# Dude, where's my NFT?
# Distributed Infrastructures for Digital Art

Leonhard Balduf
Technical University of Darmstadt
Germany

Martin Florian
Weizenbaum Institute & Humboldt-Universität zu Berlin
Germany

Björn Scheuermann
Technical University of Darmstadt
Germany

November 4, 2022

**Abstract**

We explore issues relating to the storage of digital art, based on an empirical investigation into the storage of audiovisual data referenced by non-fungible tokens (NFTs). We identify current trends in NFT data storage and highlight problems with implemented solutions. We particularly focus our investigation on the use of the Interplanetary Filesystem (IPFS), which emerges as a popular and versatile distributed storage solution for NFTs. Based on the analysis of discovered data storage techniques, we propose a set of best practices to ensure long-term storage survivability of NFT data. While helpful for forming the NFT art market into a legitimate long-term environment for digital art, our recommendations are also directly applicable for improving the availability and integrity of non-NFT digital art.

## 1 Introduction

Ensuring the availability and integrity of digital art is an inherent challenge [5], as is its monetization. Non-fungible tokens (NFTs) promise to solve these issues. This most modern form of digital art packaging attempts to create digital scarcity and immutability by securing the "ownership" and metadata of digital art copies on a blockchain. And it allegedly leverages distributed infrastructures for guaranteeing the long-term storage and availability of the actual content. In this paper we focus on the availability and integrity challenges surrounding digital art, as observable in the context of NFTs. We investigate how popular NFTs *actually* approach the storage question, and to what extent the content they reference is *actually* immutable. Counter to reasonable consumer expectations, we find that many of the most popular NFTs store data on classical centralized platforms and permit changes to both metadata and actual content. Still, our empirical analysis of 1000 top-ranking NFTs also reveals multiple prominent NFTs that successfully leverage on-chain storage and distributed storage infrastructures such as the Interplanetary Filesystem (IPFS). For content stored on IPFS, we additionally take a look at IPFS itself, quantifying the extent to which stored NFT pieces are replicated (and hence, available) within the network. Finally, we formulate a set of recommendations for NFT data storage. The lessons we extract from the innovative approaches that (some) NFT projects take to digital art preservation are directly transferable to non-tokenized digital art as well. Our key contributions are threefold: 1. a classification and empirical study of NFT data storage techniques 2. an investigation of replication of popular content on IPFS, and 3. a set of recommendations for NFTs data storage to enable long-term survivability .

# 2 NFT Storage in Current Practice

In the following, we present an empirical study into how popular digital art NFTs approach the challenge of storing metadata and assets. The choice of a storage method has a determining impact on the long-term availability and (im)mutability properties of NFTs and the associated digital art.

## 2.1 NFTs: Definition and Implementation

NFTs are digital records that cryptographically connect a distinguishable digital token to an owner. In practice, each individual token is identified by an integer referred to as token ID. They are usually stored on a blockchain (the "chain") and thus generally perceived as permanent. This perception may extend to not only the digital token, but the assets referenced by it. A highly prominent use case for NFTs is to encapsulate ownership of *digital art* and audiovisual content more generally. Other use cases include gaming (e. g., Mirandus[1]) and virtual reality — or Metaverse — (e. g., Decentraland[2]).

Most commonly, Ethereum [11] serves as the blockchain on top of which NFTs are built[3]. NFT implementations on Ethereum are typically governed by either the `ERC721` [9] or `ERC1155` [10] smart contract interfaces. These interfaces govern identification and ownership of NFTs, including transfer thereof, as well as association of storage uniform resource identifiers (URIs) to tokens, via optional metadata extensions.

[1] `https://mirandus.game/`
[2] `https://decentraland.org/`
[3] According to `https://dappradar.com/nft`

Listing 1 shows an excerpt of the commonly used `ERC721Metadata` extension interface. The interface is used to attach a URI for descriptive metadata to tokens. The standard recommends that the metadata is a JSON document with a name, a description, and an image. Commonly, the per-token metadata URI is constructed on-the-fly from a base URI and the numerical token ID, in order to save on-chain storage. Audiovisual content, if any, is referenced via a URI within the metadata document. EIP721 provides no guidelines on how or where to store referenced content. We investigate plausible techniques and the approaches that popular NFTs projects take in the upcoming Sec. 2.

```
interface ERC721Metadata {
  ...
  function tokenURI(uint256 _tokenId)
  [...] returns (string);
}
```

Listing 1: The ERC721Metadata Interface.

The more recent `ERC1155` interface allows tracking multiple distinct fungible and non-fungible tokens in one contract. For NFTs, this behaves exactly like `ERC721` by using multiple tokens that each have a total supply of one. An extension to assign metadata exists for `ERC1155` as well: `ERC1155Metadata_URI`. Here, the metadata *must* be a JSON document with a predefined set of required keys. An example is given in Listing 2.

```
{
  "name": "Asset Name",
  "description": "Lorem ipsum...",
  "image": "http://example.com/{id}.png",
  "properties": { ... }
}
```

Listing 2: An example EIP1155 Metadata File.

It is noteworthy that a single contract, regardless of whether it implements `ERC721`, `ERC1155`, or both, can track multiple *collections* of NFTs. The notion of grouping multiple tokens into a collection can exist on a contract-level, but is usually tracked on a higher level, i. e., on NFT marketplaces.

## 2.2 Discovering Popular Digital Art NFTs

We obtained a list of the top 1000 most popular NFTs on the OpenSea marketplace. We used "last sale price" as a proxy for popularity and extracted the ranking from OpenSea on August 1st 2022. OpenSea is the most popular[4] NFT marketplace and thus well-suited for our initial sample collection.

[4] by trading volume, according to DAppRadar, see `https://dappradar.com`

From our initial list of popular NFTs, we proceeded to extract all NFTs associated with digital art. In a classification approach inspired by previous works such as [7], we label each NFT as digital that art that has the primary purpose of referencing one or more audiovisual assets. We explicitly exclude gaming and metaverse NFTs from this class, as well as utility tokens more generally. 786 NFTs from our sample fit this classification[5]. They are managed by a total of 51 smart contracts, many of which represent NFT collections with multiple representatives within our sample. The subsequent discussion is based on these 786 NFTs.

## 2.3 Metadata and Asset Storage

We proceed to classify digital art NFTs in our sample based on the storage system they use for storing metadata and assets. We collected data stored in the NFT smart contracts, e. g., data exposed through the EIP721 or EIP1155 standards, and manually inspected custom contract implementations for data stored on-chain.

Motivated by the goal of digital art preservation, we first identify what data is necessary to *reconstruct* the artwork. For example, generative art projects may choose to store source code on-chain, and provide a pre-rendered version of the artwork on cloud storage for convenience. We classify this case as being stored on-chain, since reconstruction from on-chain data is possible. It is potentially necessary to also preserve a runtime environment for this data, but this falls outside the scope of this work.

We differentiate between three main storage approaches: cloud storage, on-chain storage, and decentralized storage systems — IPFS and Arweave. Their properties with regards to integrity and availability vary greatly:

- **Cloud storage** offers no hard guarantees — stored data can be mutated or (re)moved at any time.

- **On-chain storage** guarantees result directly from the security properties of the underlying blockchain system. Data availability and data integrity are among the core features of permissionless blockchain systems and it is a reasonable assumption that popular systems like Ethereum will remain intact — and effective at providing availability and integrity for stored data — for many years to come.

- **Decentralized storage systems** can have varying properties with regards to integrity and availability. The systems used by the digital art NFTs in our sample rely on cryptographically-generated content identifiers: links into these storage systems correspond to cryptographic hashes of the stored content. In effect, changes to the original content can be detected as long as the integrity of the URI of the asset is maintained.

[5]In addition, we classify 179 NFTs as metaverse-related, 19 as game-related, 8 as collectibles, and 8 as utility tokens.

Table 1: Metadata and Asset Storage Locations for Digital Art NFTs

| Metadata | Asset(s) | Count |
|---|---|---|
| on-chain | on-chain | 412 |
| ipfs | ipfs | 254 |
| cloud | cloud | 79 |
| cloud | ipfs | 17 |
| on-chain | cloud | 11 |
| ipfs-dead | | 7 |
| arweave | arweave | 3 |
| cloud | on-chain | 2 |
| cloud-dead | | 1 |

> The availability of content stored in decentralized data storage systems is highly dependent on system particularities and the popularity of the stored content. In the upcoming Sec. 4, we report on a deeper investigation into how the NFTs in our sample are inserted into the popular IPFS network, and to what extent they are replicated there.

The results of our classification are shown in Table 1. To our surprise, the metadata and assets of more than half of our sampled digital art NFTs are already stored on-chain. This is partly due to our manual inspection of contracts, which reveals that various generative art projects store their source code on-chain in addition to providing pre-rendered versions of the artwork via other storage systems. We discuss a representative example in Sec. 3.

We attempted to resolve all metadata URIs pointing to cloud storage and decentralized storage systems on August 1st 2022. We discovered one case in which a metadata URI pointing to a cloud endpoint could not be resolved (cloud-dead). Keeping in mind that our sample is based on the top 1000 most valuable (based on our data) NFTs on OpenSea and that the availability of intact metadata is crucial to reconstructing the assets embodied by an NFT, this is a troubling result. OpenSea continued to provide a cached versions of both metadata and assets for the NFT in question. Relying on the caching mechanisms of NFT marketplace platforms for availability or integrity is highly problematic, however.

## 2.4 (Im)mutability

It is a reasonable assumption that a digital artwork, and especially a digital art NFT, should be immutable and maintain its original state. In some cases, NFTs do enable mutability, however. While mutability might be a deliberate aspect of a digital artwork, digital art that is meant to be immutable should also be stored in an immutable manner.

### 2.4.1 Metadata mutability

The most evident cases in which the immutability of assets referenced by an NFT is not secured are NFTs that use cloud storage for storing metadata or assets, without any additional precautions such as recording a cryptographic hash of referenced data at a findable on-chain location It must be noted, however, that it is not a *sufficient* criterion for immutability that the metadata and assets of an NFT are stored in an integrity-securing manner. In many cases, an NFT's smart contract might allow the metadata URI itself to be modified, thereby enabling a mutability of both metadata and assets. Even though changes are visible on the blockchain, it is unreasonable to expect consumers to keep track of NFT contracts.

We examine NFTs in our sample for transactions changing the metadata URI of tokens. In general, neither EIP721 nor EIP1155 specify how URIs should be stored or updated, but we found that many contracts expose functions of the form `(set|update)[base](URI|uri)`. By processing transaction data on the Ethereum blockchain, we found that out of the 51 contracts examined, 19 changed the URI of token metadata at least once. Notably, we found one case where a base URI was changed from a cloud provider to IPFS, and later back to another cloud provider. Closer examination of this case revealed that the metadata contained a link to the project homepage, which later changed location. OpenSea still caches the IPFS URI for this project, highlighting the problem of cache inconsistency. Additionally, we found one case where the base URI of a project was temporarily set to an invalid URI, which was corrected around 4 hours later.

### 2.4.2 Deliberate mutability

Having surveyed the state of mutability via URI-changing transactions, we then explored whether NFT art projects are intentionally mutable. For that, we examined each project and decided whether, from a consumer perspective and with reasonable time invested to research them, they presented themselves as mutable. We conclude that 782 of the digital art NFTs in our sample are immutable, and only 4 are deliberately and understandably not. These projects communicate their mutability openly: One of them is an evolving art piece, which will be updated by the artist over time. Another one is the MoonCats[6] project, which allows owners to dress their virtual cats with accessories, which ultimately show up in the NFT-referenced image, visible on marketplaces and collector sites. The latter is especially interesting, since this functionality is implemented via smart contracts, including all assets necessary to construct the images. Artists can create accessories, which are (after a review) stored on-chain and made available to be bought and donned by owners of cats. Internally, accessories are positioned and layered upon the base cat image to construct the resulting image.

[6] `https://mooncat.community/`

Notably, a few of the cases that we perceived as immutable, but exhibited mutability through changing URIs, did so for marketing purposes: they changed the URI from a preview or "teaser" version of the collection to the complete data on launch. In most cases, however, mutability seems to be a byproduct of technical issues, such as the changing-homepage case laid out earlier.

# 3 Focus: On-chain Storage

NFTs that feature on-chain storage of metadata or assets are typically based on a custom smart contract implementation. Instead of generalizing, we will therefore focus on a specific example case: the popular NFT collection *CryptoPunks*.

CryptoPunks are a set of $10,000$ unique, $24 \times 24$ pixel character images. There are, historically, multiple versions of the contract deployed on Ethereum[7]. The original contract, contains a SHA256 hash of an image sprite containing all $10,000$ $24 \times 24$ pixel avatars, with transparent background. The image itself is hosted in multiple locations off-chain, and the token ID of the NFT provides an index into the correct position in the image.

In August 2021, a contract to reconstruct the image data was deployed. The contract builds images by compositing a "punk" from eight assets, which are each indexed through one byte. The assets themselves are encoded compressed as a series of three-byte tuples, specifying coordinates and colors to apply. Each tuple can paint an area of $2 \times 2$ pixels, which allows for assets of various sizes and positions to be encoded efficiently. The contract stores a total of 133 assets of various sizes, $100 \times 100 \times 8$ bytes of punk-data, a series of precomputed alpha-blended colors and a palette of base colors. The developers spent 73 million Gas on deploying the contract and data[8], resulting in a total fee of approximately 4.05 ETH $\approx$ USD 12 205 (at the time).

Overall, even though the contracts are not linked and thus some manual intervention is necessary, it is possible to retrieve images from on-chain data only. This is impressive, and sets a good example for future projects. It must be noted, of course, that the pixel-art style, composability from assets, and relatively small image size of $24 \times 24$ pixels are enablers for this feature. Even so, the deployment incurred significant costs. It still seems unlikely that large rasterized images will be stored on-chain in the future.

The CryptoPunks case exemplifies core characteristics of the on-chain storage approach for digital art NFTs. In summary:

- The integrity and availability of embodied art pieces or codes is guaranteed as long as the underlying blockchain is intact and accessible.

- The involved smart contracts are non-standard, implying higher technical hurdles, respectively higher security risks.

[7] `https://github.com/cryptopunksnotdead/punks.contracts`
[8] `https://www.larvalabs.com/blog/2021-8-18-18-0/on-chain-cryptopunks`

- Due to the high cost of storing data on public blockchain networks like Ethereum, on-chain data storage is likely viable only for digital art with low storage space requirements, such as pixel art and generative art.
- Even though core data required for the reconstruction of the artwork is obtainable from the blockchain, auxiliary dependencies might still be required. We conclude that even on-chain storage should be paired with a holistic digital art preservation concept if long-term availability is a goal.

# 4 Focus: Storage on IPFS

In the following, we examine data stored on IPFS in greater detail. IPFS is the most popular distributed storage system for NFTs at the moment, see Table 1. We are furthermore able to leverage existing methods and tooling [4, 1] for gathering detailed information about storage locations and data distribution within the IPFS network.

## 4.1 Background

IPFS is a popular peer-to-peer (P2P) data storage and retrieval system [2, 1]. It uses content addressing via so-called content identifiers (CIDs). A CID is formed based on a cryptographic hash of the addressed data, hence pointing to *immutable* data. Directories and large files are stored as directed acyclic graphs (DAGs) of hash-addressed blocks.

Downloading data is a two-step process: First, *providers* are sought via a one-hop broadcast mechanism and a distributed hash table (DHT)-search [4, 1]. Then, requests to download the data are sent to providers. Downloaded content is, by default, cached locally and (re)provided upon request. This effectively causes an increase in replication for data items that are frequently requested.

It is possible to add a whole directory of files to IPFS, so that only the CID of the directory (the root of the DAG) must be stored. Individual files, e. g., metadata for individual tokens of a collection, can then be addressed relatively to this CID. Resolving such directories on IPFS requires downloading the directory block, which contains CID links to the files, and then downloading the file. This is suboptimal for preservation since both blocks need to be downloadable.

Often, content is not requested via IPFS directly but via a HTTP-IPFS gateway, which fetches content via IPFS and returns it via HTTP. Public gateways are provided by the community. This is convenient since it does not require running an IPFS node locally, and public gateways are generally well-connected and quick to fetch content. It does, however, create dependence on centralized infrastructure [1].

Since content is downloaded directly from nodes on the network, it is available for as long as it is held by at least one nodes that is online and discoverable. Some companies provide paid *pinning* services, which store content on their nodes, making it available to the IPFS network.

## 4.2 Empirical Study

We focus on the 254 digital art NFTs in our sample that use IPFS for storing both metadata and assets. Out of those, one contract did not conform to the metadata specification and returned an image instead — we ignore the associated NFTs in the following.

### 4.2.1 CID Encoding

In principle, a CID addressing raw binary data can be stored in 32 bytes, by encoding only the hash and generating a valid (textual) CID from it on the fly. Still, we found that *all* of the 254 investigated NFTs used regular text-coded IPFS URIs requiring upwards of 50 bytes of on-chain storage. Many contracts additionally use a hard-coded uniform resource locator (URL) to a public gateway server, in the form `https://<gateway>/ipfs/<cid>`, instead of a regular `ipfs://<cid>` URI. While this makes it more convenient for users to access the referenced metadata and assets, it is wasteful in terms of on-chain storage. Additionally, when users resolve IPFS CIDs exclusively via public HTTP gateways, the expectable levels of availability and integrity degrade to the level of cloud-base storage [1].

### 4.2.2 Data Availability

We extract a list of all IPFS metadata URI from all 51 considered NFT smart contracts, which we resolve to obtain a list of all metadata and asset CIDs for the 254 digital art NFTs in our sample. We then conducted a measurement study to determine the availability and replication of these data items. Specifically, we attempt to identify all providers for the referenced data items at six-hour intervals over the course of one week. We do this through searching the DHT, using a well-connected, non-NATed node located in Germany.

Notably, the DHT is not the only discovery mechanism in IPFS — there is also a one-hop broadcast mechanism via the Bitswap protocol [1]. We plan to extend the study presented here with Bitswap discovery results in the future.

We consider unique (peer ID, CID) tuples over all measurements and investigate how often a CID appears within the set of all tuples, i. e., how many times a CID is replicated. The results, visualized as a histogram of the discrete counts as well as an ECDF, are shown in Fig. 1. Among other tings, it can be seen that 50% of CIDs were provided by less than ten nodes each. (Over the course of the whole week, not just at one instant!) Additionally, for a surprising 33 CIDs no providers could be discovered via the DHT at all.

These results are to some extent puzzling — NFT platforms such as OpenSea do not typically communicate to users (in our experience) that some NFT might be unavailable beyond the current platform's cache. It is possible that our results hint at problems on IPFS's DHT layer, i.e., that some providers for a given DHT can only be discovered via one-hop broadcasting. Our results from Sec. 2, in which we were unable to resolve only 7 CIDs (for a different time frame and not restricting the content search to the IPFS DHT), also point to this conclusion.

Further studies are necessary, involving both DHT lookups and broadcast-based searches from nodes with a large number of simultaneous peers (e.g., modified nodes as in [1]).

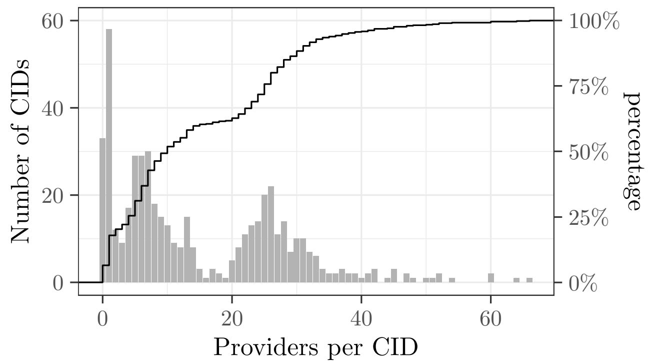


Figure 1: Number of Replicas per CID. Based on one week of DHT crawls every six hours, counting unique (CID, peer ID) tuples.

# 5 Related Work

Nadini et al. [7] investigate NFT trades from 2017 to 2021 on Ethereum and WAX. They describe trade networks and investigate visual similarity between items as indicators for price. A classification of NFTs by use case (e. g., profile pictures) is performed, however without a consideration of storage and immutability questions. Vasan, Janosov, and Barabási [8] explore NFTs from a network-scientific view. They note that, due to the inherent transparency of the underlying blockchains, novel insight into the art ecosystem is possible. The authors perform extensive analyses of trader-, collector-, and artist networks on the *Foundation* platform. Storage and immutability challenges are, again, not explicitly addressed.

Das et al. [3] showcase a wide range of security issues in the NFT ecosystem. They analyze popular NFT marketplaces and identify various security issues pertaining to, e. g., verification of authenticity, preventing fraud, etc. Most relevantly to our work, they showcase the persistence of off-chain NFT data as an unsolved issue. They analyze the storage location and availability of NFT metadata and assets based on a dataset of 12,215,650 NFTs (of all types) from OpenSea. They find that only a fraction of the referenced content is stored on IPFS, and that a large portion of data stored on other systems is inaccessible already. Interestingly, while the research question behind this part of their work is similar to ours, their results and conclusions differ from what we outline in Sec. 2. This fact can be explained by differences in the dataset used in both studies — we focus only on especially popular NFTs, for example — as well as our detailed investigation of contract implementations to identify on-chain storage solutions. We extend the work of Das et al. by looking at NFT storage approaches and their impact in greater detail. We furthermore contribute an in-depth exploration of IPFS as a storage medium for NFT, including an empirical study on the IPFS network itself that aims to determine the level of replication of surveyed NFTs. Empirical studies on IPFS have been performed in the past [4, 1] and offer a valuable methodological foundation for our work. To the best of our knowledge, however, no published works exist that evaluate the level of replication of a specific set of objectively popular content within the IPFS network.

Our work is tangentially related to the study of digital art preservation which arose with the introduction and use of digital systems in the creation and storage of art in the second half of the 20th century. Lee et al. [5] give an overview of techniques used for digital preservation. Since both NFTs and IPFS are young technologies, they might not yet be in the focus of digital art preservation efforts. With this work, we hope to give guidelines on best practices for artists and creators to make digital preservation efforts easier in the future — both of digital art NFT but also of non-tokenized art projects that want to leverage NFT-related technologies for improved availability and integrity.

# 6 Conclusions and Outlook

In this work, we surveyed the current state of NFT data storage empirically, showcased examples and concrete implementations of various storage approaches, and measured replication of data on IPFS. Based on the preceding discussion and our empirical results, we can arrive at a list of recommendations for upcoming digital art NFTs:

- On-chain storage should be the gold standard, and considered for any data small enough to allow for this. Metadata files can already be stored on-chain using `data` URIs [6]. This could further be optimizing by storing only relevant field values on-chain and constructing the metadata JSON on-the-fly. On-chain storage of assets is only feasible for a subset of digital art, such as pixel art, ASCII art, or generative art. For the latter, dependency management and reproducibility should be considered, e. g. through version-locking dependencies on-chain.

- Assets should be immutable by default, which realistically corresponds to user expectation. Mutability should be communicated explicitly. Mutations to metadata should be done transparently in a way that allows reconstruction from on-chain data.

- If on-chain storage is not an option, a content-addressed distributed storage system should be considered. For any off-chain storage solution, a hash of the referenced data should be stored on-chain. This hash can generally be stored binary and reconstructed on-the-fly. Distributed storage systems should be employed in a way that incentivizes replication of the data. Some storage systems explicitly promise long-term availability of data, which should be a goal.

Potential follow-up works include extending our study on the replication and availability of NFT-related content on IPFS, to also consider content discovery mechanisms beyond DHT lookups. Our insights on NFT storage practices might furthermore inspire the development of improved long-term storage approaches. For example, it is conceivable to devise a system that stores content on multiple storage systems in parallel. Stored hashes might be used for automatically generating, or in some way authenticating, content identifiers for multiple distributed storage systems (IPFS and BitTorrent, for example). Improved storage approaches like the sketched idea would benefit all forms of digital preservation efforts, including the preservation of non-NFT digital art.

# References


[1] Leonhard Balduf et al. *Monitoring Data Requests in Decentralized Data Storage Systems: A Case Study of IPFS*. May 23, 2022. DOI: 10.48550/arXiv.2104.09202. arXiv: 2104.09202 [cs]. URL: http://arxiv.org/abs/2104.09202 (visited on 07/28/2022).

[2] Juan Benet. "IPFS - Content Addressed, Versioned, P2P File System." July 14, 2014. arXiv: 1407.3561 [cs]. URL: http://arxiv.org/abs/1407.3561 (visited on 03/15/2020).

[3] Dipanjan Das et al. "Understanding Security Issues in the NFT Ecosystem." Apr. 27, 2022. arXiv: 2111.08893 [cs]. URL: http://arxiv.org/abs/2111.08893 (visited on 05/02/2022).

[4] Sebastian Henningsen et al. "Mapping the Interplanetary Filesystem." In: *2020 IFIP Networking Conference (Networking)*. 2020 IFIP Networking Conference (Networking). June 2020, pp. 289–297.

[5] Kyong-Ho Lee et al. "The State of the Art and Practice in Digital Preservation." In: *Journal of Research of the National Institute of Standards and Technology* 107.1 (2002), pp. 93–106. ISSN: 1044-677X. DOI: 10.6028/jres.107.010. pmid: 27446721. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4865277/ (visited on 07/28/2022).

[6] Larry M. Masinter. *The "Data" URL Scheme*. Request for Comments RFC 2397. Internet Engineering Task Force, Aug. 1998. 5 pp. DOI: 10.17487/RFC2397. URL: https://datatracker.ietf.org/doc/rfc2397 (visited on 08/15/2022).

[7] Matthieu Nadini et al. "Mapping the NFT Revolution: Market Trends, Trade Networks, and Visual Features." In: *Scientific Reports* 11.1 (1 Oct. 22, 2021), p. 20902. ISSN: 2045-2322. DOI: 10.1038/s41598-021-00053-8. URL: https://www.nature.com/articles/s41598-021-00053-8 (visited on 05/02/2022).

[8] Kishore Vasan, Milán Janosov, and Albert-László Barabási. "Quantifying NFT-driven Networks in Crypto Art." In: *Scientific Reports* 12.1 (1 Feb. 17, 2022), p. 2769. ISSN: 2045-2322. DOI: 10.1038/s41598-022-05146-6. URL: https://www.nature.com/articles/s41598-022-05146-6 (visited on 05/02/2022).

[9] William Entriken et al. *EIP-721: Non-Fungible Token Standard*. Ethereum Improvement Proposals. Jan. 24, 2018. URL: https://eips.ethereum.org/EIPS/eip-721 (visited on 08/11/2022).

[10] Witek Radomski et al. *EIP-1155: Multi Token Standard*. Ethereum Improvement Proposals. June 17, 2018. URL: https://eips.ethereum.org/EIPS/eip-1155 (visited on 08/11/2022).

[11] Dr Gavin Wood. "Ethereum: A Secure Decentralised Generalised Transaction Ledger." In: (2014), p. 32.