
Hyperledger Indy Node Documentation

Hyperledger

Aug 16, 2020

Contents

1	Transactions	1
2	Requests	41
3	Default AUTH_MAP Rules	77
4	Pool Upgrade Guideline	79
5	Create a Network and Start Nodes	83
6	Add Node to Existing Pool	87
7	Helper Scripts	89
8	Setup iptables rules (recommended)	91
9	Node Monitoring Tools for Stewards	93
10	Node State Diagnostic Tools Workflow	97
11	Workflow	101
12	Continuous Integration / Delivery	103
13	Indy File Folders Structure Guideline	111
14	Code quality requirements guideline	115
15	Dev Setup	119
16	List of breaking changes for migration from indy-node 1.3 to 1.4	125

CHAPTER 1

Transactions

- *General Information*
- *Genesis Transactions*
- *Common Structure*
- *Domain Ledger*
 - *NYM*
 - *ATTRIB*
 - *SCHEMA*
 - *CLAIM_DEF*
 - *REVOC_REG_DEF*
 - *REVOC_REG_ENTRY*
 - *JSON_LD_CONTEXT*
 - *RICH_SCHEMA*
 - *RICH_SCHEMA_ENCODING*
 - *RICH_SCHEMA_MAPPING*
 - *RICH_SCHEMA_CRED_DEF*
 - *RICH_SCHEMA_PRES_DEF*
- *Pool Ledger*
 - *NODE*
- *Config Ledger*
 - *POOL_UPGRADE*
 - *NODE_UPGRADE*
 - *POOL_CONFIG*

- *AUTH_RULE*
- *AUTH_RULES*
- *TRANSACTION_AUTHOR_AGREEMENT*
- *TRANSACTION_AUTHOR_AGREEMENT_AML*
- *TRANSACTION_AUTHOR_AGREEMENT_DISABLE*

1.1 General Information

This doc is about supported transactions and their representation on the Ledger (that is, the internal one). If you are interested in the format of a client's request (both write and read), then have a look at [requests](#).

- All transactions are stored in a distributed ledger (replicated on all nodes)
- The ledger is based on a [Merkle Tree](#)
- The ledger consists of two things:
 - transactions log as a sequence of key-value pairs where key is a sequence number of the transaction and value is the serialized transaction
 - merkle tree (where hashes for leaves and nodes are persisted)
- Each transaction has a sequence number (no gaps) - keys in transactions log
- So, this can be considered a blockchain where each block's size is equal to 1
- There are multiple ledgers by default:
 - *pool ledger*: transactions related to pool/network configuration (listing all nodes, their keys and addresses)
 - *config ledger*: transactions for pool configuration plus transactions related to pool upgrade
 - *domain ledger*: all main domain and application specific transactions (including NYM transactions for DID)
- All transactions are serialized to MsgPack format
- All transactions (both transaction log and merkle tree hash stores) are stored in a LevelDB
- One can use the `read_ledger` script to get transactions for a specified ledger in a readable format (JSON)
- See [roles and permissions](#) for a list of roles and they type of transactions they can create.

Below you can find the format and description of all supported transactions.

1.2 Genesis Transactions

As Indy is a public **permissioned** blockchain, each ledger may have a number of pre-defined transactions defining the initial pool and network.

- pool genesis transactions define initial trusted nodes in the pool
- domain genesis transactions define initial trusted trustees and stewards

1.3 Common Structure

Each transaction has the following structure consisting of metadata values (common for all transaction types) and transaction specific data:

```
{
  "ver": <...>,
  "txn": {
    "type": <...>,
    "ver": <...>,
    "protocolVersion": <...>,

    "data": {
      <txn-specific fields>
    },

    "metadata": {
      "reqId": <...>,
      "from": <...>,
      "endorser": <...>,
      "digest": <...>,
      "payloadDigest": <...>,
      "taaAcceptance": {
        "taaDigest": <...>,
        "mechanism": <...>,
        "time": <...>
      }
    },
  },
  "txnMetadata": {
    "txnTime": <...>,
    "seqNo": <...>,
    "txnId": <...>
  },
  "reqSignature": {
    "type": <...>,
    "values": [{
      "from": <...>,
      "value": <...>
    }]
  }
}
```

- **ver (string):**

Transaction version to be able to evolve content. The content of all sub-fields may depend on this version.

- **txn (dict):**

Transaction-specific payload (data)

- **type (enum number as string):**

Supported transaction types:

- * NODE = "0"
- * NYM = "1"
- * TXN_AUTHOR_AGREEMENT = "4"

- * `TXN_AUTHOR_AGREEMENT_AML` = “5”
- * `ATTRIB` = “100”
- * `SCHEMA` = “101”
- * `CLAIM_DEF` = “102”
- * `POOL_UPGRADE` = “109”
- * `NODE_UPGRADE` = “110”
- * `POOL_CONFIG` = “111”
- * `REVOC_REG_DEF` = “113”
- * `REVOC_REG_DEF` = “114”
- * `AUTH_RULE` = “120”
- * `AUTH_RULES` = “122”
- * `JSON_LD_CONTEXT` = “200”
- * `RICH_SCHEMA` = “201”
- * `RICH_SCHEMA_ENCODING` = “202”
- * `RICH_SCHEMA_MAPPING` = “203”
- * `RICH_SCHEMA_CRED_DEF` = “204”
- * `RICH_SCHEMA_PRES_DEF` = “205”

– `ver` (string)

Transaction’s payload version as defined in the input request. If the input request doesn’t have the version specified, then default one will be used. Some transactions (for example `TRANSACTION_AUTHOR_AGREEMENT`) have non-default transaction payload version defined in source code as a result of evolution of business logic and features.

– `protocolVersion` (integer; optional):

The version of client-to-node or node-to-node protocol. Each new version may introduce a new feature in requests/replies/data. Since clients and different nodes may be at different versions, we need this field to support backward compatibility between clients and nodes.

– `data` (dict):

Transaction-specific data fields (see following sections for each transaction’s description).

– `metadata` (dict):

Metadata as came from the request.

- * `from` (base58-encoded string):

Identifier (DID) of the transaction author as base58-encoded string for 16 or 32 bit DID value. It may differ from `endorser` field who submits the transaction on behalf of `identifier`. If `endorser` is absent, then the author (`identifier`) plays the role of endorser and submits request by his own. It also may differ from `dest` field for some of requests (for example `NYM`), where `dest` is a target identifier (for example, a newly created DID identifier).

Example:

- `identifier` is a DID of a transaction author who doesn’t have write permissions; `endorser` is a DID of a user with Endorser role (that is with write permissions).

- new NYM creation: `identifier` is a DID of an Endorser creating a new DID, and `dest` is a newly created DID.
- * `reqId` (integer): Unique ID number of the request with transaction.
- * `digest` (SHA256 hex digest string): SHA256 hash hex digest of all fields in the initial requests (including signatures)
- * `payloadDigest` (SHA256 hex digest string): SHA256 hash hex digest of the payload fields in the initial requests, that is all fields excluding signatures and plugins-added ones
- * `endorser` (base58-encoded string, optional): Identifier (DID) of an Endorser submitting a transaction on behalf of the original author (`identifier`) as base58-encoded string for 16 or 32 bit DID value. If `endorser` is absent, then the author (`identifier`) plays the role of endorser and submits request by his own. If `endorser` is present then the transaction must be multi-signed by the both author (`identifier`) and Endorser (`endorser`).
- * `taaAcceptance` (dict, optional): If transaction author agreement is set/enabled, then every transaction (write request) from Domain and plugins-added ledgers must include acceptance of the latest transaction author agreement.
 - `taaDigest` (SHA256 hex digest string): SHA256 hex digest of the latest Transaction Author Agreement on the ledger. The digest is calculated from concatenation of *TRANSACTION_AUTHOR_AGREEMENT*'s `version` and `text`.
 - `mechanism` (string): a mechanism used to accept the signature; must be present in the latest list of transaction author agreement acceptance mechanisms on the ledger
 - `time` (integer as POSIX timestamp): transaction author agreement acceptance time. The time needs to be rounded to date to prevent correlation of different transactions which is possible when acceptance time is too precise.
- `txnMetadata` (dict):

Metadata attached to the transaction.

 - * `version` (integer): Transaction version to be able to evolve `txnMetadata`. The content of `txnMetadata` may depend on the version.
 - * `txnTime` (integer as POSIX timestamp): The time when transaction was written to the Ledger as POSIX timestamp.
 - * `seqNo` (integer): A unique sequence number of the transaction on Ledger
 - * `txnId` (string, optional): Txn ID as State Trie key (address or descriptive data). It must be unique within the ledger. Usually present for Domain transactions only.
- `reqSignature` (dict):

Submitter's signature over request with transaction (`txn` field).

 - `type` (string enum):
 - * `ED25519`: ed25519 signature
 - * `ED25519_MULTI`: ed25519 signature in multisig case.
 - `values` (list):
 - * `from` (base58-encoded string): Identifier (DID) of signer as base58-encoded string for 16 or 32 byte DID value.
 - * `value` (base58-encoded string): signature value

Please note that all these metadata fields may be absent for genesis transactions.

1.4 Domain Ledger

1.4.1 NYM

Creates a new NYM record for a specific user, endorser, steward or trustee. Note that only trustees and stewards can create new endorsers and a trustee can be created only by other trustees (see [roles](#)).

The transaction can be used for creation of new DIDs, setting and rotation of verification key, setting and changing of roles.

- `dest` (base58-encoded string):

Target DID as base58-encoded string for 16 or 32 byte DID value. It may differ from the `from` metadata field, where `from` is the DID of the submitter. If they are equal (in permissionless case), then transaction must be signed by the newly created `verkey`.

Example: `from` is a DID of a Endorser creating a new DID, and `dest` is a newly created DID.

- `role` (enum number as integer; optional):

Role of a user that the NYM record is being created for. One of the following values

- None (common USER)
- “0” (TRUSTEE)
- “2” (STEWARD)
- “101” (ENDORSER)
- “201” (NETWORK_MONITOR)

A TRUSTEE can change any Nym’s role to None, thus stopping it from making any further writes (see [roles](#)).

- `verkey` (base58-encoded string, possibly starting with “~”; optional):

Target verification key as base58-encoded string. It can start with “~”, which means that it’s an abbreviated verkey and should be 16 bytes long when decoded, otherwise it’s a full verkey which should be 32 bytes long when decoded. If not set, then either the target identifier (`did`) is 32-bit cryptonym CID (this is deprecated), or this is a user under guardianship (doesn’t own the identifier yet). Verkey can be changed to “None” by owner, it means that this user goes back under guardianship.

- `alias` (string; optional):

NYM’s alias.

If there is no NYM transaction for the specified DID (`did`) yet, then this can be considered as the creation of a new DID.

If there is already a NYM transaction with the specified DID (`did`), then this is considered an update of that DID. In this case **only the values that need to be updated should be specified** since any specified one is treated as an update even if it matches the current value in ledger. All unspecified values remain unchanged.

So, if key rotation needs to be performed, the owner of the DID needs to send a NYM request with `did` and `verkey` only. `role` and `alias` will stay the same.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": "1",
    "ver": 1,
```

(continues on next page)

(continued from previous page)

```

    "protocolVersion":2,

    "data": {
      "dest":"GEzcdDLhCpGCYRHW82kjHd",
      "verkey":"~HmUWn928bnFT6Ephf65YXv",
      "role":101,
    },

    "metadata": {
      "reqId":1513945121191691,
      "from":"L5AD5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },
  },
  "txnMetadata": {
    "txnTime":1513945121,
    "seqNo": 10,
    "txnId": "N22KY2Dyvmuu2PyyqSFKue|01"
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  }
}

```

1.4.2 ATTRIB

Adds an attribute to a NYM record

- `dest` (base58-encoded string):

Target DID we set an attribute for as base58-encoded string for 16 or 32 byte DID value. It differs from `from` metadata field, where `from` is the DID of the submitter.

Example: `from` is a DID of a Endorser setting an attribute for a DID, and `dest` is the DID we set an attribute for.

- `raw` (sha256 hash string; mutually exclusive with `hash` and `enc`):

Hash of the raw attribute data. Raw data is represented as JSON, where the key is the attribute name and the value is the attribute value. The ledger only stores a hash of the raw data; the real (unhashed) raw data is stored in a separate attribute store.

- `hash` (sha256 hash string; mutually exclusive with `raw` and `enc`):
Hash of attribute data (as sent by the client). The ledger stores this hash; nothing is stored in an attribute store.
- `enc` (sha256 hash string; mutually exclusive with `raw` and `hash`):
Hash of encrypted attribute data. The ledger contains the hash only; the real encrypted data is stored in a separate attribute store.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": "100",
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "dest": "GEzcdDLhCpGCYRHW82kjHd",
      "raw": "3cbale3cf23c8ce24b7e08171d823fbd9a4929aafd9f27516e30699d3a42026a",
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "N22KY2Dyvmuu2PyyqSFKue|02"
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztggE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd"
↪ ""
    }]
  }
}
```

1.4.3 SCHEMA

Adds a Claim's schema.

It's not possible to update an existing schema. So, if the Schema needs to be evolved, a new Schema with a new version or new name needs to be created.

- data (dict):

Dictionary with Schema's data:

- attr_names: array of attribute name strings
- name: Schema's name string
- version: Schema's version string

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 101,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "data": {
        "attr_names": ["undergrad", "last_name", "first_name", "birth_date",
→ "postgrad", "expiry_date"],
        "name": "Degree",
        "version": "1.0"
      },
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
→ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
→ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
→ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
→ "4X3skpoEK2DRgZxQ9PwuevCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
→ "
    }]
  }
}
```

(continues on next page)

(continued from previous page)

```

    }
}

```

1.4.4 CLAIM_DEF

Adds a claim definition (in particular, public key), that Issuer creates and publishes for a particular claim schema.

- **data (dict):**

Dictionary with claim definition's data:

- **primary (dict):** primary claim public key
- **revocation (dict):** revocation claim public key

- **ref (string):**

Sequence number of a schema transaction the claim definition is created for.

- **signature_type (string):**

Type of the claim definition (that is claim signature). CL (Camenisch-Lysyanskaya) is the only supported type now.

- **tag (string, optional):**

A unique tag to have multiple public keys for the same Schema and type issued by the same DID. A default tag will be used if not specified.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 102,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "data": {
        "primary": {
          ...
        },
        "revocation": {
          ...
        }
      },
      "ref": 12,
      "signature_type": "CL",
      "tag": "some_tag"
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
→ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",

```

(continues on next page)

(continued from previous page)

```

        "payloadDigest":
→ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
        "taaAcceptance": {
            "taaDigest":
→ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
            "mechanism": "EULA",
            "time": 1513942017
        }
    },
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
        "txnId": "HHAD5g65TDQr1PPHHRoiGf2L5AD5g65TDQr1PPHHRoiGf1|Degree1|CL|key1",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
→ "
        }]
    }
}

```

1.4.5 REVOC_REG_DEF

Adds a Revocation Registry Definition, that Issuer creates and publishes for a particular Claim Definition. It contains public keys, maximum number of credentials the registry may contain, reference to the Claim Def, plus some revocation registry specific data.

- value (dict):

Dictionary with revocation registry definition's data:

- maxCredNum (integer): a maximum number of credentials the Revocation Registry can handle
- tailsHash (string): tails' file digest
- tailsLocation (string): tails' file location (URL)
- issuanceType (string enum): defines credentials revocation strategy. Can have the following values:
 - * ISSUANCE_BY_DEFAULT: all credentials are assumed to be issued initially, so that Revocation Registry needs to be updated (REVOC_REG_ENTRY txn sent) only when revoking. Revocation Registry stores only revoked credentials indices in this case. Recommended to use if expected number of revocation actions is less than expected number of issuance actions.
 - * ISSUANCE_ON_DEMAND: no credentials are issued initially, so that Revocation Registry needs to be updated (REVOC_REG_ENTRY txn sent) on every issuance and revocation. Revocation Registry stores only issued credentials indices in this case. Recommended to use if expected number of issuance actions is less than expected number of revocation actions.
- publicKey (dict): Revocation Registry's public key
- id (string): Revocation Registry Definition's unique identifier (a key from state trie is currently used)

- `credDefId` (string): The corresponding Credential Definition's unique identifier (a key from state trie is currently used)
- `revocDefType` (string enum): Revocation Type. `CL_ACCUM` (Camenisch-Lysyanskaya Accumulator) is the only supported type now.
- `tag` (string): A unique tag to have multiple Revocation Registry Definitions for the same Credential Definition and type issued by the same DID.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 113,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "id": "L5AD5g65TDQr1PPHHRoiGf:3:FC4aWomrA13YyvYC1Mxw7:3:CL:14:some_tag:CL_
↪ACCUM:tag1",
      "credDefId": "FC4aWomrA13YyvYC1Mxw7:3:CL:14:some_tag"
      "revocDefType": "CL_ACCUM",
      "tag": "tag1",
      "value": {
        "maxCredNum": 1000000,
        "tailsHash": "6619ad3cf7e02fc29931a5cdc7bb70ba4b9283bda3badae297",
        "tailsLocation": "http://tails.location.com",
        "issuanceType": "ISSUANCE_BY_DEFAULT",
        "publicKeys": {},
      },
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
↪"4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪"21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
↪"6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "L5AD5g65TDQr1PPHHRoiGf:3:FC4aWomrA13YyvYC1Mxw7:3:CL:14:some_tag:CL_
↪ACCUM:tag1",
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
```

(continues on next page)

(continued from previous page)

```

        "value":
    ↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztggE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
    ↪ "
        ]]
    }
}

```

1.4.6 REVOC_REG_ENTRY

The RevocReg entry containing the new accumulator value and issued/revoked indices. This is just a delta of indices, not the whole list. So, it can be sent each time a new claim is issued/revoked.

- value (dict):

Dictionary with revocation registry's data:

- accum (string): the current accumulator value
- prevAccum (string): the previous accumulator value; it's compared with the current value, and txn is rejected if they don't match; it's needed to avoid dirty writes and updates of accumulator.
- issued (list of integers): an array of issued indices (may be absent/empty if the type is ISSUANCE_BY_DEFAULT); this is delta; will be accumulated in state.
- revoked (list of integers): an array of revoked indices (delta; will be accumulated in state)
- revocRegDefId (string): The corresponding Revocation Registry Definition's unique identifier (a key from state trie is currently used)
- revocDefType (string enum): Revocation Type. CL_ACCUM (Camenisch-Lysyanskaya Accumulator) is the only supported type now.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 114,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "revocRegDefId":
    ↪ "L5AD5g65TDQr1PPHHRoiGf:3:FC4aWomrA13YyvYC1Mxw7:3:CL:14:some_tag:CL_ACCUM:tag1"
      "revocDefType": "CL_ACCUM",
      "value": {
        "accum": "accum_value",
        "prevAccum": "prev_acuum_value",
        "issued": [],
        "revoked": [10, 36, 3478],
      },
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
    ↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",

```

(continues on next page)

(continued from previous page)

```

        "payloadDigest":
→ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
        "taaAcceptance": {
            "taaDigest":
→ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
            "mechanism": "EULA",
            "time": 1513942017
        }
    },
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
        "txnId": "5:L5AD5g65TDQr1PPHHRoiGf:3:FC4aWomrA13YyvYC1Mxw7:3:CL:14:some_tag:CL_
→ ACCUM:tag1",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
→ "
        }]
    }
}

```

1.4.7 JSON_LD_CONTEXT

Adds a JSON LD Context as part of Rich Schema feature.

It's not possible to update an existing Context. If the Context needs to be evolved, a new Context with a new id and name-version needs to be created.

- **id (string):**
A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the content field)
- **content (json-serialized string):**
Context object as JSON serialized in canonical form. It must have @context as a top level key. The @context value must be either:
 1. a URI (it should dereference to a Context object)
 2. a Context object (a dict)
 3. an array of Context objects and/or Context URIs
- **rsType (string):**
Context's type. Currently expected to be `ctx`.
- **rsName (string):**
Context's name
- **rsVersion (string):**

Context's version

rsType, rsName and rsVersion must be unique among all rich schema objects on the ledger.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 200,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "id": "did:sov:GGAD5g65TDQr1PPHHRoiGf",
      "content": "{
        "@context": [
          {
            "@version": 1.1
          },
          "https://www.w3.org/ns/odrl.jsonld",
          {
            "ex": "https://example.org/examples#",
            "schema": "http://schema.org/",
            "rdf": "http://www.w3.org/1999/02/22-rdf-syntax-ns#"
          }
        ]
      }",
      "rsName": "SimpleContext",
      "rsVersion": "1.0",
      "rsType": "ctx"
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
        ↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
        ↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
          ↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },

    "txnMetadata": {
      "txnTime": 1513945121,
      "seqNo": 10,
      "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
    },
    "reqSignature": {
      "type": "ED25519",
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
          ↪ "4X3skpoEK2DRgzxQ9PwueVCJpL8JHdQ8X4HDDFyztggEI5DM2ZnkvrAh9bQYI6egVinZT7(continues on next page)yDgd"
      }],
    },
  },
}
```

(continued from previous page)

```

    }
  }
}

```

1.4.8 RICH_SCHEMA

Adds a Rich Schema object as part of Rich Schema feature.

It's not possible to update an existing Rich Schema. If the Rich Schema needs to be evolved, a new Rich Schema with a new id and name-version needs to be created.

- `id` (string):
A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the content field)
- `content` (json-serialized string):
Rich Schema object as JSON serialized in canonical form. This value must be a json-ld, rich schema object. json-ld supports many parameters that are optional for a rich schema txn. However, the following parameters must be there:
 - `@id`: The value of this property must be (or map to, via a context object) a URI.
 - `@type`: The value of this property must be (or map to, via a context object) a URI.
 - `@context`(optional): If present, the value of this property must be a context object or a URI which can be dereferenced to obtain a context object.
- `rsType` (string):
Rich Schema's type. Currently expected to be `sch`.
- `rsName` (string):
Rich Schema's name
- `rsVersion` (string):
Rich Schema's version

`rsType`, `rsName` and `rsVersion` must be unique among all rich schema objects on the ledger.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 201,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "id": "did:sov:HGAD5g65TDQr1PPHHRoiGf",
      "content": "{
        \"@id\": \"test_unique_id\",
        \"@context\": \"ctx:sov:2f9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD\",
        \"@type\": \"rdfs:Class\",
        \"rdfs:comment\": \"ISO18013 International Driver License\",
        \"rdfs:label\": \"Driver License\",
      }"
    }
  }
}

```

(continues on next page)

(continued from previous page)

```

        "rdfs:subClassOf": {
            "@id": "sch:Thing"
        },
        "driver": "Driver",
        "dateOfIssue": "Date",
        "dateOfExpiry": "Date",
        "issuingAuthority": "Text",
        "licenseNumber": "Text",
        "categoriesOfVehicles": {
            "vehicleType": "Text",
            "vehicleType-input": {
                "@type": "sch:PropertyValueSpecification",
                "valuePattern": "^(A|B|C|D|BE|CE|DE|AM|A1|A2|B1|C1|D1|C1E|D1E)
→ $"
            },
            "dateOfIssue": "Date",
            "dateOfExpiry": "Date",
            "restrictions": "Text",
            "restrictions-input": {
                "@type": "sch:PropertyValueSpecification",
                "valuePattern": "^[A-Z][1-9]) $"
            }
        },
        "administrativeNumber": "Text"
    },
    "rsName": "SimpleRichSchema",
    "rsVersion": "1.0",
    "rsType": "sch"
},

    "metadata": {
        "reqId": 1513945121191691,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "endorser": "D6HG5g65TDQr1PPHHRoiGf",
        "digest":
→ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "payloadDigest":
→ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
        "taaAcceptance": {
            "taaDigest":
→ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
            "mechanism": "EULA",
            "time": 1513942017
        }
    },
},
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
        "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
→ "

```

(continues on next page)

(continued from previous page)

```

    }
  }
}

```

1.4.9 RICH_SCHEMA_ENCODING

Adds an Encoding object as part of Rich Schema feature.

It's not possible to update an existing Encoding. If the Encoding needs to be evolved, a new Encoding with a new id and name-version needs to be created.

- `id` (string):

A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the content field)

- `content` (json-serialized string):

Encoding object as JSON serialized in canonical form.

- `input`: a description of the input value
- `output`: a description of the output value
- `algorithm`:

- * `documentation`: a URL which references a specific github commit of the documentation that fully describes the transformation algorithm.

- * `implementation`: a URL that links to a reference implementation of the transformation algorithm. It is not necessary to use the implementation linked to here, as long as the implementation used implements the same transformation algorithm.

- * `description`: a brief description of the transformation algorithm.

- `testVectors`: a URL which references a specific github commit of a selection of test vectors that may be used to provide assurance that a transformation algorithm implementation is correct.

- `rsType` (string):

Encoding's type. Currently expected to be `enc`.

- `rsName` (string):

Encoding's name

- `rsVersion` (string):

Encoding's version

`rsType`, `rsName` and `rsVersion` must be unique among all rich schema objects on the ledger.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 202,
    "ver": 1,
    "protocolVersion": 2,

    "data": {

```

(continues on next page)

(continued from previous page)

```

    "id": "did:sov:HGAD5g65TDQr1PPHHRoiGf",
    "content": {
      "input": {
        "id": "DateRFC3339",
        "type": "string"
      },
      "output": {
        "id": "UnixTime",
        "type": "256-bit integer"
      },
      "algorithm": {
        "description": "This encoding transforms an
          RFC3339-formatted datetime object into the number
          of seconds since January 1, 1970 (the Unix epoch).",
        "documentation": "URL to specific github commit",
        "implementation": "URL to implementation"
      },
      "test_vectors": "URL to specific github commit"
    },
    "rsName": "SimpleEncoding",
    "rsVersion": "1.0",
    "rsType": "enc"
  },

  "metadata": {
    "reqId": 1513945121191691,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
    "endorser": "D6HG5g65TDQr1PPHHRoiGf",
    "digest":
    ↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
    "payloadDigest":
    ↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    "taaAcceptance": {
      "taaDigest":
    ↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "mechanism": "EULA",
      "time": 1513942017
    }
  },

  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
    ↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
    ↪ "
    }]
  }
}

```

1.4.10 RICH_SCHEMA_MAPPING

Adds a Mapping object as part of Rich Schema feature.

It's not possible to update an existing Mapping. If the Mapping needs to be evolved, a new Mapping with a new id and name-version needs to be created.

- `id` (string):

A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the `content` field)

- `content` (json-serialized string):

Mapping object as JSON serialized in canonical form. This value must be a json-ld object. json-ld supports many parameters that are optional for a rich schema txn. However, the following parameters must be there:

- `@id`: The value of this property must be (or map to, via a context object) a URI.
- `@type`: The value of this property must be (or map to, via a context object) a URI.
- `@context`(optional): If present, the value of this property must be a context object or a URI which can be dereferenced to obtain a context object.
- `schema`: An id of the corresponding Rich Schema
- `attributes` (dict): A dict of all the schema attributes the Mapping object is going to map to encodings and use in credentials. An attribute may have nested attributes matching the schema structure. It must also contain the following default attributes required by any W3C compatible verifiable credential (plus any additional attributes that may have been included from the W3C verifiable credentials data model):

- * `issuer`
- * `issuanceDate`
- * any additional attributes

Every leaf attribute's value is an array of the following pairs:

- * `enc` (string): Encoding object (referenced by its `id`) to be used for representation of the attribute as an integer.
- * `rank` (int): Rank of the attribute to define the order in which the attribute is signed by the Issuer. It is important that no two `rank` values may be identical.

- `rsType` (string):

Mapping's type. Currently expected to be `map`.

- `rsName` (string):

Mapping's name

- `rsVersion` (string):

Mapping's version

`rsType`, `rsName` and `rsVersion` must be unique among all rich schema objects on the ledger.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 203,
    "ver": 1,
```

(continues on next page)

(continued from previous page)

```

    "protocolVersion":2,

    "data": {
      "id": "did:sov:HGAD5g65TDQr1PPHHRoiGf",
      "content": "{
        '@id': "did:sov:5e9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
        '@context': "did:sov:2f9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
        '@type': "rdfs:Class",
        "schema": "did:sov:4e9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
        "attribuities" : {
          "issuer": [{
            "enc":
↪ "did:sov:9x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 1
          }],
          "issuanceDate": [{
            "enc":
↪ "did:sov:119F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 2
          }],
          "expirationDate": [{
            "enc":
↪ "did:sov:119F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 11
          }],
          "driver": [{
            "enc":
↪ "did:sov:1x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 5
          }],
          "dateOfIssue": [{
            "enc":
↪ "did:sov:2x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 4
          }],
          "issuingAuthority": [{
            "enc":
↪ "did:sov:3x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 3
          }],
          "licenseNumber": [
            {
              "enc":
↪ "did:sov:4x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
              "rank": 9
            },
            {
              "enc":
↪ "did:sov:5x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
              "rank": 10
            },
          ],
          "categoriesOfVehicles": {
            "vehicleType": [{
              "enc":
↪ "did:sov:6x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
              "rank": 6
            }

```

(continues on next page)

(continued from previous page)

```

        },
        "dateOfIssue": [{
            "enc":
↪ "did:sov:7x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
            "rank": 7
        }],
    },
    "administrativeNumber": [{
        "enc":
↪ "did:sov:8x9F8ZmxuvDqRiqqY29x6dx9oU4qwFTkPbDpWtwGbdUsrCD",
        "rank": 8
    }]
}
},
"rsName": "SimpleMapping",
"rsVersion": "1.0",
"rsType": "map"
},

"metadata": {
    "reqId": 1513945121191691,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
    "endorser": "D6HG5g65TDQr1PPHHRoiGf",
    "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
    "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    "taaAcceptance": {
        "taaDigest":
↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
    }
},
},
"txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
},
"reqSignature": {
    "type": "ED25519",
    "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
}
}

```

1.4.11 RICH_SCHEMA_CRED_DEF

Adds a Credential Definition object as part of Rich Schema feature.

Credential Definition is considered as a mutable object as the Issuer may rotate keys present there. However, rotation

of Issuer's keys should be done carefully as it will invalidate all credentials issued for this key.

- `id` (string):

A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the content field)

- `content` (json-serialized string):

Credential Definition object as JSON serialized in canonical form.

- `signatureType` (string): Type of the ZKP signature. CL (Camenisch-Lysyanskaya) is the only supported type now.
- `mapping` (string): An id of the corresponding Mapping
- `schema` (string): An id of the corresponding Rich Schema. The mapping must reference the same Schema.
- `publicKey` (dict): Issuer's public keys. Consists of primary and revocation keys.
 - * `primary` (dict): primary key
 - * `revocation` (dict, optional): revocation key

- `rsType` (string):

Credential Definition's type. Currently expected to be `cdf`.

- `rsName` (string):

Credential Definition's name

- `rsVersion` (string):

Credential Definition's version

`rsType`, `rsName` and `rsVersion` must be unique among all rich schema objects on the ledger.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 204,
    "ver": 1,
    "protocolVersion": 2,
    "data": {
      "id": "did:sov:HGAD5g65TDQr1PPHHRoiGf",
      "content": "{
        \"signatureType\": \"CL\",
        \"mapping\": \"did:sov:UVj5w8DRzcmPVDpUMr4AZhJ\",
        \"schema\": \"did:sov:U5x5w8DRzcmPVDpUMr4AZhJ\",
        \"publicKey\": {
          \"primary\": \"...\",
          \"revocation\": \"...\"
        }
      }",
      "rsName": "SimpleCredDef",
      "rsVersion": "1.0",
      "rsType": "cdf"
    },
  },
}
```

(continues on next page)

(continued from previous page)

```

        "metadata": {
            "reqId": 1513945121191691,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "endorser": "D6HG5g65TDQr1PPHHRoiGf",
            "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
            "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bcc2e0106cd905685",
            "taaAcceptance": {
                "taaDigest":
↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
                "mechanism": "EULA",
                "time": 1513942017
            }
        },
    },
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
        "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    }
}

```

1.4.12 RICH_SCHEMA_PREDEF

Adds a Presentation Definition object as part of Rich Schema feature.

Presentation Definition is considered as a mutable object since restrictions to Issuers, Schemas and Credential Definitions to be used in proof may evolve. For example, Issuer's key for a given Credential Definition may be compromised, so Presentation Definition can be updated to exclude this Credential Definition from the list of recommended ones.

- `id` (string):
A unique ID (for example a DID with a id-string being base58 representation of the SHA2-256 hash of the content field)
- `content` (json-serialized string):
Presentation Definition object as JSON serialized in canonical form.
- `rsType` (string):
Presentation Definition's type. Currently expected to be `pdf`.
- `rsName` (string):
Presentation Definition's name
- `rsVersion` (string):

Presentation Definition's version

rsType, rsName and rsVersion must be unique among all rich schema objects on the ledger.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 205,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "id": "did:sov:HGAD5g65TDQr1PPHHRoiGf",
      "content": "{
        TBD
      }",
      "rsName": "SimplePresDef",
      "rsVersion": "1.0",
      "rsType": "pdf"
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "endorser": "D6HG5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      "taaAcceptance": {
        "taaDigest":
↪ "6sh15d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "mechanism": "EULA",
        "time": 1513942017
      }
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
    "txnId": "L5AD5g65TDQr1PPHHRoiGf1|Degree|1.0",
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  }
}
```

1.5 Pool Ledger

1.5.1 NODE

Adds a new node to the pool or updates an existing node in the pool

- `data` (dict):

Data associated with the Node:

- `alias` (string): Node's alias
- `blskey` (base58-encoded string; optional): BLS multi-signature key as base58-encoded string (it's needed for BLS signatures and state proofs support)
- `client_ip` (string; optional): Node's client listener IP address, that is the IP clients use to connect to the node when sending read and write requests (ZMQ with TCP)
- `client_port` (string; optional): Node's client listener port, that is the port clients use to connect to the node when sending read and write requests (ZMQ with TCP)
- `node_ip` (string; optional): The IP address other Nodes use to communicate with this Node; no clients are allowed here (ZMQ with TCP)
- `node_port` (string; optional): The port other Nodes use to communicate with this Node; no clients are allowed here (ZMQ with TCP)
- `services` (array of strings; optional): the service of the Node. `VALIDATOR` is the only supported one now.

- `dest` (base58-encoded string):

Target Node's verkey as base58-encoded string for 16 or 32 byte DID value. It differs from `identifier` metadata field, where `identifier` is the DID of the transaction submitter (Steward's DID).

Example: `identifier` is a DID of a Steward creating a new Node, and `dest` is the verkey of this Node.

If there is no NODE transaction with the specified Node ID (`dest`), then it can be considered as creation of a new NODE.

If there is a NODE transaction with the specified Node ID (`dest`), then this is update of existing NODE. In this case we can specify only the values we would like to override. All unspecified values remain the same. So, if a Steward wants to rotate BLS key, then it's sufficient to send a NODE transaction with `dest` and a new `blskey`. There is no need to specify all other fields, and they will remain the same.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 0,
    "protocolVersion": 2,

    "data": {
      "data": {
        "alias": "Delta",
        "blskey":
↪ "4kkk7y7NQVzcfvY4SAelHBMynFohAJ2ygLeJd3nC77SFv2mJAmebH3BGbrGPHamLZMAFWQJNHEM81P62RfZjnb5SER6cQk1MNN
↪ ",
        "client_ip": "127.0.0.1",
        "client_port": 7407,

```

(continues on next page)

(continued from previous page)

```

        "node_ip": "127.0.0.1",
        "node_port": 7406,
        "services": ["VALIDATOR"]
    },
    "dest": "4yC546FFzorLPgTNTc6V43DnpFrR8uHvtunBxb2Suaa2",
},
{
    "metadata": {
        "reqId": 1513945121191691,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
        "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bcc2e0106cd905685",
    },
},
{
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztggE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    }
}
}

```

1.6 Config Ledger

1.6.1 POOL_UPGRADE

Command to upgrade the Pool (sent by Trustee). It upgrades the specified Nodes (either all nodes in the Pool, or some specific ones).

- **name (string):**
Human-readable name for the upgrade.
- **action (enum: start or cancel):**
Starts or cancels the Upgrade.
- **version (string):**
The version of indy-node package we perform upgrade to. Must be greater than existing one (or equal if `reinstall` flag is `True`).
- **schedule (dict of node DIDs to timestamps):**
Schedule of when to perform upgrade on each node. This is a map where Node DIDs are keys, and upgrade time is a value (see example below). If `force` flag is `False`, then it's required that time difference between each Upgrade must be not less than 5 minutes (to give each Node enough time and not make the whole Pool go down during Upgrade).

- `sha256` (sha256 hash string):

sha256 hash of the package

- `force` (boolean; optional):

Whether we should apply transaction (schedule Upgrade) without waiting for consensus of this transaction. If false, then transaction is applied only after it's written to the ledger. Otherwise it's applied regardless of result of consensus, and there are no restrictions on the Upgrade `schedule` for each Node. So, we can Upgrade the whole Pool at the same time when it's set to True. False by default. Avoid setting to True without good reason.

- `reinstall` (boolean; optional):

Whether it's allowed to re-install the same version. False by default.

- `timeout` (integer; optional):

Limits upgrade time on each Node.

- `justification` (string; optional):

Optional justification string for this particular Upgrade.

Example:

```
{
  "ver": 1,
  "txn": {
    "type": 109,
    "ver": 1,
    "protocolVersion": 2,
    "data": {
      "name": "upgrade-13",
      "action": "start",
      "version": "1.3",
      "schedule": {
        "4yC546FFzorLPgTNTc6V43DnpFrR8uHvtunBxb2Suaa2": "2017-12-
↪ 25T10:25:58.271857+00:00",
        "AtDfpKFe1RPgcr5nnYBw1Wxkgyn8Zjyh5MzFoEUTeoV3": "2017-12-
↪ 25T10:26:16.271857+00:00",
        "DG5M4zFm33Shrhjj6JB7nmX9BoNJUq219UXDfvwBDPe2": "2017-12-
↪ 25T10:26:25.271857+00:00",
        "JpYerf4CssDrH76z7jyQPJLnZ1vwYgvKbvcp16AB5RQ": "2017-12-
↪ 25T10:26:07.271857+00:00"},
      "sha256": "db34a72a90d026dae49c3b3f0436c8d3963476c77468ad955845a1ccf7b03f55
↪ ",
      "force": false,
      "reinstall": false,
      "timeout": 1,
      "justification": null,
    },
    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    "reqSignature": {
      "type": "ED25519",
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
      ]}]
    }
  }
}

```

1.6.2 NODE_UPGRADE

Status of each Node's upgrade (sent by each upgraded Node)

- action (enum string):
One of in_progress, complete or fail.
- version (string):
The version of indy-node the node was upgraded to.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 110,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "action": "complete",
      "version": "1.2"
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
  }
}

```

(continues on next page)

(continued from previous page)

```

    }
  }
}

```

1.6.3 POOL_CONFIG

Command to change Pool's configuration

- `writes` (boolean):

Whether any write requests can be processed by the pool (if false, then pool goes to read-only state). True by default.

- `force` (boolean; optional):

Whether we should apply transaction (for example, move pool to read-only state) without waiting for consensus of this transaction. If false, then transaction is applied only after it's written to the ledger. Otherwise it's applied regardless of result of consensus. False by default. Avoid setting to True without good reason.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 111,
    "ver": 1,
    "protocolVersion": 2,

    "data": {
      "writes": false,
      "force": true,
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest":
↪ "4ba05d9b2c27e52aa8778708fb4b3e5d7001eecd02784d8e311d27b9090d9453",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
  },
  "txnMetadata": {
    "txnTime": 1513945121,
    "seqNo": 10,
  },
  "reqSignature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd"
↪ ""
    }]
  }
}

```

1.6.4 AUTH_RULE

A command to change authentication rules. Internally authentication rules are stored as a key-value dictionary: {action} -> {auth_constraint}.

The list of actions is static and can be found in [auth_rules.md](#). There is a default Auth Constraint for every action (defined in [auth_rules.md](#)).

The AUTH_RULE command allows to change the Auth Constraint. So, it's not possible to register new actions by this command. But it's possible to override authentication constraints (values) for the given action.

Please note, that list elements of GET_AUTH_RULE output can be used as an input (with a required changes) for AUTH_RULE.

The following input parameters must match an auth rule from the [auth_rules.md](#):

- `auth_type` (string enum)

The type of transaction to change the auth constraints to. (Example: "0", "1", ...). See transactions description to find the txn type enum value.

- `auth_action` (enum: ADD or EDIT)

Whether this is adding of a new transaction, or editing of an existing one.

- `field` (string)

Set the auth constraint of editing the given specific field. * can be used to specify that an auth rule is applied to all fields.

- `old_value` (string; optional)

Old value of a field, which can be changed to a new_value. Must be present for EDIT auth_action only. * can be used if it doesn't matter what was the old value.

- `new_value` (string)

New value that can be used to fill the field. * can be used if it doesn't matter what was the old value.

The `constraint_id` fields is where one can define the desired auth constraint for the action:

- `constraint` (dict)

- `constraint_id` (string enum)

Constraint Type. As of now, the following constraint types are supported:

```
- "ROLE": a constraint defining how many signatures of a given role are
↪required
- "OR": logical disjunction for all constraints from `auth_constraints`
- "AND": logical conjunction for all constraints from `auth_constraints`
```

- fields if "constraint_id": "OR" or "constraint_id": "AND"

- `auth_constraints` (list)

A list of constraints. Any number of nested constraints is supported recursively

- fields if "constraint_id": "ROLE":

- `role` (string enum)

Who (what role) can perform the action Please have a look at [NYM](#) transaction description for a mapping between role codes and names.

* `sig_count` (int):

The number of signatures that is needed to do the action

* `need_to_be_owner` (boolean):

Flag to check if the user must be the owner of a transaction (Example: A steward must be the owner of the node to make changes to it). The notion of the `owner` is different for every auth rule. Please reference to [auth_rules.md](#) for details.

* `off_ledger_signature` (boolean, optional, False by default):

Whether signatures against keys not present on the ledger are accepted during verification of required number of valid signatures. An example when it can be set to `True` is creation of a new DID in a permissionless mode, that is when `identifier` is not present on the ledger and a newly created `verkey` is used for signature verification. Another example is signing by cryptonyms (where `identifier` is equal to `verkey`), but this is not supported yet. If the value of this field is `False` (default), and the number of required signatures is greater than zero, then the transaction author's DID (`identifier`) must be present on the ledger (corresponding NYM txn must exist).

* `metadata` (dict; optional):

Dictionary for additional parameters of the constraint. Can be used by plugins to add additional restrictions.

– fields if "`constraint_id`": "`FORBIDDEN`":

no fields

Example:

Let's consider an example of changing a value of a NODE transaction's `service` field from `[VALIDATOR]` to `[]` (demotion of a node). We are going to set an Auth Constraint, so that the action can be only be done by two TRUSTEE or one STEWARD who is the owner (the original creator) of this transaction.

```
{
  "txn": {
    "type": "120",
    "protocolVersion": 2,
    "data": {
      "auth_type": "0",
      "auth_action": "EDIT",
      "field": "services",
      "old_value": [VALIDATOR],
      "new_value": [],
      "constraint": {
        "constraint_id": "OR",
        "auth_constraints": [
          {
            "constraint_id": "ROLE",
            "role": "0",
            "sig_count": 2,
            "need_to_be_owner": False,
            "metadata": {}
          },
          {
            "constraint_id": "ROLE",
            "role": "2",
            "sig_count": 1,
            "need_to_be_owner": True,
            "metadata": {}
          }
        ]
      }
    }
  },
}
```

(continues on next page)

(continued from previous page)

```

    "metadata":{
      "reqId":252174114,
      "from":"M9BJDuS24bqbJNvBRsoGg3",
      "digest":"6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    }
  },
  "txnMetadata":{
    "txnTime":1551785798,
    "seqNo":1
  },
  "reqSignature":{
    "type":"ED25519",
    "values":[
      {
        "value":
↪ "4wpLLAtkT6SeiKEXPVsMcCirx9KvkeKKd11Q4VsMXmSv2tnJrRw1TQKFyov4m2BuPP4C5oCiZ6RUwS9w3EPdywnz
↪ ",
        "from":"M9BJDuS24bqbJNvBRsoGg3"
      }
    ]
  },
  "ver":"1"
}

```

1.6.5 AUTH_RULES

A command to set multiple AUTH_RULES by one transaction. Transaction AUTH_RULES is not divided into a few AUTH_RULE transactions, and is written to the ledger with one transaction with the full set of rules that come in the request. Internally authentication rules are stored as a key-value dictionary: {action} -> {auth_constraint}.

The list of actions is static and can be found in [auth_rules.md](#). There is a default Auth Constraint for every action (defined in [auth_rules.md](#)).

The AUTH_RULES command allows to change the Auth Constraints. So, it's not possible to register new actions by this command. But it's possible to override authentication constraints (values) for the given action.

Please note, that list elements of GET_AUTH_RULE output can be used as an input (with a required changes) for the field rules in AUTH_RULES.

- The `rules` field contains a list of auth rules. One rule has the following list of parameters which must match an auth rule from the [auth_rules.md](#):
 - `auth_type` (string enum)

The type of transaction to change the auth constraints to. (Example: "0", "1", ...). See transactions description to find the txn type enum value.
 - `auth_action` (enum: ADD or EDIT)

Whether this is adding of a new transaction, or editing of an existing one.
 - `field` (string)

Set the auth constraint of editing the given specific field. * can be used to specify that an auth rule is applied to all fields.

- `old_value` (string; optional)

Old value of a field, which can be changed to a `new_value`. Must be present for `EDIT` `auth_action` only. `*` can be used if it doesn't matter what was the old value.

- `new_value` (string)

New value that can be used to fill the field. `*` can be used if it doesn't matter what was the old value.

The `constraint_id` fields is where one can define the desired auth constraint for the action:

- `constraint` (dict)

- * `constraint_id` (string enum)

Constraint Type. As of now, the following constraint types are supported:

```
- "ROLE": a constraint defining how many signatures of a given role_
↪are required
- "OR": logical disjunction for all constraints from `auth_constraints`
- "AND": logical conjunction for all constraints from `auth_constraints`
- "FORBIDDEN": a constraint for not allowed actions
```

- * fields if `"constraint_id": "OR"` or `"constraint_id": "AND"`

- `auth_constraints` (list)

A list of constraints. Any number of nested constraints is supported recursively

- * fields if `"constraint_id": "ROLE"`:

- `role` (string enum)

Who (what role) can perform the action Please have a look at [NYM](#) transaction description for a mapping between role codes and names.

- `sig_count` (int):

The number of signatures that is needed to do the action

- `need_to_be_owner` (boolean):

Flag to check if the user must be the owner of a transaction (Example: A steward must be the owner of the node to make changes to it). The notion of the `owner` is different for every auth rule. Please reference to [auth_rules.md](#) for details.

- `off_ledger_signature` (boolean, optional, False by default):

Whether signatures against keys not present on the ledger are accepted during verification of required number of valid signatures. An example when it can be set to `True` is creation of a new DID in a permissionless mode, that is when `identifier` is not present on the ledger and a newly created `verkey` is used for signature verification. Another example is signing by cryptonyms (where `identifier` is equal to `verkey`), but this is not supported yet. If the value of this field is `False` (default), and the number of required signatures is greater than zero, then the transaction author's DID (`identifier`) must be present on the ledger (corresponding NYM txn must exist).

- `metadata` (dict; optional):

Dictionary for additional parameters of the constraint. Can be used by plugins to add additional restrictions.

- * fields if `"constraint_id": "FORBIDDEN"`:

no fields

Example:

```

{
  "txn":{
    "type":"120",
    "protocolVersion":2,
    "data":{
      rules: [
        {
          "auth_type": "0",
          "auth_action": "EDIT",
          "field": "services",
          "old_value": [VALIDATOR],
          "new_value": []
          "constraint":{
            "constraint_id": "OR",
            "auth_constraints": [{
              "constraint_id": "ROLE",
              "role": "0",
              "sig_count": 2,
              "need_to_be_owner": False,
              "metadata": {}},
              {
                "constraint_id": "ROLE",
                "role": "2",
                "sig_count": 1,
                "need_to_be_owner": True,
                "metadata": {}
              }
            ]
          }
        },
        ...
      ]
      "metadata":{
        "reqId":252174114,
        "from":"M9BJDuS24bqbJNvBRsoGg3",
        "digest":"6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c",
        "payloadDigest":
        ↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
      }
    },
    "txnMetadata":{
      "txnTime":1551785798,
      "seqNo":1
    },
    "reqSignature":{
      "type":"ED25519",
      "values":[
        {
          "value":
          ↪ "4wpLLAtkT6SeiKEXPVsmcCirx9KvkeKKd11Q4VsMXmSv2tnJrRw1TQKFyov4m2BuPP4C5oCiZ6RUwS9w3EPdywnz
          ↪ ",
          "from":"M9BJDuS24bqbJNvBRsoGg3"
        }
      ]
    },
    "ver":"1"
  }
}

```

1.6.6 TRANSACTION_AUTHOR_AGREEMENT

Setting (enabling/disabling) a transaction author agreement for the pool.

If transaction author agreement is set, then all write requests to Domain ledger (transactions) must include additional metadata pointing to the latest transaction author agreement's digest which is signed by the transaction author. If no transaction author agreement is set, or there are no active transaction author agreements, then no additional metadata is required.

Each transaction author agreement has a unique version. If TRANSACTION_AUTHOR_AGREEMENT transaction is sent for already existing version it is considered an update (for example of retirement timestamp), in this case text and ratification timestamp should be either absent or equal to original values.

For any given version of transaction author agreement text and ratification timestamp cannot be changed once set. Ratification timestamp cannot be in future. In order to update Transaction Author Agreement TRANSACTION_AUTHOR_AGREEMENT transaction should be sent, containing new version and new text of agreement. This makes it possible to use new Transaction Author Agreement, but doesn't disable previous one automatically.

Individual transaction author agreements can be disabled by setting retirement timestamp using same transaction. Retirement timestamp can be in future, in this case deactivation of agreement won't happen immediately, it will be automatically deactivated at required time instead.

It is possible to change existing retirement timestamp of agreement by sending a TRANSACTION_AUTHOR_AGREEMENT transaction with a new retirement timestamp. This may potentially re-enable already retired Agreement. Re-enabling retired Agreement needs to be considered as an exceptional case used mostly for fixing disabling by mistake or with incorrect retirement timestamp specified.

It is possible to delete retirement timestamp of agreement by sending a TRANSACTION_AUTHOR_AGREEMENT transaction without a retirement timestamp or retirement timestamp set to `None`. This will either cancel retirement (if it hasn't occurred yet), or disable retirement of already retired transaction (re-enable the Agreement). Re-enabling retired Agreement needs to be considered as an exceptional case used mostly for fixing disabling by mistake or with incorrect retirement timestamp specified.

Latest transaction author agreement cannot be disabled using this transaction.

It is possible to disable all currently active transaction author agreements (including latest) using separate transaction *TRANSACTION_AUTHOR_AGREEMENT_DISABLE*. This will immediately set current timestamp as retirement one for all not yet retired Transaction Author Agreements.

It's not possible to re-enable an Agreement right after disabling all agreements because there is no active latest Agreement at this point. A new Agreement needs to be sent instead.

At least one *TRANSACTION_AUTHOR_AGREEMENT_AML* must be set on the ledger before submitting TRANSACTION_AUTHOR_AGREEMENT txn.

- `version` (string):

Unique version of the transaction author agreement

- `text` (string; optional):

Transaction author agreement's text. Must be specified when creating a new Agreement. Should be either omitted or equal to existing value in case of updating an existing Agreement (setting `retirement_ts`).

- `ratification_ts` (integer as POSIX timestamp; optional):

Transaction Author Agreement ratification timestamp as POSIX timestamp. May have any precision up to seconds. Must be specified when creating a new Agreement. Should be either omitted or equal to existing value in case of updating an existing Agreement (setting `retirement_ts`).

- `retirement_ts` (integer as POSIX timestamp; optional):

Transaction Author Agreement retirement timestamp as POSIX timestamp. May have any precision up to seconds. Can be any timestamp either in future or in the past (the Agreement will be retired immediately in the latter case). Must be omitted when creating a new (latest) Agreement. Should be used for updating (deactivating) non-latest Agreement on the ledger.

New Agreement Example:

```
{
  "ver": 1,
  "txn": {
    "type": 4,
    "ver": 2,
    "protocolVersion": 2,

    "data": {
      "version": "1.0",
      "text": "Please read carefully before writing anything to the ledger",
      "ratification_ts": 1514304094738044
    },

    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest": "6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c
↪ ",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
    "txnMetadata": {
      "txnTime": 1577836799,
      "seqNo": 10,
    },
    "reqSignature": {
      "type": "ED25519",
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
↪ "4X3skpoEK2DRgZxQ9PwEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
      ]
    }
  }
}
```

Retire Agreement Example:

```
{
  "ver": 2,
  "txn": {
    "type": 4,
    "protocolVersion": 2,

    "data": {
      "ver": 2,
      "version": "1.0",
      "retirement_ts": 1515415195838044
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest": "6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c
→",
      "payloadDigest":
→ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
    "txnMetadata": {
      "txnTime": 1577836799,
      "seqNo": 10,
    },
    "reqSignature": {
      "type": "ED25519",
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
→"
      }]
    }
  }
}

```

1.6.7 TRANSACTION_AUTHOR_AGREEMENT_AML

Setting a list of acceptance mechanisms for transaction author agreement.

Each write request for which a transaction author agreement needs to be accepted must point to a mechanism from the latest list on the ledger. The chosen mechanism is signed by the write request author (together with the transaction author agreement digest).

Each acceptance mechanisms list has a unique version.

- `version` (string):
Unique version of the transaction author agreement acceptance mechanisms list
- `aml` (dict):
Acceptance mechanisms list data in the form `<acceptance mechanism label>: <acceptance mechanism description>`
- `amlContext` (string, optional):
A context information about Acceptance mechanisms list (may be URL to external resource).

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 5,
    "protocolVersion": 2,

    "data": {

```

(continues on next page)

(continued from previous page)

```

        "ver": 1,
        "version": "1.0",
        "aml": {
            "EULA": "Included in the EULA for the product being used",
            "Service Agreement": "Included in the agreement with the service_
↪ provider managing the transaction",
            "Click Agreement": "Agreed through the UI at the time of submission",
            "Session Agreement": "Agreed at wallet instantiation or login"
        },
        "amlContext": "http://aml-context-descr"
    },

    "metadata": {
        "reqId": 1513945121191691,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "digest": "6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c
↪ ",
        "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bccc2e0106cd905685",
    },
    "txnMetadata": {
        "txnTime": 1513945121,
        "seqNo": 10,
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    }
}

```

1.6.8 TRANSACTION_AUTHOR_AGREEMENT_DISABLE

Immediately retires all active Transaction Author Agreements at once by setting current timestamp as a retirement one.

It's not possible to re-enable an Agreement right after disabling all agreements because there is no active latest Agreement at this point. A new Agreement needs to be sent instead.

Example:

```

{
  "ver": 1,
  "txn": {
    "type": 8,
    "protocolVersion": 2,

    "data": {
      "ver": 1,
    },
  },
}

```

(continues on next page)

(continued from previous page)

```
    "metadata": {
      "reqId": 1513945121191691,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "digest": "6cee82226c6e276c983f46d03e3b3d10436d90b67bf33dc67ce9901b44dbc97c
↪",
      "payloadDigest":
↪ "21f0f5c158ed6ad49ff855baf09a2ef9b4ed1a8015ac24bcc2e0106cd905685",
    },
    "txnMetadata": {
      "txnTime": 1577836799,
      "seqNo": 10,
    },
    "reqSignature": {
      "type": "ED25519",
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪"
      }]
    }
  }
}
```

- *Common Message Structure*
- *Signed Message Structure*
- *Common Request Structure*
- *Common Reply Structure*
- *Command Reply Structure*
- *ACK Structure*
- *NACK Structure*
- *Reject Structure*
- *Write Requests*
 - *NYM*
 - *ATTRIB*
 - *SCHEMA*
 - *CLAIM_DEF*
 - *NODE*
 - *POOL_UPGRADE*
 - *POOL_CONFIG*
- *Read Requests*
 - *GET_NYM*
 - *GET_ATTRIB*
 - *GET_SCHEMA*
 - *GET_CLAIM_DEF*
 - *GET_TXN*

This doc is about supported client's Request (both write and read ones). If you are interested in transactions and their representation on the Ledger (that is internal one), then have a look at [transactions](#).

`indy-sdk` expects the format as specified below.

See [roles and permissions](#) on the roles and who can create each type of transactions.

2.1 Base Client-to-Node and Node-to-Node serialization

The main Client-to-Node and Node-to-Node envelope is serialized in MsgPack format.

2.2 Common Message Structure

This is a common structure for ALL messages (both Node-to-Node and Client-to-Node).

```
"type": <...>,
"protocolVersion": <...>,
"ver": 1,

"data": {
  "ver": 1,
  <msg-specific fields>
},

"metadata": {
  <msg-specific fields>
},

"pluginData": {
  <plugin-specific-fields>
}
```

- `type` (enum integer):
Msg type.
- `protocolVersion` (integer; optional):
The version of client-to-node or node-to-node protocol. Each new version may introduce a new feature in Requests/Replies/Data. Since clients and different Nodes may be at different versions, we need this field to support backward compatibility between clients and nodes.
- `ver` (integer, optional):
Data/Metadata version.
- `data` (dict):
Message-specific data.
- `metadata` (dict):
Message-specific metadata.
- `pluginData` (dict):
Plugin-specific data.

2.3 Signed Message Structure

A message (see above) can be wrapped into a Signed Message envelope. All write requests must be signed.

```
{
  "type": <...>,
  "ver": <...>,

  "signature": {
    "type": <...>,
    "values": [{
      "from": <...>,
      "value": <...>,
    }]
  },

  "serialization": <...>,
  "msg": <serialized-msg>
}
```

- signature (dict):

Submitter's signature over serialized msg field.

- type (string enum):

- * ED25519: ed25519 signature
- * ED25519_MULTI: ed25519 signature in multisig case.

- values (list):

- * from (base58-encoded string): Identifier (DID) of signer as base58-encoded string for 16 or 32 bit DID value.
- * value (base58-encoded string): signature value

- serialization (string enum, optional):

Defines how the msg is serialized

- JSON: json
- MSG_PACK: msgpack

- msg (dict):

Serialized message.

2.4 Common Request Structure

Each Request (both write and read) follows the pattern as shown above.

```
{
  "type": <...>,
  "ver": <...>,
  "protocolVersion": <...>,

  "data": {
```

(continues on next page)

(continued from previous page)

```
    "ver": <...>,
    <msg-specific fields>
  },

  "metadata": {
    "reqId": <...>,
    "from": <...>,
    "endorser": <...>
  },
}
```

- Message Type `type` is one of the following values:

- `NODE` = 0
- `NYM` = 1
- `ATTRIB` = 100
- `SCHEMA` = 101
- `CLAIM_DEF` = 102
- `POOL_UPGRADE` = 109
- `NODE_UPGRADE` = 110
- `POOL_CONFIG` = 111
- `GET_TXN` = 3
- `GET_ATTR` = 104
- `GET_NYM` = 105
- `GET_SCHEMA` = 107
- `GET_CLAIM_DEF` = 108

- `metadata` (dict):

Metadata coming with the Request and saving in the transaction as is (if this is a write request).

- `from` (base58-encoded string): Identifier (DID) of the transaction author as base58-encoded string for 16 or 32 bit DID value. It may differ from `endorser` field who submits the transaction on behalf of `identifier`. If `endorser` is absent, then the `author(identifier)` plays the role of `endorser` and submits request by his own. It also may differ from `dest` field for some of requests (for example `NYM`), where `dest` is a target identifier (for example, a newly created DID identifier).

Example:

- * `identifier` is a DID of a transaction author who doesn't have write permissions; `endorser` is a DID of a user with `Endorser` role (that is with write permissions).
- * `new NYM creation`: `identifier` is a DID of an `Endorser` creating a new DID, and `dest` is a newly created DID.
- `endorser` (base58-encoded string, optional): Identifier (DID) of an `Endorser` submitting a transaction on behalf of the original author (`identifier`) as base58-encoded string for 16 or 32 bit DID value. If `endorser` is absent, then the `author(identifier)` plays the role of `endorser` and submits request by his own. If `endorser` is present then the transaction must be multi-signed by the both `author(identifier)` and `Endorser(endorser)`.
- `reqId` (integer): Unique ID number of the request with transaction.

- Please find the format of each request-specific data for each type of request below.

2.5 Common Reply Structure

Each Write/Read Reply follows the pattern as shown above.

```
{
  "type": REPLY,
  "ver": <...>,
  "protocolVersion": <...>,

  "data": {
    "ver": <...>,
    "results": [
      "result": {
        <result>
      },

      "multiSignature": {
        "type": <...>,
        "value": <...>,
        "from": <...>,
        "serialization": <...>,
        "signedState": <...>
      },

      "stateProof": <...>,
      "auditProof": <...>,
    ]
  },
  "metadata": {
    "reqId": <...>,
    "from": <...>,
  },
}
```

where multiSignature's signedState is a serialized value having the following form:

```
{
  "ledgerMetadata": {
    "ledgerId": <...>,
    "rootHash": <...>,
    "size": <...>,
  },

  "stateMetadata": {
    "timestamp": <...>,
    "poolRootHash": <...>,
    "rootHash": <...>,
  }
}
```

- type (string):
Request type as was in the corresponding Request.
- results (array):

Array of results. Each result may have either a state proof (it means that `result` is taken from state), audit proof (it means that `result` is taken from ledger), or no proofs (it means that this is some calculated data, and it's up to the client to verify it).

- `result` (dict):

The main result. It can be a transaction from state (see `state proof` then), a transaction from ledger (see `audit proof`), or any custom result. It usually includes transaction data and request metadata (as was in `Request`). It also includes `txnMetadata` and `reqSignature`.

- `multiSignature` (dict):

Nodes' Multi-signature against the given `ledgerMetadata` and `stateMetadata` (serialized to `MsgPack`)

- `value` (enum string): multi-signature type
 - * BLS: BLS multi-signature
- `value` (base58-encoded string): the value of the BLS multi-signature
- `from` (array of strings): Aliases of Nodes participated in BLS multi-signature (the number of participated nodes is not less than `n-f`)
- `serialization` (enum string, optional): serialization type of the `signedState` (`MsgPack` by default).
- `signedState` (bytes): serialized the multi-signed state the signature is calculated against
 - * `ledgerMetadata` (dict): Metadata associated with the current state of the Ledger.
 - `ledgerId` (enum integer): ID of the ledger the transaction is written to:
 - `POOL_LEDGER` = 0
 - `DOMAIN_LEDGER` = 1
 - `CONFIG_LEDGER` = 2
 - `rootHash` (base58-encoded hash string): base58-encoded ledger's merkle tree root hash
 - `size` (integer): Ledger's size (that is the last `seqNo` present at the ledger at that time).
 - * `stateMetadata` (dict): Metadata associated with the state (Patricia Merkle Trie) the returned transaction belongs to (see `ledgerId`).
 - `timestamp` (integer as POSIX timestamp): last update of the state
 - `poolRootHash` (base58-encoded hash string): pool state trie root hash to get the current state of the Pool (it can be used to get the state of the Pool at the moment the BLS multi-signature was created).
 - `rootHash` (base58-encoded string): state trie root hash for the ledger the returned transaction belongs to

- `stateProof` (base64-encoded string; optional):

Patricia Merkle Trie State proof for the returned transaction. It proves that the returned transaction belongs to the Patricia Merkle Trie with a root hash as specified in `stateMetadata`. It is present for replies to read requests only (except `GET_TXN` Reply).

- `auditProof` (array of base58-encoded hash strings; optional):

Ledger's merkle tree audit proof as array of base58-encoded hash strings. This is a cryptographic proof to verify that the returned transaction has been appended to the ledger with the root hash as specified in `ledgerMetadata`. It is present for replies to write requests only, and `GET_TXN` Reply.

- `metadata` (dict; optional):

Metadata as in Request. It may be absent for Reply to write requests as `txn` fields already contains this information as part of transaction written to the ledger.

2.6 Command Reply Structure

Each Reply to commands/actions follows the pattern as shown above.

```
{
  "type": REPLY_COMMAND,
  "protocolVersion": <...>,
  "ver": <...>,

  "data": {
    "ver": <...>,
    "results": [
      "result": {
        <result>
      },
    ]
  },
  "metadata": {
    "reqId": <...>,
    "from": <...>,
  },
}
```

2.7 ACK Structure

Each ACK follows the pattern as shown above.

```
{
  "type": REQACK,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
  },
  "metadata": {
    "reqId": 1514215425836443,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
  },
}
```

2.8 NACK Structure

Each NACK follows the pattern as shown above.

```
{
  "type": REQNACK,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "reason": <reason_str>
  },
  "metadata": {
    "reqId": 1514215425836443,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
  },
}
```

2.9 Reject Structure

Each Reject follows the pattern as shown above.

```
{
  "type": REJECT,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "reason": <reason_str>
  },
  "metadata": {
    "reqId": 1514215425836443,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
  },
}
```

2.10 Write Requests

The format of each request-specific data for each type of request.

2.10.1 NYM

Creates a new NYM record for a specific user, endorser, steward or trustee. Note that only trustees and stewards can create new endorsers and trustee can be created only by other trustees (see [roles](#)).

The request can be used for creation of new DIDs, setting and rotation of verification key, setting and changing of roles.

- `did` (base58-encoded string):

Target DID as base58-encoded string for 16 or 32 bit DID value. It differs from `from` metadata field, where `from` is the DID of the submitter.

Example: `from` is a DID of a Endorser creating a new DID, and `did` is a newly created DID.

- `role` (enum number as integer; optional):

Role of a user NYM record being created for. One of the following numbers

- None (common USER)
- 0 (TRUSTEE)
- 2 (STEWARD)
- 101 (ENDORSER)

A TRUSTEE can change any Nym's role to None, this stopping it from making any writes (see [roles](#)).

- `verkey` (base58-encoded string; optional):

Target verification key as base58-encoded string. If not set, then either the target identifier (`did`) is 32-bit cryptonym CID (this is deprecated), or this is a user under guardianship (doesn't own the identifier yet). Verkey can be changed to None by owner, it means that this user goes back under guardianship.

- `alias` (string; optional):

NYM's alias.

If there is no NYM transaction with the specified DID (`did`), then it can be considered as creation of a new DID.

If there is a NYM transaction with the specified DID (`did`), then this is update of existing DID. In this case we can specify only the values we would like to override. All unspecified values remain the same. So, if key rotation needs to be performed, the owner of the DID needs to send a NYM request with `did` and `verkey` only. `role` and `alias` will stay the same.

Request Example:

```
{
  "type": 1,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "ver": 1,
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 1,
    "ver": 1,
    "protocolVersion": 1,

    "data": {
      "ver": 1,
      "did": "N22KY2Dyvmuu2PyyqSFKue",
      "role": "101",
      "verkey": "~HmUWn928bnFT6Ephf65YXv"
    },
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
  },
}
```

Reply Example:

```

{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "did": "N22KY2Dyvmuu2PyyqSFKue",
            "role": "101",
            "verkey": "~HmUWn928bnFT6Ephf65YXv"
          },
          "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
          },
        },
        "txnMetadata": {
          "creationTime": 1514211268,
          "seqNo": 300,
          "txnId": "N22KY2Dyvmuu2PyyqSFKue|01"
        },
        "reqSignature": {
          "type": "ED25519",
          "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
          }]
        },
      },
    ],

    "multiSignature": {
      "type": "BLS",
      "value":
↪ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGaklvYAgvj2ecAKXQZXwcfos
↪ ",
      "from": ["Delta", "Gamma", "Alpha"],
      "serialization": "MsgPack",
      "signedState": {
        "ledgerMetadata": {
          "ledgerId": 1,
          "rootHash": "DqQ7G4fgDHBfdfVLE6DCdYyyED1fY5oKw76aDeFsLVr",
          "size": 300,
        },
        "stateMetadata": {
          "timestamp": 1514214795,
          "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
          "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
        }
      }
    }
  }
}

```

(continues on next page)

(continued from previous page)

```

    },
  },
  "auditProof": [ "Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKfk2eZmcySAWm2T5rnzZdEypW7A5SKi1Qt" ],
]
},
"metadata": {
  "reqId": 1514215425836443,
  "from": "L5AD5g65TDQr1PPHHRoiGf",
}
}

```

2.10.2 ATTRIB

Adds attribute to a NYM record.

- `did` (base58-encoded string):

Target DID we set an attribute for as base58-encoded string for 16 or 32 bit DID value. It differs from `from` metadata field, where `from` is the DID of the submitter.

Example: `from` is a DID of a Endorser setting an attribute for a DID, and `did` is the DID we set an attribute for.

- `raw` (json; mutually exclusive with `hash` and `enc`):

Raw data is represented as json, where key is attribute name and value is attribute value.

- `hash` (sha256 hash string; mutually exclusive with `raw` and `enc`):

Hash of attribute data.

- `enc` (string; mutually exclusive with `raw` and `hash`):

Encrypted attribute data.

Request Example:

```

{
  "type": 100,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 100,
    "ver": 1,
    "protocolVersion": 1,

```

(continues on next page)

(continued from previous page)

```

    "data": {
      "ver": 1,
      "did": "N22KY2Dyvmuu2PyyqSFKue",
      "raw": "{\"name\": \"Alice\"}"
    },
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
  },
}

```

Reply Example:

```

{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 100,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "did": "N22KY2Dyvmuu2PyyqSFKue",
            "raw": "{\"name\": \"Alice\"}"
          },
          "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
          },
        },
        "txnMetadata": {
          "creationTime": 1514211268,
          "seqNo": 300,
          "txnId": "N22KY2Dyvmuu2PyyqSFKue|02"
        },
        "reqSignature": {
          "type": "ED25519",
          "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↳ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↳ "

          }]
        },
        "multiSignature": {
          "type": "BLS",
          "value":
↳ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfo
↳ "

```

(continues on next page)

(continued from previous page)

```

    "from": ["Delta", "Gamma", "Alpha"],
    "serialization": "MsgPack",
    "signedState": {
      "ledgerMetadata": {
        "ledgerId": 1,
        "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fY5oKw76aDeFsLVr",
        "size": 300,
      },
      "stateMetadata": {
        "timestamp": 1514214795,
        "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
        "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
      },
    },
  },

  "auditProof": ["Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKFk2eZmcySAWm2T5rnzZdEypW7A5SKi1Qt"],
]
},
"metadata": {
  "reqId": 1514215425836443,
  "from": "L5AD5g65TDQr1PPHHRoiGf",
}
}

```

2.10.3 SCHEMA

Adds Claim's schema.

It's not possible to update existing Schema. So, if the Schema needs to be evolved, a new Schema with a new version or name needs to be created.

- **id (string):**
Schema's ID as State Trie key (address or descriptive data). It must be unique within the ledger. It must be equal (or be mapped to) the real key of the SCHEMA state in the State Trie.
- **name (string):**
Schema's name string.
- **version (string):**
Schema's version string
- **value (dict):**
Schema's specific data
 - **attrNames (array of strings):** Array of attribute name strings.

Request Example:

```

{
  "type": 101,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {

```

(continues on next page)

(continued from previous page)

```

        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    },
    "msg": {
        "type": 101,
        "ver": 1,
        "protocolVersion": 1,

        "data": {
            "ver": 1,
            "id": "L5AD5g65TDQr1PPHHRoiGf:Degree:1.0",
            "version": "1.0",
            "name": "Degree",
            "value": {
                "attrNames": ["undergrad", "last_name", "first_name", "birth_date",
↪ "postgrad", "expiry_date"]
            }
        },
        "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
        },
    }
}

```

Reply Example:

```

{
    "type": "WRITE_REPLY",
    "ver": 1,
    "protocolVersion": 1,

    "data": {
        "ver": 1,
        "results": [
            "result": {
                "txn": {
                    "type": 101,
                    "protocolVersion": 1,

                    "data": {
                        "ver": 1,
                        "version": "1.0",
                        "name": "Degree",
                        "value": {
                            "attrNames": ["undergrad", "last_name", "first_name",
↪ "birth_date", "postgrad", "expiry_date"]
                        }
                    },
                    "metadata": {
                        "reqId": 1514215425836443,
                        "from": "L5AD5g65TDQr1PPHHRoiGf",
                    },
                }
            }
        ]
    }
}

```

(continues on next page)

(continued from previous page)

```

    },
    "txnMetadata": {
      "creationTime": 1514211268,
      "seqNo": 300,
      "txnId": "L5AD5g65TDQr1PPHHRoiGf:Degree:1.0",
    },
    "reqSignature": {
      "type": "ED25519",
      "ver": 1,
      "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
→ "
      ]
    },
    "multiSignature": {
      "type": "BLS",
      "value":
→ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnvNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfo
→ ",
      "from": ["Delta", "Gamma", "Alpha"],
      "serialization": "MsgPack",
      "signedState": {
        "ledgerMetadata": {
          "ledgerId": 1,
          "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fy5oKw76aDeFsLVr",
          "size": 300,
        },
        "stateMetadata": {
          "timestamp": 1514214795,
          "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
          "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
        },
      },
    },
    "auditProof": ["Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
→ "3phchUcMsnKfk2eZmcySAWm2T5rnzZdEypW7A5SKi1Qt"],
  ]
},
"metadata": {
  "reqId": 1514215425836443,
  "from": "L5AD5g65TDQr1PPHHRoiGf",
}
}

```

2.10.4 CLAIM_DEF

Adds a claim definition (in particular, public key), that Issuer creates and publishes for a particular Claim Schema.

It's not possible to update data in existing Claim Def. So, if a Claim Def needs to be evolved (for example, a key needs to be rotated), then a new Claim Def needs to be created by a new Issuer DID (did).

- `id` (string):
Schema's ID as State Trie key (address or descriptive data). It must be unique within the ledger. It must be equal (or be mapped to) the real key of the SCHEMA state in the State Trie.
- `type` (string):
Type of the claim definition (that is claim signature). CL (Camenisch-Lysyanskaya) is the only supported type now.
- `tag` (string):
A unique descriptive tag of the given CRED_DEF for the given Issuer and Schema. An Issuer may have multiple CRED_DEFS for the same Schema created with different tags.
- `schemaId` (string):
ID of a Schema transaction the claim definition is created for.
- `value` (dict):
Type-specific value:
 - `publicKeys` (dict):
Dictionary with Claim Definition's public keys:
 - * `primary`: primary claim public key
 - * `revocation`: revocation claim public key

Request Example:

```
{
  "type": 102,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 102,
    "ver": 1,
    "protocolVersion": 1,

    "data": {
      "ver": 1,
      "id": "HHAD5g65TDQr1PPHHRoiGf2L:5AD5g65TDQr1PPHHRoiGf:Degree:1.0:CL:key1",
      "signatureType": "CL",
      "schemaRef": "L5AD5g65TDQr1PPHHRoiGf1Degree1",
      "publicKeys": {
        "primary": ....,
        "revocation": ....
      },
      "tag": "key1",
    },
    "metadata": {
```

(continues on next page)

(continued from previous page)

```

    "reqId": 1514215425836443,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example:

```

{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 102,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "signatureType": "CL",
            "schemaRef": "L5AD5g65TDQr1PPHHRoiGf1Degree1",
            "publicKeys": {
              "primary": ....,
              "revocation": ....
            },
            "tag": "key1",
          },
          "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
          },
          "txnMetadata": {
            "creationTime": 1514211268,
            "seqNo": 300,
            "txnId": "HHAD5g65TDQr1PPHHRoiGf2L:5AD5g65TDQr1PPHHRoiGf:Degree:1.
↪0:CL:key1",
          },
          "reqSignature": {
            "type": "ED25519",
            "values": [{
              "from": "L5AD5g65TDQr1PPHHRoiGf",
              "value":
↪"4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪"
            }]
          },
          "multiSignature": {
            "type": "BLS",
            "value":
↪"RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfos
↪",

```

(continues on next page)

(continued from previous page)

```

        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
            "ledgerMetadata": {
                "ledgerId": 1,
                "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fY5oKw76aDeFsLVr",
                "size": 300,
            },
            "stateMetadata": {
                "timestamp": 1514214795,
                "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
                "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbFJQ1HoEpK",
            },
        },
    },
    "auditProof": ["Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKfK2eZmcySAWm2T5rnzZdEypW7A5SKi1Qt"],
    ],
    "metadata": {
        "reqId": 1514215425836443,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
    }
}

```

2.10.5 NODE

Adds a new node to the pool, or updates existing node in the pool.

- `did` (base58-encoded string):

Target Node's verkey as base58-encoded string for 16 or 32 byte DID value. It differs from `from` metadata field, where `from` is the DID of the transaction submitter (Steward's DID).

- `alias` (string):

Node's alias

- `blskey` (base58-encoded string; optional):

BLS multi-signature key as base58-encoded string (it's needed for BLS signatures and state proofs support)

- `clientIp` (string; optional):

Node's client listener IP address, that is the IP clients use to connect to the node when sending read and write requests (ZMQ with TCP)

- `clientPort` (string; optional):

Node's client listener port, that is the port clients use to connect to the node when sending read and write requests (ZMQ with TCP)

- `nodeIp` (string; optional):

The IP address other Nodes use to communicate with this Node; no clients are allowed here (ZMQ with TCP)

- `nodePort` (string; optional):

The port other Nodes use to communicate with this Node; no clients are allowed here (ZMQ with TCP)

- `services` (array of strings; optional):

The service of the Node. `VALIDATOR` is the only supported one now.

If there is no `NODE` transaction with the specified Node ID (`did`), then it can be considered as creation of a new `NODE`.

If there is a `NODE` transaction with the specified Node ID (`did`), then this is update of existing `NODE`. In this case we can specify only the values we would like to override. All unspecified values remain the same. So, if a Steward wants to rotate BLS key, then it's sufficient to send a `NODE` transaction with `did` and a new `blskey`. There is no need to specify all other fields, and they will remain the same.

Request Example:

```
{
  "type": 0,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 0,
    "ver": 1,
    "protocolVersion": 1,

    "data": {
      "ver": 1,
      "did": "6HoV7DUEfNDiUP4ENnSC4yePja8w7JDQJ5uzVgyW4nL8"
      "alias": "Node1",
      "clientIp": "127.0.0.1",
      "clientPort": 7588,
      "nodeIp": "127.0.0.1",
      "nodePort": 7587,
      "blskey": "00000000000000000000000000000000",
      "services": ["VALIDATOR"]
    },
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
  },
}
```

Reply Example:

```
{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
```

(continues on next page)

(continued from previous page)

```

    "result": {
      "txn": {
        "type": 0,
        "protocolVersion": 1,

        "data": {
          "ver": 1,
          "did": "6HoV7DUEfNDiUP4ENnSC4yePja8w7JDQJ5uzVgyW4nL8"
          "alias": "Node1",
          "clientIp": "127.0.0.1",
          "clientPort": 7588,
          "nodeIp": "127.0.0.1",
          "nodePort": 7587,
          "blskey": "00000000000000000000000000000000",
          "services": ["VALIDATOR"]
        },
        "metadata": {
          "reqId": 1514215425836443,
          "from": "L5AD5g65TDQr1PPHHRoiGf",
        },
      },
      "txnMetadata": {
        "creationTime": 1514211268,
        "seqNo": 300,
      },
      "reqSignature": {
        "type": "ED25519",
        "values": [{
          "from": "L5AD5g65TDQr1PPHHRoiGf",
          "value":
→ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
→ "
        }]
      },
      "multiSignature": {
        "type": "BLS",
        "value":
→ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnvNN5rD2nXkqaVW3USGaklvyAgvj2ecAKXQZXwcfo
→ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
          "ledgerMetadata": {
            "ledgerId": 0,
            "rootHash": "DqQ7G4fgDHBfdfVLR6DCdYyyED1fY5oKw76aDeFsLVr",
            "size": 300,
          },
          "stateMetadata": {
            "timestamp": 1514214795,
            "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
            "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
          },
        },
      },
    },
  },

```

(continues on next page)

(continued from previous page)

```

    "auditProof": [ "Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
    ↪ "3phchUcMsnKfK2eZmcySAWm2T5rnzZdEypW7A5SKilQt" ],
    ],
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    }
  }
}

```

2.10.6 POOL_UPGRADE

Command to upgrade the Pool (sent by Trustee). It upgrades the specified Nodes (either all nodes in the Pool, or some specific ones).

- **name (string):**
Human-readable name for the upgrade.
- **action (enum: start or cancel):**
Starts or cancels the Upgrade.
- **version (string):**
The version of indy-node package we perform upgrade to. Must be greater than existing one (or equal if `reinstall` flag is `True`).
- **schedule (dict of node DIDs to timestamps):**
Schedule of when to perform upgrade on each node. This is a map where Node DIDs are keys, and upgrade time is a value (see example below). If `force` flag is `False`, then it's required that time difference between each Upgrade must be not less than 5 minutes (to give each Node enough time and not make the whole Pool go down during Upgrade).
- **sha256 (sha256 hash string):**
sha256 hash of the package
- **force (boolean; optional):**
Whether we should apply transaction (schedule Upgrade) without waiting for consensus of this transaction. If `false`, then transaction is applied only after it's written to the ledger. Otherwise it's applied regardless of result of consensus, and there are no restrictions on the Upgrade `schedule` for each Node. So, we can Upgrade the whole Pool at the same time when it's set to `True`. `False` by default. Avoid setting to `True` without good reason.
- **reinstall (boolean; optional):**
Whether it's allowed to re-install the same version. `False` by default.
- **timeout (integer; optional):**
Limits upgrade time on each Node.
- **justification (string; optional):**
Optional justification string for this particular Upgrade.

Request Example:

```

{
  "type": 109,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 109,
    "ver": 1,
    "protocolVersion": 1,

    "data": {
      "ver": 1,
      "action": `start`,
      "version": `1.3`,
      "schedule": { "4yC546FFzorLPgTNTc6V43DnpFrR8uHvtunBxb2Suaa2": "2017-12-
↪ 25T10:25:58.271857+00:00", "AtDfpKFe1RPgcr5nnYBw1Wxkgyn8Zjyh5MzFoEUTeoV3": "2017-12-
↪ 25T10:26:16.271857+00:00", "DG5M4zFm33Shrhjj6JB7nmX9BoNJUq219UXDfwbDPe2": "2017-12-
↪ 25T10:26:25.271857+00:00", "JpYerf4CssDrH76z7jyQPJLnZ1vwYgvKbvcp16AB5RQ": "2017-12-
↪ 25T10:26:07.271857+00:00" },
      "sha256": ↪
↪ `db34a72a90d026dae49c3b3f0436c8d3963476c77468ad955845a1ccf7b03f55`,
      "force": false,
      "reinstall": false,
      "timeout": 1
    },
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
  },
}

```

Reply Example:

```

{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 109,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "action": `start`,

```

(continues on next page)

(continued from previous page)

```

        "version": `1.3`,
        "schedule": { "4yC546FFzorLPgTNTc6V43DnpFrR8uHvtunBxb2Suaa2":
↪ "2017-12-25T10:25:58.271857+00:00", "AtDfpKFe1RPgcr5nnYBw1Wxkgyn8Zjyh5MzFoEUTeoV3":
↪ "2017-12-25T10:26:16.271857+00:00", "DG5M4zFm33Shrhjj6JB7nmX9BoNJUq219UXDfvwBDPe2":
↪ "2017-12-25T10:26:25.271857+00:00", "JpYerf4CssDrH76z7jyQPJLnZ1vwYgvKbvcpl6AB5RQ":
↪ "2017-12-25T10:26:07.271857+00:00"},
        "sha256": ↪
↪ `db34a72a90d026dae49c3b3f0436c8d3963476c77468ad955845a1ccf7b03f55`,
        "force": false,
        "reinstall": false,
        "timeout": 1
    },
    "metadata": {
        "reqId": 1514215425836443,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
    },
    "txnMetadata": {
        "creationTime": 1514211268,
        "seqNo": 300,
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    },
    },
    "multiSignature": {
        "type": "BLS",
        "value":
↪ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfos
↪ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
            "ledgerMetadata": {
                "ledgerId": 2,
                "rootHash": "DqQ7G4fgDHBfdfVLR6DCdYyyED1fY5oKw76aDeFsLVr",
                "size": 300,
            },
            "stateMetadata": {
                "timestamp": 1514214795,
                "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
                "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbFJQ1HoEpK",
            },
        },
    },
    "auditProof": ["Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKFk2eZmcySAWm2T5rnzZdEypW7A5SKi1Qt"],
    ]
},

```

(continues on next page)

(continued from previous page)

```

    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    }
  }
}

```

2.10.7 POOL_CONFIG

Command to change Pool's configuration

- `writes` (boolean):

Whether any write requests can be processed by the pool (if false, then pool goes to read-only state). True by default.

- `force` (boolean; optional):

Whether we should apply transaction (for example, move pool to read-only state) without waiting for consensus of this transaction. If false, then transaction is applied only after it's written to the ledger. Otherwise it's applied regardless of result of consensus. False by default. Avoid setting to True without good reason.

Request Example:

```

{
  "type": 111,
  "ver": 1,
  "serialization": "MSG_PACK",
  "signature": {
    "type": "ED25519",
    "ver": 1,
    "values": [{
      "from": "L5AD5g65TDQr1PPHHRoiGf",
      "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztggE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    }]
  },
  "msg": {
    "type": 111,
    "ver": 1,
    "protocolVersion": 1,

    "data": {
      "ver": 1,
      "writes": false,
      "force": true
    },
    "metadata": {
      "reqId": 1514215425836443,
      "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
  },
}

```

Reply Example:

```

{
  "type": "WRITE_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 111,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "writes": false,
            "force": true
          },
          "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
          },
        },
        "txnMetadata": {
          "creationTime": 1514211268,
          "seqNo": 300,
        },
        "reqSignature": {
          "type": "ED25519",
          "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↳ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmkaFM4jjyDgd
↳ "

          }]
        },
      ],

      "multiSignature": {
        "type": "BLS",
        "value":
↳ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfos
↳ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
          "ledgerMetadata": {
            "ledgerId": 2,
            "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fY5oKw76aDeFsLVr",
            "size": 300,
          },
          "stateMetadata": {
            "timestamp": 1514214795,
            "poolRootHash": "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
            "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
          },
        },
      },
    ],
  },
}

```

(continues on next page)

(continued from previous page)

```

    },
    "auditProof": [ "Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKFk2eZmcySAWm2T5rnzZdEypW7A5SKilQt" ],
  ],
  "metadata": {
    "reqId": 1514215425836443,
    "from": "L5AD5g65TDQr1PPHHRoiGf",
  }
}

```

2.11 Read Requests

2.11.1 GET_NYM

Gets information about a DID (NYM).

- did (base58-encoded string):

Target DID as base58-encoded string for 16 or 32 bit DID value. It differs from `from` metadata field, where `from` is the DID of the sender (may not exist on ledger at all).

Request Example:

```

{
  "type": 105,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "did": "N22KY2Dyvmuu2PyyqSFKue",
  },
  "metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example:

```

{
  "type": "READ_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

```

(continues on next page)

(continued from previous page)

```

        "data": {
            "ver": 1,
            "did": "N22KY2Dyvmuu2PyyqSFKue",
            "role": "101",
            "verkey": "~HmUWn928bnFT6Ephf65YXv"
        },
        "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
        },
    },
    "txnMetadata": {
        "creationTime": 1514211268,
        "seqNo": 300,
        "txnId": "N22KY2Dyvmuu2PyyqSFKue|01"
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztggE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        ]
    },
    },
    "multiSignature": {
        "type": "BLS",
        "value":
↪ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfo
↪ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
            "ledgerMetadata": {
                "ledgerId": 1,
                "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fY5oKw76aDeFsLVr
↪ ",
                "size": 300,
            },
            "stateMetadata": {
                "timestamp": 1514214795,
                "poolRootHash":
↪ "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
                "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
            },
        },
    },
    "state_proof": "+QHl+FGAgICg0he/hjc9t/
↪ tPFzmCrb2T+nHnN0cRwqPKqZE3pw2iCaAoAsA80p3oFwfl4dDaKkNI8z8weRsSaS9Y8n3HoardRzxgICAgICAgICAgID4naAgv
↪ CULIerYmmnnK2A6hNlu4ofU2eihKBna5MOCHiaObMfgjhsZ8KBSbC6EpTFruD02fuGKlF1q4CAGICgBk8Cpc14mIr78WguSeT7-
↪ rLT8qykKxzI4IO5ZMQwSmAoLsEwI+BkQFBiPsN8F610IjAg3+MVMbBjzugJKDo4NhYoFJ0lnlwq3FTW00iwlzoUcO3FPjSh5yvt
↪ ",
    ]

```

(continues on next page)

(continued from previous page)

```

    },
    "metadata": {
      "reqId": 1514215425836444,
      "from": "DDAD5g65TDQr1PPHHRoiGf",
    }
  }
}

```

2.11.2 GET_ATTRIB

Gets information about an Attribute for the specified DID.

NOTE: GET_ATTRIB for hash and enc attributes is something like the “proof of existence”, i.e. reply data contains requested value only.

- did (base58-encoded string):

Target DID we get an attribute for as base58-encoded string for 16 or 32 bit DID value. It differs from from metadata field, where from is the DID of the sender (may not exist on ledger at all).

- raw (string; mutually exclusive with hash and enc):

Requested attribute name.

- hash (sha256 hash string; mutually exclusive with raw and enc):

Requested attribute hash.

- enc (string; mutually exclusive with raw and hash):

Encrypted attribute.

Request Example:

```

{
  "type": 105,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "did": "AH4RRiPR78DUrCWatnCW2w",
    "raw": "dateOfBirth"
  },
  "metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example:

```

{
  "type": "READ_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {

```

(continues on next page)

(continued from previous page)

```

    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "did": "N22KY2Dyvmuu2PyyqSfKue",
            "raw": "{\"name\": \"Alice\"}"
          },
          "metadata": {
            "reqId": 1514215425836443,
            "from": "L5AD5g65TDQr1PPHHRoiGf",
          },
        },
        "txnMetadata": {
          "creationTime": 1514211268,
          "seqNo": 300,
          "txnId":
↪ "HHAD5g65TDQr1PPHHRoiGf2L:5AD5g65TDQr1PPHHRoiGf:Degree:1.0:CL:key1",
        },
        "reqSignature": {
          "type": "ED25519",
          "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
          }]
        },
        "multiSignature": {
          "type": "BLS",
          "value":
↪ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGaklvAgvj2ecAKXQZXwcfos
↪ ",
          "from": ["Delta", "Gamma", "Alpha"],
          "serialization": "MsgPack",
          "signedState": {
            "ledgerMetadata": {
              "ledgerId": 1,
              "rootHash": "DqQ7G4fgDHBfdVfVLR6DCdYyyED1fY5oKw76aDeFsLVr
↪ ",
              "size": 300,
            },
            "stateMetadata": {
              "timestamp": 1514214795,
              "poolRootHash":
↪ "TfMhX3KDjrq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
              "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbJfJQ1HoEpK",
            },
          },
        },
      },
    ],
  },

```

(continues on next page)

(continued from previous page)

```

        "state_proof": "+QHl+FGAgICg0he/hjc9t/
→tPFzmCrb2T+nHnN0cRwqPKqZEc3pw2iCaAoAsA80p3oFwfl4dDaKkNI8z8weRsSaS9Y8n3HoardRzxgICAgICAgICAgID4naAgv
→CULIerYmmnnK2A6hNlu4ofU2eihKBna5MOCHiaObMfghjsZ8KBSbC6EpTFruD02fuGKlFlq4CAgICgBk8Cpc14mIr78WguSeT7-
→rLT8qykKxzI4IO5ZMQwSmAoLsEwI+BkQFBiPsN8F610IjAg3+MVMbBjzugJKDo4NhYoFJ0lnlwq3FTW00iw1zoUcO3FPjSh5ytv
→",
    ],
    },
    "metadata": {
        "reqId": 1514215425836444,
        "from": "DDAD5g65TDQr1PPHHRoiGf",
    }
}
}

```

2.11.3 GET_SCHEMA

Gets Claim's Schema.

- id (string):

Schema's ID as State Trie key (address or descriptive data). It must be unique within the ledger. It must be equal (or be mapped to) the real key of the SCHEMA state in the State Trie.

Request Example:

```

{
  "type": 105,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "id": "L5AD5g65TDQr1PPHHRoiGf:Degree:1.0",
  },
  "metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example:

```

{
  "type": "READ_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

          "data": {
            "ver": 1,

```

(continues on next page)

(continued from previous page)

```

        "version": "1.0",
        "name": "Degree",
        "value": {
            "attrNames": ["undergrad", "last_name", "first_name",
↪ "birth_date", "postgrad", "expiry_date"]
        }
    },
    "metadata": {
        "reqId": 1514215425836443,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
    },
    "txnMetadata": {
        "creationTime": 1514211268,
        "seqNo": 300,
        "txnId": "L5AD5g65TDQr1PPHHRoiGf:Degree:1.0",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQY16egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        ]
    },
    },
    "multiSignature": {
        "type": "BLS",
        "value":
↪ "RTyxbErBLcmTHBLjlrYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGak1vyAgvj2ecAKXQZXwcfos
↪ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
            "ledgerMetadata": {
                "ledgerId": 1,
                "rootHash": "DqQ7G4fgDHBfdfVLR6DCdYyyED1fY5oKw76aDeFsLVr
↪ ",
                "size": 300,
            },
            "stateMetadata": {
                "timestamp": 1514214795,
                "poolRootHash":
↪ "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
                "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
            },
        },
    },
    "state_proof": "+QHl+FGAgICg0he/hjc9t/
↪ tPFzmCrb2T+nHnN0cRwqPKqZEc3pw2iCaAoAsA80p3oFwfl4dDaKkNI8z8weRsSaS9Y8n3HoardRzxgICAgICAgICAgID4naAgv
↪ CULIerYmmnnK2A6hNlu4ofU2eihKBna5MOCHiaObMfghjsZ8KBSbC6EpTFruD02fuGKlF1q4CAGICgBk8Cpc14mIr78WguSeT7-
↪ rLT8qykKxzI4IO5ZMQwSmAoLsEwI+BkQFBiPsN8F610IjAg3+MVMbBjzugJKDo4NhYoFJ0lnlwq3FTWO0iwlzoUcO3FPjSh5yvt
↪ ",
    ]

```

(continues on next page)

(continued from previous page)

```

    },
    "metadata": {
      "reqId": 1514215425836444,
      "from": "DDAD5g65TDQr1PPHHRoiGf",
    }
  }
}

```

2.11.4 GET_CLAIM_DEF

Gets Claim Definition.

- `id` (string):

Schema's ID as State Trie key (address or descriptive data). It must be unique within the ledger. It must be equal (or be mapped to) the real key of the SCHEMA state in the State Trie.

Request Example:

```

{
  "type": 105,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "id": "HHAD5g65TDQr1PPHHRoiGf2L:5AD5g65TDQr1PPHHRoiGf:Degree:1.0:CL:key1",
  },
  "metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example:

```

{
  "type": "READ_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "signatureType": "CL",
            "schemaRef": "L5AD5g65TDQr1PPHHRoiGf1Degree1",
            "publicKeys": {
              "primary": ....,
              "revocation": ....
            }
          }
        }
      }
    ]
  }
}

```

(continues on next page)

(continued from previous page)

```

        },
        "tag": "key1",
    },
    "metadata": {
        "reqId": 1514215425836443,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
    },
    "txnMetadata": {
        "creationTime": 1514211268,
        "seqNo": 300,
        "txnId":
↪ "HHAD5g65TDQr1PPHHRoiGf2L:5AD5g65TDQr1PPHHRoiGf:Degree:1.0:CL:key1",
    },
    "reqSignature": {
        "type": "ED25519",
        "values": [{
            "from": "L5AD5g65TDQr1PPHHRoiGf",
            "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
        }]
    },
    },
    "multiSignature": {
        "type": "BLS",
        "value":
↪ "RTyxbErBLcmTHBLj1rYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGaklvyAgvj2ecAKXQZXwcfos
↪ ",
        "from": ["Delta", "Gamma", "Alpha"],
        "serialization": "MsgPack",
        "signedState": {
            "ledgerMetadata": {
                "ledgerId": 1,
                "rootHash": "DqQ7G4fgDHBfdfVLR6DCdYyyED1fY5oKw76aDeFsLVr
↪ ",
                "size": 300,
            },
            "stateMetadata": {
                "timestamp": 1514214795,
                "poolRootHash":
↪ "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
                "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbfJQ1HoEpK",
            },
        },
    },
    },
    "state_proof": "+QH1+FGAgICg0he/hjc9t/
↪ tPFzmCrb2T+nHnN0cRwqPKqZEc3pw2iCaAoAsA80p3oFwfl4dDaKkNI8z8weRsSaS9Y8n3HoardRzxgICAgICAgICAgID4naAgv
↪ CULIerYmmnnK2A6hNlu4ofU2eihKBna5MOCHiaObMfghjsZ8KBSbC6EpTFruD02fuGKlF1q4CAGICgBk8Cpc14mIr78WguSeT7-
↪ rLT8qykKxzI4IO5ZMQwSmAoLsEwI+BkQFBiPsN8F610IjAg3+MVMbBjzugJKDo4NhYoFJ0ln1wq3FTW00iw1zoUcO3FPjSh5yt
↪ ",
    ]
    },
    "metadata": {
        "reqId": 1514215425836444,

```

(continues on next page)

(continued from previous page)

```

        "from": "DDAD5g65TDQr1PPHHRoiGf",
    }
}
}

```

2.11.5 GET_TXN

A generic request to get a transaction from Ledger by its sequence number.

- `ledgerId` (int enum):
ID of the ledger the requested transaction belongs to (Pool=0; Domain=1; Config=2).
- `seqNo` (int):
Requested transaction sequence number as it's stored on Ledger.

Request Example (requests a NYM txn with seqNo=9):

```

{
  "type": 105,
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "ledgerId": "1",
    "seqNo": 9,
  },
  "metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
  },
}

```

Reply Example (requests a NYM txn with seqNo=9):

```

{
  "type": "READ_REPLY",
  "ver": 1,
  "protocolVersion": 1,

  "data": {
    "ver": 1,
    "results": [
      "result": {
        "txn": {
          "type": 1,
          "protocolVersion": 1,

          "data": {
            "ver": 1,
            "did": "2VkbBskPNNyWrLrZq7DBhk",
            "role": "101",
            "verkey": "31V83xQnJDkZTSvm796X4MnzZFtUc96Tq6GJtuVkJQBE"
          },
        },
        "metadata": {

```

(continues on next page)

(continued from previous page)

```

        "reqId": 1514215425836443,
        "from": "L5AD5g65TDQr1PPHHRoiGf",
    },
},
"txnMetadata": {
    "creationTime": 1514211268,
    "seqNo": 9,
},
"reqSignature": {
    "type": "ED25519",
    "values": [{
        "from": "L5AD5g65TDQr1PPHHRoiGf",
        "value":
↪ "4X3skpoEK2DRgZxQ9PwuEvCJpL8JHdQ8X4HDDFyztgqE15DM2ZnkvrAh9bQYl6egVinZTzwHqznmnkaFM4jjyDgd
↪ "
    ]
},
},
"multiSignature": {
    "type": "BLS",
    "value":
↪ "RTyxbErBLcmTHBLjlYCAEpMMkLnL65kchGni2tQczqzomYWZx9QQpLvnnNN5rD2nXkqaVW3USGaklvyAgvj2ecAKXQZXwcfo
↪ ",
    "from": ["Delta", "Gamma", "Alpha"],
    "serialization": "MsgPack",
    "signedState": {
        "ledgerMetadata": {
            "ledgerId": 1,
            "rootHash": "DqQ7G4fgDHBfdfVLrE6DCdYyyED1fY5oKw76aDeFsLVr
↪ ",
            "size": 300,
        },
        "stateMetadata": {
            "timestamp": 1514214795,
            "poolRootHash":
↪ "TfMhX3KDjrqq94Wj7BHV9sZrgivZyjbHJ3cGRG4h1Zj",
            "rootHash": "7Wdj3rrMCZ1R1M78H4xK5jxikmdUUGW2kbFJQ1HoEpK",
        },
    },
},
"auditProof": ["Cdsoz17SVqPodKpe6xmY2ZgJ9UcywFDZTRgWSAYM96iA",
↪ "3phchUcMsnKFk2eZmcySAWm2T5rnzZdEypW7A5SKilQt"],
],
},
"metadata": {
    "reqId": 1514215425836444,
    "from": "DDAD5g65TDQr1PPHHRoiGf",
}
}
}

```

Default AUTH_MAP Rules

3.1 Who Is Owner

3.2 Endorser using

- Endorser is required only when the transaction is endorsed, that is signed by someone else besides the author.
- If transaction is endorsed, Endorser must sign the transaction.
- If author of txn has role `ENDORSER`, then no multi-sig is required, since he's already signed the txn.
- Endorser is required for unprivileged roles only.
- Unprivileged users cannot submit any transaction (including administrative transactions like pool upgrade or restart) without a signature from a DID with the endorser role that is specified in the endorser field.

Pool Upgrade Guideline

There is quite interesting and automated process of the pool (network) upgrade.

- The whole pool (that is each node in the pool) can be upgraded automatically without any manual actions via `POOL_UPGRADE` transaction.
- As a result of Upgrade, each Node will be at the specified version, that is a new package, for example deb package, will be installed.
- Migration scripts can also be performed during Upgrade to deal with breaking changes between the versions.

4.1 Pool Upgrade Transaction

- Pool Upgrade is done via `POOL_UPGRADE` transaction.
- The txn defines a schedule of Upgrade (upgrade time) for each node in the pool.
- Only the `TRUSTEE` can send `POOL_UPGRADE`.
- This is a common transaction (written to config ledger), so consensus is required.
- There are two main modes for `POOL_UPGRADE`: forced and non-forced (default).
 - Non-forced mode schedules upgrade only after `POOL_UPGRADE` transaction is written to the ledger, that is there was consensus. Forced upgrade schedules upgrade for each node regardless of whether `POOL_UPGRADE` transaction is actually written to the ledger, that is it can be scheduled even if the pool lost consensus.
 - Non-forced mode requires that upgrade of each node is done sequentially and not at the same time (so that a pool is still working and can reach consensus during upgrade). Forced upgrade allows upgrade of the whole pool at the same time.
- One should usually use non-forced Upgrades assuming that all changes are backward-compatible.
- If there are non-backward-compatible (breaking) changes, then one needs to use forced Upgrade and make it happen at the same time on all nodes (see below).

4.2 Node Upgrade Transaction

- Each node sends `NODE_UPGRADE` transaction twice:
 - `in_progress` action: just before start of the Upgrade (that is re-starting the node and applying a new package) to log that Upgrade started on the node.
 - `success` or `fail` action: after upgrade of the node to log the upgrade result.
- `NODE_UPGRADE` transaction is a common transaction (written to config ledger), so consensus is required.

4.3 Node Control Tool

- Upgrade is performed by a `node-control-tool`.
- See `node_control_tool.py`.
- On Ubuntu it's installed as a systemd service (`indy-node-control`) in addition to the node service (`indy-node`).
- `indy-node-control` is executed from the `root` user.
- When upgrade time for the node comes, it sends a message to `node-control-tool`.
- The `node-control-tool` then does the following:
 - stops `indy-node` service;
 - upgrades `indy-node` package (`apt-get install` on Ubuntu);
 - back-ups node data (ledger, etc.);
 - runs migration scripts (see `migration_tool.py`);
 - starts `indy-node` service;
 - restarts `indy-node-control` service.
- If upgrade failed for some reasons, `node-control-tool` tries to restore the data (ledger) from back-up and revert to the version of the code before upgrade.

4.4 Migrations

- Although we must try keeping backward-compatibility between the versions, it may be possible that there are some (for example, changes in ledger and state data format, re-branding, etc.).
- We can write migration scripts to support this kind of breaking changes and perform necessary steps for data migration and/or running some scripts.
- The migration should go to `data/migration` folder under the package name (so this is `data/migration/deb` on Ubuntu).
- Please have a look at the following [doc](#) for more information about how to write migration scripts.

4.5 When to Run Forced Upgrades

- Any changes in Ledger transactions format leading to changes in transactions root hash.
- Any changes in State transactions format (for example new fields added to State values) requiring re-creation of the state from the ledger.
- Any changes in Requests/Replies/Messages without compatibility and versioning support.

Create a Network and Start Nodes

Please be aware that recommended way of starting a pool is to use [Docker](#).

In order to run your own Network, you need to do the following for each Node:

1. Install Indy Node

- A recommended way for ubuntu is installing from deb packages

```
sudo apt-key adv --keyserver keyserver.ubuntu.com --recv-keys CE7709D068DB5E88
sudo bash -c 'echo "deb https://repo.sovrin.org/deb xenial stable" >> /etc/apt/
↪sources.list'
sudo apt-get update
sudo apt-get install indy-node
```

- It's also possible to install from pypi for test purposes: `pip install indy-node`

2. Initialize Node to be included into the Network

- if `indy-node` were installed from pypi basic directory structure should created manually with the command `create_dirs.sh`
- set Network name in config file
 - the location of the config depends on how a Node was installed. It's usually inside `/etc/indy` for Ubuntu.
 - the following needs to be added: `NETWORK_NAME={network_name}` where `{network_name}` matches the one in genesis transaction files above
- generate keys
 - ed25519 transport keys (used by ZMQ for Node-to-Node and Node-to-Client communication)
 - BLS keys for BLS multi-signature and state proofs support
- provide genesis transactions files which will be a basis of initial Pool.
 - pool transactions genesis file:
 - * The file must be named as `pool_transactions_genesis`

- * The file contains initial set of Nodes a Pool is started from (initial set of NODE transactions in the Ledger)
- * New Nodes will be added by sending new NODE txn to be written into the Ledger
- * All new Nodes and Clients will use genesis transaction file to connect to initial set of Nodes, and then catch-up all other NODE transactions to get up-to-date Ledger.
- * File must be located in `/var/lib/indy/{network_name}` folder
- domain transactions genesis file:
 - * The file must be named as `domain_transactions_genesis`
 - * The file contains initial NYM transactions (for example, Trustees, Stewards, etc.)
 - * File must be located in `/var/lib/indy/{network_name}` folder
- configure iptables to limit the number of simultaneous clients connections (recommended)

5.1 Scripts for Initialization

There are a number of scripts which can help in generation of keys and running a test network.

5.1.1 Generating keys

For deb installation

The following script should be used to generate both ed25519 and BLS keys for a node named Alpha with node port 9701 and client port 9702:

```
init_indy_node Alpha 0.0.0.0 9701 0.0.0.0 9702 [--seed_  
↪ 11111111111111111111111111111111Alpha]
```

Also this script generates indy-node environment file needed for systemd service config and indy-node iptables setup script.

For pip installation

The following script can generate both ed25519 and BLS keys for a node named Alpha:

```
init_indy_keys --name Alpha [--seed 11111111111111111111111111111111Alpha] [--force]
```

Note: Seed can be any randomly chosen 32 byte value. It does not have to be in the format 11..

Please note that these scripts must be called *after* `CURRENT_NETWORK` is set in config (see above).

5.1.2 Generating keys and test genesis transaction files for a test network

There is a script that can generate keys and corresponding test genesis files to be used with a Test network.

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 1 [--ips '191.177.  
↪ 76.26,22.185.194.102,247.81.153.79,93.125.199.45'] [--network=sandbox]
```

- `--nodes` specifies a total number of nodes in the pool

- `--clients` specifies a number of pre-configured clients in the pool (in `domain_transactions_file_{network_name}_genesis`)
- `--nodeNum` specifies a number of this particular node (from 1 to `--nodes` value), that is a number of the Node to create private keys locally for
- `--ip` specifies IP addresses for all nodes in the pool (if not specified, then `localhost` is used)
- `--network` specifies a Network generate transaction files and keys for. `sandbox` is used by default

We can run the script multiple times for different networks.

5.1.3 Setup iptables (recommended)

It is strongly recommended to add iptables (or some other firewall) rule that limits the number of simultaneous clients connections for client port. There are at least two important reasons for this:

- preventing the indy-node process from reaching of open file descriptors limit caused by clients connections
- preventing the indy-node process from large memory usage as ZeroMQ creates the separate queue for each TCP connection.

Instructions related to iptables setup can be found [here](#).

5.1.4 Running Node

The following script will start a Node process which can communicate with other Nodes and Clients:

```
start_indy_node Alpha 0.0.0.0 9701 0.0.0.0 9702
```

The node uses separate TCP channels for communicating with nodes and clients. The first IP/port pair is for the node-to-node communication channel and the second IP/port pair is for node-to-client communication channel. IP addresses may be changed according to hardware configuration. Different IP addresses for node-to-node and node-to-client communication may be used.

5.2 Local Test Network Example

If you want to try out an Indy cluster of 4 nodes with the nodes running on your local machine, then you can do the following:

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 1 2 3 4
By default node with the name Node1 will use ports 9701 and 9702 for nodestack and
↪clientstack respectively
Node2 will use ports 9703 and 9704 for nodestack and clientstack respectively
Node3 will use ports 9705 and 9706 for nodestack and clientstack respectively
Node4 will use ports 9707 and 9708 for nodestack and clientstack respectively
```

Now you can run the 4 nodes as

```
start_indy_node Node1 0.0.0.0 9701 0.0.0.0 9702
```

```
start_indy_node Node2 0.0.0.0 9703 0.0.0.0 9704
```

```
start_indy_node Node3 0.0.0.0 9705 0.0.0.0 9706
```

```
start_indy_node Node4 0.0.0.0 9707 0.0.0.0 9708
```

5.3 Remote Test Network Example

Now let's say you want to run 4 nodes on 4 different machines as

1. Node1 running on 191.177.76.26
2. Node2 running on 22.185.194.102
3. Node3 running on 247.81.153.79
4. Node4 running on 93.125.199.45

On machine with IP 191.177.76.26 you will run

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 1 --ips '191.177.
↪76.26,22.185.194.102,247.81.153.79,93.125.199.45'
This node with name Node1 will use ports 9701 and 9702 for nodestack and clientstack_
↪respectively
```

On machine with IP 22.185.194.102 you will run

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 2 --ips '191.177.
↪76.26,22.185.194.102,247.81.153.79,93.125.199.45'
This node with name Node2 will use ports 9703 and 9704 for nodestack and clientstack_
↪respectively
```

On machine with IP 247.81.153.79 you will run

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 3 --ips '191.177.
↪76.26,22.185.194.102,247.81.153.79,93.125.199.45'
This node with name Node3 will use ports 9705 and 9706 for nodestack and clientstack_
↪respectively
```

On machine with IP 93.125.199.45 you will run

```
~$ generate_indy_pool_transactions --nodes 4 --clients 5 --nodeNum 4 --ips '191.177.
↪76.26,22.185.194.102,247.81.153.79,93.125.199.45'
This node with name Node4 will use ports 9707 and 9708 for nodestack and clientstack_
↪respectively
```

Now you can run the 4 nodes as

```
start_indy_node Node1 0.0.0.0 9701 0.0.0.0 9702
```

```
start_indy_node Node2 0.0.0.0 9703 0.0.0.0 9704
```

```
start_indy_node Node3 0.0.0.0 9705 0.0.0.0 9706
```

```
start_indy_node Node4 0.0.0.0 9707 0.0.0.0 9708
```

Add Node to Existing Pool

6.1 Node preparation

1. Add this files from a running node:

```
/var/lib/indy/network_name/pool_transactions_genesis
/var/lib/indy/network_name/domain_transactions_genesis
```

1. Initialize keys, aliases and ports on the new node using `init_indy_node` script. Example:

```
sudo su - indy -c "init_indy_node NewNode 0.0.0.0 9701 0.0.0.0 9702_
↪0000000000000000000000000000NewNode"
```

1. When the node starts for the first time, it reads the content of `genesis_pool_transactions_sandbox` and `domain_transactions_sandbox` files and adds it to the ledger. The Node reads genesis transactions only once during the first start-up, so make sure the genesis files are correct before starting the service.

```
sudo systemctl start indy-node
sudo systemctl status indy-node
sudo systemctl enable indy-node
```

6.2 Add Node to the Pool

1. As Trustee add another Steward if needed (only Steward can add a new Validator Node; a Steward can add one and only one Validator Node).
2. Using Indy CLI, run the following command as Steward:

```
ledger node target=6G9QhQa3HWjRKeRmEvEkLbWWf2t7cw6KLtafzi494G4G client_port=9702_
↪client_ip=10.255.255.255 alias=NewNode node_ip=10.0.0.10.255.255.255 node_port=9701_
↪services=VALIDATOR_
↪blskey=zi65fRHZjK2R8wdJfDzeWVgcf9imXUsMSEY64LQ4HyhDMsSn3Br1vhnwXHE7NyGjxVnw4FGPqxpzY8HrQ2PnrL9tu4
↪blskey_
↪pop=RaY9xGLbQbrBh8np5gWWQAWisaxd96FtvtbXKjyzBj4fUYyPq4pkyCHTYvQz jehmUK5pNfnyhwWqGg1ahPwtWopenuRjAeCl
(continues on next page)
```

- `alias` specifies unique Node name
- `blskey` specifies BLS key from `init_indy_node` script
- `blskey_pop` specifies Proof of possession for BLS key from `init_indy_node` script
- `target` specifies base58 of the node public key ('Verification key' field in output of `init_indy_node`)

Example: Verification key is `ab78300b3a3eca0a1679e72dd1656075de9638ae79dc6469a3093ce1cc8b424f` In order to get base58 of the verkey execute in your shell (you should have `indy-plenum` installed):

```
python3 -c "from plenum.common.test_network_setup import TestNetworkSetup; print (TestNetworkSetup.getNymFromVerkey(str.encode('ab78300b3a3eca0a1679e72dd1656075de9638ae79dc6469a3093ce1cc8b424f')))"
```

Output:
`4Tn3wZMNCvhSTXPcLinQDnHyj56DTLQtL61ki4jo2Loc`

6.3 Make sure that Node is workable

Do `systemctl restart indy-node` and verify that node completed catch-up successfully.

Indy-node comes with a number of useful helper scripts:

- `init_indy_keys`
Initializes (or rotates) Node's keys (private and public ones) needed for communication with the Nodes via CurveCP protocol (CurveZMQ)
- `init_bls_keys`
Initializes (or rotates) Node's BLS keys required for BLS multi-signature and State Proofs support
- `read_ledger`
Reads transaction from a specified ledger in JSON format
- `validator-info`
Monitors the current status of the Node
- `generate_indy_pool_transactions`
Initializes a test Pool (generates keys and genesis transactions)
- `clear_node`
Clean up of all data on the Node

Setup iptables rules (recommended)

It is strongly recommended to add iptables (or some other firewall) rule that limits the number of simultaneous clients connections for client port. There are at least two important reasons for this:

- preventing the indy-node process from reaching of open file descriptors limit caused by clients connections
- preventing the indy-node process from large memory usage as ZeroMQ creates the separate queue for each TCP connection.

NOTE: limitation of the number of *simultaneous clients connections* does not mean that we limit the number of *simultaneous clients* the indy-node works with in any time. The IndySDK client does not keep connection infinitely, it uses the same connection for request-response session with some optimisations, so it's just about **connections**, **not** about **clients**.

Also iptables can be used to deal with various DoS attacks (e.g. syn flood) but rules' parameters are not estimated yet.

NOTE: you should be a root to operate with iptables.

8.1 Setting up clients connections limit

8.1.1 Using raw iptables command or iptables front-end

In case of deb installation the indy-node environment file `/etc/indy/indy.env` is created by `init_indy_node` script. This environment file contains client port (`NODE_CLIENT_PORT`) and recommended clients connections limit (`CLIENT_CONNECTIONS_LIMIT`). This parameters can be used to add the iptables rule for chain INPUT:

```
# iptables -I INPUT -p tcp --syn --dport 9702 -m connlimit --connlimit-above 500 --  
→connlimit-mask 0 -j REJECT --reject-with tcp-reset
```

Some key options:

- `--dport` - a port for which limit is set
- `--connlimit-above` - connections limit, exceeding new connections will be rejected using TCP reset
- `--connlimit-mask` - group hosts using the prefix length, 0 means “all subnets”

Corresponding fields should be set in case of some iptables front-end usage.

8.1.2 Using indy scripts

For ease of use and for people that are not familiar with iptables we've added two scripts:

- `setup_iptables`: adds a rule to iptables to limit the number of simultaneous clients connections for specified port;
- `setup_indy_node_iptables`: a wrapper for `setup_iptables` script which gets client port and recommended connections limit from indy-node environment file that is created by `init_indy_node` script.

Links to these scripts:

- https://github.com/hyperledger/indy-node/blob/master/scripts/setup_iptables
- https://github.com/hyperledger/indy-node/blob/master/scripts/setup_indy_node_iptables

NOTE: for now the iptables chain for which the rule is added is not parameterized, the rule is always added for INPUT chain, we can parameterize it in future if needed.

For deb installation

To setup the limit of the number of simultaneous clients connections it is enough to run the following script without parameters

```
# setup_indy_node_iptables
```

This script gets client port and recommended connections limit from the indy-node environment file.

NOTE: this script should be called *after* `init_indy_node` script.

For pip installation

The `setup_indy_node_iptables` script can not be used in case of pip installation as indy-node environment file does not exist, use the `setup_iptables` script instead (9702 is a client port, 500 is recommended limit for now)

```
# setup_iptables 9702 500
```

In fact, the `setup_indy_node_iptables` script is just a wrapper for the `setup_iptables` script.

Node Monitoring Tools for Stewards

- *Plugin Manager*
 - *Events Emitted*
- *Email Plugin*
 - *Prerequisites*
 - *Install*
 - *Configuration*
 - *Email delivery frequency*
- *AWS SNS Plugin*
 - *Prerequisites*
 - *Setup*
 - *Configuration*
 - *Events*
 - *Hints*
 - *Example*

9.1 Plugin Manager

Currently, indy-node emits different events via the Plugin Manager when certain criteria are met. The Plugin Manager tries to import all pip packages which names start with “indynotifier*”. Each of these packages is required to expose `send_message`; interface which is used to pass the event with the associated message to the package for further handling.

The Plugin Manager code is located at [here](#).

9.1.1 Events Emitted

- `.nodeRequestSpike` : `NodeRequestSuspiciousSpike`
- `.clusterThroughputSpike` : `ClusterThroughputSuspiciousSpike`
- `.clusterLatencyTooHigh` : `ClusterLatencyTooHigh`
- `.nodeUpgradeScheduled` : `NodeUpgradeScheduled`
- `.nodeUpgradeComplete` : `NodeUpgradeComplete`
- `.nodeUpgradeFail` : `NodeUpgradeFail`
- `.poolUpgradeCancel` : `PoolUpgradeCancel`

9.2 Email Plugin

9.2.1 Prerequisites

- SMTP server must be running on localhost.
- Install SMTP server (if you don't have one already)

The most simple way on Ubuntu is to use `sendmail`:

```
$ sudo apt-get install sendmail
```

To check that it's working execute:

```
echo "Subject: sendmail test" | sendmail -v youremail@example.com -f alert@noreply.com
```

If you get a email on your `youremail@example.com` then `sendmail` is working.

9.2.2 Install

```
# pip3 install indynotifieremail
```

9.2.3 Configuration

The spike detection and notification mechanisms should be enabled by appending of the following line to `indy_config.py` configuration file:

```
SpikeEventsEnabled=True
```

The package depends on two environment variables:

- `INDY_NOTIFIER_EMAIL_RECIPIENTS` (required)
- `INDY_NOTIFIER_EMAIL_SENDER` (optional)

Add these variables to `/etc/indy/indy.env` environment file as you are required to set such system environment variables for `indy-node` service in form described below.

INDY_NOTIFIER_EMAIL_RECIPIENTS

`INDY_NOTIFIER_EMAIL_RECIPIENTS` should be a string in a format of:

recipient1@address.com [optional list of events the recipient is going to get],
recipient2@address.com [event list]

If no list was provided the recipient is going to get notifications for all events. Example:

```
steward1@company.com event1 event2, steward2@company.com, steward3@company.com  
event3
```

This way steward1 is going to get notifications for event1 and event2, steward2 is going to get all possible notifications and steward3 is going to get notifications for event3 only.

The current list of events can be found above.

INDY_NOTIFIER_EMAIL_SENDER

By default every email notification is going to be from alert@noreply.com. You can change this by setting INDY_NOTIFIER_EMAIL_SENDER. May be useful for email filters.

9.2.4 Email delivery frequency

By default you will not get a email with the same topic more than once an hour. This is defined by SILENCE_TIMEOUT. It can be overridden by setting INDY_NOTIFIER_SILENCE_TIMEOUT environment variable in /etc/indy/indy.env file. Emails regarding update procedure are always delivered.

9.3 AWS SNS Plugin

9.3.1 Prerequisites

- .A AWS SNS topic created with permissions to publish to it.
- .A installed Sovrin Validator instance.

9.3.2 Setup

Install the python package for sovrin-notifier-awssns. This should be only be installed using pip3.

```
pip3 install sovrinnotifierawssns
```

9.3.3 Configuration

To configure AWS Credentials you will need to know the values for: aws_access_key_id and aws_secret_access_key. Follow the steps documented here [Boto3 Configuring Credentials](#).

Use either of the following ways:

- .Environment variables AWS_ACCESS_KEY_ID, AWS_SECRET_ACCESS_KEY
- .Shared credential file (~/.aws/credentials)
- .Boto2 config file (/etc/boto.cfg and ~/.boto)

Configure AWS Region you will need to know the value where the SNS Topic is hosted e.g. us-west-1, us-west-2, sa-east-1

To achieve this:

- .Set a Environment variable AWS_DEFAULT_REGION

- `.Set` region using file (`~/aws/config`)

Define environment variable `SOVRIN_NOTIFIER_AWSSNS_TOPICARN` on the Validator and set valid AWS SNS TopicARN as the value.

9.3.4 Events

Events that cause a notification:

- `NodeRequestSuspiciousSpike`
- `ClusterThroughputSuspiciousSpike`
- `ClusterLatencyTooHigh`
- `NodeUpgradeScheduled`
- `NodeUpgradeComplete`
- `NodeUpgradeFail`,
- `PoolUpgradeCancel`

9.3.5 Hints

The home directory for the account that runs `sovrin-node.service` on a Validator is `/home/sovrin/`. So the `aws` credentials/config files must be created in `/home/sovrin/.aws` folder.

To set an environment variable on the Validator you must add it to the file `/home/sovrin/.sovrin/sovrin.env` and restart the Validator. The TopicARN must be defined in this file.

To restart the Validator on a Ubuntu system you must execute the command `sudo systemctl restart sovrin-node.service` while not logged in as a `sovrin` user.

Example

This simple script will complete the setup, assuming that the `sovrinnotifierawssns` package is already installed:

```
#!/bin/bash

sudo mkdir /home/sovrin/.aws

sudo sh -c "printf \"[default]\nregion=us-west-2\" > /home/sovrin/.aws/config"

sudo sh -c "printf \".[default]\naws_access_key_id=AKIAIGKGW3CKRXKKWPZA\naws_secret_
↪access_key=<YOUR_SECRET_KEY>\" > /home/sovrin/.aws/credentials"

sudo sh -c "printf \"SOVRIN_NOTIFIER_AWSSNS_TOPICARN=arn:aws:sns:us-west-
↪2:034727365312:validator-health-monitor-STN\" >> /home/sovrin/.sovrin/sovrin.env"

sudo chown -R sovrin:sovrin /home/sovrin/.aws /home/sovrin/.sovrin/sovrin.env

sudo systemctl restart sovrin-node
```

Node State Diagnostic Tools Workflow

10.1 Introduction

The Node state diagnostic tools are for working with indy node and its state. The tools are intended to work with the recorder functionality and should help with diagnostic work in the indy-node development and QA workflows.

There are currently three scripts located in tools/diagnostics. They are nscapture, nsreplay and nsdiff. Basic functionality and intended workflow are documented below.

This documentation is high level and will lack complete detail about the parameters and options for these scripts. Please see the ‘-h’ output from the script for more details.

10.2 nscapture

nscapture has no required parameters and makes the following assumptions. These assumptions should work in most production like environments (QA environments or other Debian install environments).

Assumptions:

1. **ROOT_DIR:** nscapture will assume that the root directory is the base of the filesystem (‘/’). This should be correct for Debian installed systems. For other installation patterns, the ‘-o’ flag can be used to point to another location.
2. **NODE_NAME:** nscapture will look for nodes; starting it’s search from the root directory. If only one node configuration is found, nscapture will capture it. Otherwise, nscapture will require the user to specify which node to capture. Most nodes will only have one configuration. The ‘-n’ option can be used to specify the desired node to capture.
3. **OUTPUT_DIR:** By default, nscapture will output the archive to the current working directory of the running process. If another location is desired, the ‘-o’ option can be used.

Other Notes

- Capture file name structure is:

```
<node name>.<pool name>.<capture date>.<file extension>
```

ex:

```
Node1.sandbox.20180517163819.tar.gz
```

This format is expected by the other scripts.

10.3 nsreplay

If a node's nscapture archive contains a recording, nsreplay starts a single instance of indy-node and replays captured events (mostly messages) using replay logic in Plenum.

Assumptions

1. Node state was captured using nscapture.
2. HAS RECORDING: The replay requires a recording to replay. See doc/recorder.md in indy-plenum for details on how to configure the recorder. nscapture captures the recording in data/<node name>/recorder/*.

Other Notes

10.4 nsdiff

nsdiff will compare two nscapture archives. The result will highlight what parts of the Node state are different and how they differ in *nix diff format. The contents of the nscapture archives dictates what is being compared. What nscapture captures for comparison may change over time. There are comments in the nsdiff script that describe what is being compared and how to get a comprehensive list of things being compared.

nsdiff is intended to help understand how node state differs between different nodes in the same cluster or between a node and its replayed state. It is not intended to lessen/reduce the need for in-depth understanding of the Indy codebase.

Assumptions

1. Node state was captured using nscapture.

Other Notes

When using nsdiff to compare a node and its replayed state, you will be required to do the following:

1. Run nscapture on a node. This will produce a *.tar.gz file.

```
$ nscapture
```

1. Run nsreplay on the Node State Archive (*.tar.gz) captured in step 1, explicitly defining an output directory in which to write replayed state. Note that nsreplay does not have to be run on the node where the Node State Archive was captured. An equivalent development environment with compatible versions of indy-node, indy-plenum, etc. is recommended, because you will also have visual debug tools available (i.e. PyCharm)

```
$ nsreplay -o <OUTPUT_DIR> <step 1 *.tar.gz>
```

1. Run nscapture with step 2 <OUTPUT_DIR> as the ROOT_DIR to capture the replayed state. This will create another *.tar.gz

```
$ nscapture -r <step 2 OUTPUT_DIR>
```

1. nsdiff the two Node State Archives

```
$ nsdiff <step 1 *.tar.gz> <step 3 *.tar.gz>
```


The following is a simple but hopefully a common workflow for using these tools.

11.1 Steps

11.1.1 Step #1 - Configure recording

Add `STACK_COMPANION` to `/etc/indy/indy_config.py` if using `indy-node` version `>= 1.4.485` `/etc/indy/plenum_config.py` otherwise and set to the value to 1 (numeric value, NOT a String).

The value 0 will DISABLE recording.

The value 1 will ENABLE recording.

Indy will require a restart for the configuration to take effect.

11.1.2 Step #2 - Get indy into the desired state

Run load or other experiments to get the node into the desired state. This step is open-ended and up to the desires/objectives of the person using the workflow.

11.1.3 Step #3 - Capture Node state and recording

In the environment (on the server, VM, docker container, etc) where Indy is running, stop Indy (`indy-node` service) and then run `nscapture`. Doing this will produce an archive. This archive could be large depending on the state of the node. Every message for every instance will be contained in the recording. The archive must be moved to the environment where the replay will be performed. That can be done via SCP or similar method.

11.1.4 Step #4 - Replay recording

In a development environment, use `nsreplay` to replay the recording. Currently, the replay will take a similar amount of time compared to the recording. During the replay, debugging and other development tools should be usable.

11.1.5 Step #5 - Capture replayed Node state

In the environment (on the server, VM, docker container, etc) where `nsreplay` was run, execute `nscapture` to capture the replayed state.

11.1.6 Step #6 - Diff captured Node state with captured replay state

If desired, the `nsdiff` can be used to check that the replay was able to reproduce the same state as the recording. Use `nsdiff` to check the difference between the states. Ideally, they will be identical but if they are different, insight might still be gained.

12.1 Branches

- `master` branch contains the latest changes. All PRs usually need to be sent to `master`.
- `stable` branch contains latest releases (<https://github.com/hyperledger/indy-node/releases>). Hotfixes need to be sent to both `stable` and `master`.
- `release-*` branches hold release candidates during release workflow
- `hotfix-*` branches hold release candidates during hotfix workflow

12.2 Pull Requests

- Each PR needs to be reviewed.
- PR can be merged only after all tests pass and code is reviewed.

12.3 Continuous Integration

- for each PR we execute:
 - static code validation
 - Unit/Integration tests
- We use pipeline in code approach and Jenkins as our main CI/CD server.
- CI part of the pipeline (running tests for each PR) is defined in `Jenkinsfile.ci` file.
- CI part is run on Hyperledger and Sovrin Foundation Jenkins servers, so they are public and open as every contributor needs to see results of the tests run for his or her PR.

12.3.1 Static Code Validation

- We use flake8 for static code validation.
- It's run against every PR. PR fails if there are some static code validation errors.
- Not all checks are enabled (have a look at `.flake8` file at the project root)
- You can run static code validation locally:
 - Install flake8: `pip install flake8`
 - Run validation on the root folder of the project: `flake8 .`

12.4 Continuous Delivery

- CD part of the pipeline is defined in `Jenkinsfile.cd` file.
- CD part is run on Sovrin Foundation Jenkins server dealing with issuing and uploading new builds.

12.4.1 Builds

What artifacts are produced after each push

- to `master` branch:
 - all artifacts include developmental release segment `devN` in their version
 - indy-plenum:
 - * indy-plenum in `pypi`
 - * indy-plenum deb package in `https://repo.sovrin.org/deb xenial master-latest`
 - indy-node:
 - * indy-node in `pypi`
 - * indy-node deb package in `https://repo.sovrin.org/deb xenial master-latest`
 - * indy-node deb package in `https://repo.sovrin.org/deb xenial master` (copied from `master-latest`)
 - * indy-plenum deb package in `https://repo.sovrin.org/deb xenial master` (copied from `master-latest`)
- to `release-*` and `hotfix-*` branches:
 - all artifacts include pre-release segment `rcN` in their version
 - indy-plenum:
 - * indy-plenum in `pypi`
 - * indy-plenum deb package in `https://repo.sovrin.org/deb xenial rc-latest`
 - indy-node:
 - * indy-node in `pypi`
 - * indy-node deb package in `https://repo.sovrin.org/deb xenial rc-latest`

- * indy-node deb package in <https://repo.sovrin.org/deb xenial rc> (copied from rc-latest)
- * indy-plenum deb package in <https://repo.sovrin.org/deb xenial rc> (copied from rc-latest)
- to stable branch:
 - indy-plenum:
 - * indy-plenum in [pypi](#)
 - * indy-plenum deb package in <https://repo.sovrin.org/deb xenial stable-latest>
 - * indy-plenum release tag (<https://github.com/hyperledger/indy-plenum/releases>)
 - indy-node:
 - * indy-node in [pypi](#)
 - * indy-node deb package in <https://repo.sovrin.org/deb xenial stable-latest> (re-packed from rc-latest)
 - * indy-node deb package in <https://repo.sovrin.org/deb xenial stable> (copied from rc-latest)
 - * indy-plenum deb package in <https://repo.sovrin.org/deb xenial stable> (copied from stable-latest)
 - * indy-node release tag (<https://github.com/hyperledger/indy-node/releases>)

Use cases for artifacts

- PyPI artifacts can be used for development experiments, but not intended to be used for production.
- Using deb packages is recommended way to be used for a test/production pool on Ubuntu.
 - indy-node deb package from <https://repo.sovrin.org/deb xenial stable> is one and the only official stable release that can be used for production (stable version).
 - indy-node deb package from <https://repo.sovrin.org/deb xenial master> contains the latest changes (from master branch). It's not guaranteed that that this code is stable enough.

12.4.2 Packaging

Supported platforms and OSes

- Ubuntu 16.04 on x86_64

Build scripts

We use [fpm](#) for packaging python code into deb packages. Build scripts are placed in `build-scripts` folders:

- <https://github.com/hyperledger/indy-node/blob/master/build-scripts>
- <https://github.com/hyperledger/indy-plenum/blob/master/build-scripts>

We also pack some 3rd parties dependencies which are not presented in canonical ubuntu repositories:

- <https://github.com/hyperledger/indy-node/blob/master/build-scripts/ubuntu-1604/build-3rd-parties.sh>
- <https://github.com/hyperledger/indy-plenum/blob/master/build-scripts/ubuntu-1604/build-3rd-parties.sh>

Each `build-scripts` folder includes `Readme.md`. Please check them for more details.

12.4.3 Versioning

- Please note, that we are using versioning that satisfies [PEP 440](#) with release segment as `MAJOR.MINOR.PATCH` that satisfies [SemVer](#) as well.
- Version is set in the code (see `__version__.json`).
- Version is bumped for new releases / hotfixes either manually or using `bump_version.sh` script. The latter is preferred.
- During development phase version includes developmental segment `devN`, where `N` is set for CD pipeline artifacts as incremented build number of build server jobs. In the source code it is just equal to 0 always.
- During release preparation phase (release / hotfix workflows) version includes pre-release segment `rcN`, where `N` ≥ 1 and set in the source code by developers.
- Each dependency (including indy-plenum) has a strict version (see `setup.py`)
- If you install indy-node (either from pypi, or from deb package), the specified in `setup.py` version of indy-plenum is installed.
- Master and Stable share the same versioning scheme.
- Differences in master and stable code:
 - `setup.py`: different versions of indy-plenum dependency
 - different versions in migrations scripts

For releases < 1.7.0 (deprecated)

- Please note, that we are using semver-like approach for versioning (major, minor, build) for each of the components.
- Major and minor parts are set in the code (see `__metadata__.py`). They must be incremented for new releases manually from code if needed.
- Build part is incremented with each build on Jenkins (so it always increases, but may be not sequentially)
- Each dependency (including indy-plenum) has a strict version (see `setup.py`)
- If you install indy-node (either from pypi, or from deb package), the specified in `setup.py` version of indy-plenum is installed.
- Master and Stable builds usually have different versions.
- Differences in master and stable code:
 - `setup.py`:
 - * dev suffix in project names and indy-plenum dependency in master; no suffixes in stable
 - * different versions of indy-plenum dependency
 - different versions in migrations scripts

12.5 Release workflow

12.5.1 Feature Release

1. Release Candidate Preparation

1. [Maintainer]

- Create `release-X.Y.Z` branch from `stable` (during the first RC preparation only).

2. [Contributor]

- Create `rc-X.Y.Z.rcN` branch from `release-X.Y.Z` (N starts from 1 and is incremented for each new RC).
- Apply necessary changes from `master` (either merge or cherry-pick).
- *(optional)* [indy-node] Set **stable** (just X.Y.Z) `indy-plenum` version in `setup.py`.
- Set the package version `./bump_version.sh X.Y.Z.rcN`.
- Commit, push and create a PR to `release-X.Y.Z`.

3. Until PR is merged:

1. [build server]

- Run CI for the PR and notifies GitHub.

2. [Maintainer]

- Review the PR.
- Either ask for changes or merge.

3. [Contributor]

- *(optional)* Update the PR if either CI failed or reviewer asked for changes.
- *(optional)* [indy-node] Bump `indy-plenum` version in `setup.py` if changes require new `indy-plenum` release.

2. Release Candidate Acceptance

Note If any of the following steps fails new release candidate should be prepared.

1. [Maintainer]

- **Start release candidate pipeline manually.**

2. [build server]

- Checkout the repository.
- Publish to PyPI as `X.Y.Z.rcN`.
- Bump version locally to `X.Y.Z`, commit and push as the `release` commit to remote.
- Build debian packages:
 - for the project: source code version would be `X.Y.Z`, debian package version `X.Y.Z~rcN`;
 - for the 3rd party dependencies missed in the official debian repositories.
- Publish the packages to `rc-latest` debian channel.

- [indy-node] Copy the package along with its dependencies (including `indy-plenum`) from `rc-latest` to `rc` channel.
 - [indy-node] Run system tests for the `rc` channel.
 - Create **release PR** from `release-X.Y.Z` (that points to `release commit`) branch to `stable`.
 - Notify maintainers.
 - Wait for an approval to proceed. **It shouldn't be provided until release PR passes all necessary checks** (e.g. DCO, CI testing, maintainers reviews etc.).
3. [build server]
 - Run CI for the PR and notify GitHub.
 4. [QA]
 - *(optional)* Perform additional testing.
 5. [Maintainer]
 - Review the PR but **do not merge it**.
 - If approved: let build server to proceed.
 - Otherwise: stop the pipeline.
 6. [build server]
 - If approved:
 - perform fast-forward merge;
 - create and push tag `vX.Y.Z`;
 - Notify maintainers.
 - Otherwise rollback `release commit` by moving `release-X.Y.Z` to its parent.

3. Publishing

1. [build server] triggered once the `release PR` is merged
 - Publish to PyPI as `X.Y.Z`.
 - Download and re-pack debian package `X.Y.Z~rcN` (from `rc-latest` channel) to `X.Y.Z` changing only the package name.
 - Publish the package to `rc-latest` debian channel.
 - Copy the package along with its dependencies from `rc-latest` to `stable-latest` channel.
 - [indy-node] Copy the package along with its dependencies (including `indy-plenum`) from `stable-latest` to `stable` channel.
 - [indy-node] Run system tests for the `stable` channel.
 - Notify maintainers.

4. New Development Cycle Start

1. [Contributor]:

- Create PR to `master` with version bump to `X' . Y' . Z' . dev0`, where `X' . Y' . Z'` is next target release version. Usually it increments one of `X`, `Y` or `Z` and resets lower parts (check [SemVer](#) for more details), e.g.:
 - `X . Y . Z+1` - bugfix release
 - `X . Y+1 . 0` - feature release, backwards compatible API additions/changes
 - `X+1 . 0 . 0` - major release, backwards incompatible API changes

12.5.2 Hotfix Release

Hotfix release is quite similar except the following difference:

- hotfix branches named `hotfix-X.Y.Z`;
- `master` usually is not merged since hotfixes (as a rule) should include only fixes for stable code.

Indy File Folders Structure Guideline

Indy-node service works with some files and folders on the file system. We need to be careful when selecting this files and folders or adding new ones.

13.1 1. Use system-specific files and folder for indy-node service

As of now, we support indy-node service (using systemctl) on Ubuntu only. But in future we will support more platforms (CentOS, Windows Server, etc.)

So, we should follow the following principles:

- Use system-specific folder for storing config files
 - Indy-config files
 - other config files (such as service config)

Ubuntu: /etc/indy

- Use system-specific folder for storing data, such as
 - ledger (transaction log, states)
 - Genesis transaction files
 - Node keys (transport and BLS)

Ubuntu: /var/lib/indy

- Use system-specific folder to store log files

Ubuntu: /var/log/indy

- Avoid using /home folder for indy-node service

13.2 2. Organize file folders to support possibility to work with multiple networks (live, test, local, etc.)

We may have multiple Networks (identified by different genesis transaction files) installed for the same indy-node service. The file structure should support it.

- The current Network to work with is specified explicitly in the main config file (NETWORK_NAME=):

Ubuntu: `/etc/indy`

- Separate config files for each Network

Ubuntu: `/var/lib/indy/{network_name}`

- Separate data for each Network

- Separate ledgers (transaction log, states)
- Separate genesis transaction files
- Separate Node keys (transport and BLS)

Ubuntu: `/var/lib/indy/{network_name}`

- Separate log files for each Network

Ubuntu: `/var/log/indy/{network_name}`

13.3 3. Set proper permissions for files and folders

Make sure that all data, and, especially, keys have proper permissions. Private keys can only be read by the user the service is run for (indy user usually)

13.4 4. Provide a way to override config and other data for different networks

Each Network may have its own config extending base config.

Ubuntu:

- `/etc/indy` - base config
- `/etc/indy/{network_name}` - config extension

13.5 5. Client should use /home folder

Client doesn't need a service, and should use its own home directory for files with proper permissions.

Indy-sdk uses `~/indy_client`

13.6 6. Separate node and client folders

Client and Node should be two independent products not sharing any common files/folders

13.7 7. It should be possible to work with both node and client (libindy) on the same machine

We may install and work with both node and client on the same machine independently

13.8 8. It should be possible to run client (libindy) for multiple users

We may have multiple users working with client code on the same machine. Each user must have separate data and key files with proper permissions.

Code quality requirements guideline

Please make sure that you take into account the following items before sending a PR with the new code:

14.1 General items

- Consider sending a design doc into `design` folder (as markdown or PlantUML diagram) for a new feature before implementing it.
- Make sure that a new feature or fix is covered by tests (try following TDD)
- Make sure that documentation is updated according to your changes (see `docs` folder). In particular, update *transactions* and *requests* if you change any transactions or requests format.
- Follow incremental re-factoring approach:
 - do not hesitate to improve the code
 - put TODO and FIXME comments if you see an issue in the code
 - log tickets in *Indy Jira* if you see an issue.

14.2 Code quality items

The current code is not perfect, so feel free to improve it.

- Follow PEP 8 and keep same code style in all files.
 - Naming. Snake case for functions, camel case for classes, etc.
 - Give names starting with the underscore to class members that are not public by their intention. Don't use these members outside of the class which they belong to.
 - Docstrings for all public modules and functions/methods. Pay special attention to describing argument list

- Annotate types for arguments of public functions
 - Annotate return type for public functions
 - Use flake8 in developer environment and CI (fail a build if the flake8 check does not pass). (you can run `flake8` on the project root to check it; you can install flake8 from pypi: `pip install flake8`)
- Inheritance and polymorphism
 - Use ABC when creating abstract class to ensure that all its fields implemented in successors.
 - Prefer composition instead of multiple inheritance
 - Avoid creating deep hierarchies
 - Avoid type checks (`type(instance) == specific_subclass`), we have them in some places and it is awful, do polymorphism instead
 - If there is an interface or abstract class then it should be used everywhere instead of specific instance
 - Make sure that methods of subclass are fully compatible with their declarations in interface/abstract class
- Separation of concerns
 - Follow separation of concerns principles
 - Make components as independent as possible
- Avoid high coupling
 - Some classes have references to each other, for example Node and Replica, so we get a kind of “spaghetti code”
- Use classes instead of parallel arrays
- Clear and not duplicated Utilities
- Decomposition
 - There are functions and classes which are too long. For example class Node has more than 2000 lines of code, this complicates understanding of logic.
 - This can be solved by destructuring of such classes or functions on a smaller one by aggregation, composition and inheritance
- No multiple enclosed if-elif-else statements
 - It’s hard to read the code containing multiple enclosed if-else statements
 - Consider checking some conditions at the beginning of a method and return immediately
- Choose good variable and class names
 - Make variable and name classes to be clear, especially for Public API
 - Avoid unclear abbreviations in public API
- Use asyncio instead of callbacks
 - It looks and behaves just like a method call - you know that further code is not executed until coroutine completed.
 - It requires no callback preparation - code is cleaner
 - It does not break context - exception raised in a coroutine has the same stack trace as an exception raised in a method (while the one raised in callback don’t)
 - It is easy to chain coroutine calls

- If it is needed you can easily wrap coroutine in a Future and execute it in non-blocking fashion
- Consider using Actor models (this is our long-term goal for re-factoring to achieve better code quality and performance)
- Think about performance of the pool

14.3 Test quality items

- Write good tests
 - Not all of our tests are clean and clear enough
 - Follow TDD in writing the tests first
 - Have Unit tests where possible/necessary
 - Avoid using ‘module’ level fixtures if possible
- Use indy-sdk in all tests where possible. If it’s not possible to use indy-sdk for some reasons (for example, indy-sdk doesn’t have a required feature), then log a ticket in [Indy-sdk Jira](#) and put a TODO comment in the test.
- Use `txnPoolNodeSet`, not `nodeSet` fixture. `nodeSet` will be deprecated soon.
- Importing fixtures
 - Fixture imports are not highlighted by IDEs, that’s why they can be accidentally removed by someone. To avoid this try to move fixture to `conftest` of the containing package or `conftest` of some of higher-level package - this allows fixture to be resolved automatically. If you still import fixtures, mark such the imports with the following hint for IDE not to mark them as unused and not to remove them when optimizing imports:

```
# noinspection PyUnresolvedReferences
```
- Try to use the same config and file folder structure for integration tests as in real environment. As of now, all tests follow the same file folder structure (see [indy-file-structure-guideline](#)) as Indy-node service, but the folders are created inside `tmp` test folder.

There are scripts that can help in setting up environment and project for developers. The scripts are in [dev-setup](#) folder.

Note: as of now, we provide scripts for Ubuntu only. It's not guaranteed that the code is working on Windows.

- One needs Python 3.5 to work with the code
- We recommend using Python virtual environment for development
- We use pytest for unit and integration testing
- There are some dependencies that must be installed before being able to run the code

15.1 Quick Setup on Ubuntu 16.04:

This is a Quick Setup for development on a clean Ubuntu 16.04 machine. You can also have a look at the scripts mentioned below to follow them and perform setup manually.

1. Get scripts from [dev-setup-ubuntu](#)
2. Run `setup-dev-python.sh` to setup Python3.5, pip and virtualenv
3. Run `source ~/.bashrc` to apply virtual environment wrapper installation
4. Run `setup-dev-depend-ubuntu16.sh` to setup dependencies (libindy, libindy-crypto, libsodium)
5. Fork [indy-plenum](#) and [indy-node](#)
6. Go to the destination folder for the project
7. Run `init-dev-project.sh <github-name> <new-virtualenv-name>` to clone indy-plenum and indy-node projects and create a virtualenv to work in
8. Activate new virtualenv workon `<new-virtualenv-name>`
9. [Optionally] Install Pycharm
10. [Optionally] Open and configure projects in Pycharm:

- Open both indy-plenum and indy-node in one window
- Go to File -> Settings
- Configure Project Interpreter to use just created virtualenv
 - Go to Project: <name> -> Project Interpreter
 - You’ll see indy-plenum and indy-node projects on the right side tab. For each of them:
 - * Click on the project just beside “Project Interpreter” drop down, you’ll see one setting icon, click on it.
 - * Select “Add Local”
 - * Select existing virtualenv path as below: /bin/python3.5 For example: /home/user_name/.virtualenvs/new-virtualenv-name>/bin/python3.5
- Configure Project Dependency
 - Go to Project: <name> -> Project Dependencies
 - Mark each project to be dependent on another one
- Configure pytest
 - Go to Configure Tools -> Python Integrated tools
 - You’ll see indy-plenum and indy-node projects on the right side tab. For each of them:
 - * Select Py.test from the ‘Default test runner’
- Press Apply

15.2 Detailed Setup

15.2.1 Setup Python

One needs Python 3.5 to work with the code. You can use `dev-setup/ubuntu/setup_dev_python.sh` script for quick installation of Python 3.5, pip and virtual environment on Ubuntu, or follow the detailed instructions below.

Ubuntu

1. Run `sudo add-apt-repository ppa:deadsnakes/ppa`
2. Run `sudo apt-get update`
3. On Ubuntu 14, run `sudo apt-get install python3.5` (python3.5 is pre-installed on most Ubuntu 16 systems; if not, do it there as well.)

CentOS/Redhat

Run `sudo yum install python3.5`

Mac

1. Go to python.org and from the “Downloads” menu, download the Python 3.5.0 package (python-3.5.0-macosx10.6.pkg) or later.
2. Open the downloaded file to install it.
3. If you are a homebrew fan, you can install it using this brew command: `brew install python3`
4. To install homebrew package manager, see: [brew.sh](#)

Windows

Download the latest build (pywin32-220.win-amd64-py3.5.exe is the latest build as of this writing) from [here](#) and run the downloaded executable.

15.2.2 Setup Libsodium

Indy also depends on libsodium, an awesome crypto library. These need to be installed separately.

Ubuntu

1. We need to install libsodium with the package manager. This typically requires a package repo that's not active by default. Inspect `/etc/apt/sources.list` file with your favorite editor (using `sudo`). On ubuntu 16, you are looking for a line that says `deb http://us.archive.ubuntu.com/ubuntu xenial main universe`. On ubuntu 14, look for or add: `deb http://ppa.launchpad.net/chris-lea/libsodium/ubuntu trusty main` and `deb-src http://ppa.launchpad.net/chris-lea/libsodium/ubuntu trusty main`.
2. Run `sudo apt-get update`. On ubuntu 14, if you get a GPG error about public key not available, run this command and then, after, retry `apt-get update`: `sudo apt-key adv --keyserver keyserver.ubuntu.com --recv-keys B9316A7BC7917B12`
3. Install libsodium; the version depends on your distro version. On Ubuntu 14, run `sudo apt-get install libsodium13`; on Ubuntu 16, run `sudo apt-get install libsodium18`

CentOS/Redhat

Run `sudo yum install libsodium-devel`

Mac

Once you have homebrew installed, run `brew install libsodium` to install libsodium.

Windows

1. Go to <https://download.libsodium.org/libsodium/releases/> and download the latest libsodium package (libsodium-1.0.8-mingw.tar.gz is the latest version as of this writing).
2. When you extract the contents of the downloaded tar file, you will see 2 folders with the names libsodium-win32 and libsodium-win64.

3. As the name suggests, use the libsodium-win32 if you are using 32-bit machine or libsodium-win64 if you are using a 64-bit operating system.
4. Copy the libsodium-x.dll from libsodium-win32\bin or libsodium-win64\bin to C:\Windows\System or System32 and rename it to libsodium.dll.

15.2.3 Setup Indy-Crypto

Indy depends on [Indy-Crypto](#).

There is a deb package of libindy-crypto that can be used on Ubuntu:

```
sudo apt-get update
sudo apt-get install apt-transport-https ca-certificates
apt-key adv --keyserver keyserver.ubuntu.com --recv-keys CE7709D068DB5E88
sudo add-apt-repository "deb https://repo.sovrin.org/deb xenial master"
sudo apt-get update
sudo apt-get install libindy-crypto
```

See [Indy-Crypto](#) on how it can be installed on other platforms.

15.2.4 Setup RocksDB

Indy depends on RocksDB, an embeddable persistent key-value store for fast storage.

Currently Indy requires RocksDB version 5.8.8 or higher. There is a deb package of RocksDB-5.8.8 and related stuff that can be used on Ubuntu 16.04 (repository configuration steps may be skipped if Indy-Crypto installation steps have been done):

```
# Start of repository configuration steps
sudo apt-get update
sudo apt-get install apt-transport-https ca-certificates
apt-key adv --keyserver keyserver.ubuntu.com --recv-keys CE7709D068DB5E88
sudo add-apt-repository "deb https://repo.sovrin.org/deb xenial master"
# End of repository configuration steps
sudo apt-get update
sudo apt-get install libbz2-dev \
    zlib1g-dev \
    liblz4-dev \
    libsnappy-dev \
    rocksdb=5.8.8
```

See [RocksDB](#) on how it can be installed on other platforms.

15.2.5 Setup Libindy

Indy needs [Libindy](#) as a test dependency.

There is a deb package of libindy that can be used on Ubuntu:

```
sudo add-apt-repository "deb https://repo.sovrin.org/sdk/deb xenial stable"
sudo apt-get update
sudo apt-get install -y libindy
```

See [Libindy](#) on how it can be installed on other platforms.

15.2.6 Using a virtual environment (recommended)

We recommend creating a new Python virtual environment for trying out Indy. A virtual environment is a Python environment which is isolated from the system's default Python environment (you can change that) and any other virtual environment you create.

You can create a new virtual environment by:

```
virtualenv -p python3.5 <name of virtual environment>
```

And activate it by:

```
source <name of virtual environment>/bin/activate
```

Optionally, you can install virtual environment wrapper as follows:

```
pip3 install virtualenvwrapper
echo ' ' >> ~/.bashrc
echo '# Python virtual environment wrapper' >> ~/.bashrc
echo 'export VIRTUALENVWRAPPER_PYTHON=/usr/bin/python3' >> ~/.bashrc
echo 'export WORKON_HOME=$HOME/.virtualenvs' >> ~/.bashrc
echo 'source /usr/local/bin/virtualenvwrapper.sh' >> ~/.bashrc
source ~/.bashrc
```

It allows to create a new virtual environment and activate it by using

```
mkvirtualenv -p python3.5 <env_name>
workon <env_name>
```

15.2.7 Installing code and running tests

Activate the virtual environment.

Navigate to the root directory of the source (for each project) and install required packages by

```
pip install -e .[tests]
```

If you are working with both indy-plenum and indy-node, then please make sure that both projects are installed with -e option, and not from pypi (have a look at the sequence at `init-dev-project.sh`).

Go to the folder with tests (either indy-plenum, indy-node/indy_node, indy-node/indy_client or indy-node/indy_common) and run tests

```
pytest .
```


CHAPTER 16

List of breaking changes for migration from indy-node 1.3 to 1.4

-**Another format of REPLY for “WRITE” transactions.** New REPLY format you can see in: [REPLY for “WRITE” txns](#). The main goal for this changes is to divide data and metadata.

-**Change key of ReqIdrToTxn stores.** For now, ReqIdrToTxn store request from client as a map of: (digest) -> (ledger_id<some_delimiter>txn seq_no). Digest is a sha256 hash of (reqId, DID, operation). Also, digest was moved to transaction metadata section.

-**Added tag for CLAIM_DEF.** ‘TAG’ is a part of key for state tree. If this field exists in incoming transaction, then it will be used. If not, then default will be used. Full CLAIM_DEF format description can be found in [CLAIM_DEF txn format](#) and [CLAIM_DEF request format](#) All of other request specifications are in: [Request](#)

-**Migration fom levelDB to RocksDB.** Rocksdb has a native ‘seek’ methods, more flexible, has snapshots support and works well on Windows. For now, all of internal node’s storage use this db as a default key-value storages.

Migration scripts, which would be called automatically during pool upgrade procedure (while sending POOL_UPGRADE transaction):

- **Change REPLY format for WRITE transactions and using digest as a key:** [1_3_428_to_1_3_429.py](#)
- **From levelDb to Rocksdb:** [1_3_396_to_1_3_397.py](#)

16.1 Significant notes about upgrade steps:

- Please mention that Pool upgrade should be performed simultaneously for all nodes due to txn format changes.
- All indy-cli pools should be recreated with actual genesis files.

16.2 CLI Upgrading:

Old CLI (indy):

upgrade from 1.3 to 1.4 version delete `~.indy-cli/networks/<network_name>/data` folder replace both old genesis files by new ones (from 1.4 node or from sovryn repo)

New CLI (`indy-cli`):

upgrade from 1.4 to 1.5 version recreate indy-cli pool using 1.4 pool genesis file (from 1.4 node or from sovrin repo)