ETSI GS PDL 015 V1.1.1 (2023-01)

Permissioned Distributed Ledger (PDL); Reputation management

Disclaimer

The present document has been produced and approved by the Permissioned Distributed Ledger (PDL) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference
DGS/PDL-0015_Reputation

Keywords
algorithm, keyword, PDL

ETSI

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

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

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

Contents

Intellectual Property Rights

Essential patents

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

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

Trademarks

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

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

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Permissioned Distributed Ledger (PDL).

Modal verbs terminology

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

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

Executive summary

The present document discusses the use and application of reputation in PDL. The aspects of reputation discussed include:

a) The meaning of reputation.

b) Representation of reputation and the use of a normalized score.

c) Types of reputation with specific focus on:

   i) Reputation based on technical performance and adherence to service level commitments.

   ii) Reputation based on behaviour and conformance with standards and regulations.

d) The use of reputation when conducting PDL related activities.

The present document also defines methods for deriving reputation based on heuristics and measurement of performance levels.

Introduction

The present document discusses the use and applicability of reputation in PDL. The main content is broken down to four clauses as described herewith:

a) **Clause 4** defines the term and discusses the different types of reputation with respect to technology and ETSI deliverables. Focus is given to assignment of reputation to objects of different types and methods of presentation.

b) **Clause 5** defines the use of reputation with focus on indicators such as:

   i) *Quality of Service*, indicating a score based on performance of service against defined targets.

   ii) *Trustworthiness*, indicating the involvement of the object in fraudulent activities.

   iii) *Commercial reliability or stability* indicating the object's solidity when it comes to financial matters.

   iv) This clause also discusses *objective scores*, based on measurable attributes, and *subjective scores* based on perception and unmeasurable attributes.

c) **Clause 6** discusses and defines the mathematical formulas used for calculating reputation based on actual performance with focus on the duration historical events have effect on current reputation score. Such as everlasting, linear decay and logarithmic decay.

d) **Clause 7** discusses GDPR aspects of reputation and the way to ensure compliance with such requirements.

The present document is a Group Specification and as such each of the clauses includes requirements (mandatory, recommended, optional) that need to be fulfilled for an ETSI compliant PDL reputation to be defined and managed.

1 Scope

The present document discusses and specifies:

a) The meaning of reputation.

b) Representation of reputation and the use of a normalized score.

c) Types of reputation with specific focus on:

   i) Reputation based on technical performance and adherence to service level commitments.

   ii) Reputation based on behaviour and conformance with standards and regulations.

d) The use of reputation when conducting PDL related activities.

2 References

2.1 Normative references

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

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference/.

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

The following referenced documents are necessary for the application of the present document.

Not applicable.

2.2 Informative references

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

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

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

[i.1] Cambridge Dictionary definition of the term "reputation".

NOTE: Available at https://dictionary.cambridge.org/dictionary/english/reputation.

[i.2] Merriam-Webster Dictionary definition of the term "reputation".

NOTE: Available at https://www.merriam-webster.com/dictionary/reputation.

[i.3] "Reputation and its risks", Robert G. Eccles, Scott C. Newquist, and Roland Schatz, Harvard Business review, February 2007.

NOTE: Available at https://hbr.org/search?search_type=search-all&term=reputation+and+its+risks.

[i.4] Recommendation ITU-T G.107 (June 2015): "The E-model: a computational model for use in transmission planning".

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

object: device, an entity or a functionality that can be identified and defined

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

API	Application Programming Interface
CA	Certification Authority
D&B	Dun & Bradstreet ™
CIBIL	Credit Information Bureau (India) Limited
CQE	Conversational Quality Estimation
DLT	Distributed Ledger Technology
ETSI	European Telecommunications Standards Institute
EWMA	Exponential Weighted Moving Average
GDPR	General Data Protection Regulation
GoB	Good or Better
ISG	Industry Specification Group
ITU	International Telecommunication Union
MOS	Mean Opinion Score
PDL	Permissioned Distributed Ledger
PoW	Poor or Worse
SLA	Service Level Agreement
SLS	Service Level Specifications
SP	Service Provider
VoIP	Voice over IP
WMA	Weighted Moving Average
ZKP	Zero Knowledge Proof

4 Definition of Reputation

4.1 Introduction

In most PDL ecosystems quantifiable and verifiable reputation represents significant economic and operational value to the ecosystem participants/members or their delegates. While in human interactions a person would have more trust in another person or an entity with higher reputation, when it comes to digital systems such trust needs to be represented in a manner readable and usable by a machine so an algorithm of some sort can use such representation when making reputation-related decisions. Such decisions may include selection of vendors (where the algorithm may prefer a vendor with higher reputation or may consider reputation as one of multiple weighted factors such as price, delivery timelines, SLA, etc.). Reputation may be presented as a single metric but may also represent different metrics. E.g. an object may be assessed by its SLA reputation, financial stability reputation and trustworthiness reputation where each may have a certain effect on the final score calculated by an algorithm.

Certain entities offer reputation scores of various types. Banks typically define credit scores to their customers. Other companies provide scores to entities such as commercia companies and even countries. One of the main drawbacks of scores issued by such entities is the lack of transparency into the algorithms used to derive the score, and uncertainty about the motives or trustworthiness of the issuer of such scores.

PDL-based reputation scores offer a way to overcome both the issue of transparency and the uncertainty related to trust. PDL based algorithms are transparent and trust is embedded in PDL.

4.2 Etymology

4.2.1 Definition of Reputation

The Cambridge Dictionary defines reputation as *"the opinion that people in general have about someone or something, or how much respect or admiration someone or something receives, based on past behaviour or character"* [i.1].

The Meriam-Webster dictionary [i.2] gives the term Reputation three similar meanings with slightly different contexts:

a) overall quality or character as seen or judged by people in general;

b) recognition by other people of some characteristic or ability;

c) a place in public esteem or regard: good name.

Combining both sources reputation can be described as: Judgement of the behaviour of an object (object A) by another object (object B) which may affect decisions made by object B with regards to object A.

4.2.2 Branding and Reputation

Discussion of the difference between Reputation and Branding. A brand may have a reputation, meaning reputation is one of the attributes describing a brand. It is not uncommon that companies re-brand themselves and as a result their reputation may or may not follow to the new brand. e.g. an ill-repute brand may re-brand itself in an attempt to get rid of its poor reputation and start off with a clean slate. Another example could be two brands that merge into one and choose to retain the brand that has the better reputation. It is also not uncommon to see a company offering different brands in different geographies based on the reputation such brands have established in said geographies.

4.3 The value of reputation

In most PDL ecosystems quantifiable and verifiable reputation represents significant economic and operational value to the ecosystem participants/members or their delegates. The present document defines how ETSI ISG PDL manages a quantifiable and verifiable reputation framework.

[O1] The reputation of an ETSI ISG PDL entity **MAY** reflect:

a) The quality of its products and services.

b) Its trustworthiness as business/operational entity.

c) Its level of engagement in the PDL operations.

d) Additional factors not listed above.

There may be relationships among these different types of reputations, but there are important distinctions that determine how they are to be derived and used. As an example, an object (such as a node or a user) may be trustworthy when it comes to fraudulent activities, but perform poorly when it comes to its ability to process data on a timely manner. As such it will have a low Quality of Service score and a high Trustworthiness score. However – there may be scenarios where an object's slow performance or communication errors may cause it to appear as if it is injecting errors on purpose and that may also reduce its Trustworthiness score.

By leveraging the data integrity and time linearization properties of PDLs, it is possible to automate the collection, organization, and use of certain reputation information. This, in turn, can be used to make the PDL platform operations more transparent and reward those participants who offer a high-quality service and conduct their business fairly and according to governance rules.

In the context of a PDL platform the perception or opinion discussed in clause 4.1 is not necessarily held by people but may rather be held by systems or machines. The past behaviour or character is then measured against specific criteria and is expressed in comparable and readable numerical terms. The present document discusses both the methods by which such behaviour and character can be defined and measured, as well as possible ways to represent reputation in a normalized and comparable format.

The use of reputation as a measurable attribute of an object or an entity that can then determine actions related to such object or entity is the core value of reputation in a PDL platform.

Robert G. Eccles, Scott C. Newquist, and Roland Schatz discuss reputation in an article in the Harvard Business review [i.3] stating that "… *strong positive reputations attract better people. … perceived as providing more value.... organizations are especially vulnerable to anything that damages their reputations*". Leading to a conclusion that a good reputation may be an indicator of both the ability to perform well as well as the ability to sustain such performance over time. On the other hand, it also leads to the conclusion that damage to reputation not only indicates that an entity or object does not perform well, but it may also restrict such object or entity from being included in certain activities.

The resulting conclusion is that reputation offers value to both the objects or entities being measured and the objects or entities using those measurements in order to take decisions or actions.

4.4 Assignment of reputation to objects in a PDL platform

4.4.1 Assignment of Reputation to a PDL node

4.4.1.1 Service Level related reputation

A PDL node can be assigned a reputation score representing certain Service Level related attributes which measure its adherence with defined/expected operational behaviour. Examples would be:

a) Uptime of the node.

b) Responsiveness of the node (e.g. time of data processing calculations and sending a response).

c) Additional attributes defined in a service level agreement defined by the governance.

These are discussed in detail in clause 5.1.2 herewith.

4.4.1.2 Trustworthiness related reputation

A PDL node can be assigned a reputation score based on its trustworthiness within the context of PDL consensus operations and general calculations. Examples would be:

a) Involvement of the node in fraudulent activity.

b) Ability of the node to maintain proper security measures against external fraudulent activity.

These are discussed in detail in clause 5.1.3 herewith.

4.4.1.3 Commercial reputation

A PDL node can be assigned a reputation score representing its payment and financial stability and behaviour. Such score may be linked to external credit score rating entities.

In the context of PDL payment and financial stability are related to:

a) Cryptographic transactions performed using a blockchain based crypto-currency.

b) Token based fiat transactions where certain details of the transaction are recorded on-chain.

These are discussed in detail in clause 5.1.4 herewith.

4.4.2 Assignment of Reputation to entities

Entities, in the context of a PDL platform, could be node operators, external storage providers, virtual or physical hosts and others. Such entities may operate one or more device that is involved directly or indirectly in PDL operations.

Assignment of reputation to such entities can be broken down in a similar manner to that of nodes: Service Level related and Trustworthiness related. The main difference is that entities may operate more than one device and the reputation of an entity affects all devices/nodes it operates, or all objects included in it.

4.4.3 The significance of reputation of objects

As per the etymological definition of reputation earlier in this clause, reputation may impact the opinion and behaviour of certain objects towards other objects. As a result, reputation may impact the likelihood of specific objects to be used/selected for certain tasks. Examples could be:

a) A node with poor trustworthiness reputation may be banned from taking part in certain consensus operations.

b) An entity with poor commercial reputation may be less likely to receive orders from potential customers who may consider them as a financial risk.

4.5 Disengagement of Reputation from Commercial/Monetary value

4.5.1 Representation of reputation as a metric

Measuring reputation should preferably yield a score with a numerical value. Even in environments with a binary behaviour (e.g. "Operational" vs. "Non-operational") a reputation score can be achieved over time by comparing the number or duration of the binary options thus yielding a score of "290 out of 300 samples were operational". In other environments, for example temperature-controlled environments, a score can represent the average temperature and the number of times or duration the temperature exceeded the min/max thresholds.

4.5.2 Binary Reputation vs. Score-based reputation

In a binary reputation scenario, an object can be tagged as "reputable" or "irreputable" and will then be considered for inclusion in or exclusion from key operations (consensus votes, hash calculations, etc.).

In a score-based reputation scenario an object has a reputation which is somewhere between a minimum and a maximum value and may then be considered for inclusion or exclusion from key operations based on its score. E.g. in the case of a platform with, say, 8 nodes and a governance rule stating that a minimum of 5 nodes is required for a vote to be valid, the governance may select the 5 nodes with the highest reputation score.

Typically, a lower score represents lower reputation, and a higher score represents a higher reputation. However, the governance has the prerogative to define the opposite. This is useful when the score, as a numerical value, is used for calculations related to the eligibility of an object to participate in key activities.

[D1] In a Score Based reputation scenario the governance **SHOULD** define the lower limit and the upper limit of the reputation score and the meaning of such limits related to the use of the score.

4.5.3 Normalized reputation score/Metric

The examples presented in clause 4.4.1 represent the need for normalization of the metrics representing the score. When comparing samples of different populations (e.g. 300 samples in one population vs. 500 samples in another) the mere number of samples does not represent the true behaviour of one population compared to the other. A normalized value, such as percentage, will be more useful. It is thus recommended that reputation is represented in a normalized manner, be that percentage, a range between 0 % to 10 %, between 0 % to 1 % or any other convention agreed upon or decided by the governance.

[D2] Score Based reputation **SHOULD** use a Normalized Metric.

[R1] In a platform using Score Based reputation all nodes **SHALL** use the same Normalized Metric.

When defining **[D1]** a normalized metric, the following factors need to be agreed:

a) The value representing the lowest reputation. That will typically be Zero.

b) The value representing the maximum reputation. That would typically be 1 or a representation of a "Whole unit" in the respective numerical system (e.g. 100 %).

c) The resolution of details. That will typically be represented by the number of decimal positions to be captured and calculated. E.g. In a percentage representation between 0 % to 100 % it can be agreed that values are represented down to a resolution of 1 % (meaning two decimal positions: e.g. 0,23, 0,58, 0,99) or a resolution of 0,1 % (meaning a resolution of three decimal positions: e.g. 0,231, 0,578, 0,989).

5 Use of Reputation

5.1 Reputation as an indicator of performance metrics

5.1.1 Types of Quantifiable and Verifiable Reputation

[O2] The reputation of an ETSI ISG PDL entity **MAY** reflect:

a) The quality of its products and services. This is defined as Service Quality Reputation.

b) Its trustworthiness as business/operational entity. This is defined as Trustworthiness Reputation.

c) Its level of engagement in the PDL operations.

d) Additional factors not listed above.

5.1.2 Service Quality Reputation

Focused on the service provided by PDL participants to their respective customers. Enables informed decisions based on objective performance data.

Service Quality is specified by the Governance and defines measurable target performance levels for certain attributes and the methods by which they are measured. Table 1 defines some of the attributes that may be measured and rated. The present document proposes targets for demonstration purposes, but the governance may use other targets as it sees fit.

Table 1: Service Quality Reputation Attributes

Attribute	Definition	Method of measurement	Proposed target
Downtime	The percentage of time over a certain period during which the object being measured was not operational.	Measuring the availability of the object in pre-defined intervals by performing a transaction. When an object does not respond withing a pre-defined timeframe it is considered not operational. See note 1.	Proposed interval: Lower than half the duration of a typical transaction on said platform. Proposed allowed downtime: equal or lower than 0,001 of total time. Proposed period of measurement: Monthly.
Uptime	The percentage of time over a certain period during which the object being measured was operational.	Measuring the availability of the object in pre-defined intervals by performing a transaction. When an object responds withing a pre-defined timeframe it is considered operational. See notes 2 and 3.	Proposed interval: Lower than half the duration of a typical transaction on said platform. Proposed allowed downtime: equal or lower than 0,001 of total time. Proposed period of measurement: Monthly.
Object Responsiveness	The time difference between the moment an object receives a task and the moment it completes performing that task.	Measuring To (the time when the object has received a task). Measuring Tc (the time when the object has completed the task). Calculating Responsiveness = Tc-To.	Proposed responsiveness < (Period of measurement) divided by (number of transactions expected to occur during that period) e.g. 1 (minute) / 6 (transactions per minute).
System Responsiveness	The time difference between the moment a task is sent by a requester to an object and the moment the results of that task are received by the requester. See note 4.	Measuring Ts (the time when a request sends a task to an object). Measuring Tr (the time when the response from the object has been received by the requestor). Calculating Responsiveness = Tr-Ts.	Proposed responsiveness < (Period of measurement) divided by (number of transactions expected to occur during that period) e.g. 1 (minute) / 6 (transactions per minute).
Transaction Loss	The number of transactions that were not correctly completed during a period of time.	Measuring Nd (the number of transactions distributed for processing during a certain period). Measuring Np (the number of transacted completed during a certain period) Calculating Transaction Loss = Nd - Np.	This is a discrete value and is it proposed that a normalized value (Transaction Loss Rate) is used instead.
Transaction Loss Rate	A normalized representation of Transaction Loss as a fraction.	Transaction Loss Rate = 1 - (Nd - Np)/Nd.	Proposed representation using percentage. Target value depends on criticality of service and tolerance to transaction loss.

NOTE 1: The transaction may be as simple as a "ping" or a more complex API call requesting the object to report its status.
NOTE 2: In a normalized system Uptime = 1 – Downtime.
NOTE 3: The transaction may be as simple as a "ping" or a more complex API call requesting the object to report its status.
NOTE 4: This includes both the Object Responsiveness and transmission induced delays due to geography or communication network congestion (latency).

[D3] The governance of a platform and an application developer **SHOULD** jointly decide which attributes need to be included (values, units, measurement methods, data model) in the measured and claimable service quality reputation for the respective application.

[R2] An application's reputation data **SHALL** be made available (while maintaining GDPR compliance) to platform members using the respective application.

[R3] The service quality reputation score for each application **SHALL** be standardized and derived from the data processed in such application.

[O3] The service quality reputation score **MAY** be application specific.

5.1.3 Trustworthiness Reputation

In day-to-day life Trustworthiness represents the level of trust a party or a person has in another party/person. In a network environment such as a PDL such party may be any object whose behaviour can be measured in a manner representing its ability or tendency to be truthful.

The reasons reputation is associated with trustworthiness include:

- Encouragement if continued and increasing participation in honest network activities.
- Provides an objective assessment of how participants are interacting with the PDL platform.
- Provides an objective assessment of an object's ability to maintain proper security measures against external fraudulent activity.

The attributes and objectives representing Trustworthiness are discussed in table 2.

Table 2: Trustworthiness Reputation Attributes

Objective/Attribute	Description	Representation	Examples
Involvement in fraudulent activities	The level of involvement of an object in fraudulent activities.	The number of incidents where an object had been involved in fraudulent activity. See note 1.	A PDL node had participated in a "51 % attack" attempt. A PDL is using fraudulent identity.
Ability to meet advertised/claimed capabilities	The ability of an object to meet/deliver capabilities such object claims to be able to meet/deliver.	The number of incidents where an object was not able to meet/deliver capabilities it has claimed capable of meeting/delivering. See note 2.	A PDL node that claims to be able to process 10 transactions per second fails to do so and is only able to process 8 transactions per second. A PDL does not deploy the required security measures.
Longevity	The duration which an object had been known or active.	The longer an object had been known or active would typically indicate it is trusted by its users and peers.	A known brand participating in a PDL may be considered more trustworthy than an unknown brand with limited or no historical references.
Transparency	The level of transparency offered by the object.	The more transparent an object is (through use of open source and proper documentation) the more trusted it will be.	Certain objects may act as a "black-box" using proprietary code and as such reduce the ability to identify back-doors or data leakage.
NOTE 1: May also be represented as the fraction of fraudulent activity incidents from the total number of transactions in a certain period. NOTE 2: The extent of the outcome of such incidents may be counted as well, where incidents with a severe impact may count higher than incidents with minor impact.			

In some instances, an object may perform in a certain manner that may affect its reputation in more than one reputation metric. However - if the source of such score results from the behaviour of the same attribute, the object should not be penalized twice. E.g. if low computation resources cause a delay in processing of transactions as well as missing consensus calculations - such attribute (low processing power) may influence both the Service Level reputation score and the Trustworthiness reputation Score.

[D4] An object **SHOULD NOT** be double penalized for the same attribute in more than one reputation metrics.

5.1.4 Commercial Reputation

The Commercial Reputation score represents the payment and financial stability of an entity.

The Commercial Reputation score is calculated based on payment history and credit score of an entity such as a Service Provider (SP), an enterprise customer or a consumer. The score is derived from timeliness of payments, accuracy of payments, duration and effectivity of reconciliation process compared to a target performance defined in agreements signed between the entity and its suppliers.

The score may also take into account information received from external sources such as analyst reports and publicly available financial records.

The calculation formula for the Commercial Reputation score is out of scope of the present document and is for further study.

Table 3 lists the respective objectives and attributes used to define and measure commercial reputation:

Table 3: Commercial Reputation Attributes

Objective/Attribute	Description	Representation	Examples
Payment History	The history of payments including timeliness, accuracy, number of disputes and time to resolve them.	A list of payment related events and incidents and a list of incidents where such events were out of agreed-upon norms and standards.	An entity is paying later than contracted.
Credit score	The ability of an entity to make payments.	Credit score represents the ability of an entity to make timely payments. See note.	An entity does not have sufficient balance in the bank to pay for a large volume of goods and is limited by the credit score to order a limited amount until the balance is replenished.
External scores	Commercial scores issued by external firms and entities.	This score is similar in nature to a CA. It is based on trusted external sources accepted by the PDL platform participants. Different Credit-rating entities have different representations.	E.g. D&B™, Experian™, TransUnion®, Equifax™. Some are country-specific (e.g. CIBIL™ is only valid in India). Some are operated by financial institutions (e.g. CapitalOne™).
NOTE: Credit score would typically affect the amount of down payment/deposit required by a seller from a buyer prior to delivering goods.			

5.1.5 Discussion of SLS and SLA

Service Quality is measured against a Service Level Specification (SLS) that defines the metrics and required levels of such metrics that meet such specifications. As an example, an SLS may require Transaction Loss Rate of less than 0,001 %. In such case the service quality is defined by measuring the actual transaction loss rate and comparing it to the target. Should the actual loss rate be below the target (e.g. 0,00094 %) then the SLS had been met and service quality is good (within the SLS). Should it be higher than the target (e.g. 0,0014 %) the SLS had not been met and service quality is not as good as it should be.

[D5] The performance of a PDL platform **SHOULD** be measured against an SLS.

When the SLS has not been met, there may be commercial implications based on the contractual commitment between the user and the service provider. Such contractual commitment is defined in a Service Level Agreement (SLA). The SLA defines penalties (paid by the service provider) or credits (received by the user) and other types of commercial or operational actions that need to take place in the event the SLA had not been met.

[D6] The penalties/credits **SHOULD** be calculated based on an SLA signed between the user and service provider.

NOTE: The Service Provider in requirement **[D6]** may refer to the operator of the PDL platform or to a third party offering service implemented on a PDL platform. There may be scenarios of supply chain management where one service provider operates a PDL platform (SP1 for the purpose of this example) and another Service Provider (SP2) is offering a service that is implemented on that PDL platform. SP1 and SP2 may have an SLA between them and SP2 may have a back-to-back SLA with the user. When the PDL platform does not meet the SLS SP1 pays credit to SP2, who then pays credit to the user. The commercial details of such back-to-back arrangement are out of scope of the present document.

An example of the relations between an SLS and an SLA may be that an SLS defines Uptime target levels of, say, 99,999 % (considered "Gold"), 99,99 % (considered "Silver") and 99,9 % (considered "Bronze"). Suppose the user had signed up for "Gold" quality service. The SLA defines the method by which the credit is calculated based on actual performance. Should the actual performance (uptime) be higher than 99,999 % then the user is not eligible to any credit. Should the actual uptime be lower than the target, the user is eligible to credit based on a formula that considers:

a) The duration of downtime.

b) The time of day during which the outage occurred. This is an optional metric with the rationale that an outage during a busy time of day may have a larger negative effect compared to an outage where the system has little use.

c) The difference between the target and the actual performance. This is an optional metric with the rationale that longer downtime causes significantly more harm than a short one. Thus the credit will be significantly (non-linear) higher than for a short downtime.

d) Additional factors such as history (frequency of downtime).

In practicality the SLA defines the point in time (relative to the start of an event) when the service is considered "below the SLS". For example - if the service provider offers an SLA of 100 % uptime that means their customer is eligible to SLA credits starting from the instance service went down, while if the SLA offered was 99,999 % uptime then the customer would have been eligible for credit after 0,0001 % of the measurement period had elapsed (which is about 2,6 seconds in a month) and the service provider would not have to pay penalties for shorter outages.

5.1.6 Discussion of objective and subjective scores

A score can be calculated based on objective measurements and subjective measurements or a mix of both.

An example of a subjective measurement would be a satisfaction survey where a user may express their satisfaction by selecting one of a few options from a list (e.g. Extremely satisfied, Satisfied, somehow satisfied, etc.) or entering a numerical score (e.g. 1 = extremely dissatisfied, 5 = extremely satisfied). Such surveys may cover many aspects of products and services (e.g. airline service, banking service, food quality, healthcare). The results of such surveys are collected and represented as an average score. In an automated platform the calculation is performed in an on-going manner and the average score may change with time and serve as an indication to both the providers of the service and potential users. Providers may use the score to find out if they need to improve certain aspects of their service, and users may use the score to compare service providers and select those with higher scores. An example of an objective score is the Recommendation ITU-T G.107 [i.4] R-value and MOS_{CQE} that represent a normalized score (R-value ranging from 50 to 100) based on a Mean Opinion Score (MOS) related to quality of VoIP telephone calls. The formulas are using objective measurements (e.g. delay, packet loss) as well as subjective judgement such as GoB (Good or Better) and PoW (Poor or Worse) to calculate the MOS and R-value. The main drawback of subjective scores is that they are subject to users' opinions and are at risk of being manipulated by an organized campaign to lower the score by encouraging users to assign low ratings. PDLs may be used to identify such attempts through use of heuristics and trend analysis, as well as identifying conflicting subjective scores and taking proper actions accordingly (e.g. two users are rating the quality of the same telephone call where one rates it as "excellent" and the other as "horrible". This will raise suspicion that one of them is trying to manipulate the score. While if one rates the quality as "good" and the other as "adequate" the rating will be considered unbiased, and the subjective score would be estimated as somewhere in between those values).

Objective scores are based on objective measurements only and do not take users' opinions into account. As such they use measurable attributes and well-defined measurement methods. Such scores are almost impossible to manipulate (manipulation will require meddling with the measurements) hence they would be considered more reliable.

[R4] Objective scores **SHALL** be based on measurable attributes.

[R5] Objective scores **SHALL NOT** use subjective measurements.

6 Reputation Management

6.1 Introduction to Reputation Management

Reputation of the SP is an indicator of its capabilities and reliability and may play a factor in being selected for participation in a Supply Chain. The present document describes multiple forms of reputation and how the features of a PDL can be utilized to generate trustworthy and useful reputations.

6.2 Reputation management over time

6.2.1 Introduction

The governance defines the method of reputation management.

[R6] The governance **SHALL** address the following questions when defining the reputation scoring algorithms:

a) What actions will allow an entity to earn reputation and how much?

b) What actions will force an entity to lose reputation and by how much?

c) What can an entity do with its reputation?

d) What shall reputation not be used for?

e) What is the (dis)incentive model for reputation?

f) How secure is the reputation system against malicious entities?

An example of answers to these questions is listed in annex A herewith.

6.2.2 Everlasting cumulative reputation

a) Each reputation related incident is recorded and used for calculating score regardless of the time that had elapsed.

b) Old entries have same weight as new entries.

c) The effect of each individual entry decreases as entries accumulate.

6.2.3 Time dependent reputation

6.2.3.1 General discussion and introduction

Reputation related incidents are recorded but their effect on the score diminishes with time based on number of scores recorded per unit of time, the mean half-life of the operations being measured and the time that had elapsed since the event occurred.

[R7] A participating party **SHALL** derive Objective performance from measurement of Objective service attributes and comparison to SLA commitments.

[D7] The SLA reputation **SHOULD** be a normalized score that spans between 0,0 (lowest score, complete failure to meet the SLA) and 9,9 (highest score, no SLA violations).

[R8] An entity **SHALL** generate, using the WMA(t) formula described in this clause, a normalized SLA Reputation score between 0,0 to 9,9 representing their ability to meet the metrics defined in the SLA and upload such score to the DLT through consensus.

The SLA score is calculated based on a Weighted Moving Average (WMA) algorithm that gives higher value to recent performance records compared to historical performance records. C(t) is the performance of record "t" where "t" represents the latest record that was recorded, "(t-1)" is the previous record and so on.

The WMA may apply linear or logarithmic decay formulas to represent higher or lower emphasis to recent events.

[R9] The governance **SHALL** define the type of decay used by the platform.

6.2.3.2 Logarithmic decay

All records are included in the calculation but the contribution of each entry towards the score decreases logarithmically so that recent entries have a very strong influence while old records have infinitesimal influence. This type of WMA is also known as EWMA (Exponential WMA).

$$C_t = \begin{cases} Y_0, & t = 0 \\ \alpha Y_t \quad +(1-\alpha)\cdot C_{t-1}, & t > 0 \end{cases}$$

Where:

a) The coefficient α represents the degree of weighting decrease, a constant smoothing factor between 0 and 1. A higher α discounts older observations faster.

b) *Yt* is the value at a time period *t*.

c) *Ct* is the value of the *EWMA* at any time period *t*.

6.2.3.3 Linear decay

The contribution of each entry decays in a linear fashion so that at a certain point entries that have reached a certain age are not counted any longer.

$$WMA(t) = \frac{n\cdot C(t)+(n-1)\cdot C(t-1)+\ \cdots\ +2\cdot C(t-n+2)+1\cdot C(t-n+1)}{n+(n-1)+\cdots+2+1}$$

Or in shorter form:

$$WMA(t) = \frac{\sum_{j=0}^{n-1}(n-j)\ C(t-j)}{\sum_{i=1}^{n} i} \quad ; \quad Where\ n \geq 1$$

Where *n* is the duration of the longest instance of service to date, expressed as a discrete number of the time-measurement units used to measure the service (e.g. a "by-the-second" service will be measured by the number of seconds duration of the longest instance, while a "by-the-hour" service will be measured by the number of hours duration of the longest instance, not by the number of seconds). Meaning that if the longest duration of instance of service to date has been 60,1 seconds, *n* will equal 60 and the *WMA* will be based on the last 60 records (previous records are ignored).

In the event that the number of performance records is smaller than n then n will equal the number of performance records.

[D8] An entity **SHOULD** use the number of performance records that equals the number of time measurement units in the longest instance to date.

[CR1] < [D8] An entity **SHALL** use all available records in the event that the number of available records is less than the number defined in **[D8]**.

C(t) is the performance of record *t*. It is measured and normalized to a value between 0,0 - 9,9 at the end of the service instance or at the time of measurement (which is applicable for long-term services that span longer than an agreed-upon measurement interval). The method of normalization is application specific, and the per-application definition is beyond the scope of the present document.

[R10] An entity **SHALL** make performance calculations at an interval that is less than or equal to average duration of the service type being measured.

[O4] An entity **MAY** synchronize the timing of the performance calculations and SLA measurements.

[O5] An entity **MAY** make performance calculations based on past performance information that was recorded in the PDL.

The overall score is calculated as the average of all the per-application scores.

[D9] An entity **SHOULD** calculate the average of all its per-application Reputation scores and upload to the PDL through consensus.

[O6] An entity **MAY** upload its per-application Reputation score to the PDL through consensus.

The score shall be accompanied by the number of performance measurements it is based on in the following format: *x.y:n* where *x.y* is the score and *n* is the number of performance measurements. E.g. 8.9:132.

[CR2]<{[D9] OR [O6]} SP **SHALL** specify the Reputation score together with the number of performance measurements on which the score is based using the format "*x.y:n*" where *x.y* is the score and n is the number of performance measurements.

Additional technical performance metrics and scores are out of scope of the present document and are for further study.

7 Compliance

7.1 GDPR Compliance

Although the score itself may not contain any personal data, the related attributes may contain information that cannot be revealed without violating GDPR. Therefore, when processing, storing and sharing reputation scores care should be taken to avoid violating GDPR.

[R11] All reputation scores stored on the PDL and any publicly accessible storage as well as any user information such scores are derived from **SHALL** be GDPR compliant.

[O7] ZKP **MAY** be used to selectively reveal/hide information of an entity/object in scenarios where exposure of information violates GDPR.

[R12] When an entity/object is removed from a platform/application it **SHALL** have the ability to exercise its GDPR right to be forgotten.

Annex A (normative): Example of criteria for calculating reputation

This annex provides an example of the criteria defined by the governance for calculating reputation. Other criteria may be used for other example which may imply other requirements (mandatory, recommended or optional).

Table A.1: Example of reputation calculation criteria

Questions	Answer
What actions will allow an entity to earn reputation and how much?	**[R13]** Eligible Reputation Actions on a Network **SHALL** be making a reputation affecting or affected claim: a) How: A valid claim of reputation is defined as a claim that has not been disputed until the dispute period of the claim has elapsed or any raised disputes against the claim during the dispute period have been resolved successfully for the claimant. b) Action Frequency: Any time. Amount earned: 1 unit of reputation per valid claim.
What actions will force an entity to lose reputation and by how much?	**[R14]** An entity **SHALL** lose reputation when their claim of reputation is successfully disputed during the dispute period: a) How: Validated dispute against a claim of reputation. b) Action Frequency: Any time during the dispute period of a claim. Amount lost: varying percentage up to 100 % of the quantified reputation score of an entity dependent on the severity of the malicious action.
What can an entity do with its service quality reputation?	**[D10]** Reputation **SHOULD** be indicative to other entities of the quality of service an entity member typically offers. **[O8]** Application providers **MAY** require that an entity holds a certain minimal reputation level before it may use an application. **[O9]** An entity **MAY** stake its reputation.
What shall reputation not be used for?	**[R15]** Reputation **SHALL NOT** be transferred between entities.

Questions	Answer
What is the (dis)incentive model for reputation?	Incentive models need to be carefully constructed such that the incentives do not lead to: a) crowding-out effects of desired behaviour; b) game-theoretic attacks that extract value from or locks-up value in the network. Disincentive models need to also be carefully calibrated in order to avoid significant penalties to honest entities which might be unavoidable in order to ensure that all dishonest entities also known as maximizing the grievance factor defined as the ratio of the sum of penalties of malicious actors to the sum of penalties of honest actors in case of a byzantine failure. **[R16]** For Reputation the following *Non-Programmatic* Incentive model **SHALL** be utilized*:* The more an entity delivers high-quality service, the more service quality reputation the entity will earn which in turn will improve its business position by being a trustworthy business partner. **[R17]** For Reputation the following *Programmatic* Incentive model **SHALL** be utilized: a) Positive Bonding Curve: A positive bonding curve should incentivize early adoption by rewarding even a small amount of activity, reduce the reward for each additional activity as the number of activities by an entity increases until a threshold has been crossed and the rewards increase again. In addition, the curve resets after a predefined period. The effect of such a periodic parabola shaped curve is that both low volume and high-volume adoption are incentivized, and the reset allows everyone to obtain the same rewards again for a certain period of time. Such a construct ensures more equity in network usage, does not lead to discouragement when incentives flatten, and it incentivizes high volume. b) Decay curve: "How has a certain individual done lately?" has its programmatic equivalent in a decay function which reduces the amount of reputation over time through for example an exponential function. This ensures that entities are incentivized to deliver as much high-quality service in as short a time span as possible. It is not advisable to reduce the value to zero or close to zero but rather think of the decay as a "half-life" of reputation - the time it takes for the value to be cut in half. The exact time period for such a halving need to be significantly longer than typical network cycles/periods on the one hand, however, not too long in order to not be a big enough incentive to engage more with the network. c) Static long-term reputation accounting: In order to introduce a long term measure of reputation, and similar to the EEA Trusted Reward Token, any difference between the starting and ending value of an entity's reputation over an agreed-upon time period, is added or subtracted from a long term reputation value. This long term reputation cannot be used for any form of network activity, but rather functions as a long term service quality reputation ranking parameter. This allows entities or their delegates to evaluate an entity beyond any short term effort. **[R18]** For Reputation the following *Programmatic* Disincentive model **SHALL** be utilized: a) *Programmatic*: Governance a dispute arbiter is required to be a party to all application transactions utilized to make a claim of reputation on the PDLs. Consequently, the claim proof can be readily and programmatically validated by the governance if a dispute is filed by an entity by simply recomputing the proof based on the available PDL data in the application. In addition, the governance randomly audits the reputation claims. Considering four scenarios: i) The claim proof computed by the governance does not match the claim proof submitted by the claiming entity. Outcome: The service quality reputation AND the trustworthiness reputation of the claiming entity member is slashed. ii) The required data to compute the claim proof is not available on the application/PDL or is incomplete. Outcome: The service quality reputation AND the trustworthiness reputation of the claiming entity is slashed. iii) The computed claim proof by the governance does not match the claim proof filed as part of a dispute. Outcome: the entity that filed the dispute will have their service quality AND trustworthiness reputation slashed. iv) The reputation computed claim proof does not match either the claim proof of the claiming or the disputing entity. Outcome: Both entities will have their service quality AND trustworthiness reputation slashed. *Non-Programmatic Disincentive Models are n*ot required as there is a programmatic disincentive model.

Questions	Answer
How secure is the reputation system against malicious entities?	There are four major types of economic attacks that **SHALL** be addressed: **[R19] Collusion Attacks:** Two or more entities jointly lie about the business outcome used in making a claim of service quality. Mitigations: The governance as an arbiter of disputes and performing random, programmatic audits of service quality reputation claims strongly mitigate against these types of attacks. **[R20] Discouragement Attacks:** Can be an issue at the application level, when a larger entity refuses to certify a certain level of achieved quality of a smaller entity. The only mitigation for the smaller entity is to file a formal claim of a malicious action within the governance. **[R21] Extortion Attacks:** See comments under Discouragement attacks. **[R22] Value Extraction Attacks:** Such attacks can take many forms, such as "hording" reputation, in other words not use of reputation. Given that the supply of reputation is not limited and that reputation decays over time, such an attack does not pose a threat either short or longer term. See note.
NOTE: The issues of consensus attacks on the underlying PDLs are outside the scope of the present document.	

History

Document history		
V1.1.1	January 2023	Publication