

RESEARCH PAPER

Distributed Persistent Identifiers System Design

Pavel Golodoniuc¹, Nicholas N. J. Car² and Jens Klump¹¹ CSIRO Mineral Resources, Perth, WA, AU² Geoscience Australia, Canberra, ACT, AUCorresponding author: Pavel Golodoniuc (pavel.golodoniuc@csiro.au)

The need to identify both digital and physical objects is ubiquitous in our society. Past and present persistent identifier (PID) systems, of which there is a great variety in terms of technical and social implementation, have evolved with the advent of the Internet, which has allowed for globally unique and globally resolvable identifiers. PID systems have, by in large, catered for identifier uniqueness, integrity, and persistence, regardless of the identifier's application domain. Trustworthiness of these systems has been measured by the criteria first defined by Bütikofer (2009) and further elaborated by Golodoniuc *et al.* (2016) and Car *et al.* (2017).

Since many PID systems have been largely conceived and developed by a single organisation they faced challenges for widespread adoption and, most importantly, the ability to survive change of technology. We believe that a cause of PID systems that were once successful fading away is the centralisation of support infrastructure – both organisational and computing and data storage systems.

In this paper, we propose a PID system design that implements the pillars of a trustworthy system – ensuring identifiers' independence of any particular technology or organisation, implementation of core PID system functions, separation from data delivery, and enabling the system to adapt for future change. We propose decentralisation at all levels – persistent identifiers and information objects registration, resolution, and data delivery – using Distributed Hash Tables and traditional peer-to-peer networks with information replication and caching mechanisms, thus eliminating the need for a central PID data store. This will increase overall system fault tolerance thus ensuring its trustworthiness. We also discuss important aspects of the distributed system's governance, such as the notion of the authoritative source and data integrity.

Keywords: Identifier systems; persistent identifiers; distributed systems; Distributed Hash Tables; peer-to-peer networks; PID; P2P

Introduction

Persistent identifier systems have evolved significantly in the past two decades. Applications of these systems have expanded into new domains and new technologies have been introduced. Many PID systems were developed by various communities and, for different reasons, have failed to withstand the test of time, eventually sliding into paralysis and a 'zombie' stage (Beck *et al.* 2016; Huber & Klump 2016), where identifiers continue to exist but the PID system loses its resolution service.

In our investigation of what contributed to the failure of, once successful, PID systems, we observed a common trait: the reliance of PID systems on a centralised technical infrastructure or governing authority (Huber & Klump 2016). The factors that led to the failure of PID systems included: (i) complexity of the PID systems; (ii) lack of long-term financial support from hosting organisations; (iii) narrow target community; and (iv) reliance on a single governance authority or hosting organisation. In order to create identifiers that are not reliant on a centralised system, PID systems with a decentralised resolution infrastructure have been designed where nodes in the infrastructure can use any technology as long as a protocol for operations is

adhered to. Digital Object Identifiers (DOI)¹ and the Handle² system that it is based on are good examples of distributed resolver systems.

Golodoniuc *et al.* (2016) proposed an approach to development of PID systems that combines the use of (a) the Handle system as a distributed system for the registration and initial resolution of persistent identifiers, and (b) the PID Service – a particular PID resolver system implementation – to enable fine-grained resolution of PIDs to different information object representations. The proposed approach solved the problem of guaranteed initial resolution of identifiers, but left fine-grained resolution and information delivery under the control of a single authoritative source, potentially posing risk to the long-term availability of information resources. In this paper, we develop a further approach and explore the potential of large-scale decentralisation at all levels: (i) persistent identifiers and information objects registration; (ii) identifier resolution; and (iii) data delivery. To do these things we propose using Distributed Hash Tables (DHT), Peer Exchange networks (PEX), and peer-to-peer networks.

Recent advancements in the development of Internet-based Distributed Hash Tables (Loewenstern & Norberg 2009), Peer Exchange networks (Wu *et al.* 2010), and Magnet Links³ (Farrell *et al.* 2013) have led to the development of new approaches to the identification of information and this enables their use in relation to PID resolution. This proposed system architecture is based on concepts implemented by these technologies and, although in-depth knowledge is not required, a basic understanding of their principles is essential.

To better define our vision of a future PID system design, we first review the best-known implementation of distributed technologies and then propose a PID system architecture with resolution mechanisms, linkage to information objects and distributed data delivery.

Distributed Identifier Resolution Networks Overview

The balance between overall reliability of the system components and the governance model strongly influences trustworthiness of a PID system and its common acceptance by a community of practice. Whereas reliability can generally be strengthened by added redundancy via the use of distributed networks, the choice of an appropriate governance model requires a thought since there are down sides to managing networks: at the very least, integrity of data must be assured over the multiple copies in a network. Later in this paper, we focus on the technical aspects of the solution, whereas the role governance models is discussed in more details in Car *et al.* (2017). Here, we take a deeper look at different aspects of Digital Object Identifiers, Handle System, and Peer-to-peer networks technologies.

Digital Object Identifiers and Handle System

DOI identifiers exhibit many important traits of a successful PID system and many aspects of its architecture and governance fit into the proposed four PID ‘pillars’ outlined by Golodoniuc *et al.* (2016) and further discussed in (Car *et al.* 2017). In the DOI system, identifiers are not tied to any particular underlying technology or organisation names and identifier resolution is via the Handle network. Identifier resolution is often achieved by using a known DOI web-based resolver, e.g., <http://dx.doi.org/>, but DOIs are able to be resolved by any Handle identifier, including the generic Handle web-based resolver, <http://hdl.handle.net/>. To DOI consortium also concerns itself with the information member users must store alongside the identifiers (such as the authors of identified scientific papers) in addition to the metadata of the identifiers themselves used for resolution and managements (such as the identifier creator).

The Handle System, run by Corporation for National Research Initiatives (CNRI) and authorised by the DONA Foundation⁴, allocates to and manages prefix namespaces for members who then provide a distributed network of networks for identifier resolution. The DONA Foundation only operates the Global Handle Registry (GHR) which allocates identifier resolution to a particular, distributed, subnet of resolution nodes. This semi-distributed architecture adds redundancy and enables the system to run smoothly and uninterrupted when individual resolver nodes are temporarily unavailable or depart permanently. Unlike DOI, Handle members do not have to adhere to policies concerning the identified object’s metadata.

¹ The DOI System, <http://www.doi.org/>.

² The Handle System, <http://handle.net/>.

³ Magnet URI Scheme, https://en.wikipedia.org/wiki/Magnet_URI_scheme.

⁴ The DONA Foundation, <https://www.dona.net/>.

Despite their distributed approach to final identifier resolution, both DOI and Handle are prone to organisational challenges. The allocation of identifier prefixes and the maintenance of the GHR is still within the governance scope of a single organisation (albeit an international, multi-stakeholder organisation) that is then a potential single point of failure. Therefore, both systems ultimately rely on a single authoritative source for the resolution of an identifier to an information object even while removing dependence on other single systems (one of the purposes of Handle) such as the global Domain Name System (DNS)⁵.

Peer-to-peer networks technologies

The Magnet URI scheme, on which Magnet Links are based, became a de-facto standard for the unique identification of files in peer-to-peer (P2P) networks by using a cryptographic hashing function of their content, rather than their name or location. They have gained popularity particularly in P2P file sharing networks, such as those using the BitTorrent protocol⁶, where it is beneficial for access speed to have multiple copies of the same resource and for those copies to be available from multiple sources. Despite the distributed nature of BitTorrent file download, users used to initially download a torrent file⁷ from a single discoverable host, such as a well-known website, that informed them of ‘tracker’ nodes⁸, a special type of servers that assists in the communication between peers using the BitTorrent protocol. Trackers store up-to-date information about peers known to contain copies of the desired resource in order to initiate resource download from those peers. With Magnet Links, generated from a unique hash (i.e. signature) of the torrent file, a user can now obtain the torrent file by resolving the link using the same P2P network that serves to deliver the resource. This means the reliance on tracker nodes – potentially a point of failure – is bypassed and the P2P network serves to resolve both the Magnet Link to the torrent file and then the details in the torrent file to the resource content. In our proposal, due to the changing nature of the content and persistent nature of identifiers, we suggest applying hashing function to the identifier itself and then using it in a Magnet Link, e.g.:

```
magnet:?xt=urn:sha1:82353b74f78c5c8a70e751dd20e1f7e2488bfcad
```

An extension to the Distributed Hash Tables (DHT) (Loewenstern & Norberg, 2009), P2P technology used by BitTorrent for both Magnet Link and content link resolution, is the Peer Exchange (PEX) network (Wu *et al.* 2010). It allows active peers in a network to update each other as to the existence of other peers. This allows download requests to be put to active peers rather than a pure DHT list of potential active peers, many of whom will certainly be inactive a short time after list generation.

The use of decentralised architectures, such as DHTs and PEX, dramatically increase the resistance of services against a denial of service attacks⁹ (Awerbuch & Scheideler, 2007; Saad *et al.* 2008; Sit & Morris 2002) and general robustness due to the bypassing of a single point of failure. In 2009, The Pirate Bay, a well-known torrent file search and access website, announced (Sar, 2009) that it would shut down its torrent file tracker system permanently in favour of the use of DHT, PEX, and Magnet Links, which it did three years later (Sar, 2012) by which time it provided access to torrent files via Magnet Links only. This has set a precedent in distributed file sharing networks for the use of multiple levels of redundancy and resulting high content availability.

File sharing P2P networks rely on having multiple copies of the same file co-exist in various locations. Magnet Links provide a unique PID for a file as long as its content remains intact, thus eliminating the possibility of having multiple representations of the same information object as any change to the object’s content will inevitably change its unique hash signature and thus the Magnet Link that is based on it. In the next section, we propose a new approach to a PID systems design based on P2P file sharing protocols, which allow for important facets of a PID system design – a notion of an authoritative source as well as availability of information object replication mechanisms. The proposed PID system’s design not only increases the availability of the identifier resolution services, but also of data delivery services.

⁵ https://en.wikipedia.org/wiki/Domain_Name_System.

⁶ <https://en.wikipedia.org/wiki/BitTorrent>.

⁷ https://en.wikipedia.org/wiki/Torrent_file.

⁸ https://en.wikipedia.org/wiki/BitTorrent_tracker.

⁹ Denial-of-service attack, https://en.wikipedia.org/wiki/Denial-of-service_attack.

Blockchain

Blockchain¹⁰ technology represents another successful application of distributed network infrastructures and one that has gained a lot of media attention recently, mainly due to the rise of the Bitcoin cryptocurrency¹¹. Although there are multiple implementations of Blockchain technology they are all built upon the fundamental idea of a Blockchain database – to distribute an ever-growing list of transactions, known as a ledger, over a large network of participating nodes. High levels of data integrity are maintained by sequential hashing of the ledger after each new transaction is added. The use of hashing functions acts as digital signatures, which provide a reliable mechanism to detect even minor discrepancies in data to ensure overall integrity. This ordered list of hash-based links creates a monolithic record of all confirmed transactions, which guarantees that all copies of the distributed ledger are both bit-level identical and immutable.

Blockchain technology has been proposed for use by PID systems by Bolikowski *et al.* (2015) and while there is certainly the potential for its use in recording some form of transaction, such as changes in PID ownership, the main functions supplied PID systems (see Car *et al.* 2017 for a general summary of PID system functions) are not sequential transactions in nature and, thus, not generally amenable to Blockchain technology use. For this reason, we do not further address Blockchain technology in this paper.

An Approach to Decentralised Identifier Resolution and Data Distribution

PID systems commonly rely on a single or multi-node resolution service and a single data distribution service provided by a data custodian, whether for file-based information objects or dynamically generated data provided via web services. With the use of Handle, the reliability of the identifier resolution process is robust due to Handle's distributed nature. However, once the user request is resolved, the user is then redirected to a single source of data that acts as a single point of failure, thus data delivery is not guaranteed.

We propose a new approach to both the resolution of persistent identifiers and their data delivery, leveraging advancements in the development of decentralised distributed systems and peer-to-peer (P2P) file sharing networks. Despite taking a decentralised approach to PID system design, we maintain the notion of an authoritative information source and allow for circumstances prohibiting the use of data distribution networks or caching of information objects at location other than authoritative source. In the following sections, we also distinguish between PIDs for file-based information objects and those for abstract entities or dynamic datasets that can only be retrieved via a (web) data service. The use of a term 'data service' rather than 'web service' is intentional as we allow that in n -years' time web- or HTTP-based approaches to the resolution of persistent identifiers may be replaced with new technologies, and data will be delivered by other protocols. This level of abstraction from current Internet technologies (i.e. HTTP, HTTP URIs, etc.) is intended to future-proof a PID system from inevitable technological changes.

Persistent identifiers registration and resolution

Most PID systems, e.g., PURL, LSID (Orme *et al.* 2008), PID Service (Golodoniuc *et al.* 2015b), implement their own PID registration mechanisms and interfaces for resolution of those PIDs. The metadata associated with the PIDs – metadata required for PID resolution, not metadata about the object resolved – is usually stored in a database internal to the PID system (**Figure 1**) and can normally be accessed via an application programming interface (API) or a user interface. The *PID Service* (Golodoniuc *et al.* 2015a), a recent implementation of a PID system, is now deployed and operating in various environments ranging from Semantic Web applications (Golodoniuc *et al.* 2015b) to government departments delivering multiple data catalogues¹² to Big Data management applications at the Australian National Computing Infrastructure (Wang *et al.* 2016). The respective organisations maintain PID Services instances and their associated databases at level of availability, security and fault tolerance appropriate to their organisation's needs. Various solutions, including pre-configured virtual machine environments that run PID Service registration and resolution services, have been implemented by different organisations. Although these implementations deliver high service availability, a single registration and resolution service instance still remain a single point of failure.

We propose a solution whereby we move away from the necessity to maintain an internal PID Service database for storing PID descriptors and make use of Distributed Hash Tables (DHT) for that purpose. While,

¹⁰ Blockchain, [https://en.wikipedia.org/wiki/Blockchain_\(database\)](https://en.wikipedia.org/wiki/Blockchain_(database)).

¹¹ See the Bitcoin community and technology website: <https://bitcoin.org/en/> for a description of the system and a recent online newspaper's news article about Bitcoin's potential adoption by mainstream financial futures trading: <http://www.wsj.com/articles/bitcoin-futures-might-be-coming-soon-1479143252>.

¹² See <http://pid.geoscience.gov.au/>, which is the PID home page for Geoscience Australia.

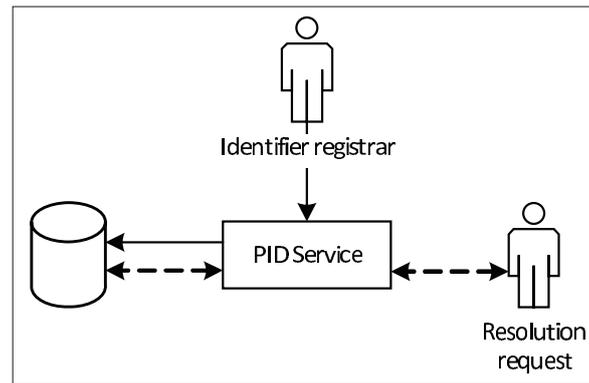


Figure 1: Close coupling between the PID Service registration and resolution service and its internal storage database. PID Service acts as a single interface between users and the database.

from the user's perspective, the PID registration and resolution processes appear the same the underlying data storage mechanisms are fundamentally different.

To create a globally unique PID, one needs to have an identifier issued following an agreed naming convention and then to have it stored in a DHT. For a given PID, a unique hash value, known as its key k , is produced and the key and PID descriptor information are then registered with any node participating in the DHT via a message call $put(k, data)$. Note the 'data' in the put function is the PID's metadata – who created it and where it resolves to, not the data of the item it identifies. The message is then forwarded from node to node through the DHT overlay network until it is finally stored in appropriate nodes according to DHT topology and keyspace partitioning policy. To retrieve the data, a requester will have to apply a hashing function to a persistent identifier and ask any participating node in DHT to find the data associated with key k . The message once again will be routed through the DHT network to an appropriate node, which will then reply with the stored *data*. Although hashing function may differ between DHT implementations, SHA-1 is most commonly used. The SHA-1 hash generates 160-bit keys that are large enough to avoid hash table key collisions. The use of hashing function also alleviates the identifiers length problem. The requester is not required to perform identifier hashing to resolve a human-readable persistent identifier and may instead may resort to using organisational or a public well-known resolver as an intermediary as shown in **Figure 2**.

The **Figure 2** schematically illustrates registration and resolution mechanisms of PIDs via a DHT. The notable aspect of the diagram is that it lacks a single storage database containing PID information at the organisational level in favour of a decentralised network of nodes in a DHT. An organisation may choose to participate in a DHT and maintain its own node as shown by the grey square node N_o . However, participation in a DHT is voluntary. Other registration/resolution services ('Public resolver' in **Figure 2**) may choose to use a known DHT node with a certain degree of confidence¹³ that the information passed into or requested from a DHT is propagated or routed through the overlay network accordingly. The implementation of DHT allows individual nodes to arrive and depart continuously with no interruption to the rest of the network assuming an adequate number of nodes and reasonable latency between neighbouring nodes in the DHT topology is maintained. Modern implementations of DHTs benefit from load balancing (Byers *et al.* 2003; Godfrey & Stoica 2005; Li *et al.* 2004), optimised routing (Manku, 2003), and other efficiency algorithms.

The data registered in a DHT and associated with the key k represents a PID descriptor containing information according to a baseline PID data model. A PID descriptor contains essential information on identifier resolution rules or, in case of a file-based information object location of the file, where 'location' is used in a very broad sense as it might be represented by a file handle in a peer-to-peer file sharing network (e.g., Magnet Link) or perhaps something else. A PID descriptor may take form of a serialised PID Service mapping

¹³ The degree of confidence could be determined by assessing reputational factors of all the participants in the network, but that is a governance/social aspect to PID system design and not dealt with here. It is also possible that determining information about the participants is a very difficult task given that participants perhaps need only adhere to a technical protocol and don't need to adhere to social or governance arrangements. Again, discussion of this aspect of DHTs is outside this paper's scope.

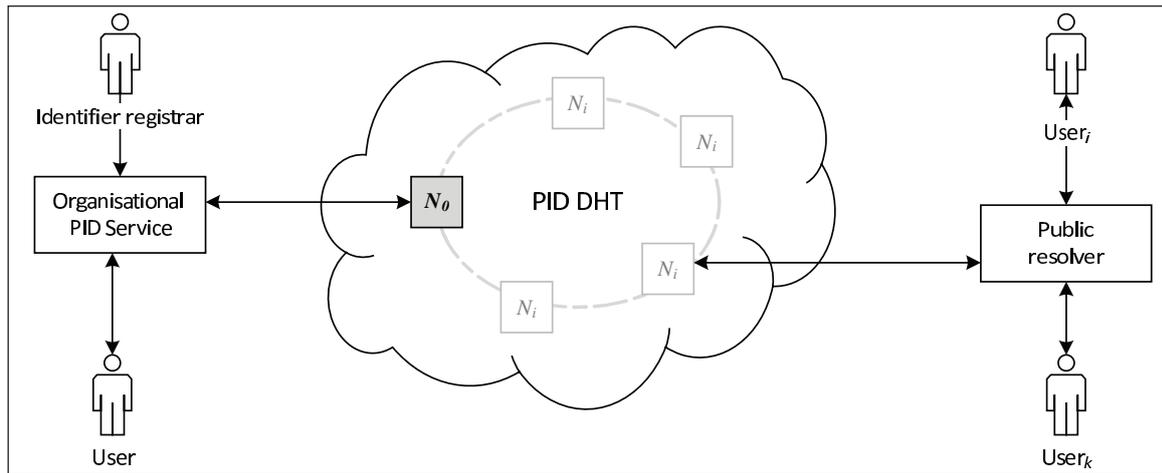


Figure 2: Registration and resolution of persistent identifiers and associated metadata via a Persistent Identifiers Network DHT that avoids maintaining a single database at the organisational level, where persistent identifiers are issued.

rule (Golodoniuc 2015a; Golodoniuc *et al.* 2015b) – a common approach to encode complex data structures into a single flat file format, e.g., XML in this case. In its serialised form, a mapping rule for a single PID, an example of which is given in **Figure 3**, is a very small file that can be transmitted between DHT nodes with little impact to overall network availability.

Information object registration

As part of the initial persistent identifier registration process, the registrar is also responsible for registering the information object identified by the identifier unless the identifier is issued for an abstract entity, such as a vocabulary term. We distinguish three types of information objects, whose types affect the registration processes needed for them as well as the way information objects are later retrieved from a distributed network:

1. **Intangible information object** – abstract entity, that neither has a digital representation nor an accessible location, e.g., vocabulary term;
2. **File-based information object** – where an identifier either has a single file representation (e.g., a specific file, DOI-tagged report, etc.) or a specific representation of an identifier, which is a file. For example, a CSV-encoded data sheet;
3. **Multi-faceted information object** – where an identifier represents an entity that has multiple facets/representations, e.g., Lake Argyle can be represented by a geometry encoded in Geography Markup Language (GML) or an ontology in a Semantic Web while being identified by the same identifier that belongs to the lake as an entity in the real world. The information object of a specific representation may potentially be cached at various levels assuming object caching and retaining policy is permitting it by the original registrar.

Due to the conditions and complexities associated with each information object type, we will further elaborate on the information object registration aspects for each type.

Registration of a persistent identifier for an intangible information object is the simplest scenario where registrar provides simple identifier metadata information according to a baseline information model and registers the identifier without the need to store one or more representations of any other data. This process creates a PID descriptor in a PID DHT and therefore marks the existence of an identifier. The original registrar has the right to update the PID descriptor, if need be, as described later in this paper.

For file-based information objects the registration process, apart from issuing an identifier, also involves a description of the file location and its retrieval mechanisms. Depending on the organisational, data security and privacy policies, which may either restrict file distribution to authoritative source only or permit creation of multiple copies in data distribution networks, the registrar may choose to provide a direct link to web-accessible resource or create a Magnet Link to a file registered in a P2P DHT and distributed through

```

<Mapping>
  <mappingInstance
    date_start="2016-10-05T22:24...
    date_end="2016-10-05T22:35:1...
    ...
  </mappingInstance>
  <mappingInstance date_start="2...
    <!-- default Condition -->
    <path>/org/(GAlga) (.ttl)?$/...
    <type>Regex</type>
    <title>GA org</title>
    <creator>car-nj</creator>
    <!-- default Action -->
    <Action>
      <type>303</type>
      <name>location</name>
      <value>http://52.62.134...
    </Action>
    <!-- end default Action -->
    <!-- end default Condition -->
    ...
  </mappingInstance>
  <Conditions>
    <Condition>
      <type>Comparator</type>
      <match>${2}=.ttl</match>
      <Actions>
        <Action>
          <type>303</type>
          <name>location</name>
          <value>http://52.62.134...
        </Action>
      </Actions>
    </Condition>
    <Condition>
      <type>HTTPHeader</type>
      <match>Accept=text/turtle</...
      <Actions>
        <Action>
          <type>303</type>
          <name>location</name>
          <value>http://52.62.134...
        </Action>
      </Actions>
    </Condition>
  </Conditions>
</mappingInstance>
</Mapping>

```

Figure 3: An XML serialised instance of a PID Service's mapping rule. Shown are the default actions (redirection to an HTML page) as well as pattern-based conditional redirects to certain Internet media types.

P2P networks. In the case of file distribution from a single authoritative source, the registrar solely is responsible for ensuring availability of the data delivery infrastructure as in common PID systems (**Figure 4**).

In the case of file distribution via P2P file sharing networks, the registrar is also responsible for creation of a file torrent and its registration in a P2P DHT. It is critically important to maintain file availability in early stages of a PID existence until it is fully replicated by multiple P2P nodes and an adequate number of seeders¹⁴ (nodes delivering parts of the file data) becomes available in a P2P network. **Figure 5** illustrates the two-phase registration processes that involves registration of an identifier for a file-based information object in a PID Distributed Hash Table in conjunction with a torrent registration in a traditional P2P file sharing network, which registrar is a participating member of (shown by the grey zoom-in outline on **Figure 5**).

The registration of an identifier for a multi-faceted information object requires the creation of a more sophisticated PID descriptor that contains one or many representations of the information object, e.g., a geographic feature can have a GML, imagery, and ontology representations at the same time. Each representation itself may be either a file distributed via a P2P network, a reference to a web resource, or request delegated to a web service, etc. In a simple scenario, this type of identifier would not benefit from the data redundancy mechanisms described above but they may be used to create identifiers to static web resources (e.g., web pages) creating permanent links to these resources known as aliases. **Figure 6** illustrates the use of a persistent identifier registered in the PID DHT for a multi-faceted information object.

Resolution and information object retrieval

The first stage of the resolution of a persistent identifier involves hashing the identifier to a key k and resolution of that key to a PID descriptor through a PID DHT. The PID descriptor itself does not contain the information object identified. It only contains metadata that enables an implementation of a PID resolver to

¹⁴ A term defined in the Glossary of BitTorrent terms, https://en.wikipedia.org/wiki/Glossary_of_BitTorrent_terms#Seed_.2F_seeding.

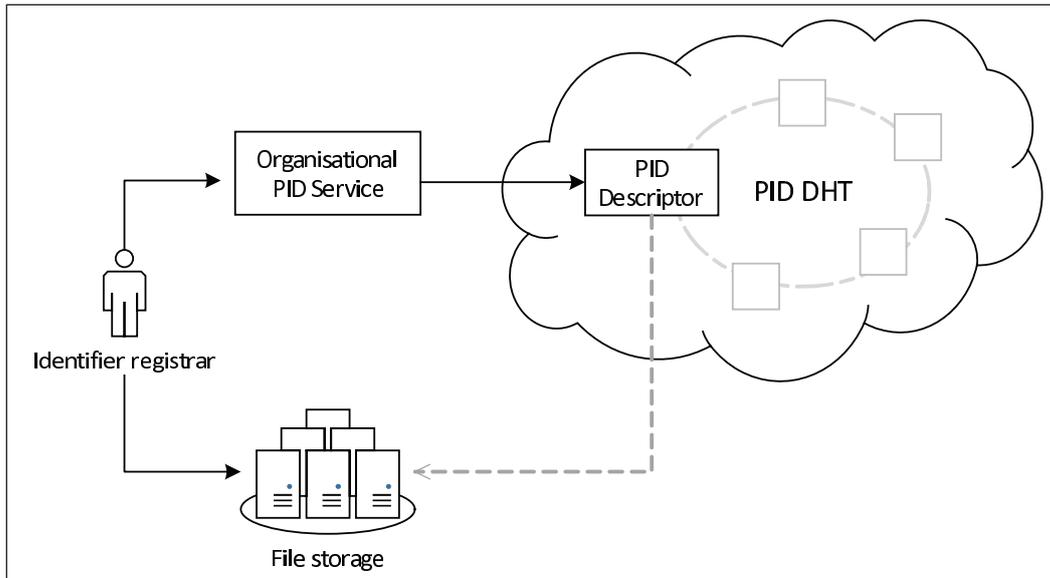


Figure 4: Persistent identifier registration for a file-based information object delivered from a single authoritative source. A single data delivery node puts additional strain on organisational network and essentially is a single point of failure.

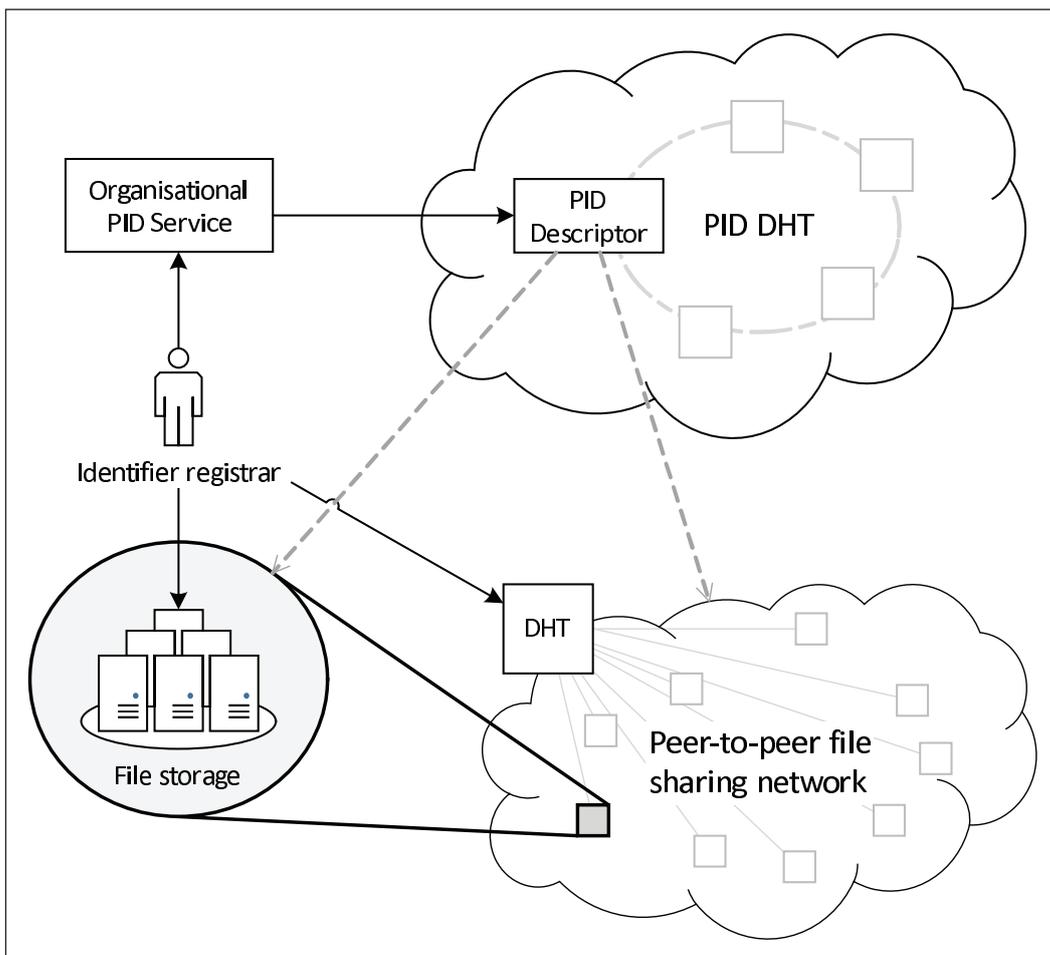


Figure 5: Persistent identifier registration for a file-based information object in a PID Distributed Hash Table and parallel registration of the file torrent in a traditional P2P file sharing network, which registrar's organisation is a participating node of.

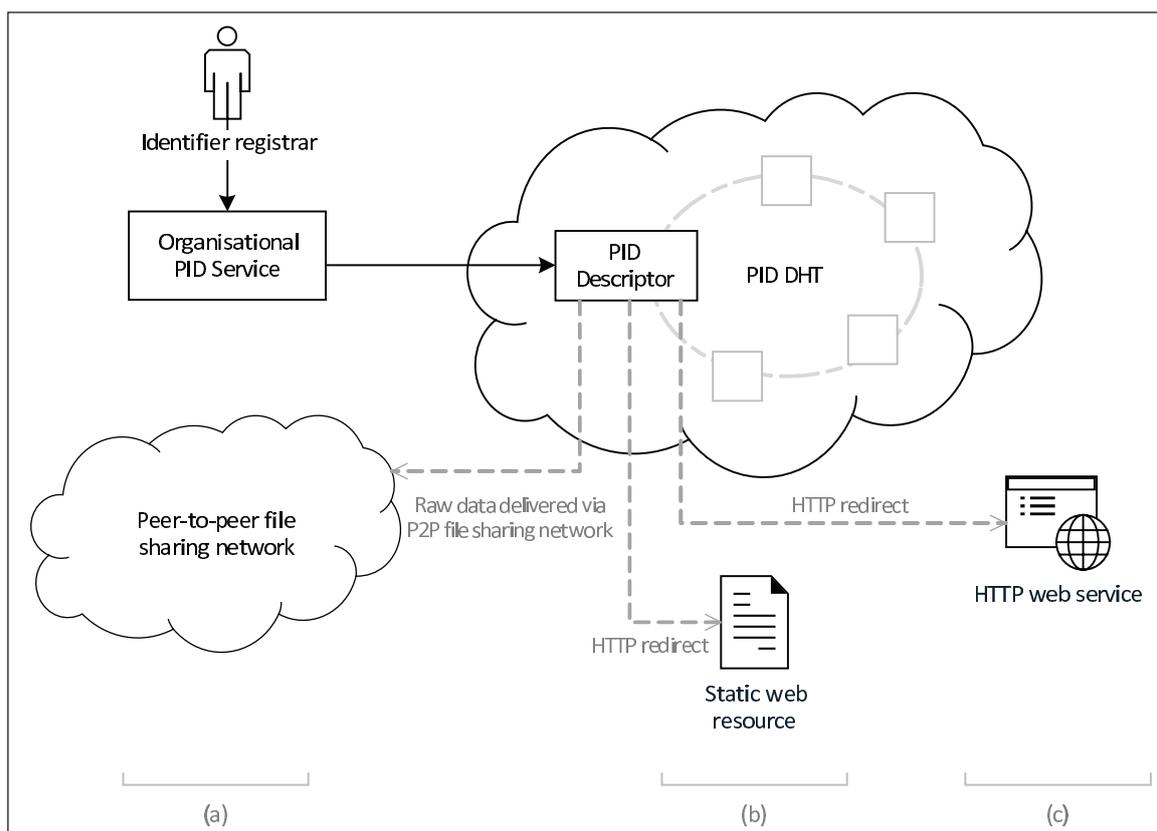


Figure 6: Persistent identifier registration for a multi-faceted information object with different types of views – **(a)** file distributed via P2P file sharing network, **(b)** a reference to a static web resource (e.g., HTML page, report, etc.), and **(c)** a delegated web request to a web service.

retrieve the identified object's content and exactly what is contained will depend on the identified information object type and its data distribution policy, as described below.

A PID for an intangible information object resolves to a PID descriptor, as in all other cases. The descriptor provides essential registration metadata details of the registrar, registration date, status, etc. Successful resolution of the PID descriptor indicates that the identifier is valid and this is the final resolution step as there is no content associated with the identifier.

For a file-based information object, the retrieved PID descriptor would either provide a direct link to the file distributed solely from a single authoritative source or a reference to the file in a P2P file-sharing network. In the latter scenario, an implementation of a PID resolver should minimally provide a Magnet Link that can then be fetched using P2P torrent client applications. However, a specific implementation of a PID resolver (whether public or run by an organisation) may provide an optional torrent leeching¹⁵ service (a node that downloads content from the seeding peer nodes) that would retrieve the file and provide it as an HTTP download. When leeching files from a P2P network, a torrent-enabled PID resolver would automatically cache the file, if data distribution policy permits, and become a new seeder node thus further improving data availability in the distribution network.

The resolution of a PID for a multi-faceted information object depends on the object representation. For file representation, the resolution will replicate the resolution process described above, whereas other views may resolve to an HTTP redirect response (as shown in **Figure 3**), or content of the information object delivered directly to the user, or perhaps by any other means specific to the technology of the day. This technological agnosticism allows for flexibility in the resolution of persistent identifiers and does not couple the proposed PID system to a specific technology stack.

The PID descriptor metadata should also contain a description of the caching and data replication policy used for its object in file sharing and content delivery distributed networks. The policy specified by

¹⁵ A term defined in the Glossary of BitTorrent terms, https://en.wikipedia.org/wiki/Glossary_of_BitTorrent_terms#Leech.

the registrar sets out restrictions: whether or not the content is allowed to be cached by intermediary PID resolver nodes; whether or not the caching is permitted within a specific geographic extent, e.g., outside of a specific country; temporal restrictions on cached information objects by setting a time-to-live (TTL) parameter, etc. The TTL is particularly important and must be finite. Even though it may improve the performance of caching, it should also be balanced in order to avoid outdated cached copies of information objects being trapped in PID resolver nodes for extended periods. The PID DHT's overarching caching policy may also impose restrictions on the upper TTL limit and reject requests that exceed that limit.

Update of PID descriptors in DHT

PID descriptors being distributed via PID DHT may be replicated over multiple nodes in a DHT and thus a robust governance model is required to ensure integrity of PID descriptors as well as security mechanisms and protection from malicious interference. For integrity, each message (i.e. PID descriptor) that is registered in a PID DHT should be digitally signed by the registrar using a private key of an asymmetric encryption algorithm and provide a public key that will be preserved in a PID DHT record. The PID DHT can then use the public key to verify authenticity of the request sender when modifications to an existing PID descriptor are made. For existing records requiring updating, only original registrar, or its trusted successors that possess the private key, would be able to update the PID descriptor in the PID DHT. Since the public key stored in the PID DHT is only used to verify the authenticity of the update request, the submission of the public key along with the PID descriptor is only required during original registration.

Persistent identifiers for collections

In some circumstances, the registration of persistent identifiers for individual items within a very large set of information objects might be impractical. Existing PID systems have implemented various approaches to deal with this including some which identify the collection as a whole and then route users' requests for a specific item within it to a delegated subsystem. Golodoniuc *et al.* (2015b) reviewed various technological solutions that support persistent identifiers for collections and hierarchies of identifier pattern matching. The PID Service (Golodoniuc *et al.* 2015b), developed by CSIRO¹⁶, uses regular expressions-based patterns to match HTTP URIs. The patterns, and dependent hierarchies of them, may also be complex enough to describe even the most difficult matching scenarios, but are usually simple enough for non-technical users to manage. The tool implements a Graphical User Interfaces (GUI) for non-technical users to perform management with and has proved itself a viable tool for collection identifier management.

With the introduction of hashed identifiers (Farrell *et al.* 2013) registered in a PID Distributed Hash Table, text-based pattern recognition in the traditional sense is no longer applicable. We have looked at other options in the context of distributed PID resolution and data delivery networks and propose the use of the Magnet URI scheme (i.e. Magnet Links) for identifiers registered in a PID DHT. A Magnet Link consists of a magnet protocol identifier, 'magnet', a protocol/parameter separator, ':?', and a series of one or more parameters. The design allows additional, vendor-specific, parameters to be introduced at any stage to extend the feature set of a PID system. Each parameter in a Magnet Link is separated by an ampersand, '&', and the parameters themselves are URL-encoded. The most important and the only required parameter is 'xt' ("exact topic"), which is a hash of the identifier and the key *k* registered in a PID DHT, e.g.:

```
magnet:?xt=urn:sha1:82353b74f78c5c8a70e751dd20e1f7e2488bf cad
```

In the case of a collection identifier, the exact topic identifies the collection itself and, like all other identifiers, may provide different representations of the collection. However, an additional parameter 'x.xo' ('x' – supplement parameter indicator, 'xo' – 'exact object' in the collection) can be supplied when resolving a persistent identifier to an exact information object within a collection, e.g.:

```
magnet:?xt=urn:sha1:82353b74f78c5c8a70e751dd20e1f7e2488bf cad
&x.xo=urn:sha1:3b45325cae761a1893b1651607c635dd15c9e8fb
```

The PID DHT will then use the 'xt' parameter to retrieve the PID descriptor and pass it onto a PID resolver that will further process user's request to resolve 'x.xo' to a specific information object in the collection.

¹⁶ Commonwealth Scientific and Industrial Research Organisation, Australian Government, <http://www.csiro.au>.

This proposed approach allows for the simple registration of persistent identifiers for collections of similarly typed information objects, which are delivered in a unified manner without the need to register identifiers for each individual object separately. This approach does not restrict the creation of persistent identifiers for individual information objects in a collection: any individual information object may also be identified by one or many persistent identifier and be part of one or many collections.

Conclusion

In this paper, we have explored the potential use of certain distributed system technologies for persistent identifier (PID). This work follows on from our previous research into the factors that make PID systems trustworthy and our analyses of the failures of once successful identifiers systems, e.g., PURL and LSID. Large-scale decentralisation, where there is no central governing authority but instead a common protocol that all participating adhere to, is the core of this proposed PID system design. We aimed to decentralise not only the persistent identifiers resolution mechanisms but also information object storage by employing advancements in peer-to-peer file sharing network technologies. The proposed approach uses a PID Distributed Hash Tables (DHT) for the registration of PID metadata and tolerates the continual arrival and departure of participants thus it lacks a single point of failure. For the distribution of information objects, caching for performance is supported while the notion of an authoritative source, critical for persistent identifiers, is maintained. The underlying technology, DHT and peer-to-peer (P2P) file sharing protocols, for this proposal is mature, however, the application to PID systems is novel. We believe that this proposed system overcomes some shortfalls present in even the most recent PID systems, namely single points of failure and technology dependence. The entry barrier for individual participants is projected to be low once a network of PID nodes in a DHT is established. The lack of a central database for metadata storage relieves data custodians from the necessity to commit to long-term maintenance of support infrastructure. Participation in a DHT or in a P2P file sharing network is voluntary and depends on the specific requirements of a participating organisation. The intrinsic replication of the PID metadata simplifies the process of PID custodianship handover and allows resolution mechanisms to remain functional in the event of departure of the PID issuing organisation or supporting infrastructure.

Competing Interests

The authors have no competing interests to declare.

References

- Awerbuch, B** and **Scheideler, A** 2007 A denial-of-service resistant DHT. In: *The 26th annual ACM symposium on Principles of distributed computing (PODC '07)*. New York, NY, USA, pp. 370–371. DOI: <https://doi.org/10.1145/1281100.1281178>
- Beck, K, Ritz, R** and **Wittenburg, P** 2016 Towards a Global Digital Object Cloud – Report from the Views on PID Systems training course and workshop. In: *RDA Europe Workshop August–September 2016, Max Planck Compute and Data Facility (MPCDF)*. Garching-Munich, Germany. Available at: https://www.rd-alliance.org/sites/default/files/attachment/20160901_RDA_PID_event_Garching_report_final.pdf (Last accessed 17 May 2017).
- Bolikowski, Ł, Nowiński, A** and **Sylwestrzak, W** 2015 A System for Distributed Minting and Management of Persistent Identifiers. *International Journal of Digital Curation*, 10(1): 280–286. DOI: <https://doi.org/10.2218/ijdc.v10i1.368>
- Bütikofer, N** 2009 *Catalogue of criteria for assessing the trustworthiness of PI systems, nestor-Materialien, Niedersächsische Staats und Universitätsbibliothek Göttingen*. Göttingen, Germany. Available at: <http://nbn-resolving.de/urn:nbn:de:0008-20080710227> (Last accessed 11 November 2016).
- Byers, J W, Considine, J** and **Mitzenmacher, M** 2003 Simple load balancing for distributed hash tables. In: *IPTPS*. Springer Berlin Heidelberg, pp. 80–87. DOI: https://doi.org/10.1007/978-3-540-45172-3_7
- Car, N J, Golodoniuc, P** and **Klump, J** 2017 The challenge of ensuring persistency of identifier systems in the world of ever-changing technology. *Data Science Journal*, 16(13): 1–18. DOI: <https://doi.org/10.5334/dsj-2017-013>
- Farrell, S, Kutscher, D, Dannewitz, C, Ohlman, B, Keranen, A** and **Hallam-Baker, P** 2013 Naming Things with Hashes. In: *RFC 6920*. DOI: <https://doi.org/10.17487/rfc6920>
- Godfrey, P B** and **Stoica, I** 2005 Heterogeneity and load balance in distributed hash tables. In: *The IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies* (vol. 1), pp. 596–606. DOI: <https://doi.org/10.1109/INFCOM.2005.1497926>

- Golodoniuc, P** 2015a Persistent Identifier Service (PID Service). Available at: <https://www.seegrid.csiro.au/wiki/Siss/PIDService> (Last accessed 3 November 2016).
- Golodoniuc, P, Car, N J, Cox, S J D** and **Atkinson, R A** 2015b PID Service – an advanced persistent identifier management service for the Semantic Web. In: *The 21st International Congress on Modelling and Simulation (MODSIM2015)*. Modelling and Simulation Society of Australia and New Zealand, Broadbeach, Australia, December 2015, pp. 767–773. ISBN: 978-0-9872143-5-5. Available at: <http://mssanz.org.au/modsim2015/C8/golodoniuc.pdf> (Last accessed 26 October 2016).
- Golodoniuc, P, Klump, J** and **Car, N J** 2016 Trustworthy persistent identifier systems of the future. *Geophysical Research Abstracts*, 18: EGU2016-1506-2, Copernicus Society. Available at: <http://meetingorganizer.copernicus.org/EGU2016/EGU2016-1506-2.pdf> (Last accessed 26 October 2016).
- Huber, R** and **Klump, J** 2016 How dead is dead in the PID Zombie Zoo? In: *RDA Europe Workshop August-September 2016, Max Planck Compute and Data Facility (MPCDF)*. Garching-Munich, Germany. Available at: https://www.rd-alliance.org/sites/default/files/attachment/20160902-RDA_EU_View_on_PID_Systems_Garching-Robert_Huber-Jens_Klump-How_dead_is_dead_in_the_PID_Zombie_zoo.pdf (Last accessed 17 May 2017).
- Li, J, Stribling, J, Gil, T M, Morris, R** and **Kaashoek, M F** 2004 Comparing the performance of distributed hash tables under churn. In: *The International Workshop on Peer-to-Peer Systems*. Springer Berlin Heidelberg, pp. 87–99.
- Loewenstern, A** and **Norberg, A** 2009 DHT Protocol. BitTorrent.org standard BEP5. Available at: http://www.bittorrent.org/beps/bep_0005.html (Last accessed 26 October 2016).
- Manku, G S** 2003 Routing networks for distributed hash tables. In: *The 22nd Annual Symposium on Principles of Distributed Computing*, pp. 133–142. DOI: <https://doi.org/10.1145/872035.872054>
- Orme, E R, Jones, A C** and **White, R J** 2008 *LSID Deployment in the Catalogue of Life, BNCOD 2008 Biodiversity Informatics Workshop*. Cardiff University: Wales, UK.
- Saad, R, Nait-Abdesselam, F** and **Serhrouchni, A** 2008 A collaborative peer-to-peer architecture to defend against DDoS attacks. In: The 33rd IEEE Conference on Local Computer Networks (LCN), Montreal, Quebec, pp. 427-434. DOI: <https://doi.org/10.1109/LCN.2008.4664200>
- Sar, E v d** 2009 The Pirate Bay tracker shuts down for good. Available at: <https://torrentfreak.com/the-pirate-bay-tracker-shuts-down-for-good-091117/> (Last accessed 26 October 2016).
- Sar, E v d** 2012 The Pirate Bay will stop serving torrents. Available at: <https://torrentfreak.com/the-pirate-bay-will-stop-serving-torrents-120112/> (Last accessed 26 October 2016).
- Sit, E** and **Morris, R** 2002 Security Considerations for Peer-to-Peer Distributed Hash Tables. In: Druschel, P, Kaashoek, M F and Rowstron, A I T (Eds.) *Revised papers from the First International Workshop on Peer-to-Peer Systems (IPTPS '01)*. Springer-Verlag: London, UK, pp. 261–269. DOI: https://doi.org/10.1007/3-540-45748-8_25
- Wang, J, Si, W, Car, N J** and **Evans, B** 2016 Persistent Identifier Practice for Big Data Management at NCI. In: *eResearch Australiasia 2016*, Melbourne, Australia. Available at: https://eresearchau.files.wordpress.com/2016/03/eresau2016_paper_90.pdf (Last accessed 26 October 2016).
- Wu, D, Dhungel, P, Hei, X, Zhang, C** and **Ross, K W** 2010 Understanding Peer Exchange in BitTorrent Systems. In: IEEE 10th International Conference on Peer-to-Peer Computing (P2P), Delft, Netherlands, pp. 1–8. DOI: <https://doi.org/10.1109/P2P.2010.5569967>

How to cite this article: Golodoniuc, P, Car, N N J and Klump, J 2017 Distributed Persistent Identifiers System Design. *Data Science Journal*, 16: 34, pp. 1–12, DOI: <https://doi.org/10.5334/dsj-2017-034>

Submitted: 16 November 2016 **Accepted:** 13 June 2017 **Published:** 28 June 2017

Copyright: © 2017 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

 *Data Science Journal* is a peer-reviewed open access journal published by Ubiquity Press.

OPEN ACCESS 