoch n affect *only* Local Trusts calculated in the epoch n . That fact is not emphasized explicitly further.

Assessment of interactions with a node is a binary value, i.e. equals either 0 (false) or 1 (true).

Let $sat(i, j)$ be a number of positive assessments of interactions with node j by node i .

Let $unsat(i, j)$ be a number of negative assessments of interactions with node j by node i .

Therefore, the total number of interactions between node i and node j :

$$all(i, j) := sat(i, j) + unsat(i, j).$$

Then:

$$S_{ij} = \frac{sat(i, j)}{all(i, j)} \in [0, 1]$$

is the averaged assessment of interactions with node j by node i .

Let us define $C_{i,j}$ as following:

$$C_{i,j} := \begin{cases} \frac{S_{i,j}}{\sum_j S_{i,j}}, \sum_j S_{i,j} \neq 0 \\ \frac{1}{N-1}, otherwise \end{cases}, \in [0, 1],$$

where N is the number of network participants. $C_{i,j}$ according to the chapter 4.5 of the EigenTrust article²⁷ is normalized trust of node i to node j . This value is taken as the final Local reputation of node j for node i .

Transport

At the end of each epoch, all child nodes must announce the accumulated statistics — the set of all Local Trusts for the epoch that has just finished — to their managers.

²⁷<http://ilpubs.stanford.edu:8090/562/1/2002-56.pdf>

The transfer is made by calling manager's `AnnounceLocalTrust` gRPC method²⁸.

Each node should not only collect and transmit its local Trusts, but also be ready to accept (act as a server) and correctly process similar data from another node, i.e. act as a manager.

Global Trust

Global Trust the result of the EigenTrust algorithm is the trust in the network participant, which has been obtained regarding *all* Local Trusts of *all* nodes.

Subject and Object of a trust

The Subject of the assessment is the entire system. Any node in the network is an Object of trust.

The evaluation criterion is the aggregate of Local Trust of *all* network participants to the current node, but adjusted for similar Local Trust to each participant from the rest of the network.

Calculating trust

Each manager retrieves the Local Trust value of its child at the end of each epoch and stores it in the `DaughterStorage`. The further task of each manager is to calculate the Global Trust of the system in its child node, and also to transfer all the information it has to another managers so that they, in turn, can perform a similar task.

Global Trust is calculated iteratively. While calculating `GlobalTrust`, according to the EigenTrust algorithm, managers should exchange “intermediate” trust values, the so-called `IntermediateTrust`. With the help of these values, managers at each iteration approximate (recalculate the results of the previous iteration) the value of the Global Trust to the common Global Value Limit: $T = [t_0, \dots, t_n]$, where t_i is the limiting Global Trust value in node i .

Using all these formulations, terms and statements, we recognize that Local Trust is only applicable in the context of a certain epoch. It means that peer-to-peer interactions of nodes for an epoch n affect *only* Global Trusts calculated in the epoch n . `IntermediateTrust` values, likewise, only make sense when used in the context of some epoch **and iteration number**. That fact is not emphasized explicitly further.

The iterative formula:

²⁸<https://github.com/nspcc-dev/neofs-api/blob/master/reputation/service.proto#L18>

$$t_j^{k+1} = (1 - \alpha) \sum_{i=0}^n C_{ij} t_i^k + \alpha t_j^0$$

where C_{ij} is Local Trust of node i to node j ; α is a value that determines the influence of the t_i^0 ("blind" trust in node i) - on the final result; $C_{i,j} t_i^k$, in the introduced terminology, is `IntermediateTrust` which has been transferred by the manager of node i to the manager of node j . In order for the algorithm to work synchronously, the number of iterations and α is read by all nodes from the Global Configuration of the `Netmap` contract at the beginning of each epoch.

Actually, blind trust t_i^0 can be different for different nodes. It can take into account how long ago a node joined the network, whether the node belongs to the developers of the network, etc. In the current implementation $t_i^0 = \frac{1}{N}$, $\forall i$, where N is the number of network participants in the current epoch.

Transport

In order for the manager of node j to be able to calculate t_j^{k+1} , which is the intermediate Global Trust (this is *not* `IntermediateTrust`) to its child node, the following is essential:

1. The manager must have all the existing `IntermediateTrusts` - $C_{ij} t_i^k, \forall i \neq j$.
2. The manager must calculate and transfer all `IntermediateTrusts`, according to the information received from its child node, that is, calculate and transfer all existing $C_{ji} t_j^k, \forall i \neq j$.

To do this, the manager must act both as a client and a server. The transfer itself is run by calling the `AnnounceIntermediateResult` gRPC method²⁹.

The manager should store intermediate `IntermediateTrust` values in `ConsumerStorage` considering both the epoch and the iteration numbers. Each manager must transmit them based on the number of iterations and the number of blocks in the current epoch.

²⁹<https://github.com/nspcc-dev/neofs-api/blob/master/reputation/service.proto#L22>

Incentive model

Economic of NeoFS uses GAS stored in NeoFS contract in the main chain. It increases with user deposits and decreases after withdrawals. Therefore, assets don't appear out of nowhere and don't go nowhere. All internal accounts and settlements are managed by Balance contract in FS chain. Every epoch many settlements take place, which we can divide into two groups:

- data storage payments,
- service fees.

Data storage payments

Data storage payments are made once an epoch. Epochs are measured in seconds based on FS chain blocks' timestamps, and can change their duration depending on network settings. Thus, estimated profit and expense for an hour, week or month may vary depending on epoch duration. Payments are done based on load reports that storage nodes do every epoch (see [Storage metrics](#)). Storage node must be part of container to report its storage in it and receive incoming payments. Storing data without reporting about it leads to no balance increase.

Data storage payments flow:

- accumulating storage reports in the Balance contract and calculate normalized averages for epochs based on the number of reports and their announcement time in an epoch,
- initiating transfers to storage nodes by the Alphabet node consensus for every actual container for an epoch,
- atomic payment distribution on the Balance contract side without any external interference, fully based on the actual FS chain state.

Basic income

Basic income provides basic asset flow from data owners to storage node owners. Basic income settlements are calculated per container. Exact payment price is calculated based on the normalized container size retrieved from the storage node's reports and the basic income rate in NeoFS network configuration. Note that all the data is paid, not only the uploaded one: objects may be duplicated based on container policies (and, in fact, are expected to have additional replicas in most of the cases), container size in such cases is calculated as the sum of every stored object on storage nodes, including duplicated data.

Basic income rate is a NeoFS network configuration value managed by Alphabet nodes of the Inner Ring. It is stored as GAS per GiB value. Once an epoch, Storage Nodes send reports to the Container

contract about containers they serve, single report per every container. This data is normalized inside Container contract taking into account time of the previous report. Storage nodes receive payments independently based on the reports they have sent previously.

Normalized billing size for j node in i container:

$$R_{i,j} = \sum_{k=0}^N r_k \frac{(n_{k+1} - n_k)}{n_N - n_0}$$

where N is number of reports in this epoch - 1, r_i is reported size, and n_i is timestamp of the report (n_N is epoch end timestamp). Time is defined based on the block's timestamp at which report was announced, this is based on FS chain's execution state. If there are no changes in node's storage related to a container, it may not update its report in the contract. In such cases, the latest reported value is considered as storage size for all the following epochs and no normalization is required.

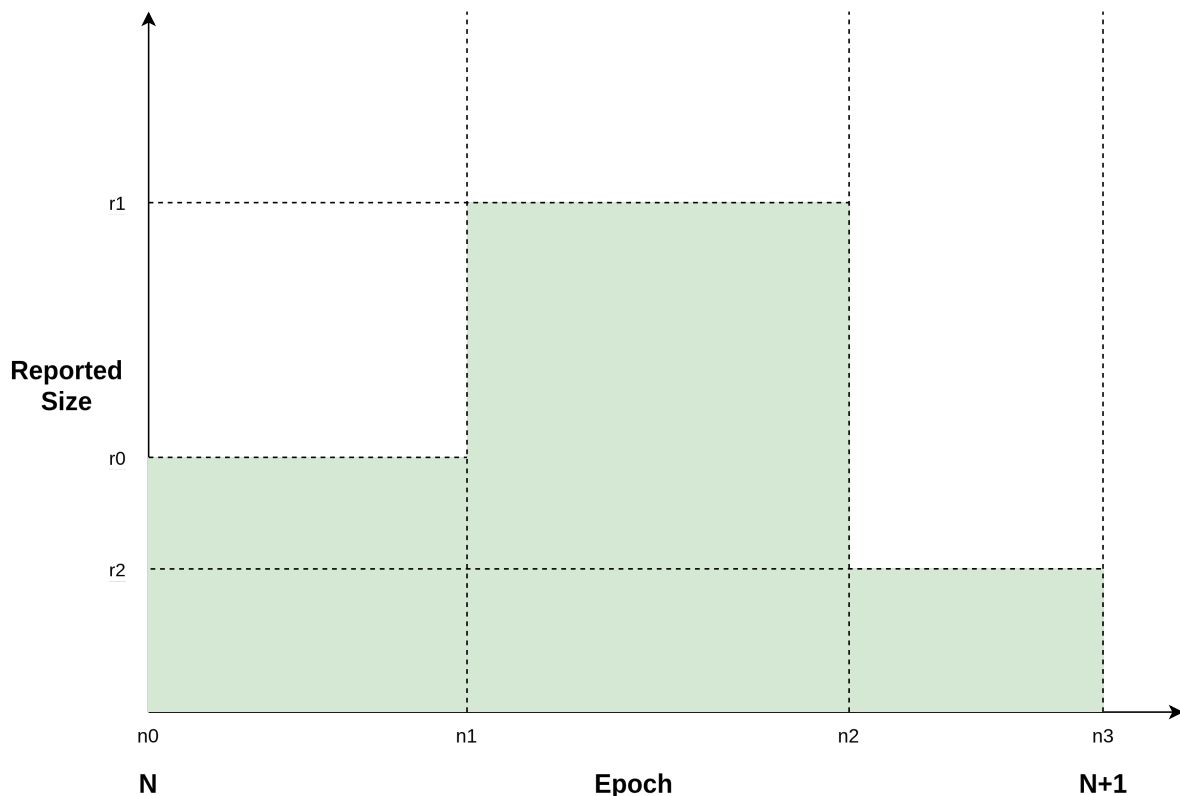


Figure 21: Report size normalization

When epoch N starts, Inner Ring nodes initiate payments for N-1 epoch, sending consensus payment transaction to FS chain. Balance contract distribute payments based on billing information stored in Container contract which in turn is based on storage node reports.

Payment that every data owner must pay every epoch:

$$\sum_i \sum_j \frac{R_{i,j} \cdot Rate}{2^{30}}$$

GAS, where $R_{i,j}$ is normalized container size report from the j -th node in i -th container the owner owns, $Rate$ is the basic income rate. Every storage node receives GAS equals the sum above with corresponded fixed j and only i s that correspond containers storage node belongs to.

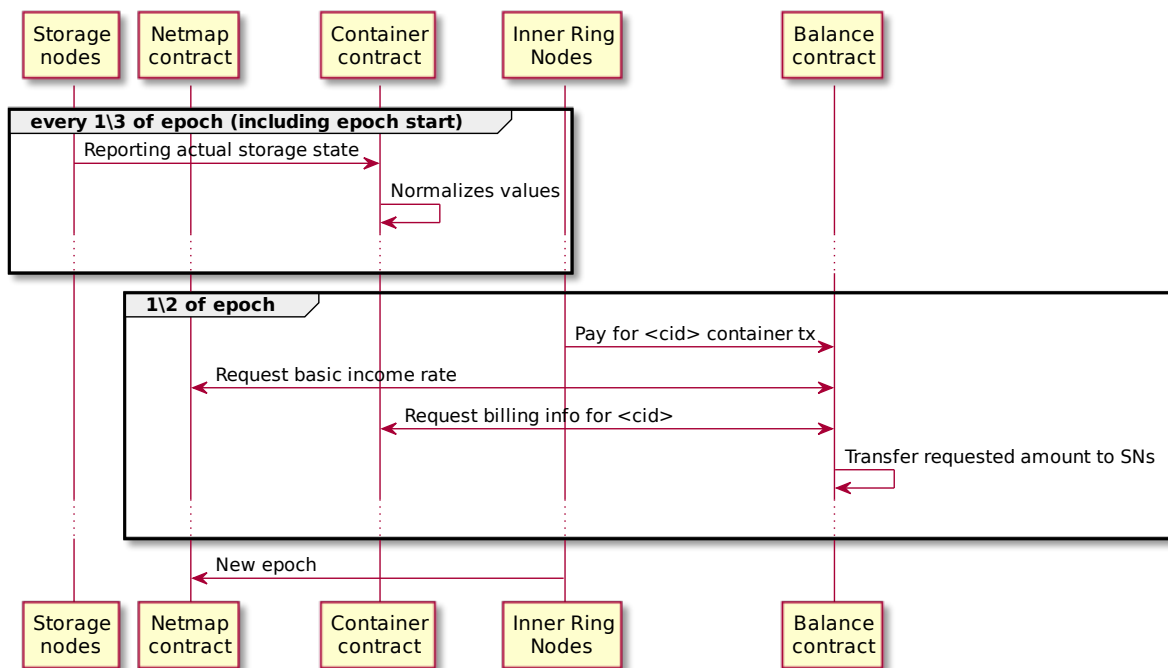


Figure 22: Basic income collection

d## Service fees

Container creation fee

To create a container, a data owner should pay fee. It is calculated as $7 \cdot fee$, where fee is a value from NeoFS network configuration (Container Fee). Each Alphabet node gets fee GAS during this operation.

Inner Ring candidate fee

To become a part of the Inner Ring list, an Inner Ring candidate must register its key in NeoFS contract. This operation transfers to NeoFS contract fee value that is set in NeoFS network configuration (InnerRingCandidateFee).

Withdraw fee

To withdraw assets, Alphabet nodes need mainchain GAS to invoke Cheque method of NeoFS contract that transfers assets back to the user. This GAS is paid by the user at Withdraw invocation. In notary enabled environment, the user pays fee value set in NeoFS network configuration (WithdrawFee). In notary disabled environment, the user pays $7 \cdot fee$.

NeoFS API v2

NeoFS API v2 is focusing on simplification of previous versions of API used during early stages of NeoFS development. The new structure makes it easier to implement NeoFS API in different languages and for different platforms.

All data structures are defined in protobuf format and grouped together with corresponding services into comparatively independent packages. This allows to significantly simplify development (automatically generate for most parts) of a library for working in NeoFS. One can start with a required minimum instead of implementing the whole package right in the beginning. We tried to make the packages as independent as possible, and thus minimizing horizontal dependency.

For transport layer, by default we assume that gRPC will be used. These are popular, simple, and time-tested tools relevant for most languages and platforms. Although gRPC is used now, we have everything to transfer structures through other protocols, e.g. JSON-RPC.

The API described below corresponds to the version 2.22.0³⁰.

Nodes and their identification

NeoFS API does not differentiate client and server. All members of the network communicate using the same protocol. The protocol does not distinguish between a small **Command Line Interface (CLI)** utility and a full-featured storage node. Both are nodes of the same p2p network.

A **NeoFS Node** is identified by a pair of keys for encryption and decryption.

Public key is encoded in compressed form according to ANSI x9.62³¹ (section 4.3.6).

Elliptic curve	secp256r1
Private key size	32 bytes
Public key size	33 bytes
Example keys:	
private key	6af2b8b41ad2e78f19aa0bc4fb5cb746d61ad44ebf9ba2a43b6e5cc3e46715a6
public key	03065e513fdaccc4556e7de010bf3d5445552357fb17928f3bd8cea33e092a64eb
Neo 3.0 address	Na6DELLB6dtnPmsD7y1HFjVZNZp8S5BdCJ

³⁰<https://github.com/nspcc-dev/neofs-api/releases/tag/v2.22.0>

³¹<https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.202.2977&rep=rep1&type=pdf>

The key format is compatible with the Neo 3.0 Wallets keys. It lets smart contracts verify the sender by public keys. Thus, each member may have an internal NEP-5 balance, where NEO Wallet address is formed from a public key, as it happens in Neo 3.0.

NeoFS is a peer-to-peer network, which means that the clients are equal and each of them needs a pair of asymmetric keys. To create a container and place an object, one should have **GAS Utility Token** on FS chain internal balance. When a user makes a deposit from his NEO Wallet, operations of container creation and object placement should be carried out with the same key (if no other keys have been associated with that OwnerID).

Requests and Responses

All request and responses in NeoFS API v2 have the same structure. Only their `.body` fields are different. Any request consists of body, meta headers, and verification headers. See the corresponding paragraph in the relevant package description section to learn about the structure of particular parts of a message in detail.

The `.body` field delivers the structure with the data making up the request or a response to it.

The `.meta_header` contains metadata to the request.

The `.verify_header` carries cryptographic signatures for `.body` and `.meta_header`. It allows to check if the message is authentic, see if it has been correctly transferred between two nodes, and provide an assumed route of the message where each intermediate node has left its signature.

Signing RPC messages and data structures

The messages exchanged between the users of the network involves ECDSA signatures. These signatures is defined in `refs.Signature` type structure. The `.sign` field keeps byte representation of a signature, while `.key` field contains public key to verify the signature.

Stable serialization

To sign messages or structures, one should first turn them into a byte series. NeoFS protocol is described in protocol buffers v3 format. Protocol Buffers v3 defines the format of serialization for all messages, but specification³² allows serialization process to be unstable: the same message can be encoded correctly by various methods. The field order in the encoded message may be changed randomly.

³²<https://developers.google.com/protocol-buffers/docs/encoding#order>

When a message is serialized, there is no guaranteed order for how its known or unknown fields should be written. Serialization order is an implementation detail and the details of any particular implementation may change in the future. Therefore, protocol buffer parsers must be able to parse fields in any order.

To generate unified signatures for messages, all NeoFS nodes **must stably serialize structures and messages described in the protocol.**

Serialization is considered stable when it does not change the order of the fields in an encoded message. All fields are encoded in ascending order regarding their numbers specified in protocol description.

```
message Foo {
  bytes one = 1; // C0 FF EE
  bytes two = 2; // BE EF
}

StableSerialize(Foo) =
[0A 03 C0 FF EE][12 02 BE EF] // OK for GRPC
                               // OK for NeoFS Signature

UnstableSerialize(Foo) =
[12 02 BE EF][0A 03 C0 FF EE] // OK for GRPC
                               // FAIL for NeoFS Signature
```

Most auto-generated serializers behave occasionally stable, but there is no guarantee that they will remain the same in future. Clients exploiting auto-generated serializers should be aware of such risk.

Signature generation format

The signature of a message or a structure should be stably serialized. The serialized byte array is hashed with SHA-512. The obtained signature (R,S) is enciphered uncompressed according to ANSI x9.62 (section 4.3.6)

Elliptic curve	secp256r1
Hash function	SHA-512
Signature size	65 bytes

```
message Foo { bytes one = 1; // C0 FF EE bytes two = 2; // BE EF }
```

```
private key = 6af2b8b41ad2e78f19aa0bc4fb5cb746d61ad44ebf9ba2a43b6e5cc3e46715a6
StableSerialize(Foo) = 0a03coffee1202beaf
SHA-512(StableSerialize(Foo)) =
5efec2432616ca322824a7140d5ac332c6a3a388d2746f8cff6e48909d36829f9cb8586718d457c9
```

```
R,S = Sign(private key, SHA-512(StableSerialize(Foo)), secp256r1)
R = e13f3e71db728b85acc4cea688d3dae6b01453d2bff1b5ebc2695cedfef7fdd5
S = 2ecbc0cc0ae4f70696682b4e358a4b698d74f9b708c13470e5c808fe04f526e5
```

```
Signature =
04e13f3e71db728b85acc4cea688d3dae6b01453d2bff1b5ebc2695cedfef7fdd52ecbc0cc0ae4f7
```

Signature function uses random number generator, which grants different signatures to the same message with each signing cycle.

Signature chaining in requests and responses

The structure and the order of signatures for Request and Response messages are the same. Therefore, everything about Request message written below is true for Response on corresponding structures.

All signatures are kept in RequestVerificationHeader structure which is filled before a message is sent. The **NeoFS Nodes** involved in the transmission and resigning of the messages fill in their own structure copies and put them in origin field recursively.

If a user sends a request directly to a receiver, the chain consists of one RequestVerificationHeader only.

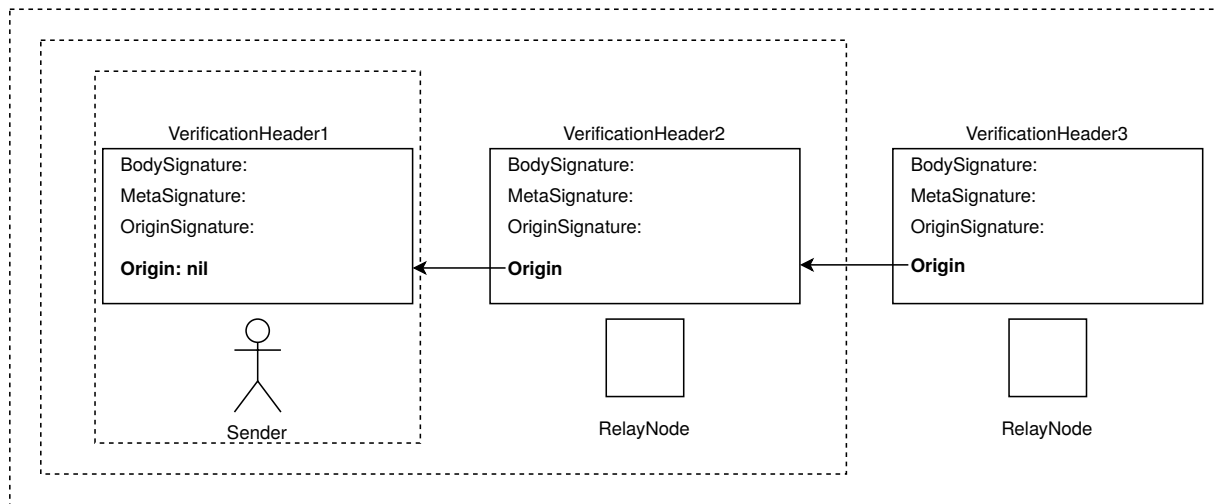


Figure 23: Signature chain verification

The RequestVerificationHeader always carries three signatures:

- Message body signature
- Meta header signature
- Verification header signature

Message body signature

Each RPC message has a field called Body. This structure is serialized stably; it is signed by the sender only. If a message is retransmitted, this signature is never set in the new copies of VerificationHeader.

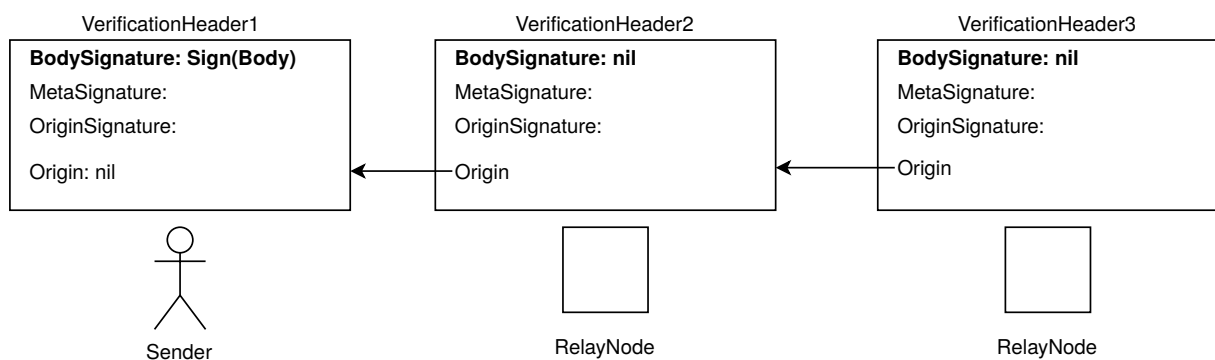


Figure 24: Body signature verification

Meta header signature

Each RPC message has a field called meta_header. Meta headers are changed for every retransmission (e.g. TTL is reduced) and form an equivalent chain. meta_signature field contains the signature for an already organized structure of the meta header.

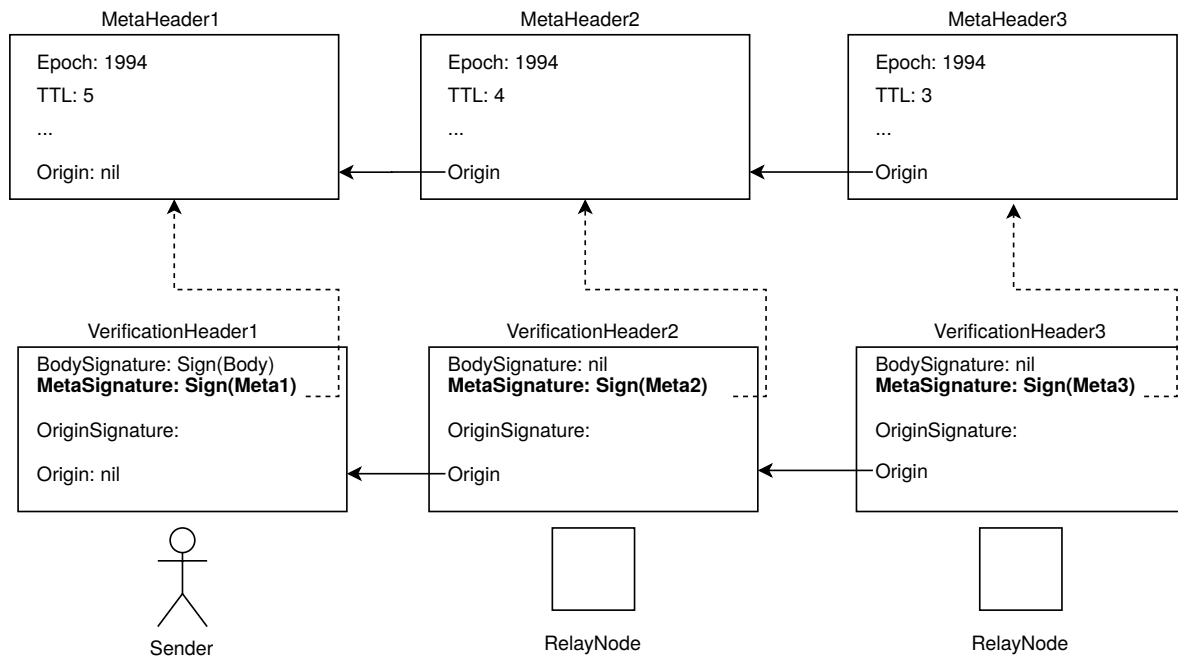


Figure 25: Meta header signature verification

Verification header signature

While putting a new verification header, intermediate nodes should sign the preceding verification header and put the signature in the origin_signature field. The requestor does not set this signature.

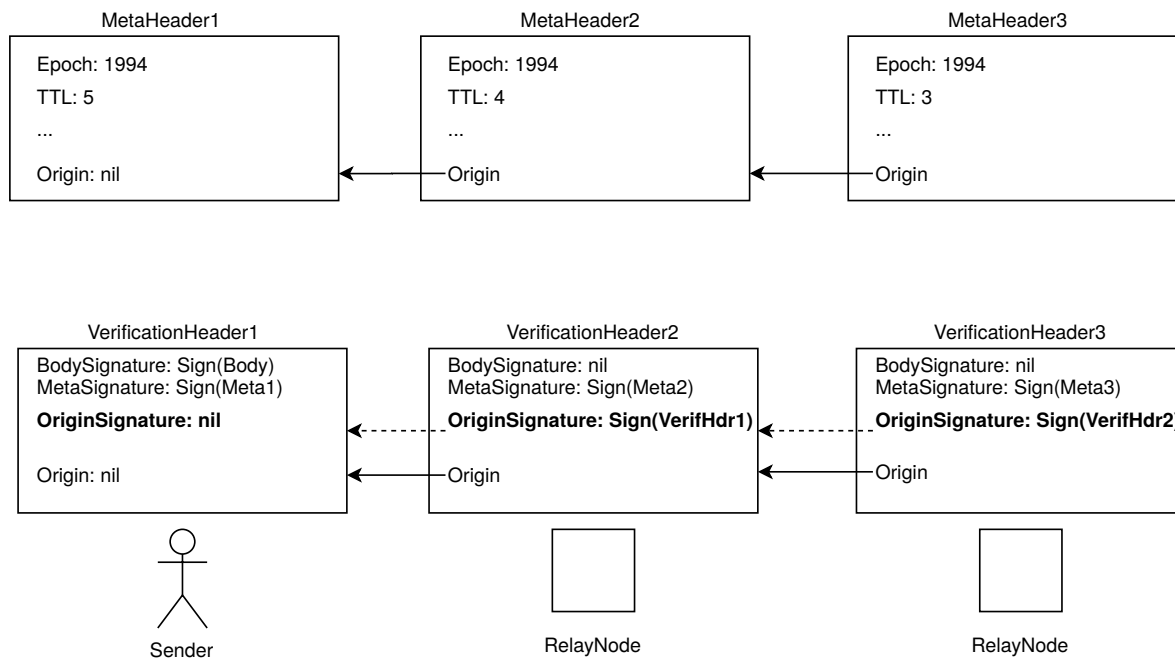


Figure 26: Verification header signature verification

Container service signatures

In addition to the RPC requests themselves there is a need to sign following structures:

- Container in `container.PutRequest.Body` message
- Container ID in `container.DeleteRequest.Body` message
- Extended ACL table structure in `container.SetExtendedACLRequest.Body` message

Those structures' signature is verified by smart contracts, hence it must be compatible with **Neo Virtual Machine**. The signature format supported by **Neo Virtual Machine** is different from the format described in previous sections.

A stably serialized message is hashed using SHA-256 algorithm. The resulting signature (R,S) is encoded uncompressed as a concatenation of the 32-byte sized coordinates of R and S.

Elliptic curve	secp256r1
Hash function	SHA-256
Signature size	64 bytes

```
message ContainerID {
```

```
bytes value = 1; // 29fe85bb8c36f5cb676e256113193235a2ba0c0abe6a71f84654afa92801c
}
```

```
private key = 6af2b8b41ad2e78f19aa0bc4fb5cb746d61ad44ebf9ba2a43b6e5cc3e46715a6
StableSerialize(Foo) = 0a2029fe85bb8c36f5cb676e256113193235a2ba0c0abe6a71f84654af
SHA-256(StableSerialize(Foo)) = a086acbc03862c01bdff3f850b8254f1be9a6d56ec5661c6e
```

```
R,S = Sign(private key, SHA-256(StableSerialize(Foo)), secp256r1)
R = 1233d0e5c87a24c5a56c518596da64b1ceb8d667723b0030c4888b524229ff8a
S = d4e42952d516c2959ba1825e2768cbfe3f4336e7a14c635236ae2ea95fa50435
```

Signature =

```
1233d0e5c87a24c5a56c518596da64b1ceb8d667723b0030c4888b524229ff8ad4e42952d516c295
```

Object service and Session signatures

Object service and the rest of the services and structures in NeoFS API v2 use the same signature format as in RPC messages signing.

Elliptic curve	secp256r1
Hash function	SHA-512
Signature size	65 bytes

neo.fs.v2.accounting

Service “AccountingService”

Accounting service provides methods for interaction with FS chain via other NeoFS nodes to get information about the account balance. Deposit and Withdraw operations can't be implemented here, as they require Mainnet NeoFS smart contract invocation. Transfer operations between internal NeoFS accounts are possible if both use the same token type.

Method Balance

Returns the amount of funds in GAS token for the requested NeoFS account.

Statuses: - **OK** (0, SECTION_SUCCESS): balance has been successfully read; - Common failures (SECTION_FAILURE_COMMON).

Request Body: BalanceRequest.Body

To indicate the account for which the balance is requested, its identifier is used. It can be any existing account in FS chain Balance smart contract. If omitted, client implementation **MUST** set it to the request's signer Owner ID.

Field	Type	Description
owner_id	OwnerID	Valid user identifier in OwnerID format for which the balance is requested. Required field.

Response Body BalanceResponse.Body

The amount of funds in GAS token for the Owner ID's account requested. Balance is given in the Decimal format to avoid precision issues with rounding.

Field	Type	Description
balance	Decimal	Amount of funds in GAS token for the requested account.

Message Decimal

Standard floating point data type can't be used in NeoFS due to inexactness of the result when doing lots of small number operations. To solve the lost precision issue, special `Decimal` format is used for monetary computations.

Please see [The General Decimal Arithmetic Specification](#)³³ for detailed problem description.

Field	Type	Description
value	int64	Number in the smallest Token fractions.
precision	uint32	Precision value indicating how many smallest fractions can be in one integer.

neo.fs.v2.acl

Message BearerToken

`BearerToken` allows to attach signed Extended ACL rules to the request in `RequestMetaHeader`. If container's Basic ACL rules allow, the attached rule set will be checked instead of one attached to the container itself. Just like `JWT`³⁴, it has a limited lifetime and scope, hence can be used in the similar use cases, like providing authorisation to externally authenticated party.

`BearerToken` can be issued only by the container's owner and must be signed using the key associated with the container's Owner ID.

Field	Type	Description
body	Body	Bearer Token body
signature	Signature	Signature of BearerToken body

Message BearerToken.Body

Bearer Token body structure contains Extended ACL table issued by the container owner with additional information preventing token abuse.

³³<http://speleotrove.com/decimal/>

³⁴<https://jwt.io>

Field	Type	Description
eacl_table	EACLTable	Table of Extended ACL rules to use instead of the ones attached to the container. If it contains container_id field, bearer token is only valid for this specific container. Otherwise, any container of the same owner is allowed.
owner_id	OwnerID	OwnerID defines to whom the token was issued. It must match the request originator's OwnerID. If empty, any token bearer will be accepted.
lifetime	TokenLifetime	Token expiration and valid time period parameters
issuer	OwnerID	Token issuer's user ID in NeoFS. It must equal to the related container's owner.

Message BearerToken.Body.TokenLifetime

Lifetime parameters of the token. Field names taken from rfc7519³⁵.

Field	Type	Description
exp	uint64	Expiration epoch, the last epoch when token is valid.
nbf	uint64	Not valid before epoch, the first epoch when token is valid.
iat	uint64	Issued at Epoch

Message EACLRecord

Describes a single eACL rule.

Field	Type	Description
operation	Operation	NeoFS request Verb to match
action	Action	Rule execution result. Either allows or denies access if filters match.

³⁵<https://tools.ietf.org/html/rfc7519>

Field	Type	Description
filters	Filter	List of filters to match and see if rule is applicable
targets	Target	List of target subjects to apply ACL rule to

Message **EACLRecord.Filter**

Filter to check particular properties of the request or the object.

The `value` field must be empty if `match_type` is an unary operator (e.g. `NOT_PRESENT`). If `match_type` field is numeric (e.g. `NUM_GT`), the `value` field must be a base-10 integer.

By default key field refers to the corresponding object's `Attribute`. Some Object's header fields can also be accessed by adding `$Object:` prefix to the name. For such attributes, field 'match_type' must not be 'NOT_PRESENT'. Here is the list of fields available via this prefix:

- `$Object:version`
version
- `$Object:objectID`
object_id
- `$Object:containerID`
container_id
- `$Object:ownerID`
owner_id
- `$Object:creationEpoch`
creation_epoch
- `$Object:payloadLength`
payload_length
- `$Object:payloadHash`
payload_hash
- `$Object:objectType`
object_type
- `$Object:homomorphicHash`
homomorphic_hash

Numeric `match_type` field can only be used with `$Object:creationEpoch` and `$Object:payloadLength` system attributes.

Please note, that if request or response does not have object's headers of full object (Range, Range-Hash, Search, Delete), it will not be possible to filter by object header fields or user attributes. From

the well-known list only `$Object:objectID` and `$Object:containerID` will be available, as it's possible to take that information from the requested address.

Field	Type	Description
<code>header_type</code>	HeaderType	Define if Object or Request header will be used
<code>match_type</code>	MatchType	Match operation type
<code>key</code>	string	Name of the Header to use
<code>value</code>	string	Expected Header Value or pattern to match

Message `EACLRecord.Target`

Target to apply ACL rule. Can be a subject's role class or a list of public keys to match.

Field	Type	Description
<code>role</code>	Role	Target subject's role class
<code>keys</code>	bytes	List of 25-byte accounts to identify target subjects. 33-byte public keys are also supported, however, they are deprecated and script hashes should be derived from them.

Message `EACLTable`

Extended ACL rules table. A list of ACL rules defined additionally to Basic ACL. Extended ACL rules can be attached to a container and can be updated or may be defined in `BearerToken` structure. Please see the corresponding NeoFS Technical Specification section for detailed description.

Field	Type	Description
<code>version</code>	Version	eACL format version. Effectively, the version of API library used to create eACL Table.
<code>container_id</code>	ContainerID	Identifier of the container that should use given access control rules
<code>records</code>	EACLRecord	List of Extended ACL rules

Emun Action

Rule execution result action. Either allows or denies access if the rule's filters match.

Number	Name	Description
0	ACTION_UNSPECIFIED	Unspecified action, default value
1	ALLOW	Allow action
2	DENY	Deny action

Emun HeaderType

Enumeration of possible sources of Headers to apply filters.

Number	Name	Description
0	HEADER_UNSPECI	Unspecified header, default value.
1	REQUEST	Filter request headers
2	OBJECT	Filter object headers
3	SERVICE	Filter service headers. These are not processed by NeoFS nodes and exist for service use only.

Emun MatchType

MatchType is an enumeration of match types.

Number	Name	Description
0	MATCH_TYPE_UNSPECIFIED	Unspecified match type, default value.
1	STRING_EQUAL	Return true if strings are equal
2	STRING_NOT_EQUAL	Return true if strings are different
3	NOT_PRESENT	Absence of attribute
4	NUM_GT	Numeric 'greater than'
5	NUM_GE	Numeric 'greater or equal than'

Number	Name	Description
6	NUM_LT	Numeric 'less than'
7	NUM_LE	Numeric 'less or equal than'

Emun Operation

Request's operation type to match if the rule is applicable to a particular request.

Number	Name	Description
0	OPERATION_UNSPECIFIED	Unspecified operation, default value
1	GET	Get
2	HEAD	Head
3	PUT	Put
4	DELETE	Delete
5	SEARCH	Search
6	GETRANGE	GetRange
7	GETRANGEHASH	GetRangeHash

Emun Role

Target role of the access control rule in access control list.

Number	Name	Description
0	ROLE_UNSPECIFIED	Unspecified role, default value
1	USER	User target rule is applied if sender is the owner of the container
2	SYSTEM	System target rule is applied if sender is a storage node within the container or an inner ring node

Number	Name	Description
3	OTHERS	Others target rule is applied if sender is neither a user nor a system target

neo.fs.v2.audit

Message DataAuditResult

DataAuditResult keeps record of conducted Data Audits. The detailed report is generated separately.

DEPRECATED: This audit mechanism is no longer supported (2.18+).

Field	Type	Description
version	Version	Data Audit Result format version. Effectively, the version of API library used to report DataAuditResult structure.
audit_epoch	fixed64	Epoch number when the Data Audit was conducted
container_id	ContainerID	Container under audit
public_key	bytes	Public key of the auditing InnerRing node in a binary format
complete	bool	Shows if Data Audit process was complete in time or if it was cancelled
requests	uint32	Number of request done at PoR stage
retries	uint32	Number of retries done at PoR stage
pass_sg	ObjectID	List of Storage Groups that passed audit PoR stage
fail_sg	ObjectID	List of Storage Groups that failed audit PoR stage
hit	uint32	Number of sampled objects under the audit placed in an optimal way according to the containers placement policy when checking PoP

Field	Type	Description
miss	uint32	Number of sampled objects under the audit placed in suboptimal way according to the containers placement policy, but still at a satisfactory level when checking PoP
fail	uint32	Number of sampled objects under the audit stored inconsistently with the placement policy or not found at all when checking PoP
pass_nodes	bytes	List of storage node public keys that passed at least one PDP
fail_nodes	bytes	List of storage node public keys that failed at least one PDP

neo.fs.v2.container

Service “ContainerService”

ContainerService provides API to interact with Container smart contract in FS chain via other NeoFS nodes. All of those actions can be done equivalently by directly issuing transactions and RPC calls to FS chain nodes.

Method Put

Sends transaction calling contract method to create container, and waits for the transaction to be executed. Deadline is determined by the transport protocol (e.g. `grpc-timeout` header). If the deadline is not set, server waits 15s after submitting the transaction.

Statuses: - **OK** (0, SECTION_SUCCESS):

container successfully created; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_AWAIT_TIMEOUT** (3075, SECTION_CONTAINER):

transaction was sent but not executed within the deadline.

Request Body: PutRequest.Body

Container creation request has container structure’s signature as a separate field. It’s not stored in FS chain, just verified on container creation by Container smart contract. Container ID is a SHA256

hash of the stable-marshalled container structure, hence there is no need for additional signature checks.

Field	Type	Description
container	Container	Container structure to register in NeoFS
signature	SignatureRFC6979	Signature of a stable-marshalled container according to RFC-6979.

Response Body PutResponse.Body

Container put response body contains information about the newly registered container as seen by Container smart contract. Container ID can be calculated beforehand from the container structure and compared to the one returned here to make sure everything has been done as expected.

Field	Type	Description
container_id	ContainerID	Unique identifier of the newly created container

Method Delete

Sends transaction calling contract method to delete container, and waits for the transaction to be executed. Deadline is determined by the transport protocol (e.g. `grpc-timeout` header). If the deadline is not set, server waits 15s after submitting the transaction. NOTE: a container deletion leads to the removal of every object in that container, regardless of any restrictions on the object removal (e.g. lock/locked object would be also removed).

Statuses: - **OK** (0, SECTION_SUCCESS):

container successfully removed; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_LOCKED** (3074, SECTION_CONTAINER):

deleting a locked container is prohibited; - **CONTAINER_AWAIT_TIMEOUT** (3075, SECTION_CONTAINER): transaction was sent but not executed within the deadline.

Request Body: DeleteRequest.Body

Container removal request body has signed Container ID as a proof of the container owner's intent. The signature will be verified by Container smart contract, so signing algorithm must be supported by NeoVM.

Field	Type	Description
container_id	ContainerID	Identifier of the container to delete from NeoFS
signature	SignatureRFC6979	ContainerID signed with the container owner's key according to RFC-6979.

Response Body DeleteResponse.Body

DeleteResponse has an empty body because delete operation is asynchronous and done via consensus in Inner Ring nodes.

Method Get

Returns container structure from Container smart contract storage.

Statuses: - **OK** (0, SECTION_SUCCESS):

container has been successfully read; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

requested container not found.

Request Body: GetRequest.Body

Get container structure request body.

Field	Type	Description
container_id	ContainerID	Identifier of the container to get

Response Body GetResponse.Body

Get container response body does not have container structure signature. It has been already verified upon container creation.

Field	Type	Description
container	Container	Requested container structure
signature	SignatureRFC6979	Signature of a stable-marshalled container according to RFC-6979.

Field	Type	Description
session_token	SessionToken	Session token if the container has been created within the session

Method List

Returns all owner's containers from 'Container' smart contract' storage.

Statuses: - **OK** (0, SECTION_SUCCESS):

container list has been successfully read; - Common failures (SECTION_FAILURE_COMMON).

Request Body: ListRequest.Body

List containers request body.

Field	Type	Description
owner_id	OwnerID	Identifier of the container owner

Response Body ListResponse.Body

List containers response body.

Field	Type	Description
container_ids	ContainerID	List of ContainerIDs belonging to the requested OwnerID

Method SetExtendedACL

Sends transaction calling contract method to set container's extended ACL, and waits for the transaction to be executed. Deadline is determined by the transport protocol (e.g. `grpc-timeout` header). If the deadline is not set, server waits 15s after submitting the transaction.

Statuses: - **OK** (0, SECTION_SUCCESS):

container eACL successfully set; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_AWAIT_TIMEOUT** (3075, SECTION_CONTAINER):

transaction was sent but not executed within the deadline.

Request Body: SetExtendedACLRequest.Body

Set Extended ACL request body does not have separate ContainerID reference. It will be taken from EACLTable.container_id field.

Field	Type	Description
eacl	EACLTable	Extended ACL table to set for the container
signature	SignatureRFC6979	Signature of stable-marshalled Extended ACL table according to RFC-6979.

Response Body SetExtendedACLResponse.Body

SetExtendedACLResponse has an empty body because the operation is asynchronous and the update should be reflected in Container smart contract's storage after next block is issued in FS chain.

Method GetExtendedACL

Returns Extended ACL table and signature from Container smart contract storage.

Statuses: - **OK** (0, SECTION_SUCCESS):

container eACL has been successfully read; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

container not found; - **EACL_NOT_FOUND** (3073, SECTION_CONTAINER):
eACL table not found.

Request Body: GetExtendedACLRequest.Body

Get Extended ACL request body

Field	Type	Description
container_id	ContainerID	Identifier of the container having Extended ACL

Response Body GetExtendedACLResponse.Body

Get Extended ACL Response body can be empty if the requested container does not have Extended ACL Table attached or Extended ACL has not been allowed at the time of container creation.

Field	Type	Description
eacl	EACLTable	Extended ACL requested, if available
signature	SignatureRFC6979	Signature of stable-marshalled Extended ACL according to RFC-6979.
session_token	SessionToken	Session token if Extended ACL was set within a session

Method `AnnounceUsedSpace`

Announces the space values used by the container for P2P synchronization.

Statuses: - **OK** (0, SECTION_SUCCESS):

estimation of used space has been successfully announced; - Common failures (SECTION_FAILURE_COMMON).

DEPRECATED: every storage node must send storage load directly to container contract.

Request Body: `AnnounceUsedSpaceRequest.Body`

Container used space announcement body.

Field	Type	Description
announcements	Announcement	List of announcements. If nodes share several containers, announcements are transferred in a batch.

Response Body `AnnounceUsedSpaceResponse.Body`

`AnnounceUsedSpaceResponse` has an empty body because announcements are one way communication.

Method `SetAttribute`

Sends transaction calling contract method to set container attribute, and waits for the transaction to be executed. Deadline is determined by the transport protocol (e.g. `grpc-timeout` header). If the deadline is not set, server waits 15s after submitting the transaction.

Statuses: - **OK** (0, SECTION_SUCCESS):

attribute successfully set; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_AWAIT_TIMEOUT**

(3075, SECTION_CONTAINER):

transaction was sent but not executed within the deadline.

Request Body: SetAttributeRequest.Body

Request payload message.

Field	Type	Description
parameters	Parameters	Op parameters.
signature	SignatureRFC6979	N3 witness of stable-marshalled parameters field. The signature must authenticate either container owner or one of subjects in the session_token field if any. Signature according to ECDSA_RFC6979_SHA256 scheme is also supported.
session_token	SessionTokenV2	Optional session token. The token must be issued by the container owner. The token must have at least one subject authenticated by signature field. The token must have at least one context with this container and CONTAINER_SETATTRIBUTE verb.
session_token_v1	SessionToken	Optional session token (V1). It must not be set together with session_token field that is highly recommended to be used instead. Requirements are the same for both.

Method RemoveAttribute

Sends transaction calling contract method to remove container attribute, and waits for the transaction to be executed. Deadline is determined by the transport protocol (e.g. grpc-timeout header). If the deadline is not set, server waits 15s after submitting the transaction.

Statuses: - **OK** (0, SECTION_SUCCESS):

attribute successfully removed; - Common failures (SECTION_FAILURE_COMMON); - **CONTAINER_AWAIT_TIMEOUT** (3075, SECTION_CONTAINER):

transaction was sent but not executed within the deadline.

Request Body: RemoveAttributeRequest.Body

Request payload message.

Field	Type	Description
parameters	Parameters	Op parameters.
signature	SignatureRFC6979	N3 witness of stable-marshalled parameters field. The signature must authenticate either container owner or one of subjects in the session_token field if any. Signature according to ECDSA_RFC6979_SHA256 scheme is also supported.
session_token	SessionTokenV2	Optional session token. The token must be issued by the container owner. The token must have at least one subject authenticated by signature field. The token must have at least one context with this container and CONTAINER_REMOVEATTRIBUTE verb.
session_token_v1	SessionToken	Optional session token (V1). It must not be set together with session_token field that is highly recommended to be used instead. Requirements are the same for both.

Message AnnounceUsedSpaceRequest.Body.Announcement

Announcement contains used space information for a single container.

Field	Type	Description
epoch	uint64	Epoch number for which the container size estimation was produced.
container_id	ContainerID	Identifier of the container.
used_space	uint64	Used space is a sum of object payload sizes of a specified container, stored in the node. It must not include inhumed objects.

Message RemoveAttributeRequest.Body.Parameters

Op parameters message.

If container does not have the `attribute`, nothing is done and status OK is returned.

`valid_until` is required request expiration time in Unix Timestamp format.

`attribute` must be one of: - `CORS`; - `S3_TAGS`; - `S3_SETTINGS`; - `S3_NOTIFICATIONS`; - `__NEOFS__LOCK_UNTIL`.

Attribute-specific requirements: - `__NEOFS__LOCK_UNTIL`: current timestamp must have already passed if any

Field	Type	Description
<code>container_id</code>	ContainerID	Identifier of the container to remove attribute from.
<code>attribute</code>	string	Attribute to be removed.
<code>valid_until</code>	uint64	Request expiration time.

Message SetAttributeRequest.Body.Parameters

Op parameters message.

If container does not have the `attribute`, it is added. Otherwise, its value is swapped.

`valid_until` is required request expiration time in Unix Timestamp format.

`attribute` must be one of: - `CORS`; - `S3_TAGS`; - `S3_SETTINGS`; - `S3_NOTIFICATIONS`; - `__NEOFS__LOCK_UNTIL`.

In general, requirements for `value` are the same as for container creation. Attribute-specific requirements: - `__NEOFS__LOCK_UNTIL`: new timestamp must be after the current one if any - `S3_TAGS`: must be a valid JSON object - `S3_SETTINGS`: must be a valid JSON object - `S3_NOTIFICATIONS`: must be a valid JSON object

Field	Type	Description
<code>container_id</code>	ContainerID	Identifier of the container to set attribute for.
<code>attribute</code>	string	Attribute to be set.
<code>value</code>	string	New attribute value.

Field	Type	Description
valid_until	uint64	Request expiration time.

Message Container

Container is a structure that defines object placement behaviour. Objects can be stored only within containers. They define placement rule, attributes and access control information. An ID of a container is a 32 byte long SHA256 hash of stable-marshalled container message.

Field	Type	Description
version	Version	Container format version. Effectively, the version of API library used to create the container.
owner_id	OwnerID	Identifier of the container owner
nonce	bytes	Nonce is a 16 byte UUIDv4, used to avoid collisions of ContainerIDs
basic_acl	uint32	BasicACL contains access control rules for the owner, system and others groups, as well as permission bits for BearerToken and Extended ACL
attributes	Attribute	Attributes represent immutable container's meta data
placement_policy	PlacementPolicy	Placement policy for the object inside the container

Message Container.Attribute

Attribute is a user-defined Key-Value metadata pair attached to the container.

Key name must be a container-unique valid UTF-8 string. Value can't be empty. Containers with duplicated attribute names or attributes with empty values will be considered invalid. Zero byte is also forbidden in UTF-8 strings.

There are some “well-known” attributes affecting system behaviour:

- `__NEOFS__SUBNET`

DEPRECATED. Was used for a string ID of a container's storage subnet. Currently doesn't affect anything.

- `__NEOFS__NAME`
String of a human-friendly container name registered as a domain in NNS contract.
- `__NEOFS__ZONE`
String of a zone for `__NEOFS__NAME`. Used as a TLD of a domain name in NNS contract. If no zone is specified, use default zone: `container`.
- `__NEOFS__DISABLE_HOMOMORPHIC_HASHING`
Disables homomorphic hashing for the container if the value equals "true" string. Any other values are interpreted as missing attribute. Container could be accepted in a NeoFS network only if the global network hashing configuration value corresponds with that attribute's value. After container inclusion, network setting is ignored.
- `__NEOFS__METAINFO_CONSISTENCY`
Policy rule that defines the condition under which an object is considered processed. Acceptable values and meanings:
 1. "strict": SN processes objects' meta information, it is validated, indexed and signed accordingly by a required minimal number of nodes that are included in the container, a corresponding object inclusion notification can be caught
 2. "optimistic": the same as "strict" but a successful PUT operation does not mean objects' meta information has been multi signed and indexed correctly, however, SNs will try to do it asynchronously; in general PUT operations are expected to be faster than in the "strict" case
 3. : SN does not process objects' meta information, it is not indexed and object presence/number of copies is not proven after a successful object PUT operation; the behavior is the same as it was before this attribute introduction
- `__NEOFS__LOCK_UNTIL`
Timestamp until which the container cannot be deleted in Unix Timestamp format

And some well-known attributes used by applications only:

- Name
Human-friendly name
- Timestamp
User-defined local time of container creation in Unix Timestamp format
- CORS
It is used to configure your container to allow cross-origin requests (CORS). The rules are represented as a JSON array of objects with the following fields:
 1. "AllowedMethods": In this element, you specify allowed HTTP methods: GET, PUT, POST,

- DELETE, HEAD.
2. “AllowedOrigins”: In this element, you specify the origins that you want to allow cross-domain requests from, for example, `http://www.example.com`. The origin string can contain only one “” *wildcard character, such as* `http://.example.com`. You can optionally specify “*” as the origin to enable all the origins to send cross-origin requests. You can also specify `https` to enable only secure origins.
 3. “AllowedHeaders”: The element specifies which headers are allowed in a preflight request through the “Access-Control-Request-Headers” request header. Each AllowedHeaders string in your configuration can contain at most one “” *wildcard character. For example,* `x-app-`.
 4. “ExposeHeaders”: Each ExposeHeader element identifies a header in the response that you want customers to be able to access from their applications (for example, from a JavaScript XMLHttpRequest object).
 5. “MaxAgeSeconds”: The element specifies the time in seconds that your browser can cache the response for a preflight request as identified by the resource, the HTTP method, and the origin.

The CORS schema is based on Amazon S3 CORS (<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cors-configuration>).

- S3_TAGS

It is used to store S3 gate-specific container tags. The value is controlled by the gate itself. Despite it, the value must be valid JSON object.

- S3_SETTINGS

It is used to store S3 gate-specific container settings. The value is controlled by the gate itself. Despite it, the value must be valid JSON object.

- S3_NOTIFICATIONS

It is used to store S3 gate-specific container notification settings. The value is controlled by the gate itself. Despite it, the value must be valid JSON object.

Field	Type	Description
key	string	Attribute name key
value	string	Attribute value

neo.fs.v2.link

Message Link

Link is a payload of helper objects that contain the full list of the split chain objects' IDs. It is created only after the whole split chain is known and signed. This object is the only object that refers to every "child object" ID. It is NOT required for the original object assembling. It MUST have ALL the "child objects" IDs. Child objects MUST be ordered according to the original payload split, meaning the first payload part holder MUST be placed at the first place in the corresponding link object. Sizes MUST NOT be omitted and MUST be a real object payload size in bytes.

Field	Type	Description
children	MeasuredObject	Full list of the "child" object descriptors.

Message Link.MeasuredObject

Object ID with its object's payload size.

Field	Type	Description
id	ObjectID	Object ID.
size	uint32	Object size in bytes.

neo.fs.v2.lock

Message Lock

DEPRECATED: lock objects 1-to-1 with `__NEOFS__ASSOCIATE` object attribute in a LOCK-typed object with no payload. For objects of 2.18+ API version, it is prohibited to have LOCK objects with payload.

Lock objects protects a list of objects from being deleted. The lifetime of a lock object is limited similar to regular objects in `__NEOFS__EXPIRATION_EPOCH` attribute. Lock object MUST have expiration epoch. It is impossible to delete a lock object via `ObjectService.Delete` RPC call. Deleting a container containing lock/locked objects results in their removal too, regardless of their expiration epochs.

Field	Type	Description
members	ObjectID	List of objects to lock. Must not be empty or carry empty IDs. All members must be of the REGULAR type.

neo.fs.v2.netmap

Service “NetmapService”

NetmapService provides methods to work with Network Map and the information required to build it. The resulting Network Map is stored in FS chain Netmap smart contract, while related information can be obtained from other NeoFS nodes.

Method LocalNodeInfo

Get NodeInfo structure from the particular node directly. Node information can be taken from Netmap smart contract. In some cases, though, one may want to get recent information directly or to talk to the node not yet present in the Network Map to find out what API version can be used for further communication. This can be also used to check if a node is up and running.

Statuses: - **OK** (0, SECTION_SUCCESS): information about the server has been successfully read; - Common failures (SECTION_FAILURE_COMMON).

Request Body: LocalNodeInfoRequest.Body

LocalNodeInfo request body is empty.

Response Body LocalNodeInfoResponse.Body

Local Node Info, including API Version in use.

Field	Type	Description
version	Version	Latest NeoFS API version in use
node_info	NodeInfo	NodeInfo structure with recent information from node itself

Method NetworkInfo

Read recent information about the NeoFS network.

Statuses: - **OK** (0, SECTION_SUCCESS): information about the current network state has been successfully read; - Common failures (SECTION_FAILURE_COMMON).

Request Body: NetworkInfoRequest.Body

NetworkInfo request body is empty.

Response Body NetworkInfoResponse.Body

Information about the network.

Field	Type	Description
network_info	NetworkInfo	NetworkInfo structure with recent information.

Method NetmapSnapshot

Returns network map snapshot of the current NeoFS epoch.

Statuses: - **OK** (0, SECTION_SUCCESS): information about the current network map has been successfully read; - Common failures (SECTION_FAILURE_COMMON).

Request Body: NetmapSnapshotRequest.Body

Get netmap snapshot request body.

Response Body NetmapSnapshotResponse.Body

Get netmap snapshot response body

Field	Type	Description
netmap	Netmap	Structure of the requested network map.

Message Filter

This filter will return the subset of nodes from NetworkMap or another filter's results that will satisfy filter's conditions.

Field	Type	Description
name	string	Name of the filter or a reference to a named filter. '*' means application to the whole unfiltered NetworkMap. At top level it's used as a filter name. At lower levels it's considered to be a reference to another named filter
key	string	Key to filter
op	Operation	Filtering operation
value	string	Value to match
filters	Filter	List of inner filters. Top level operation will be applied to the whole list.

Message Netmap

Network map structure

Field	Type	Description
epoch	uint64	Network map revision number.
nodes	NodeInfo	Nodes presented in network.

Message NetworkConfig

NeoFS network configuration

Field	Type	Description
parameters	Parameter	List of parameter values

Message NetworkConfig.Parameter

Single configuration parameter. Key MUST be network-unique.

System parameters: - **AuditFee**

Fee paid by the storage group owner to the Inner Ring member. Value: little-endian integer. Default: 0. - **BasicIncomeRate**

Cost of storing one gigabyte of data for a period of one epoch. Paid by container owner to container nodes. Value: little-endian integer. Default: 0. - **ContainerAliasFee**

Fee paid for named container's creation by the container owner. Value: little-endian integer. Default: 0. - **ContainerFee**

Fee paid for container creation by the container owner. Value: little-endian integer. Default: 0. - **EigenTrustAlpha**

Alpha parameter of EigenTrust algorithm used in the Reputation system. Value: decimal floating-point number in UTF-8 string representation. Default: 0. - **EigenTrustIterations**

Number of EigenTrust algorithm iterations to pass in the Reputation system. Value: little-endian integer. Default: 0. - **EpochDuration**

NeoFS epoch duration measured in seconds. Value: little-endian integer. Default: 0. - **HomomorphicHashingDisabled**

Flag of disabling the homomorphic hashing of objects' payload. Value: true if any byte != 0. Default: false. - **InnerRingCandidateFee**

Fee for entrance to the Inner Ring paid by the candidate. Value: little-endian integer. Default: 0. - **MaintenanceModeAllowed**

Flag allowing setting the MAINTENANCE state to storage nodes. Value: true if any byte != 0. Default: false. - **MaxObjectSize**

Maximum size of physically stored NeoFS object measured in bytes. Value: little-endian integer. Default: 0. - **WithdrawFee**

Fee paid for withdrawal of funds paid by the account owner. Value: little-endian integer. Default: 0.

Field	Type	Description
key	bytes	Parameter key. UTF-8 encoded string (with no zero bytes).
value	bytes	Parameter value

Message NetworkInfo

Information about NeoFS network

Field	Type	Description
current_epoch	uint64	Number of the current epoch in the NeoFS network

Field	Type	Description
magic_number	uint64	Magic number of FS chain of the NeoFS network
ms_per_block	int64	MillisecondsPerBlock network parameter of FS chain of the NeoFS network
network_config	NetworkConfig	NeoFS network configuration

Message NodeInfo

NeoFS node description

Field	Type	Description
public_key	bytes	Public key of the NeoFS node in a binary format
addresses	string	Ways to connect to a node
attributes	Attribute	Carries list of the NeoFS node attributes in a key-value form. Key name must be a node-unique valid UTF-8 string (without zero bytes). Value can't be empty. NodeInfo structures with duplicated attribute names or attributes with empty values will be considered invalid.
state	State	Carries state of the NeoFS node

Message NodeInfo.Attribute

Administrator-defined Attributes of the NeoFS Storage Node.

`Attribute` is a Key-Value metadata pair. Key name must be a valid UTF-8 string (without zero bytes that are forbidden). Value can't be empty.

Attributes can be constructed into a chain of attributes: any attribute can have a parent attribute and a child attribute (except the first and the last one). A string representation of the chain of attributes in NeoFS Storage Node configuration uses ":" and "/" symbols, e.g.:

```
`NEOFS_NODE_ATTRIBUTE_1=key1:val1/key2:val2`
```

Therefore the string attribute representation in the Node configuration must use “:”, “/” and “\” escaped symbols if any of them appears in an attribute’s key or value.

Node’s attributes are mostly used during Storage Policy evaluation to calculate object’s placement and find a set of nodes satisfying policy requirements. There are some “well-known” node attributes common to all the Storage Nodes in the network and used implicitly with default values if not explicitly set:

- Capacity
Total available disk space in Gigabytes.
- Price
Price in GAS tokens for storing one GB of data during one Epoch. In node attributes it’s a string presenting floating point number with comma or point delimiter for decimal part. In the Network Map it will be saved as 64-bit unsigned integer representing number of minimal token fractions.
- `__NEOFS_SUBNET%s`
DEPRECATED. Defined if the node is included in the %s subnetwork or not. Currently ignored.
- UN-LOCODE
Node’s geographic location in UN/LOCODE³⁶ format approximated to the nearest point defined in the standard.
- CountryCode
Country code in ISO 3166-1_alpha-2³⁷ format. Calculated automatically from UN-LOCODE attribute.
- Country
Country short name in English, as defined in ISO-3166³⁸. Calculated automatically from UN-LOCODE attribute.
- Location
Place names are given, whenever possible, in their national language versions as expressed in the Roman alphabet using the 26 characters of the character set adopted for international trade data interchange, written without diacritics. Calculated automatically from UN-LOCODE attribute.
- SubDivCode
Country’s administrative subdivision where node is located. Calculated automatically from UN-LOCODE attribute based on SubDiv field. Presented in ISO 3166-2³⁹ format.
- SubDiv

³⁶https://www.unece.org/cefact/codesfortrade/codes_index.html

³⁷https://en.wikipedia.org/wiki/ISO_3166-1_alpha-2

³⁸<https://www.iso.org/obp/ui/#search>

³⁹https://en.wikipedia.org/wiki/ISO_3166-2

Country's administrative subdivision name, as defined in ISO 3166-2⁴⁰. Calculated automatically from UN-LOCODE attribute.

- **Continent**

Node's continent name according to the [Seven-Continent model] (<https://en.wikipedia.org/wiki/Continent#Nu>). Calculated automatically from UN-LOCODE attribute.

- **ExternalAddr** Node's preferred way for communications with external clients. Clients SHOULD use these addresses if possible. Must contain a comma-separated list of multi-addresses. DEPRECATED. Use 'addresses' field instead.
- **Version** Node implementation's version in a free string form.
- **VerifiedNodesDomain** Confirmation of admission to a group of storage nodes. The value is the domain name registered in the NeoFS NNS. If attribute is specified, the storage node requesting entry into the NeoFS network map with this attribute must be included in the access list located on the specified domain. The access list is represented by a set of TXT records: Neo addresses resolved from public keys. To be admitted to the network, Neo address of the node's public key declared in 'public_key' field must be present in domain records. Otherwise, registration will be denied. Value must be a valid NeoFS NNS domain name. Note that if this attribute is absent, this check is not carried out.

For detailed description of each well-known attribute please see the corresponding section in NeoFS Technical Specification.

Field	Type	Description
key	string	Key of the node attribute
value	string	Value of the node attribute
parents	string	Parent keys, if any. For example for City it could be Region and Country.

Message PlacementPolicy

Set of rules to select a subset of nodes from NetworkMap able to store container's objects. The format is simple enough to transpile from different storage policy definition languages.

⁴⁰https://en.wikipedia.org/wiki/ISO_3166-2

Field	Type	Description
replicas	Replica	Rules to set number of object replicas and place each one into a named bucket. Limited to 256 items.
container_backup_factor	uint32	Container backup factor (CBF) controls how deep NeoFS will search for alternative nodes to include into container's nodes subset. In total, the number of container nodes is Selector (used by Replica) count times CBF. This number is limited to 64 per-Replica and 512 overall.
selectors	Selector	Set of Selectors to form the container's nodes subset
filters	Filter	List of named filters to reference in selectors
subnet_id	SubnetID	DEPRECATED. Was used for subnetwork ID to select nodes from, currently ignored.
ec_rules	ECRule	Erasur coding rules. Limited to 4 items.
initial	Initial	Initial placement rules.

Message PlacementPolicy.ECRule

Erasur coding rule for container objects.

For each original object, the payload is split into `data_part_num` data and `parity_part_num` parity parts. Each part is the same size. Data parts contain the original payload. If its length is not divisible by `data_part_num`, the last part is aligned with zero bytes. Both `data_part_num` and `parity_part_num` MUST NOT be zero or exceed 64, including in total. Hashes from all parts are written in the `__NEOFS__EC_PART_HASHES` attribute of the original object's header.

For each payload part, a part object is created. Original object's ID, signature and header is written in `header.split.parent`, `header.split.parent_signature` and `header.split.parent_header` fields correspondingly. Part index is written in the `__NEOFS__EC_PART_IDX` attribute as base-10 integer. Rule index in `PlacementPolicy.ec_rules` list is written in the `__NEOFS__EC_RULE_IDX` attribute as base-10 integer.

Each part object is stored in the container in one copy. Storage nodes are selected from the network map similar to `PlacementPolicy.replicas` rules. Optional `selector` acts the same way. The

object for the i -th part is placed in the i -th node. If it is unavailable, the backup nodes with indexes $m * n + i$ ($n = \text{data_part_num} + \text{parity_part_num}, m = 1, \dots, \text{CBF} - 1$). If all nodes for the i -th part are unavailable, nodes for the $i+1$ -th (0 for the last) part are tried, and so on.

Once part objects are stored in the container, the original object remains available if at least data_part_num of any part objects are available. In other words, unavailability (including complete loss) of any of parity_part_num part objects does not violate availability of the original one.

Objects of TOMBSTONE and LOCK types are not encoded and stored as they are because they have no payload.

Field	Type	Description
<code>data_part_num</code>	uint32	Number of data parts
<code>parity_part_num</code>	uint32	Number of parity parts
<code>selector</code>	string	Name of the linked selector

Message PlacementPolicy.Initial

Rules applied during initial data placement.

`replica_limits` allows to override `Replica.count` and EC partitions. If set, `replica_limits` must have a length equal to the sum of `replicas (RN)` and `ec_rules` length. Each of first `RN` elements of `replica_limits` must be less than or equal to corresponding `Replica.count`. The remaining elements must be either 0 (corresponding EC rule is skipped) or 1 (done). At least one `replica_limits` element must be non-zero.

`max_replicas` allows to limit total number of replicas and EC partitions for successful operation. If set, `max_replicas` must not overflow total replica limit (`replica_limits` or main ones).

`prefer_local` allows to tell server to try to store `MaxReplicas` replicas in locations that include this server. `prefer_local` must be set along with `max_replicas` only.

Either `replica_limits` or `max_replicas` must be specified.

Initial policy must not repeat the main one. In particular, policy with `replica_limits` equal to main ones only is invalid.

Field	Type	Description
replica_limits	uint32	Limits on the number of replicas and EC partitions
max_replicas	uint32	Maximum total number of replicas
prefer_local	bool	Flag to prefer local placement over regular one

Message Replica

Number of object replicas in a set of nodes from the defined selector. If no selector set, the root bucket containing all possible nodes will be used by default.

Field	Type	Description
count	uint32	How many object replicas to put. Limited to 8.
selector	string	Named selector bucket to put replicas

Message Selector

Selector chooses a number of nodes from the bucket taking the nearest nodes to the provided ContainerID by hash distance.

Field	Type	Description
name	string	Selector name to reference in object placement section
count	uint32	How many nodes to select from the bucket
clause	Clause	Selector modifier showing how to form a bucket
attribute	string	Bucket attribute to select from
filter	string	Filter reference to select from

Emun Clause

Selector modifier shows how the node set will be formed. By default selector just groups nodes into a bucket by attribute, selecting nodes only by their hash distance.

Number	Name	Description
0	CLAUSE_UNSPECIFIED	No modifier defined. Nodes will be selected from the bucket randomly
1	SAME	SAME will select only nodes having the same value of bucket attribute
2	DISTINCT	DISTINCT will select nodes having different values of bucket attribute

Emun NodeInfo.State

Represents the enumeration of various states of the NeoFS node.

Number	Name	Description
0	UNSPECIFIED	Unknown state
1	ONLINE	Active state in the network
2	OFFLINE	Network unavailable state
3	MAINTENANCE	Maintenance state

Emun Operation

Operations on filters

Number	Name	Description
0	OPERATION_UNSPECIFIED	No Operation defined
1	EQ	Equal
2	NE	Not Equal
3	GT	Greater then
4	GE	Greater or equal
5	LT	Less then
6	LE	Less or equal

Number	Name	Description
7	OR	Logical OR
8	AND	Logical AND

neo.fs.v2.object

Service “ObjectService”

ObjectService provides API for manipulating objects. Object operations do not affect FS chain and are only served by nodes in P2P style.

Method Get

Receive full object structure, including Headers and payload. Response uses gRPC stream. First response message carries the object with the requested address. Chunk messages are parts of the object’s payload if it is needed. All messages, except the first one, carry payload chunks. The requested object can be restored by concatenation of object message payload and all chunks keeping the receiving order.

Extended headers can change Get behaviour: * `__NEOFS__NETMAP_EPOCH`

Will use the requested version of Network Map for object placement calculation. DEPRECATED: header ignored by servers. * `__NEOFS__NETMAP_LOOKUP_DEPTH`

Will try older versions (starting from `__NEOFS__NETMAP_EPOCH` if specified or the latest one otherwise) of Network Map to find an object until the depth limit is reached. DEPRECATED: header ignored by servers.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

object has been successfully read; - Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT):

read access to the object is denied; - **OBJECT_NOT_FOUND** (2049, SECTION_OBJECT):

object not found in container; - **OBJECT_ALREADY_REMOVED** (2052, SECTION_OBJECT):

the requested object has been marked as deleted; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

object container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):

provided session token has expired.

Request Body: GetRequest.Body

GET Object request body

Field	Type	Description
address	Address	Address of the requested object
raw	bool	If raw flag is set, request will work only with objects that are physically stored on the peer node

Response Body GetResponse.Body

GET Object Response body

Field	Type	Description
init	Init	Initial part of the object stream
chunk	bytes	Chunked object payload
split_info	SplitInfo	Meta information of split hierarchy for object assembly.

Method Put

Put the object into container. Request uses gRPC stream. First message SHOULD be of PutHeader type. ContainerID and OwnerID of an object SHOULD be set. Session token SHOULD be obtained before PUT operation (see session package). Chunk messages are considered by server as a part of an object payload. All messages, except first one, SHOULD be payload chunks. Chunk messages SHOULD be sent in the direct order of fragmentation.

Extended headers can change Put behaviour: * __NEOFS__NETMAP_EPOCH

Will use the requested version of Network Map for object placement calculation. DEPRECATED: header ignored by servers.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

object has been successfully saved in the container; - **INCOMPLETE** (1, SECTION_SUCCESS):

object was put to some nodes, but the number of replicas is not sufficient to satisfy placement policy;

- Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT): write access to the container is denied; - **LOCKED** (2050, SECTION_OBJECT): placement of an object of type TOMBSTONE that includes at least one locked object is prohibited; - **LOCK_NON_REGULAR_OBJECT** (2051, SECTION_OBJECT): placement of an object of type LOCK that includes at least one object of type other than REGULAR is prohibited; - **QUOTA_EXCEEDED** (2054, SECTION_OBJECT): size quota set by user was exceeded; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER): object storage container not found; - **TOKEN_NOT_FOUND** (4096, SECTION_SESSION): (for trusted object preparation) session private key does not exist or has been deleted; - **TOKEN_EXPIRED** (4097, SECTION_SESSION): provided session token has expired.

Request Body: PutRequest.Body

PUT request body

Field	Type	Description
init	Init	Initial part of the object stream
chunk	bytes	Chunked object payload

Response Body PutResponse.Body

PUT Object response body

Field	Type	Description
object_id	ObjectID	Identifier of the saved object

Method Delete

Delete the object from a container. There is no immediate removal guarantee. Object will be marked for removal and deleted eventually. Notice that some types of objects (see ObjectType) can not be removed, currently it's Tombstone and Lock.

Extended headers can change Delete behaviour: * __NEOFS__NETMAP_EPOCH

Will use the requested version of Network Map for object placement calculation. DEPRECATED: header ignored by servers.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

object has been successfully marked to be removed from the container; - **INCOMPLETE** (1, SECTION_SUCCESS):

some nodes have accepted the deletion mark, but some may still store the object; - Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT):

delete access to the object is denied; - **LOCKED** (2050, SECTION_OBJECT):

deleting a locked object is prohibited; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

object container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):

provided session token has expired.

Request Body: DeleteRequest.Body

Object DELETE request body

Field	Type	Description
address	Address	Address of the object to be deleted

Response Body DeleteResponse.Body

Object DELETE Response has an empty body.

Field	Type	Description
tombstone	Address	Address of the tombstone created for the deleted object

Method Head

Returns the object Headers without data payload. By default full header is returned. If `main_only` request field is set, the short header with only the very minimal information will be returned instead.

Extended headers can change Head behaviour: * `__NEOFS__NETMAP_EPOCH`

Will use the requested version of Network Map for object placement calculation.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

object header has been successfully read; - Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT):

access to operation HEAD of the object is denied; - **OBJECT_NOT_FOUND** (2049, SECTION_OBJECT):

object not found in container; - **OBJECT_ALREADY_REMOVED** (2052, SECTION_OBJECT):

the requested object has been marked as deleted; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

object container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):
provided session token has expired.

Request Body: HeadRequest.Body

Object HEAD request body

Field	Type	Description
address	Address	Address of the object with the requested Header
main_only	bool	Return only minimal header subset.

DEPRECATED. This field is ignored. || raw | bool | If raw flag is set, request will work only with objects that are physically stored on the peer node |

Response Body HeadResponse.Body

Object HEAD response body

Field	Type	Description
header	HeaderWithSignature	Full object's Header with ObjectID signature
short_header	ShortHeader	Short object header.

DEPRECATED. Use HeaderWithSignature instead. || split_info | SplitInfo | Meta information of split hierarchy. |

Method Search

Search objects in container. Search query allows to match by Object Header's filed values. Please see the corresponding NeoFS Technical Specification section for more details.

DEPRECATED: please use SearchV2.

Extended headers can change Search behaviour: * __NEOFS__NETMAP_EPOCH
Will use the requested version of Network Map for object placement calculation.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

objects have been successfully selected; - **INCOMPLETE** (1, SECTION_SUCCESS):

some nodes were unable to process the request, so the result may not contain all data; - Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT):

access to operation SEARCH of the object is denied; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

search container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):

provided session token has expired.

Request Body: SearchRequest.Body

Object Search request body

Field	Type	Description
container_id	ContainerID	Container identifier were to search
version	uint32	Version of the Query Language used
filters	SearchFilter	List of search expressions

Response Body SearchResponse.Body

Object Search response body

Field	Type	Description
id_list	ObjectID	List of ObjectIDs that match the search query

Method SearchV2

Search for objects in a container. Similar to Search, but: * sorted * limited in amount of returned data * single message * allows for additional header fields to be returned

Result is ordered by the 1st requested attribute (if any) and object ID.

Request Body: SearchV2Request.Body

Object Search request body

Field	Type	Description
container_id	ContainerID	Container where the search is being performed.
version	uint32	Version of the Query Language used.
filters	SearchFilter	List of search expressions. Limited to 8. If additional attributes are requested (see attributes below) then the first filter's key MUST be the first requested attribute. <i>'Object : containerID' and 'Object:objectID'</i> filters are prohibited. Numeric filters' values MUST be in range [-MaxUint256, MaxUint256].
cursor	string	Cursor to continue search. Can be omitted or empty for the new search.
count	uint32	Limits the number of responses to the specified number. Can't be more than 1000.
attributes	string	List of attribute names (including special ones as defined by SearchFilter key) to include into the reply. Limited to 8, these attributes also affect result ordering (result is ordered by the 1st one and then by OID). If additional attributes are requested, then the first filter's key (see filters above) MUST be the first requested attribute. <i>'Object : containerID' and 'Object:objectID'</i> attributes are prohibited. If meta_header.ttl = 1 and the first filter is not STRING_EQUAL, values of the first filtered attribute are requested automatically.

Response Body SearchV2Response.Body

Main result structure.

Field	Type	Description
result	OIDWithMeta	List of object IDs with additional requested attributes.

Field	Type	Description
cursor	string	Cursor that can be used for subsequent requests. For users it's an opaque string that is omitted or empty when there are no more results to list. For nodes to interoperate this is defined as the latest OID for queries without filters and primary (first) attribute value plus OID. Values are encoded in base64.

Method GetRange

Get byte range of data payload. Range is set as an (offset, length) tuple. Like in Get method, the response uses gRPC stream. Requested range can be restored by concatenation of all received payload chunks keeping the receiving order.

Extended headers can change GetRange behaviour: * `__NEOFS__NETMAP_EPOCH`

Will use the requested version of Network Map for object placement calculation. DEPRECATED: header ignored by servers. * `__NEOFS__NETMAP_LOOKUP_DEPTH`

Will try older versions of Network Map to find an object until the depth limit is reached. DEPRECATED: header ignored by servers.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

data range of the object payload has been successfully read; - Common failures (SECTION_FAILURE_COMMON);

- **ACCESS_DENIED** (2048, SECTION_OBJECT):

access to operation RANGE of the object is denied; - **OBJECT_NOT_FOUND** (2049, SECTION_OBJECT):

object not found in container; - **OBJECT_ALREADY_REMOVED** (2052, SECTION_OBJECT):

the requested object has been marked as deleted. - **OUT_OF_RANGE** (2053, SECTION_OBJECT):

the requested range is out of bounds; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):

object container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):

provided session token has expired.

Request Body: GetRangeRequest.Body

Byte range of object's payload request body

Field	Type	Description
address	Address	Address of the object containing the requested payload range
range	Range	Requested payload range
raw	bool	If raw flag is set, request will work only with objects that are physically stored on the peer node.

Response Body `GetRangeResponse.Body`

Get Range response body uses streams to transfer the response. Because object payload considered a byte sequence, there is no need to have some initial preamble message. The requested byte range is sent as a series chunks.

Field	Type	Description
chunk	bytes	Chunked object payload's range.
split_info	SplitInfo	Meta information of split hierarchy.

Method `GetRangeHash`

Returns homomorphic or regular hash of object's payload range after applying XOR operation with the provided salt. Ranges are set of (offset, length) tuples. Hashes order in response corresponds to the ranges order in the request. Note that hash is calculated for XORed data.

Extended headers can change `GetRangeHash` behaviour: * `__NEOFS__NETMAP_EPOCH`

Will use the requested version of Network Map for object placement calculation. DEPRECATED: header ignored by servers. * `__NEOFS__NETMAP_LOOKUP_DEPTH`

Will try older versions of Network Map to find an object until the depth limit is reached. DEPRECATED: header ignored by servers.

Please refer to detailed XHeader description.

Statuses: - **OK** (0, SECTION_SUCCESS):

data range of the object payload has been successfully hashed; - Common failures (SECTION_FAILURE_COMMON); - **ACCESS_DENIED** (2048, SECTION_OBJECT):

access to operation RANGEHASH of the object is denied; - **OBJECT_NOT_FOUND** (2049, SECTION_OBJECT):

object not found in container; - **OUT_OF_RANGE** (2053, SECTION_OBJECT):
 the requested range is out of bounds; - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER):
 object container not found; - **TOKEN_EXPIRED** (4097, SECTION_SESSION):
 provided session token has expired.

Request Body: GetRangeHashRequest.Body

Get hash of object's payload part request body.

Field	Type	Description
address	Address	Address of the object that containing the requested payload range
ranges	Range	List of object's payload ranges to calculate homomorphic hash
salt	bytes	Binary salt to XOR object's payload ranges before hash calculation
type	ChecksumType	Checksum algorithm type

Response Body GetRangeHashResponse.Body

Get hash of object's payload part response body.

Field	Type	Description
type	ChecksumType	Checksum algorithm type
hash_list	bytes	List of range hashes in a binary format

Method Replicate

Save replica of the object on the NeoFS storage node. Both client and server must authenticate NeoFS storage nodes matching storage policy of the container referenced by the replicated object. Thus, this operation is purely system: regular users should not pay attention to it but use Put.

Statuses: - **OK** (0, SECTION_SUCCESS):

the object has been successfully replicated; - **INTERNAL_SERVER_ERROR** (1024, SECTION_FAILURE_COMMON):
 internal server error described in the text message; - **ACCESS_DENIED** (2048, SECTION_OBJECT):
 the client does not authenticate any NeoFS storage node matching storage policy of the container

referenced by the replicated object - **CONTAINER_NOT_FOUND** (3072, SECTION_CONTAINER): the container to which the replicated object is associated was not found.

Message **GetResponse.Body.Init**

Initial part of the `Object` structure stream. Technically it's a set of all `Object` structure's fields except `payload`.

Field	Type	Description
<code>object_id</code>	<code>ObjectID</code>	Object's unique identifier.
<code>signature</code>	<code>Signature</code>	Signed <code>ObjectID</code>
<code>header</code>	<code>Header</code>	Object metadata headers

Message **HeaderWithSignature**

Tuple of a full object header and signature of an `ObjectID`.

Signed `ObjectID` is present to verify full header's authenticity through the following steps:

1. Calculate SHA-256 of the marshalled `Header` structure
2. Check if the resulting hash matches `ObjectID`
3. Check if `ObjectID` signature in `signature` field is correct

Field	Type	Description
<code>header</code>	<code>Header</code>	Full object header
<code>signature</code>	<code>Signature</code>	Signed <code>ObjectID</code> to verify full header's authenticity

Message **PutRequest.Body.Init**

Newly created object structure parameters. If some optional parameters are not set, they will be calculated by a peer node.

Field	Type	Description
object_id	ObjectID	ObjectID if available.
signature	Signature	Object signature if available
header	Header	Object's Header. The maximum length is 16KB.
copies_number	uint32	Number of the object copies to store within the RPC call. By default object is processed according to the container's placement policy. DEPRECATED: use PlacementPolicy.Initial.max_replicas instead. Servers ignore this field.

Message Range

Object payload range. Ranges of zero length SHOULD be considered as invalid except for the special 0:0 request which is interpreted as “get whole payload” and allows to receive payload only when its size is not known.

Field	Type	Description
offset	uint64	Offset of the range from the object payload start
length	uint64	Length in bytes of the object payload range

Message SearchV2Response.OIDWithMeta

OID with additional requested metadata.

Field	Type	Description
id	ObjectID	Object ID that matches search criteria.
attributes	string	List of attribute data from the respective object, fields strictly follow requested ones. The only exception are attribute-less requests with TTL = 1 and the first non-STRING_EQUAL filter: this field carry value of the first filtered attribute.

Message Header

Object Header

Field	Type	Description
version	Version	Object format version. Effectively, the version of API library used to create particular object
container_id	ContainerID	Object's container
owner_id	OwnerID	Object's owner
creation_epoch	uint64	Object creation Epoch
payload_length	uint64	Size of payload in bytes. 0xFFFFFFFFFFFFFFFF means payload_length is unknown.
payload_hash	Checksum	SHA256 hash of payload bytes
object_type	ObjectType	Type of the object payload content
homomorphic_hash	Checksum	Homomorphic hash of the object payload (Tillich-Zemor).
session_token	SessionToken	Session token, if it was used during Object creation. Need it to verify integrity and authenticity out of Request scope. Only one of session_token or session_token_v2 can be set.
attributes	Attribute	User-defined object attributes. Attributes vary in length from object to object, so keep an eye on the entire Header limit depending on the context.
split	Split	Position of the object in the split hierarchy
session_token_v2	SessionTokenV2	Session token V2, if it was used during Object creation. Need it to verify integrity and authenticity out of Request scope. Only one of session_token or session_token_v2 can be set.

Message Header.Attribute

`Attribute` is a user-defined Key-Value metadata pair attached to an object.

Key name must be an object-unique valid UTF-8 string. Value can't be empty. Objects with duplicated attribute names or attributes with empty values will be considered invalid. Keys and values can't contain zero bytes as well.

There are some “well-known” attributes starting with `__NEOFS__` prefix that affect system behaviour:

- `__NEOFS__EXPIRATION_EPOCH`
Tells GC to delete object after that epoch (but object is available throughout the epoch specified in this attribute).
- `__NEOFS__ASSOCIATE`
Associated object. For TOMBSTONE, LOCK object types it defines object to delete and to lock accordingly. For objects of 2.18+ API version, it is the only way to delete/lock objects. It MUST be a single stringified (according to [refs.ObjectID] message) object ID with no leading or trailing spaces.
- `__NEOFS__TICK_EPOCH`
Decimal number that defines what epoch must produce object notification with UTF-8 object address in a body (0 value produces notification right after object put). DEPRECATED: attribute ignored by servers.
- `__NEOFS__TICK_TOPIC`
UTF-8 string topic ID that is used for object notification. DEPRECATED: attribute ignored by servers.
- `__NEOFS__EC_RULE_IDX`
Index of EC rule in container's `PlacementPolicy.ec_rules` according to which the part was created. Base-10 integer.
- `__NEOFS__EC_PART_IDX`
Index in the EC parts into which the parent object is divided according to `__NEOFS__EC_RULE_IDX` EC rule. Base-10 integer.
- `__NEOFS__EC_PART_HASHES`
Ordered list of payload hashes of EC parts into which this object is divided. Hash function is SHA-256. Items are comma-separated, each item is hex-encoded.

And some well-known attributes used by applications only:

- Name
Human-friendly name

- **FileName**
File name to be associated with the object on saving
- **FilePath**
Full path to be associated with the object on saving. Should start with a '/' and use '/' as a delimiting symbol. Trailing '/' should be interpreted as a virtual directory marker. If an object has conflicting FilePath and FileName, FilePath should have higher priority, because it is used to construct the directory tree. FilePath with trailing '/' and non-empty FileName attribute should not be used together.
- **Timestamp**
User-defined local time of object creation in Unix Timestamp format
- **Content-Type**
MIME Content Type of object's payload

For detailed description of each well-known attribute please see the corresponding section in NeoFS Technical Specification.

Field	Type	Description
key	string	string key to the object attribute
value	string	string value of the object attribute

Message Header.Split

Bigger objects can be split into a chain of smaller objects. Information about inter-dependencies between spawned objects and how to re-construct the original one is in the `Split` headers. Parent and children objects must be within the same container.

Field	Type	Description
parent	ObjectID	Identifier of the origin object. If the origin object is split to comply with the object size limit, parent is known only to the minor child.
previous	ObjectID	Identifier of the left split neighbor
parent_signature	Signature	signature field of the parent object. Used to reconstruct parent.
parent_header	Header	header field of the parent object. Used to reconstruct parent.

Field	Type	Description
children	ObjectID	DEPRECATED. Was used before creating the separate LINK object type. Keep child objects list in the LINK object's payload. List of identifiers of the objects generated by splitting current one.
split_id	bytes	DEPRECATED. Was used as an identifier of a split chain. Use the first part ID instead. 16 byte UUIDv4 used to identify the split object hierarchy parts. Must be unique inside container. All objects participating in the split must have the same <code>split_id</code> value.
first	ObjectID	Identifier of the first part of the origin object. Known to all the split parts except the first one. Identifies the split and allows to differ them.

Message Object

Object structure. Object is immutable and content-addressed. It means `ObjectID` will change if the header or the payload changes. It's calculated as a hash of header field which contains hash of the object's payload.

For non-regular object types payload format depends on object type specified in the header.

Field	Type	Description
object_id	ObjectID	Object's unique identifier.
signature	Signature	Signed <code>object_id</code>
header	Header	Object metadata headers
payload	bytes	Payload bytes

Message SearchFilter

Filter structure checks if the object header field or the attribute content matches a value.

If no filters are set, search request will return all objects of the container, including Regular object,

Tombstones and Storage Group objects. Most human users expect to get only object they can directly work with. In that case, `$Object:ROOT` filter should be used.

If `match_type` field is numerical, both `value` field and object attribute MUST be base-10 integers.

By default key field refers to the corresponding object's `Attribute`. Some Object's header fields can also be accessed by adding `$Object:` prefix to the name. Here is the list of fields available via this prefix:

- `$Object:version`
version
- `$Object:ownerID`
owner_id
- `$Object:creationEpoch`
creation_epoch
- `$Object:payloadLength`
payload_length
- `$Object:payloadHash`
payload_hash
- `$Object:objectType`
object_type
- `$Object:homomorphicHash`
homomorphic_hash
- `$Object:split.parent`
object_id of parent
- `$Object:split.splitID`
16 byte UUIDv4 used to identify the split object hierarchy parts
- `$Object:split.first`
object_id of the first part in split chain; non-acceptable for deprecated V1 split scheme

There are some well-known filter aliases to match objects by certain properties:

- `$Object:ROOT`
Returns only REGULAR type objects that are not split or that are the top level root objects in a split hierarchy. This includes objects not present physically, like large objects split into smaller objects without a separate top-level root object. Objects of other types like StorageGroups and Tombstones will not be shown. This filter may be useful for listing objects like `ls` command of some virtual file system. This filter is activated if the key exists, disregarding the value and matcher type.
- `$Object:PHY`
Returns only objects physically stored in the system. This filter is activated if the key exists,

disregarding the value and matcher type.

Following filters are deprecated:

- \$Object:objectID
object_id
- \$Object:containerID
container_id

Note: using filters with a key with prefix \$Object: and match type NOT_PRESENT is not recommended since this is not a cross-version approach. Behavior when processing this kind of filters is undefined.

Field	Type	Description
match_type	MatchType	Match type to use
key	string	Attribute or Header fields to match
value	string	Value to match

Message ShortHeader

Short header fields

DEPRECATED. It is not supported in practice and should not be used. Use full Header instead.

Field	Type	Description
version	Version	Object format version. Effectively, the version of API library used to create particular object.
creation_epoch	uint64	Epoch when the object was created
owner_id	OwnerID	Object's owner
object_type	ObjectType	Type of the object payload content
payload_length	uint64	Size of payload in bytes. 0xFFFFFFFFFFFFFFFF means payload_length is unknown
payload_hash	Checksum	SHA256 hash of payload bytes.

Field	Type	Description
homomorphic_hash	Checksum	Homomorphic hash of the object payload (Tillich-Zemor).

Message SplitInfo

Meta information of split hierarchy for object assembly. With the last part one can traverse linked list of split hierarchy back to the first part and assemble the original object. With a linking object one can assemble an object right from the object parts.

Field	Type	Description
split_id	bytes	DEPRECATED. Was used as an identifier of a split chain. Use the first part ID instead. 16 byte UUID used to identify the split object hierarchy parts.
last_part	ObjectID	The identifier of the last object in split hierarchy parts. It contains split header with the original object header.
link	ObjectID	The identifier of a linking object for split hierarchy parts. It contains split header with the original object header and a sorted list of object parts.
first_part	ObjectID	Identifier of the first part of the origin object. Known to all the split parts except the first one. Identifies the split and allows to differ them.

Emun MatchType

Type of match expression

Number	Name	Description
0	MATCH_TYPE_UNSPECIFIED	Unknown. Not used
1	STRING_EQUAL	Full string match
2	STRING_NOT_EQUAL	Full string mismatch

Number	Name	Description
3	NOT_PRESENT	Lack of key
4	COMMON_PREFIX	String prefix match
5	NUM_GT	Numerical 'greater than'
6	NUM_GE	Numerical 'greater or equal than'
7	NUM_LT	Numerical 'less than'
8	NUM_LE	Numerical 'less or equal than'

Emun ObjectType

Type of the object payload content. Only REGULAR type objects can be split, hence TOMBSTONE, STORAGE_GROUP and LOCK payload is limited by the maximum object size.

String presentation of object type is the same as definition: * REGULAR * TOMBSTONE * STORAGE_GROUP * LOCK * LINK

Number	Name	Description
0	REGULAR	Just a normal object
1	TOMBSTONE	Used internally to identify deleted objects
2	STORAGE_GROUP	StorageGroup information. DEPRECATED: no longer used for audit since 2.18.
3	LOCK	Object lock
4	LINK	Object that stores child object IDs for the split objects.

neo.fs.v2.refs

Message Address

Objects in NeoFS are addressed by their ContainerID and ObjectID.

String presentation of Address is a concatenation of string encoded ContainerID and ObjectID delimited by '/' character.

Field	Type	Description
container_id	ContainerID	Container identifier
object_id	ObjectID	Object identifier

Message Checksum

Checksum message. Depending on checksum algorithm type, the string presentation may vary:

- TZ
Hex encoded string without 0x prefix
- SHA256
Hex encoded string without 0x prefix

Field	Type	Description
type	ChecksumType	Checksum algorithm type
sum	bytes	Checksum itself

Message ContainerID

NeoFS container identifier.

ContainerID is a 32 byte long SHA256⁴¹ hash of stable-marshalled container message used for creation. ID consisting of all zero bytes is reserved for undefined value and must not be specified as a field.

String presentation is a base58⁴² encoded string.

JSON value will be data encoded as a string using standard base64 encoding with paddings. Either standard⁴³ or URL-safe⁴⁴ base64 encoding with/without paddings are accepted.

⁴¹<https://csrc.nist.gov/publications/detail/fips/180/4/final>

⁴²<https://tools.ietf.org/html/draft-msporny-base58-02>

⁴³<https://tools.ietf.org/html/rfc4648#section-4>

⁴⁴<https://tools.ietf.org/html/rfc4648#section-5>

Field	Type	Description
value	bytes	Container identifier in a binary format.

Message ObjectID

NeoFS Object unique identifier. Objects are immutable and content-addressed. It means ObjectID will change if the header or the payload changes.

ObjectID is a 32 byte long SHA256⁴⁵ hash of the object's header field, which, in its turn, contains the hash of the object's payload. ID consisting of all zero bytes is reserved for undefined value and must not be specified as a field.

String presentation is a base58⁴⁶ encoded string.

JSON value will be data encoded as a string using standard base64 encoding with paddings. Either standard⁴⁷ or URL-safe⁴⁸ base64 encoding with/without paddings are accepted.

Field	Type	Description
value	bytes	Object identifier in a binary format

Message OwnerID

OwnerID is a derivative of a user's main public key. The transformation algorithm is the same as for Neo3 wallet addresses. Neo3 wallet address can be directly used as OwnerID.

OwnerID is a 25 bytes sequence starting with Neo version prefix byte followed by 20 bytes of ScrptHash and 4 bytes of checksum.

String presentation is a Base58 Check⁴⁹ Encoded string.

JSON value will be data encoded as a string using standard base64 encoding with paddings. Either standard⁵⁰ or URL-safe⁵¹ base64 encoding with/without paddings are accepted.

⁴⁵<https://csrc.nist.gov/publications/detail/fips/180/4/final>

⁴⁶<https://tools.ietf.org/html/draft-msporny-base58-02>

⁴⁷<https://tools.ietf.org/html/rfc4648#section-4>

⁴⁸<https://tools.ietf.org/html/rfc4648#section-5>

⁴⁹https://en.bitcoin.it/wiki/Base58Check_encoding

⁵⁰<https://tools.ietf.org/html/rfc4648#section-4>

⁵¹<https://tools.ietf.org/html/rfc4648#section-5>

Field	Type	Description
value	bytes	Identifier of the container owner in a binary format

Message Signature

Signature of something in NeoFS.

Field	Type	Description
key	bytes	Public key used for signing. For N3 scheme, the field represents a verification script. The maximum allowed length is 1024 bytes.
sign	bytes	Signature. For N3 scheme, the field represents an invocation script. The maximum allowed length is 1024 bytes.
scheme	SignatureScheme	Scheme contains digital signature scheme identifier

Message SignatureRFC6979

RFC 6979 signature.

Field	Type	Description
key	bytes	Public key used for signing. For N3 auth scheme, the field represents a verification script. The maximum allowed length is 1024 bytes.
sign	bytes	Deterministic ECDSA with SHA-256 hashing. For N3 auth scheme, the field represents an invocation script. The maximum allowed length is 1024 bytes.

Message SubnetID

NeoFS subnetwork identifier.

String representation of a value is base-10 integer.

JSON representation is an object containing a single value number field.

DEPRECATED. Kept for compatibility only.

Field	Type	Description
value	fixed32	4-byte integer subnetwork identifier.

Message Version

API version used by a node.

String presentation is a Semantic Versioning 2.0.0 compatible version string with 'v' prefix. i.e. vX.Y, where X is the major number, Y is the minor number.

Field	Type	Description
major	uint32	Major API version
minor	uint32	Minor API version

Emun ChecksumType

Checksum algorithm type.

Number	Name	Description
0	CHECKSUM_TYPE_UNSPECIFIED	Unknown. Not used
1	TZ	Tillich-Zemor homomorphic hash function
2	SHA256	SHA-256

Emun SignatureScheme

Signature scheme describes digital signing scheme used for (key, signature) pair.

Number	Name	Description
0	ECDSA_SHA512	ECDSA with SHA-512 hashing (FIPS 186-3)
1	ECDSA_RFC6979_SHA256	Deterministic ECDSA with SHA-256 hashing (RFC 6979)
2	ECDSA_RFC6979_WC	Deterministic ECDSA with SHA-256 hashing using WalletConnect API. Here the algorithm is the same, but the message format differs.
3	N3	Neo N3 witness.

neo.fs.v2.reputation

Service “ReputationService”

ReputationService provides mechanisms for exchanging trust values with other NeoFS nodes. Nodes rate each other’s reputation based on how good they process requests and set a trust level based on that rating. The trust information is passed to the next nodes to check and aggregate unless the final result is recorded.

Method AnnounceLocalTrust

Announce local client trust information to any node in NeoFS network.

Statuses: - **OK** (0, SECTION_SUCCESS): local trust has been successfully announced; - Common failures (SECTION_FAILURE_COMMON).

Request Body: AnnounceLocalTrustRequest.Body

Announce node’s local trust information.

Field	Type	Description
epoch	uint64	Trust assessment Epoch number

Field	Type	Description
trusts	Trust	List of normalized local trust values to other NeoFS nodes. The value is calculated according to EigenTrust++ algorithm and must be a floating point number in [0;1] range.

Response Body `AnnounceLocalTrustResponse.Body`

Response to the node's local trust information announcement has an empty body because the trust exchange operation is asynchronous. If Trust information does not pass sanity checks, it is silently ignored.

Method `AnnounceIntermediateResult`

Announce the intermediate result of the iterative algorithm for calculating the global reputation of the node in NeoFS network.

Statuses: - **OK** (0, SECTION_SUCCESS): intermediate trust estimation has been successfully announced; - Common failures (SECTION_FAILURE_COMMON).

Request Body: `AnnounceIntermediateResultRequest.Body`

Announce intermediate global trust information.

Field	Type	Description
epoch	uint64	Iteration execution Epoch number
iteration	uint32	Iteration sequence number
trust	PeerToPeerTrust	Current global trust value calculated at the specified iteration

Response Body `AnnounceIntermediateResultResponse.Body`

Response to the node's intermediate global trust information announcement has an empty body because the trust exchange operation is asynchronous. If Trust information does not pass sanity checks, it is silently ignored.

Message GlobalTrust

Global trust level to NeoFS node.

Field	Type	Description
version	Version	Message format version. Effectively, the version of API library used to create the message.
body	Body	Message body
signature	Signature	Signature of the binary body field by the manager.

Message GlobalTrust.Body

Message body structure.

Field	Type	Description
manager	PeerID	Node manager ID
trust	Trust	Global trust level

Message PeerID

NeoFS unique peer identifier is a 33 byte long compressed public key of the node, the same as the one stored in the network map.

String presentation is a base58⁵² encoded string.

JSON value will be data encoded as a string using standard base64 encoding with paddings. Either standard⁵³ or URL-safe⁵⁴ base64 encoding with/without paddings are accepted.

Field	Type	Description
public_key	bytes	Peer node's public key

⁵²<https://tools.ietf.org/html/draft-msporny-base58-02>

⁵³<https://tools.ietf.org/html/rfc4648#section-4>

⁵⁴<https://tools.ietf.org/html/rfc4648#section-5>

Message PeerToPeerTrust

Trust level of a peer to a peer.

Field	Type	Description
trusting_peer	PeerID	Identifier of the trusting peer
trust	Trust	Trust level

Message Trust

Trust level to a NeoFS network peer.

Field	Type	Description
peer	PeerID	Identifier of the trusted peer
value	double	Trust level in [0:1] range

neo.fs.v2.session

Service “SessionService”

SessionService allows to establish a temporary trust relationship between two peer nodes and generate a SessionToken as the proof of trust to be attached in requests for further verification. Please see corresponding section of NeoFS Technical Specification for details.

Method Create

Open a new session between two peers.

Statuses: - **OK** (0, SECTION_SUCCESS): session has been successfully opened; - Common failures (SECTION_FAILURE_COMMON).

Request Body: CreateRequest.Body

Session creation request body

Field	Type	Description
owner_id	OwnerID	Session initiating user's or node's key derived Owner ID
expiration	uint64	Session expiration epoch, the last epoch when session is valid.

Response Body CreateResponse.Body

Session creation response body

Field	Type	Description
id	bytes	Identifier of a newly created session
session_key	bytes	Public key used for session

Message ContainerSessionContext

Context information for Session Tokens related to ContainerService requests.

Field	Type	Description
verb	Verb	Type of request for which the token is issued
wildcard	bool	Spreads the action to all owner containers. If set, container_id field is ignored.
container_id	ContainerID	Particular container to which the action applies. Ignored if wildcard flag is set.

Message ObjectSessionContext

Context information for Session Tokens related to ObjectService requests

Field	Type	Description
verb	Verb	Type of request for which the token is issued

Field	Type	Description
target	Target	Object session target. MUST be correctly formed and set. If objects field is not empty, then the session applies only to these elements, otherwise, to all objects from the specified container.

Message ObjectSessionContext.Target

Carries objects involved in the object session.

Field	Type	Description
container	ContainerID	Indicates which container the session is spread to. Field MUST be set and correct.
objects	ObjectID	Indicates which objects the session is spread to. Objects are expected to be stored in the NeoFS container referenced by container field. Each element MUST have correct format.

Message RequestMetaHeader

Meta information attached to the request. When forwarded between peers, request meta headers are folded in matryoshka style.

Field	Type	Description
version	Version	Peer's API version used
epoch	uint64	Peer's local epoch number. Set to 0 if unknown.
ttl	uint32	Maximum number of intermediate nodes in the request route
x_headers	XHeader	Request X-Headers
session_token	SessionToken	Session token within which the request is sent

Field	Type	Description
session_token_v2	SessionTokenV2	Session token v2 with delegation chain support. Requests are invalid if both session_token and session_token_v2 are set.
bearer_token	BearerToken	BearerToken with eACL overrides for the request
origin	RequestMetaHeader	RequestMetaHeader of the origin request
magic_number	uint64	NeoFS network magic. Must match the value for the network that the server belongs to.

Message RequestVerificationHeader

Verification info for the request signed by all intermediate nodes.

Field	Type	Description
body_signature	Signature	Request Body signature. Should be generated once by the request initiator.
meta_signature	Signature	Request Meta signature is added and signed by each intermediate node
origin_signature	Signature	Signature of previous hops
origin	RequestVerificationHeader	Chain of previous hops signatures

Message ResponseMetaHeader

Information about the response

Field	Type	Description
version	Version	Peer's API version used
epoch	uint64	Peer's local epoch number
ttl	uint32	Maximum number of intermediate nodes in the request route
x_headers	XHeader	Response X-Headers

Field	Type	Description
origin	ResponseMetaHeader	ResponseMetaHeader of the origin request
status	Status	Status return

Message ResponseVerificationHeader

Verification info for the response signed by all intermediate nodes

DEPRECATED: was eliminated from the protocol starting from version v2 . 22.

Field	Type	Description
body_signature	Signature	Response Body signature. Should be generated once by an answering node.
meta_signature	Signature	Response Meta signature is added and signed by each intermediate node
origin_signature	Signature	Signature of previous hops
origin	ResponseVerificationHeader	Chain of previous hops signatures

Message SessionContextV2

SessionContextV2 carries unified context for both ObjectService and ContainerService requests.

Field	Type	Description
container	ContainerID	Container where operation is allowed. For container operations, this is the container being operated on. For object operations, this is the container holding the objects. Empty container ID means wildcard (applies to all containers).
verbs	Verb	Operations authorized for this context. Must contain at least one verb (empty list is invalid). Verbs must be sorted in ascending order. Maximum number of verbs: 12.

Message SessionToken

NeoFS Session Token.

Field	Type	Description
body	Body	Session Token contains the proof of trust between peers to be attached in requests for further verification. Please see corresponding section of NeoFS Technical Specification for details.
signature	Signature	Signature of SessionToken information

Message SessionToken.Body

Session Token body

Field	Type	Description
id	bytes	Token identifier is a valid UUIDv4 in binary form
owner_id	OwnerID	Identifier of the session initiator
lifetime	TokenLifetime	Lifetime of the session
session_key	bytes	Public key used in session
object	ObjectSessionContext	ObjectService session context
container	ContainerSessionContext	ContainerService session context

Message SessionToken.Body.TokenLifetime

Lifetime parameters of the token. Field names taken from rfc7519.

Field	Type	Description
exp	uint64	Expiration epoch, the last epoch when token is valid.
nbf	uint64	Not valid before epoch, the first epoch when token is valid.
iat	uint64	Issued at Epoch

Message SessionTokenV2

SessionTokenV2 represents NeoFS Session Token with delegation support.

Field	Type	Description
body	Body	Session token body.
signature	Signature	Signature of the body by the issuer.
origin	SessionTokenV2	Origin token that was delegated to create this token. This creates a chain of trust through token embedding. When B receives a token from A and wants to delegate to C, B creates a new SessionTokenV2 and embeds A's token here.

Delegation validation rules: 1. Lifetime must be within origin's lifetime ($exp \geq origin.exp$, $nbf \leq origin.nbf$). 2. Contexts must be narrowed (see contexts field validation rules). 3. If origin.final is true, delegation is forbidden. 4. Maximum chain depth: 4 tokens. |

Message Target

Target account for SessionTokenV2. It can be either direct (OwnerID) or indirect (NNS domain).

Field	Type	Description
owner_id	OwnerID	Direct account reference via OwnerID (hash of verification script).
nns_name	string	Indirect account reference via NeoFS Name Service. NNS name is a domain name that resolves to an OwnerID through the NeoFS Name Service. The name must be a valid DNS-like domain name (e.g., "example.neofs") that is registered in the NNS contract on the Neo blockchain. The NNS record should contain a string record with the corresponding OwnerID value in NEP-18 address format.

Message TokenLifetime

Lifetime parameters of the token v2. Field names taken from rfc7519.

Field	Type	Description
exp	uint64	Expiration, the last valid Unix timestamp.
nbf	uint64	Not valid before, the first valid Unix timestamp.
iat	uint64	Issued at, the Unix timestamp when the token was issued.

Message XHeader

Extended headers for Request/Response. They may contain any user-defined headers to be interpreted on application level.

Key name must be a unique valid UTF-8 string. Value can't be empty. Requests or Responses with duplicated header names or headers with empty values will be considered invalid.

There are some “well-known” headers starting with `__NEOFS__` prefix that affect system behaviour:

- `__NEOFS__NETMAP_EPOCH`
Netmap epoch to use for object placement calculation. The value is string encoded `uint64` in decimal presentation. If set to '0' or not set, the current epoch only will be used. DEPRECATED: header ignored by servers.
- `__NEOFS__NETMAP_LOOKUP_DEPTH`
If object can't be found using current epoch's netmap, this header limits how many past epochs the node can look up through. The value is string encoded `uint64` in decimal presentation. If set to '0' or not set, only the current epoch will be used. DEPRECATED: header ignored by servers.

Field	Type	Description
key	string	Key of the X-Header
value	string	Value of the X-Header

Emun ContainerSessionContext.Verb

Container request verbs

Number	Name	Description
0	VERB_UNSPECIFIED	Unknown verb
1	PUT	Refers to container.Put RPC call
2	DELETE	Refers to container.Delete RPC call
3	SETEACL	Refers to container.SetExtendedACL RPC call
4	SETATTRIBUTE	Refers to container.SetAttribute RPC call
5	REMOVEATTRIBUTE	Refers to container.RemoveAttribute RPC call

Emun ObjectSessionContext.Verb

Object request verbs

Number	Name	Description
0	VERB_UNSPECIFIED	Unknown verb
1	PUT	Refers to object.Put RPC call
2	GET	Refers to object.Get RPC call
3	HEAD	Refers to object.Head RPC call
4	SEARCH	Refers to object.Search RPC call
5	DELETE	Refers to object.Delete RPC call
6	RANGE	Refers to object.GetRange RPC call
7	RANGEHASH	Refers to object.GetRangeHash RPC call

Emun Verb

Verb represents all possible operations in NeoFS that can be authorized via session tokens or delegation chains. This enum covers both object and container service operations.

Number	Name	Description
0	VERB_UNSPECIFIED	Unknown verb
1	OBJECT_PUT	Refers to object.Put RPC call
2	OBJECT_GET	Refers to object.Get RPC call
3	OBJECT_HEAD	Refers to object.Head RPC call
4	OBJECT_SEARCH	Refers to object.Search RPC call
5	OBJECT_DELETE	Refers to object.Delete RPC call
6	OBJECT_RANGE	Refers to object.GetRange RPC call
7	OBJECT_RANGEHASH	Refers to object.GetRangeHash RPC call
8	CONTAINER_PUT	Refers to container.Put RPC call
9	CONTAINER_DELETE	Refers to container.Delete RPC call
10	CONTAINER_SETEACL	Refers to container.SetExtendedACL RPC call
11	CONTAINER_SETATTRIBUTE	Refers to container.SetAttribute RPC call
12	CONTAINER_REMOVEATTRIBUTE	Refers to container.RemoveAttribute RPC call

neo.fs.v2.status

Message Status

Declares the general format of the status returns of the NeoFS RPC protocol. Status is present in all response messages. Each RPC of NeoFS protocol describes the possible outcomes and details of the operation.

Each status is assigned a one-to-one numeric code. Any unique result of an operation in NeoFS is unambiguously associated with the code value.

Numerical set of codes is split into 1024-element sections. An enumeration is defined for each section. Values can be referred to in the following ways:

- numerical value ranging from 0 to 4,294,967,295 (global code);
- values from enumeration (local code). The formula for the ratio of the local code (L) of a defined section (S) to the global one (G): $G = 1024 * S + L$.

All outcomes are divided into successful and failed, which corresponds to the success or failure of the

operation. The definition of success follows the semantics of RPC and the description of its purpose. The server must not attach code that is the opposite of the outcome type.

See the set of return codes in the description for calls.

Each status can carry a developer-facing error message. It should be a human readable text in English. The server should not transmit (and the client should not expect) useful information in the message. Field details should make the return more detailed.

Field	Type	Description
code	uint32	The status code
message	string	Developer-facing error message
details	Detail	Data detailing the outcome of the operation. Must be unique by ID.

Message Status.Detail

Return detail. It contains additional information that can be used to analyze the response. Each code defines a set of details that can be attached to a status. Client should not handle details that are not covered by the code.

Field	Type	Description
id	uint32	Detail ID. The identifier is required to determine the binary format of the detail and how to decode it.
value	bytes	Binary status detail. Must follow the format associated with ID. The possibility of missing a value must be explicitly allowed.

Emun CommonFail

Section of failed statuses independent of the operation.

Number	Name	Description
0	INTERNAL	[1024] Internal server error, default failure. Not detailed. If the server cannot match failed outcome to the code, it should use this code.
1	WRONG_MAGIC_NUMBER	[1025] Wrong magic of the NeoFS network. Details: - [0] Magic number of the served NeoFS network (big-endian 64-bit unsigned integer).
2	SIGNATURE_VERIFICATION	[1026] Signature verification failure.
3	NODE_UNDER_MAINTENANCE	[1027] Node is under maintenance.
4	BAD_REQUEST	[1028] Malformed, syntactically or semantically incorrect request, client-side error.
5	BUSY	[1029] Some node resources are exhausted and request can't be processed. This is likely to be a transient state and request can be retried in future.

Emun Container

Section of statuses for container-related operations.

Number	Name	Description
0	CONTAINER_NOT_FOUND	[3072] Container not found.
1	EACL_NOT_FOUND	[3073] eACL table not found.
2	CONTAINER_LOCKED	[3074] Operation rejected by the container lock.
3	CONTAINER_AWAIT_TIMEOUT	[3075] Async container operation timed out.

Emun Object

Section of statuses for object-related operations.

Number	Name	Description
0	ACCESS_DENIED	[2048] Access denied by ACL. Details: - [0] Human-readable description (UTF-8 encoded string).
1	OBJECT_NOT_FOUND	[2049] Object not found.
2	LOCKED	[2050] Operation rejected by the object lock.
3	LOCK_NON_REGULAR_OBJECT	[2051] Locking an object with a non-REGULAR type rejected.
4	OBJECT_ALREADY_DELETED	[2052] Object has been marked deleted.
5	OUT_OF_RANGE	[2053] Invalid range has been requested for an object.
6	QUOTA_EXCEEDED	[2054] Quota was exceeded; no further PUTs will be allowed within current settings.

Emun Section

Section identifiers.

Number	Name	Description
0	SECTION_SUCCESS	Successful return codes.
1	SECTION_FAILURE_COMMON	Failure codes regardless of the operation.
2	SECTION_OBJECT	Object service-specific errors.
3	SECTION_CONTAINER	Container service-specific errors.
4	SECTION_SESSION	Session service-specific errors.

Emun Session

Section of statuses for session-related operations.

Number	Name	Description
0	TOKEN_NOT_FOUND	[4096] Token not found.
1	TOKEN_EXPIRED	[4097] Token has expired.

Emun Success

Section of NeoFS successful return codes.

Number	Name	Description
0	OK	[0] Default success. Not detailed. If the server cannot match successful outcome to the code, it should use this code.
1	INCOMPLETE	[1] Incomplete success, only a part of the request done which may or may not be sufficient for the client. Examples are object PUT making 2 out of 3 required replicas or distributed SEARCH getting result from 4 out of 5 nodes.

neo.fs.v2.storagegroup

Message StorageGroup

StorageGroup keeps verification information for Data Audit sessions. Objects that require paid storage guarantees are gathered in StorageGroups with additional information used for the proof of storage. StorageGroup only contains objects from the same container.

Being an object payload, StorageGroup may have expiration Epoch set with `__NEOFS__EXPIRATION_EPOCH` well-known attribute. When expired, StorageGroup will be ignored by InnerRing nodes during Data Audit cycles and will be deleted by Storage Nodes.

DEPRECATED: storage groups are no longer used for audit since 2.18.

Field	Type	Description
validation_data_size	uint64	Total size of the payloads of objects in the storage group
validation_hash	Checksum	Homomorphic hash from the concatenation of the payloads of the storage group members. The order of concatenation is the same as the order of the members in the members field.
expiration_epoch	uint64	DEPRECATED. Last NeoFS epoch number of the storage group lifetime
members	ObjectID	Strictly ordered list of storage group member objects. Members MUST be unique

neo.fs.v2.subnet

Message SubnetInfo

NeoFS subnetwork description

DEPRECATED. Ignored and kept for compatibility only.

Field	Type	Description
id	SubnetID	Unique subnet identifier. Missing ID is equivalent to zero (default subnetwork) ID.
owner	OwnerID	Identifier of the subnetwork owner

neo.fs.v2.tombstone

Message Tombstone

DEPRECATED: delete objects 1-to-1 with `__NEOFS__ASSOCIATE` object attribute in a TOMBSTONE-typed object with no payload. For objects of 2.18+ API version, it is prohibited to have TOMBSTONE objects with payload.

Tombstone keeps record of deleted objects for a few epochs until they are purged from the NeoFS network. It is impossible to delete a tombstone object via `ObjectService.Delete` RPC call.

Field	Type	Description
expiration_epoch	uint64	Last NeoFS epoch number of the tombstone lifetime. It's set by the tombstone creator depending on the current NeoFS network settings. DEPRECATED. Field ignored by servers, set corresponding object attribute <code>__NEOFS__EXPIRATION_EPOCH</code> only.
split_id	bytes	16 byte UUID used to identify the split object hierarchy parts. Must be unique inside a container. All objects participating in the split must have the same <code>split_id</code> value. DEPRECATED. The field is ignored by servers.
members	ObjectID	List of objects to be deleted. IDs should be either: 1. Root object IDs (objects that are not split OR parent objects) 2. Children IDs for unfinished objects that does not have LINK objects (garbage collecting).

Terms and definitions

Object An immutable piece of data with metadata in the form of a set of key-value headers. Object has a globally unique identifier.

Glossary

ACL Access Control List. [12](#), [23](#)

Alphabet nodes Inner Ring nodes that share key with FS chain consensus nodes and send all NeoFS related multisigned transactions. [49](#)

API Application Programming Interface. [9](#)

CLI Command Line Interface. [103](#)

dApp Decentralized Application. [9](#)

EigenTrust EigenTrust algorithm is a reputation management algorithm⁵⁵ for peer-to-peer networks. [35](#), [36](#), [94](#), [186](#)

GAS Utility Token Neo Blockchain's utility token. [12](#), [48](#), [50](#), [52](#), [53](#), [104](#)

Global Trust the result of the **EigenTrust** algorithm is the trust in the network participant, which has been obtained regarding *all Local Trusts* of *all* nodes. [36](#)

HRW HRW stands for Rendezvous hashing⁵⁶. It helps to achieve 3 goals:

1. Select nodes uniformly from the whole netmap. This means that every node has a chance to be included in container nodes set.
2. Select nodes deterministically. Identical (netmap, storage policy) pair results in the same placement set on every storage node.
3. Prioritize nodes providing better conditions.

Nodes having more space, better price or better rating are to be selected with higher probability. Specific weighting algorithm is defined for NeoFS network as a whole and is beyond scope of this document. See NeoFS HRW implementation⁵⁷ for details. [95](#)

⁵⁵<http://ilpubs.stanford.edu:8090/562/1/2002-56.pdf>

⁵⁶https://en.wikipedia.org/wiki/Rendezvous_hashing

⁵⁷<https://github.com/nspcc-dev/hrw>

Local Trust trust of one node to another, calculated using *only* statistical information of their peer-to-peer network interactions. The Subject and Object of such a trust are peer-to-peer nodes .
36, 186

Multiaddress Multiaddress⁵⁸ is a format for encoding addresses from various well-established network protocols. It has two forms:

1. a human-readable version to be used when printing to the user (UTF-8);
2. a binary-packed version to be used in storage, transmissions on the wire, and as a primitive in other formats

. 42

N3 mainnet Main N3 blockchain network. See Neo Documentation⁵⁹ . 11, 49

NEO Token Token representing a share of ownership in the NEO blockchain . 50, 52

Neo Virtual Machine NeoVM is a lightweight virtual machine for executing Neo smart contracts. As the core component of Neo, NeoVM has Turing completeness and high consistency, which can implement arbitrary execution logic and ensure consistent execution results of any node in distributed network, providing strong support for decentralized applications. See Neo Documentation⁶⁰ . 109

NeoFS NeoFS Exaggerative Object File Storage. Also known as **Neo File Storage** . 9

NeoFS Node Computer system with at least following properties:

- Running relevant version of NeoFS software
- Has a NeoFS Network-wide unique identifier and a key pair
- Has good enough connectivity with other nodes
- Serves requests using NeoFS API protocol

. 9, 103, 106

RoleManagement contract native N3 contract that manages list of public keys for specific roles such as *StateValidator*, *Oracle*, *NeoFSAlphabet* . 49

validator node In the NEO network, NEO holders can enroll themselves to be validators (consensus node candidates), and then be voted as consensus nodes. The voting status of validators and number of consensus nodes are stored in blockchain . 50

⁵⁸<https://multiformats.io/multiaddr/>

⁵⁹<https://docs.neo.org/v3/docs/en-us/network/testnet.html>

⁶⁰<https://docs.neo.org/docs/en-us/basic/neovm.html>