

A Public Dataset For the ZKsync Rollup*

Maria Inês Silva[†]
Matter Labs

Johnnatan Messias[†]
Matter Labs

Benjamin Livshits
Matter Labs
Imperial College London

Abstract

Despite blockchain data being publicly available, practical challenges and high costs often hinder its effective use by researchers, thus limiting data-driven research and exploration in the blockchain space. This is especially true when it comes to Layer 2 (L2) ecosystems, and ZKsync, in particular.

To address these issues, we have curated a dataset from 1 year of activity extracted from a ZKsync Era archive node and made it freely available to external parties. In this paper, we provide details on this dataset and how it was created, showcase a few example analyses that can be performed with it, and discuss some future research directions. We also publish and share the code used in our analysis on GitHub to promote reproducibility and to support further research.

1 Introduction

The blockchain ecosystem is built on the principles of decentralization and transparency. However, there remains a significant challenge for end users and non-technical researchers: *blockchain data is not easily accessible*. This challenge is hindering the widespread adoption of blockchain technology, which should be more easily accessible.

Currently, blockchain data can be obtained by deploying an archive node in the case of chains based on Ethereum Virtual Machine (EVM) [9, 14, 34, 37], or a full node for Bitcoin [13]. However, this requires a machine with high bandwidth and high-speed storage to keep the blockchain fully synced. In other words, individual deployment of archive nodes is not a practical solution. Therefore, it would be advantageous if we could offer an alternative that improves this situation. Additionally, users can gather the data through Remote Procedure Calls (RPCs) providers. However, gathering data this way can be challenging for some non-technical users and researchers and is usually costly. Alternatively, users can depend on external data sources such as Etherscan [18], Arbiscan [10], Dune [17], ZettaBlock [49], Nansen [6], and similar platforms. Although they may provide an interesting solution, they may prove to be costly in the long run, and do not meet the requirements of certain end-users, particularly those engaged in research that relies on easily and cheaply accessible data.

We believe that anyone requiring blockchain data should have a straightforward and hassle-free way to access them, without worrying about infrastructure or hardware. This data has significant value for research purposes, such as alerting [29], Airdrop designs [33], analyzing Maximal-Extractable Value (MEV) [21, 22, 31], Automated

Market Maker (AMM) [19], and potentially improving the adoption of particular chains such as ZKsync and other Layer-2 blockchains (L2s) rollups within the research scope [23, 30]. In addition, it can contribute to the growth of these ecosystems as more researchers become involved in the field.

In a recent talk, a company called Paradigm emphasized that to empower the blockchain community with the necessary data analysis capabilities, we must make blockchain data more accessible [7]. By ensuring that these data are available affordably, quickly, and with minimal effort, we can effectively address the challenge of data unavailability in our ecosystem. They introduced a novel EVM blockchain node developed in Rust, known as *Reth* that can be applied to EVM-compatible blockchains. Paradigm provided an endpoint to their archive node allowing anyone to make requests to their node. Similar initiatives have been provided by Matter Labs and other companies interested in making blockchain data more accessible [9, 34, 37]. Nevertheless, users would need to build high-efficiency code to gather all the data they need, which could invalidate this process for some users due to network delays or lack of coding skills.

More recently, Paradigm [39, 40], BigQuery [16] and other research groups [28]) have made blockchain data accessible and available by providing them in an easy-to-download and-load schema. In that sense, we followed through and decided to make our ZKsync fully available and accessible to any user or researcher. Therefore, we provide 1 year of ZKsync data covering the period of February 14th, 2023, (block 1) and March 24th, 2024, (block 29,710,983). More details of the dataset are available in Section 2.

1.1 Why ZKsync Era Data?

Launched in March 2023, ZKsync Era is a L2 scaling solution for the Ethereum blockchain that utilizes Zero-Knowledge Proof (ZKP) for efficient transaction processing. ZKsync Era is also an EVM-based rollup, ensuring compatibility with existing Ethereum smart contracts. As of July 2024, ZKsync Era ranks among the top five ZKP chains with a Total Value Locked (TVL) estimated at 1.23 billion USD [26]. This scaling solution improves Ethereum by processing transactions in batches using ZKP, which helps maintain low transaction fees and encourages user participation.

Currently, rollups are a key strategy to scale the Ethereum ecosystem [15]. In fact, in the last few years, these L2 chains have become an important part of the ecosystem by pushing new innovations and absorbing a significant portion of user activity. At the same time, there are still many unexplored questions when it comes to these blockchains, and existing research is fairly limited.

Given the growing role of L2 chains in the Ethereum ecosystem, we believe that there is an opportunity for researchers to conduct research and expand our understanding of ZKsync. By making our ZKsync Era data easily available to external groups, we hope to

*This research article is a preliminary work of scholarship and reflects the authors' own views and opinions. It does not necessarily reflect the views or opinions of any other person or organization, including the authors' employer. Readers should not rely on this article for making strategic or commercial decisions, and the authors are not responsible for any losses that may result from such use. The authors acknowledge the help of Igor Borodin from Matter Labs and his assistance in hosting the dataset.

[†]Both authors contributed equally to the paper.

advance ZKsync-related research and knowledge about L2 chains, more generally.

1.2 Contributions

► **Public Availability of ZKsync Data.** To facilitate scientific use of our collected ZKsync Era dataset, we have made it available on a public GitHub repository [4]. This consists of a one-year dataset containing information regarding blocks, transactions, receipts, and logs. Dataset details are presented in Section 2.

► **Facilitating Research and Analysis.** We demonstrated potential applications of this dataset in research and advocating for the adoption of specific blockchain technologies like ZKsync and L2 rollups in Section 3.

► **Practical Implications for Users and Researchers.** We address challenges associated with data gathering via RPCs and external platforms like Etherscan [18], Arbiscan [10], and others, which may be slow and costly. We provide an easy-to-download and easy-to-load dataset of ZKsync Era, along with the Jupyter notebooks code to facilitate this onboarding process. Table 1 describes our code utilized to load and process our dataset. All the analyses conducted in this paper can be easily reproduced using our code, which is also available in our GitHub repository [4]. See Sections 2 and 3 for details.

1.3 Paper Organization

This paper is organized as follows. Section 2 details our dataset and its data schema, together with the necessary background to understand the context of blockchain data in general. Section 3 provides examples of analyses that can be achieved using this dataset such as regarding transaction fees and gas usage, events derived from transaction logs, token swaps, and more. Furthermore, Section 4 discusses open problems and future directions for which this dataset could be useful. Finally, we present the conclusion of our work in Section 5.

2 Data Schema and Processing

In this section, we present our ZKsync dataset in detail. The dataset covers the period of February 14th, 2023, (block number 1) and March 24th, 2024, (block number 29,710,983) corresponding to 1 year of data. It contains 327,174,035 transactions and 1,631,772 contracts deployed during this time period which triggered 2,044,221,151 events on-chain. Transactions were issued by 7,322,502 unique users. This dataset enables researchers or blockchain enthusiasts to explore all the activities that occurred in the ZKsync Era since its deployment. Table 2 summarizes our ZKsync Era dataset.

We gathered our dataset from our ZKsync Era archive node as raw data. This data consists of all the information regarding blocks, transactions, receipts, and logs. Then, to allow anyone to use the dataset, we conducted a pre-processing step. This step consists of formatting the data in a parquet format that can be easily accessible through well-known libraries available in Python, for example, Pandas [38] and Polars [42]. Due to the high volume of data (i.e., around 200 GB), we focus on using Polars for better processing, memory management, and *Lazy evaluation* [24, 41]. Next, we discuss each of the different data types separately.




Figure 1: Illustration of blocks connected to each other forming the blockchain.

2.1 Blocks

Blocks are sequential units of data within a blockchain, each identified by a unique hash. They contain a list of transactions, metadata such as timestamps, and the hash of the previous block (*parentHash*), which links them in a chain back to the genesis block (block number 0) as shown in Figure 1. This chain of blocks forms the blockchain. Blocks ensure transaction security, network consensus, and efficient data storage and processing within blockchain networks. We list the attributes of the blocks presented in our dataset in Table 3 in the Appendix. In our dataset, blocks should be sorted according to the attribute *block number* to maintain the correct order of the blocks on the blockchain.

2.2 Transactions

Transactions are digital interactions that involve the transfer of assets, the recording of data, or the execution of smart contracts between parties on blockchains like those based on EVM. Each transaction is initiated by a user, authenticated through cryptographic signatures, and sent to a decentralized network of nodes for validation. Once verified, transactions are grouped into blocks and added to the blockchain via a consensus mechanism, ensuring that they are secure, immutable, transparent, and free of intermediaries, forming the core of the blockchain system. We list the attributes of the transactions presented in our dataset in Table 4 in the Appendix.

In rollups, such as ZKsync, transactions are aggregated and processed off the underlying blockchain (e.g., Ethereum, a Layer-1 blockchain (L1)) to enhance scalability and reduce costs. Rollups bundle multiple transactions into a single batch, which is then submitted to the underlying blockchain as one transaction. This method reduces the load on the underlying chain while ensuring transaction security and finality through cryptographic proofs, such as ZKP used by ZKsync or validity checks. By processing transactions off-chain and periodically committing the results to the underlying chain, rollups improve throughput and efficiency without compromising the security and decentralization of the blockchain.

Transactions are identified by a unique transaction hash. When issuing a transaction, the user needs to specify parameters such as the recipient address (which can also be a smart contract and the functions the user wants to call), the number of tokens to transfer, the gas price, and the gas limit. The gas price represents the fee the user is willing to pay per unit of gas, while the gas limit is the maximum amount of gas the user is willing to consume for the transaction, a mechanism introduced to prevent infinite loops or excessive resource consumption.

<i>Notebook file</i>	<i>Description</i>
01-zksync-data.ipynb	It computes the basic statistics of the dataset and provides analyses used in Section 2. We use four main sources of data: <i>blocks</i> , <i>transactions</i> , <i>transaction receipts</i> , and <i>logs</i> .
02-data-exploration-fees.ipynb	It provides the analyses regarding gas usage and transaction fees for ZKsync used in Section 3.1. We use two main sources of data: <i>blocks</i> and <i>transaction receipts</i> .
03-data-exploration-contracts.ipynb	It analyzes the contract deployment and events triggered on ZKsync described in Section 3.2. We use one main source of data: <i>transaction logs</i> but also load <i>blocks</i> data to extract timestamps information.
04-data-exploration-swaps.ipynb	It analyzes the swap events on ZKsync described in Section 3.3. We use one main source of data: <i>transaction logs</i> but also load <i>blocks</i> data to extract timestamps information.

Table 1: Description of the code used for analysis in this paper. These notebook files are available on our GitHub repository [4] in the directory `./zksync-data-dump/notebooks/` and provide examples of how to interact and process our ZKsync public dataset.

<i>Chain</i>	<i>Start date</i>	<i>End date</i>	<i># of Blocks</i>	<i># of Transactions</i>	<i># of Issuers</i>	<i># of Contracts</i>	<i># of Logs</i>
ZKsync Era	February 14 th , 2023	March 24 th , 2024	29,710,983	327,174,035	7,322,502	1,631,772	2,044,221,151

Table 2: Description of our 1-year ZKsync dataset.

In our dataset, transactions should be sorted by the *blockNumber* and *transactionIndex* attributes to maintain the correct order of the transactions on the blockchain.

2.3 Transactions Receipts

Besides the raw transaction data described above, our dataset also contains a table of transaction receipts.

Transaction receipts provide a comprehensive summary of the outcome and effects of a transaction once it is processed and included in a block. They include the transaction hash, block number, and block hash to identify and verify the transaction, along with the sender (*from*) and recipient (*to*) addresses.

The receipts also detail the cumulative gas used by the transaction and all preceding transactions in the block, the actual gas used by the specific transaction, and the final gas price paid by the user. By multiplying the actual gas used by the gas price, we have the actual transaction fee the user paid.

Additional information includes *logs* for event logging (discussed next), the transaction status (success or failure), and the effective gas price paid. These receipts are crucial for users and developers to understand, audit, and interact with transactions and smart contracts on the blockchain. For example, they provide essential elements for the analysis, monitoring, and verification of transaction fees spent by users. We list the attributes of the transactions presented in our dataset in Table 5 in the Appendix.

Similarly to transactions, in our dataset, transaction receipt data should be sorted by attributes *blockNumber* and *transactionIndex* to maintain the correct order of the transactions on the blockchain.

2.4 Transactions Logs

In this section, we discuss the attributes of the transaction logs data in the ZKsync dataset in detail. Transaction logs are systematic records of events generated during the execution of transactions, particularly in interactions involving smart contracts. Each log entry contains *log index*, *data*, and *topics*, crucial to identifying

and categorizing specific events such as token transfers, approvals, swaps, minting, and voting.

These logs are emitted using the *emit* keyword within the smart contract code and play a pivotal role in monitoring activities, triggering actions within decentralized applications, and enabling event-driven programming. For instance, Decentralized Exchanges (DEXs) emit events upon trade executions, enabling user interfaces to update displays with current trade information.

These logs are stored in transaction receipts, offering a gas-efficient method to capture transient event data without permanently altering the blockchain’s state. Among the vast array of data accessible on EVM-based blockchains, transaction logs stand out as crucial sources of information for researchers, developers, and users. They facilitate analyses of various token transfer patterns and support blockchain analysis research, the focus of our contributions.

In our dataset, transaction logs should be sorted by the *blockNumber*, *transactionIndex*, and *logIndex* attributes to maintain the correct order in which they are stored on the blockchain. This is particularly important when analyzing the different states of a blockchain before and after the execution of a transaction that triggers a smart contract function.

2.4.1 Topics Attributes. The interpretation of the *topics* attributes (*topics₀*, *topics₁*, *topics₂*, and *topics₃*) depends on the implementation details of the invoked function within a smart contract. Typically, *topics₀* represents the event name, while subsequent topics represent indexed parameters of the event. The “data” attribute contains non-indexed event parameters. For example, in the context of a token transfer event, *topics₀* might indicate the event name *Transfer*, *topics₁* and *topics₂* could respectively denote sender and receiver addresses, and *data* would typically represent the number of tokens transferred.

2.4.2 Hashing and Signatures. *topics₀* corresponds to the hashed function signature using *keccak256* [12]. This signature consists of the function name followed by its parameter types. For

example, the signature of a typical *Transfer* event is *Transfer(address,address,uint256)*. After hashing it with keccak256, the result becomes *0xddf2... b3ef*, which is the *topics₀*. Below is a Python code snippet demonstrating how to verify if a given signature matches *topics₀*:

```
1 import web3
2 def check_sig(sig, topics_0):
3     return web3.Web3.keccak(text=sig).hex() == topics_0
```

2.4.3 Event Mapping. We provide a mapping of the most frequently invoked events within the ZKsync dataset in our GitHub repository [4] under */src/utls.py#events_dict*. This mapping facilitates the parsing of the majority of events in our dataset. The mapping is structured as a dictionary where the *topics₀* hex value serves as the key, and the corresponding value is a dictionary containing the parsed event name and its function signature. For instance, the *Transfer* event is represented as follows within the map:¹

```
1 events_dict['0xddf252ad ... f523b3ef'] = {
2     'name': 'Transfer',
3     'signature': 'Transfer(address,address,uint256)'}

```

2.5 L2 to L1 Logs

L2 to L1 logs are messages emitted by the ZKsync L2 network and transmitted to the Ethereum L1 network. They are essential for maintaining communication between the two layers, ensuring the security and integrity of transactions and data transfers. In ZKsync, the L1 smart contract verifies these communications by checking the messages alongside the ZKP. The only *provable* part of the communication from L2 to L1 is the native L2 to L1 logs emitted by the Virtual Machine (VM). These logs can be generated using the *to_l1 opcode* [8]. We refer the reader to the ZKsync documentation for more details [5]. We list the attributes of these logs in Table 7 in §A in the appendix.

3 Example Analyses

This section demonstrates some analyses that can be performed using our ZKsync dataset. Each subsection concentrates on a specific topic and utilizes their respective data (e.g., blocks, transactions, receipts, and logs). We consider all activity that starts with block 561,367, the first block on April 1st, until block 29,710,983, the last block in the ZKsync dataset.

This section should serve as a starting point for other researchers wishing to use this dataset for their investigations. The code for generating this analysis is discussed in Table 1 and can be accessed on our [GitHub repository](#) [4].

3.1 Gas Usage and Transaction Fees

To start, we look into transactions, gas usage, and fees. All results are generated from two main sources of data, namely blocks and transaction receipts. Note that we are using the transaction receipts instead of the transaction data because the receipts are the source of the actual units of gas used and the final transaction fee paid by users.

Figure 2 shows the daily transactions executed in ZKsync Era over the period analyzed. The network processed an average of

¹We shortened *topics₀* to *0xddf252ad... f523b3ef* for better visualization in the paper.

905,194 transactions daily, with a significant spike in December 2023 that 5,362,921 transactions in a single day. This spike was due to a boom in inscriptions, which became very popular for memecoin traders around this time. Messias et al. [30] did a comprehensive review of this phenomenon across various rollups. Since the spike stabilized, we have seen a slight increase in daily transactions, which are now hovering around 1.2 million transactions per day.

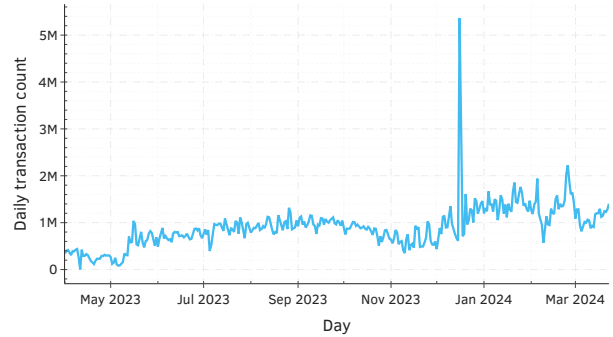


Figure 2: Total transactions executed per day.

Now, looking at gas usage, Figure 3 illustrates the total number of gas units used each day in the ZKsync Era. On average, the network has processed 567,111,979,658 gas units per day in the last year. However, gas usage has been fairly volatile. Particularly, it has experienced three significant spikes of more than 1.5 trillion gas units per day, which we detail next.

May 2023. 30% of all gas used during the month of May 2023 can be attributed to the zkApes airdrop. The address receiving the most gas units was the SyncSwap router contract, which is used to execute swaps in the largest DEX in ZKsync Era. During this time, we see more than 10% of all swaps trading ZAT, which does not occur when we look at all swaps during the year under analysis. The second address receiving the most gas units was the contract used to claim the zkApes token (ZAT) in the context of the airdrop.

December 2023. This was caused by the inscriptions boom we discussed before.

March 2024. The gas utilization during this spike is much more distributed among transaction receivers, with the top receivers being DEX routers, such as the SyncSwap V2 router and the Mute.io router (which recently re-branded to Koi), and token contracts, such as USDC.e and SOUL. This dispersed distribution of gas usage coupled with the observed steady growth in both gas usage and transaction count up to the spike suggests that it is due to the increased trading activity on ZKsync Era.

We can further contextualize gas utilization by looking at Figure 4, which shows the average number of gas units used per transaction. We see that in May 2023, transactions were using on average a significantly higher number of gas units (reaching 7,265,323 units per transaction) compared to the average over the 1-year period (784,149 units).

Finally, we can look at the transaction fees. Figure 5 shows the average transaction fees over time. In other words, this is the cost that users have paid, on average, to submit one unit of gas to ZKsync




Figure 3: Total gas units used per day.

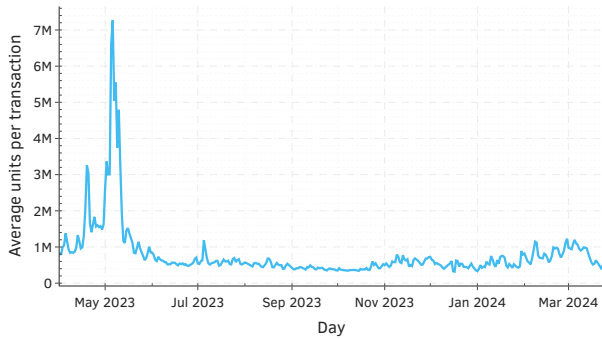


Figure 4: Average gas units used per transaction.

Era. Transaction fees were fairly stable during 2023, with an average of 0.25 gwei per unit of gas.

However, in 2024, a key upgrade led to a significant decrease in ZKsync transaction fees. Concretely, the implementation of the new prover, Boojum [27], marked a significantly reduced hardware requirement to run a Zero-Knowledge (ZK) prover and thus allowed a reduction of transaction fees to 0.1 gwei per unit of gas.

Then, in March 2024, the Dencun upgrade was implemented and deployed on the Ethereum mainnet. This upgrade, brought in EIP-4844, introduced a new type of transaction that can store “blobs” of data in the beacon node for 14 days [47]. Blobs have their independent fee model and submitting data as blobs is much cheaper than the previously used call-data [36]. ZKsync Era was one of the first rollups using blobs to publish the data related to its state changes, thus further reducing transaction fees to 0.025 gwei per unit of gas.

3.2 Events and Contract Deployments

Events are an important source of additional data about activity on EVM chains. They correspond to data emitted by smart contracts and stored on-chain. In this subsection, we analyze this data and report on some specific events. Recall that events can be obtained from the logs data in the ZKsync dataset as discussed in Section 2.

Figure 6 shows the top 15 event types with the most emitted events during the period under analysis. Different contracts may




Figure 5: Average transaction fee in gwei per unit of gas by day.

emit the same event type, so what distinguishes them is the event’s function signature, and thus its respective hash.

There is a significant overrepresentation of the top 4 event types, with *Transfer* events being by far the largest type (70.9% of all events are Transfer events). This is not unexpected as this event is emitted every time an ERC-20 token is transferred between two addresses. This occurs in both simple token transfers, but also as part of other common contract operations such as swaps in DEXs.

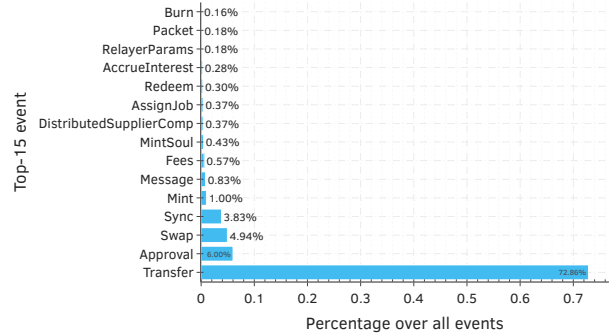


Figure 6: Top 15 event types with the highest number of events emitted between April 1st, 2023 and March 24th, 2024, and their percentage over all events emitted. Percentages include ETH transfers generated from transaction fee payments.

Before looking at *Transfer* events in more detail, we should highlight a particularity in how ZKsync Era implements transaction fee collection. In this chain, every transaction generates two additional ETH transfers — one for the initial payment of the transaction fees from the transaction submitter and another with a transaction fee refund (after all L1 and proving costs are accounted for). These transaction fee transfers always appear as ETH transfers from or to the address *0x8001*² and thus generate *Transfer* events.

These specific events account for 38.2% of all events emitted and, if we exclude them entirely from the event dataset, *Transfer* events only account for 56.1% of these filtered events.

²We shortened this address for better visualization in the paper.

After filtering the Transfer events generated by fee management ETH transfers, we can explore the top tokens transferred by examining the contract address that emitted the event. Figure 7 displays the 15 tokens most frequently transferred during the period.

Native ETH, Wrapped ETH, and Bridged USDC (USDC.e) are the most transfers, accounting collectively for 78.4%. We can also see some other stablecoins such as Tether USD or Binance USD, and Liquidity Provider (LP) tokens associated with DEXs on ZKsync Era, such as SyncSwap, and SpaceSwap.

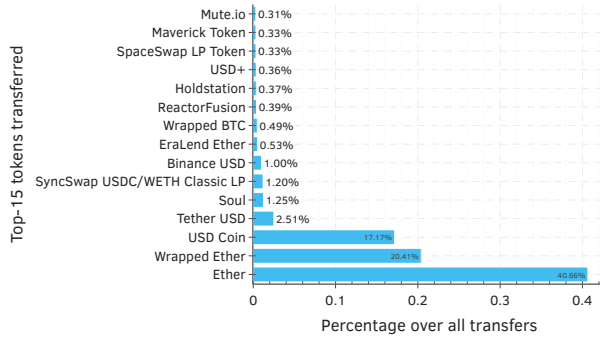


Figure 7: Top 15 ERC-20 tokens with the highest number of Transfer events emitted between April 1st, 2023 and March 24th, 2024, and their percentage over all Transfer events emitted. Percentages exclude ETH transfers generated from transaction fee payments.

Approvals are the second-largest event type, with 6% of events emitted. Similarly to the Transfer events, we can see which contracts are emitting these events. Figure 8 shows these top 15 emitters.

The top emitters are also ERC-20 token contracts, and these events represent an owner “approving” a spender to transfer a predefined amount of tokens they hold. This is common in bridged assets, for example, as is the case of Bridged USDC (USDC.e) and Wrapped ETH, the top 2 emitters.




Figure 8: Top 15 contracts with the highest number of Approval events emitted between April 1st, 2023 and March 24th, 2024, and their percentage over all Approval events emitted.

Swaps and Syncs are the third and fourth most frequently emitted event types, accounting for 4.9% and 3.8% of all events emitted,

respectively. These events are a key source of data for analyzing DEXs. However, we will leave a more in-depth exploration of these in the next subsection.

Finally, we examine contract deployments. Every time a contract is deployed, a specific event type is emitted, which allows us to easily track this metric. Although this event type is not among the 15 most emitted event types, it is still an important metric for network activity. In ZKsync, the contract deployment event is named *ContractDeployed*.

Figure 9 shows the number of contracts deployed on ZKsync each day. Before September 2023, developers averaged 2510 contract deployments per day. However, there were two major spikes during this time, the first reaching 29,114 and the second reaching 17,602 contract deployments in a single day. Then, after September 2023, contract deployments increased to a daily average of 6672.

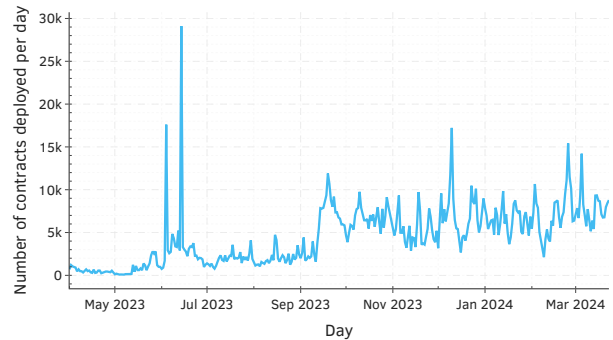


Figure 9: Total number of contracts deployed on ZKsync Era per day.

3.3 Swaps

After a high-level look at these events, we now focus on a specific event that is relevant to understanding activity on DEXs — the *Swap* event. Recall that swap events account for 4.9% of all events emitted on ZKsync.

Swap events are emitted every time a successful swap is performed in a DEX. These events contain information about the contract that emits the event, the amounts of each token being traded, and the wallets involved in the trade.

Figure 10 shows the number of swap events emitted each day during the analyzed period, which is equivalent to the number of swaps performed each day on ZKsync DEXs. We have a long-term trend of increasing the number of daily swaps, with a peak around March 2024 (which reached 531,819 swaps in a single day). During this year, users performed an average of 192,009 swaps per day.

We also see which LP contract was involved in the swap by looking at the contract address that emitted the event. Depending on the DEXs protocol, this may be the actual liquidity pool involved in the swap (e.g., SyncSwap and Koi) or a generic LP contract managing all the pools in the DEX (e.g., SpaceSwap and PancakeSwap). To understand which tokens were being traded in these generic LP contracts, we would have to process the transfer events emitted in the same transaction of these Swap events.

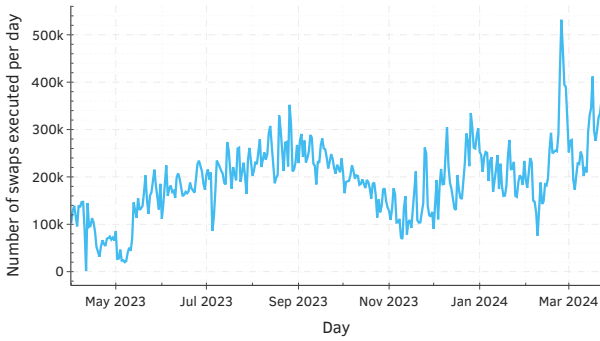


Figure 10: Number of swap events emitted per day. This corresponds to the number of swaps performed each day.

Figure 11 shows the top 15 contracts that emitted the most swaps even during the period under analysis. The largest emitter is, by a significant margin, the USDC/WETH pool on SyncSwap, which accounts for almost a third of all emitted Swap events. All combined swap events from SpaceSwap represent 11% of all swaps, followed by two USDC/WETH pools from other DEX.

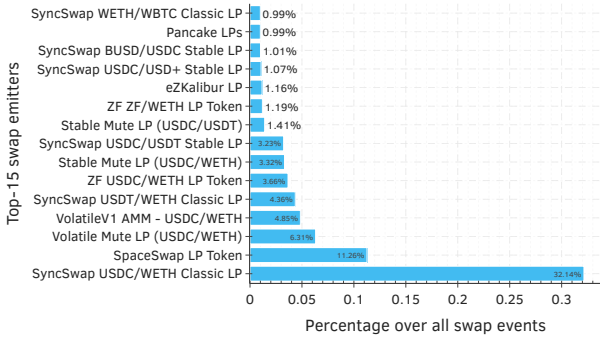


Figure 11: Top 15 contracts emitting the most Swap events between April 1st, 2023 and March 24th, 2024, and their percentage over all Swap events emitted.

Finally, we can use the “receivers” field to analyze swappers. In other words, this field provides the wallet address that is receiving the swapped token. Figure 12 displays the distribution of swaps made by unique wallet addresses.

We note that most of the users performed less than 50 swaps during the period under analysis. Concretely, the percentile 95% of this distribution is 42 swap events. However, the distribution has a significant skew to the right, with a few addresses generating a large number of swaps in the year examined. These large “traders” are usually routers and other protocols that interact with DEXs; thus, they represent many end-users. Examples include the top swappers, which are the Mute.io router, the SpaceFi router, and the Odos V2 router, respectively.

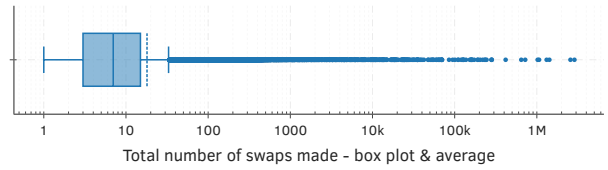


Figure 12: Distribution of swaps per unique wallet addresses.

4 Future Directions

Among the data analysis presented in this paper, our ZKsync dataset can be applied to different studies. We list some of them below, where this dataset can be valuable.

MEV and Arbitrage. MEV and arbitrage have been extensively studied in L1 blockchains [31, 43–45, 48]. However, only recently have new studies shifted their focus to L2s [11, 20, 21, 46], including ZKsync Era. Our dataset can contribute to this type of research by enabling further analysis of MEV on ZKsync Era. For example, one type of MEV known as *backrunning* involves arbitrageurs ensuring their transactions are included immediately after a target transaction. This can be particularly useful in scenarios like liquidation-MEV, where arbitrageurs take advantage of opportunities that arise right after an oracle update [31, 43, 48]. Another form of arbitrage worth analyzing is CEX-DEX MEV, where arbitrageurs exploit price deviations between Centralized Exchange (CEX) and DEX platforms. Additionally, studying cross-rollup MEV could provide insights into opportunities where arbitrageurs benefit from differences across two or more rollups.

Analysis of user activity on chain. Analyzing user activity in the ZKsync chain provides valuable insights into user behavior. This includes examining the transactions issued by users and their interactions with smart contracts. Such an analysis can also help identify airdrop farmers who create multiple accounts, known as Sybils, although this task can be challenging [33]. Another important area of analysis is measuring the impact of users in social media networks driving activities on the blockchain. For example, examining the recent boom in inscriptions in L2s chains can reveal how social networks can influence blockchain activity [30, 35]. Additionally, since our dataset includes all event logs, it allows for the analysis of decentralized governance of protocols deployed on ZKsync. This involves studying proposals to amend smart contracts and their voting processes. By filtering data related to governance contracts, such as how each user voted and the distribution of these governance tokens among users, it can shed light on the implications of token concentration on decentralized governance, raising concerns about fairness [32].

Data Science Analytics. This dataset can also be valuable for users or non-researchers interested in exploring blockchain data. It offers an opportunity for those who wish to gain a deeper understanding of blockchains. For example, data scientists can utilize this dataset to explore and analyze blockchain data on public platforms such as Kaggle [25]. Data scientists widely use it to demonstrate their skills in data analytics, machine learning, and data science in general. Therefore, this dataset can help data scientists acquire new

skills, giving them a competitive edge in the market or improving their prospects of securing a job in a blockchain company.

We hope that this dataset proves to be a valuable resource for research groups interested in conducting studies on L2s blockchains and ZKsync Era. To facilitate this, we have made the dataset available in a GitHub repository [4].

5 Conclusion

At a high level, this paper addresses a critical concern within the scientific community: *the availability and accessibility of blockchain data*. Data-driven research is dependent on high-quality datasets to make meaningful findings. Although blockchain data are theoretically publicly available, practical challenges often prevent researchers, especially those without technical expertise, from obtaining them. In addition, the cost of using external services or deploying infrastructure to run archive nodes can be prohibitive.

To address these challenges, we have collected, pre-processed, and made available our ZKsync dataset to facilitate access for researchers. We also provide a detailed background on blockchains, including blocks, transactions, receipts, and logs, designed to help nonexperts understand and utilize the data effectively. To illustrate the potential of this dataset, we offer example analyses and discuss future research directions.

We believe that our contribution will be valuable to researchers studying blockchain L2 ecosystems, particularly those interested on ZKsync. This dataset also adds to the existing body of L2 research and is designed for ease of use, using the Python library Polars to allow straightforward data processing on a local laptop.

Finally, to promote scientific reproducibility and to support further research, we have made our dataset publicly available on GitHub [4].

References

- [1] 2017. EIP-712: Typed structured data hashing and signing. <https://eips.ethereum.org/EIPS/eip-712>.
- [2] 2019. EIP-1559: Fee market change for ETH 1.0 chain. <https://eips.ethereum.org/EIPS/eip-1559>.
- [3] 2020. EIP-2930: Optional access lists. <https://eips.ethereum.org/EIPS/eip-2930>.
- [4] 2024. Data sets and scripts used to analyze the ZKsync Era blockchain. <https://github.com/matter-labs/zksync-data-dump>.
- [5] 2024. L1 <-> L2 Communication. https://docs.zksync.io/zk-stack/concepts/l1_l2_communication.
- [6] 2024. NANSEN Query: Empowering Crypto Teams With In-Depth Blockchain Data. <https://www.nansen.ai/query>.
- [7] 2024. Reth Book. <https://reth.rs>.
- [8] 2024. System contracts/bootloader description (VM v1.4.0). <https://github.com/code-423n4/2023-10-zksync/blob/main/docs/SmartcontractSection/Systemcontractsbootloaderdescription.md>.
- [9] 2024. Transaction Lifecycle. <https://docs.zksync.io/zk-stack/concepts/transaction-lifecycle#transaction-types>.
- [10] Arbitrum. 2024. Arbitrum (ETH) Blockchain Explorer. <https://arbiscan.io>. Accessed on June 5, 2024.
- [11] Arthur Bagourd and Luca Georges Francois. 2023. Quantifying MEV On Layer 2 Networks. *arXiv preprint arXiv:2309.00629* (2023).
- [12] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. 2009. Keccak sponge function family main document. *Submission to NIST (Round 2)* (2009).
- [13] bitcoin.org. 2024. Bitcoin Core. <https://bitcoin.org/en/bitcoin-core>.
- [14] bnbchain.org. 2024. BNB Smart Chain. <https://www.bnbchain.org/en/bnb-smart-chain>.
- [15] Vitalik Buterin. 2020. A rollup-centric ethereum roadmap. <https://ethereum-magicians.org/t/a-rollup-centric-ethereum-roadmap/4698>.
- [16] Alen Day and Evgeny Medvedev. 2019. Introducing six new cryptocurrencies in BigQuery Public Datasets—and how to analyze them. <https://cloud.google.com/blog/products/data-analytics/introducing-six-new-cryptocurrencies-in-bigquery-public-datasets-and-how-to-analyze-them>
- [17] Dune Analytics. 2024. Dune – Crypto Analytics Powered by Community. <https://dune.com/home>. Accessed on June 5, 2024.
- [18] Etherscan. 2024. Etherscan (ETH) Blockchain Explorer. <https://etherscan.io>. Accessed on June 5, 2024.
- [19] Krzysztof Gogol, Robin Fritsch, Johnnatan Messias, Malte Schlosser, Benjamin Kraner, and Claudio Tessone. 2024. Liquid Staking Tokens in Automated Market Makers. (2024).
- [20] Krzysztof Gogol, Johnnatan Messias, Deborah Miori, Claudio Tessone, and Benjamin Livshits. 2024. Layer-2 Arbitrage: An Empirical Analysis of Swap Dynamics and Price Disparities on Rollups. *arXiv preprint arXiv:2406.02172* (2024).
- [21] Krzysztof Gogol, Johnnatan Messias, Deborah Miori, Claudio Tessone, and Benjamin Livshits. 2024. Quantifying Arbitrage in Automated Market Makers: An Empirical Study of Ethereum ZK Rollups. *arXiv preprint arXiv:2403.16083* (2024).
- [22] Krzysztof Gogol, Johnnatan Messias, Malte Schlosser, Benjamin Kraner, and Claudio Tessone. 2023. Cross-border Exchange of CBDCs using Layer-2 Blockchain. *arXiv preprint arXiv:2312.16193* (2023).
- [23] Krzysztof Gogol, Johnnatan Messias, Maria Ines Silva, and Benjamin Livshits. 2024. The Writing is on the Wall: Analyzing the Boom of Inscriptions and its Impact on Rollup Performance and Cost Efficiency. *arXiv preprint arXiv:2404.11189* (2024).
- [24] Thomas Johnsson. 1984. Efficient compilation of lazy evaluation. In *Proceedings of the 1984 SIGPLAN symposium on Compiler construction*.
- [25] Kaggle. 2024. Kaggle: Your Machine Learning and Data Science Community. <https://www.kaggle.com>.
- [26] L2Beat. 2024. L2Beat: The state of the layer two ecosystem. <https://l2beat.com/scaling/summary?sort-by=total&sort-order=desc#layer2s>.
- [27] Matter Labs. 2024. Boojum Upgrade: zkSync Era’s New High-performance Proof System for Radical Decentralization. <https://zksync.mirror.xyz/HJ2Pj45EJkRdt5Pau-ZXwkV2ctPx8qFL19STM5jdYhc>.
- [28] Johnnatan Messias, Mohamed Alzayat, Balakrishnan Chandrasekaran, Krishna P. Gummadi, Patrick Loiseau, and Alan Mislove. 2021. Data sets and scripts used for analyzing “norm” violations in Bitcoin. <https://github.com/johnnatan-messias/blockchain-transaction-ordering>.
- [29] Johnnatan Messias, Mohamed Alzayat, Balakrishnan Chandrasekaran, Krishna P. Gummadi, Patrick Loiseau, and Alan Mislove. 2021. Selfish & Opaque Transaction Ordering in the Bitcoin Blockchain: The Case for Chain Neutrality. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC ’21)*.
- [30] Johnnatan Messias, Krzysztof Gogol, Maria Inês Silva, and Benjamin Livshits. 2024. The Writing is on the Wall: Analyzing the Boom of Inscriptions and its Impact on EVM-compatible Blockchains. *arXiv preprint arXiv:2405.15288* (2024).
- [31] Johnnatan Messias, Vabuk Pahari, Balakrishnan Chandrasekaran, Krishna P. Gummadi, and Patrick Loiseau. 2023. Dissecting Bitcoin and Ethereum Transactions: On the Lack of Transaction Contention and Prioritization Transparency in Blockchains. In *Proceedings of the Financial Cryptography and Data Security (FC’23)*.
- [32] Johnnatan Messias, Vabuk Pahari, Balakrishnan Chandrasekaran, Krishna P. Gummadi, and Patrick Loiseau. 2024. Understanding Blockchain Governance: Analyzing Decentralized Voting to Amend DeFi Smart Contracts. *arXiv:2305.17655 [cs.CR]* <https://arxiv.org/abs/2305.17655>
- [33] Johnnatan Messias, Aviv Yaish, and Benjamin Livshits. 2023. Airdrops: Giving money away is harder than it seems. *arXiv preprint arXiv:2312.02752* (2023).
- [34] Offchain Labs. 2024. A gentle introduction to Arbitrum. <https://docs.arbitrum.io/intro>.
- [35] Janna Joceli Omena, Johnnatan Messias, Fabio Gouveia, and Riccardo Ventura. 2024. Digital Methods for Blockchain Research. https://www.researchgate.net/publication/382511569_Digital_Methods_for_Blockchain_Research.
- [36] Omkar Godbole. 2024. Layer 2 Blockchains Become Cheaper After Ethereum’s Dencun Upgrade. <https://www.coindesk.com/markets/2024/03/14/layer-2-blockchains-become-cheaper-after-ethereums-dencun-upgrade>.
- [37] Optimism Foundation. 2024. Optimism. <https://www.optimism.io>.
- [38] Pandas. 2024. Pandas: Python data analysis library. <https://pandas.pydata.org>.
- [39] Paradigm. 2024. Paradigm Data Portal. <https://www.paradigm.xyz/oss/portal>.
- [40] Paradigm. 2024. Paradigm Data Portal. <https://github.com/paradigmxyz/paradigm-data-portal>.
- [41] Polars. 2024. Lazy / Eager API: Polars user guide. <https://docs.pola.rs/user-guide/concepts/lazy-vs-eager>.
- [42] Polars. 2024. Polars: Dataframes for the new era. <https://pola.rs>.
- [43] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying blockchain extractable value: How dark is the forest?. In *2022 IEEE Symposium on Security and Privacy (SP)*.
- [44] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. 2021. Attacking the defi ecosystem with flash loans for fun and profit. In *International conference on financial cryptography and data security*.
- [45] Christof Ferreira Torres, Ramiro Camino, et al. 2021. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*.

- [46] Christof Ferreira Torres, Albin Mamuti, Ben Weintraub, Cristina Nita-Rotaru, and Shweta Shinde. 2024. Rolling in the Shadows: Analyzing the Extraction of MEV Across Layer-2 Rollups. *arXiv preprint arXiv:2405.00138* (2024).
- [47] Vitalik Buterin and Dankrad Feist and Diederik Loerakker and George Kadianakis and Matt Garnett and Mofi Taiwo and Ansgar Dietrichs. 2022. EIP-4844: Shard Blob Transactions. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-4844.md>.
- [48] Ben Weintraub, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. 2022. A flash (bot) in the pan: measuring maximal extractable value in private pools. In *Proceedings of the 22nd ACM Internet Measurement Conference*.
- [49] ZettaBlock Inc. 2024. Full Stack Data Platform To Build Web3 Apps in Minutes. <https://zettablock.com>. Accessed on June 5, 2024.

A Additional Information for Data Schemas

In this section, we present the attributes available in our dataset, including their data types and brief descriptions. Our dataset consists of one year of ZKsync data, containing all blocks, transactions, receipts, and logs included in the ZKsync blockchain.

Table 3 describes the attributes of each block added to the ZKsync blockchain. Table 4 details the transaction data, including their data types and descriptions. Table 5 provides information on transaction receipts, which are generated immediately after a transaction is executed and added to a block. This is useful for computing the actual gas spent on a transaction and the corresponding transaction fees paid by users.

Next, Table 6 contains information about transaction logs, which is crucial for analyzing changes in smart contract states such as

account balances, token transfers, votes cast, and swaps. Finally, Table 7 presents the attributes related to the messages emitted by ZKsync to the Ethereum mainnet.

B Glossary

Following is a list of important notations used in this paper.

Acronyms

AMM	Automated Market Maker
CEX	Centralized Exchange
DEX	Decentralized Exchange
EVM	Ethereum Virtual Machine
L1	Layer-1 blockchain
L2	Layer-2 blockchain
LP	Liquidity Provider
MEV	Maximal-Extractable Value
RPC	Remote Procedure Call
TVL	Total Value Locked
VM	Virtual Machine
ZK	Zero-Knowledge
ZKP	Zero-Knowledge Proof

<i>Attribute</i>	<i>Type</i>	<i>Description</i>
hash	str	Unique identifier for the block.
parentHash	str	Unique identifier of the parent block.
sha3Uncles	str	SHA-3 hash of the uncles' block headers. On ZKsync it is set to $0x1dcc \dots 9347$ since there are no uncle blocks.
miner	str	Address of the miner who mined the block. This is set as Null address ($0x0$) in all ZKsync blocks since it does not have miners or block validators.
stateRoot	str	Root hash of the state trie. Set to Null address ($0x0$).
transactionsRoot	str	Root hash of the transaction trie. Set to Null address ($0x0$).
receiptsRoot	str	Root hash of the receipts trie. Set to Null address ($0x0$).
number	i64	Block number or height.
l1BatchNumber	str	L1 batch number, related to the sequence of batches submitted to Layer 1 on zkRollup systems.
gasUsed	i64	Total amount of gas used by all transactions in the block.
gasLimit	i64	Maximum amount of gas that can be used by all transactions in the block. Set to 4,294,967,296 (2^{32}) units of gas.
baseFeePerGas	i64	Base fee, in Wei ($10^{-18} ETH$), per unit of gas.
extraData	str	Extra data included by the miner in the block. Set to $0x$ since there are no miners.
logsBloom	str	Bloom filter for the logs of the block. Set to $0x0 \dots 0$.
timestamp	i64	Unix timestamp of when the block was collated.
l1BatchTimestamp	str	L1 batch timestamp (in HEX format) associated with the block.
difficulty	i64	Difficulty target for mining the block. Set to 0 since there is no mining.
totalDifficulty	i64	Cumulative difficulty of the blockchain up to and including this block. Set to 0 since there is no mining.
sealFields	list[null]	Seal fields containing proof-of-work or proof-of-stake information. List containing null as value.
uncles	list[null]	Uncle blocks that were mined but not included in the main chain. List with Null value since there is no mining.
size	i64	Size of the block in bytes. Set to 0.
mixHash	str	Hash used in the mining process to prove that enough computational work has been performed. Set to $0x0 \dots 0$ since there is no mining.
nonce	str	Value used in the mining process to find a valid block hash. Set to $0x0 \dots 0$ since there is no mining.

Table 3: Description of block attributes on ZKsync.

<i>Attribute</i>	<i>Type</i>	<i>Description</i>
blockHash	str	Unique identifier of the block containing the transaction.
blockNumber	i64	Block number or height containing the transaction.
chainId	i64	Identifier of the blockchain network. Set to 324 that represents ZKsync chain.
from	str	Address of the sender of the transaction.
gas	i64	Amount of gas provided as <i>gasLimit</i> for the transaction.
gasPrice	i64	Price per unit of gas the sender is willing to pay.
hash	str	Unique identifier for the transaction.
input	str	Data sent along with the transaction. In HEX code.
l1BatchNumber	str	L1 batch number related to the transaction in zkRollup systems.
l1BatchTxIndex	str	Index of the transaction in the L1 batch.
maxFeePerGas	i64	Maximum fee, in Wei (10^{-18} ETH), per unit of gas.
maxPriorityFeePerGas	i64	Maximum priority fee, in Wei (10^{-18} ETH), per unit of gas.
nonce	i64	Number of transactions sent by the sender prior to this one.
r	str	First part of the ECDSA signature.
s	str	Second part of the ECDSA signature.
to	str	Address of the receiver of the transaction.
transactionIndex	i64	Index of the transaction within the block.
type	i64	Type of transaction divided into 5 categories: <i>Legacy</i> (0 or <i>0x0</i>), <i>EIP-2930</i> (1 or <i>0x1</i>) [3], <i>EIP-1559</i> (2 or <i>0x2</i>) [2], <i>EIP-712</i> (113 or <i>0x71</i>) [1], and <i>Priority</i> (255 or <i>0xff</i>). See details in ZKsync documentation [9].
v	f64	Recovery id of the ECDSA signature.
value	str	Amount of tokens (in Wei and in HEX format) that are transferred in the transaction to the recipient address (<i>to</i>).

Table 4: Description of transaction attributes on ZKsync.

<i>Attribute</i>	<i>Type</i>	<i>Description</i>
blockHash	str	Unique identifier of the block containing the transaction.
blockNumber	i64	Block number or height containing the transaction.
contractAddress	str	Address of the contract created by the transaction, if applicable.
cumulativeGasUsed	i64	Total amount of gas used when the transaction was executed in the block. Set to 0 on ZKsync.
effectiveGasPrice	i64	Actual price per unit of gas, in Wei (10^{-18} ETH), paid.
from	str	Address of the sender of the transaction.
gasUsed	i64	Amount of gas used by the transaction.
l1BatchNumber	str	L1 batch number related to the transaction in zkRollup systems.
l1BatchTxIndex	str	Index of the transaction in the L1 batch.
logsBloom	str	Bloom filter for the logs of the transaction. Set to <i>0x0 ... 0</i> on ZKsync.
root	str	State root after the transaction is executed.
status	i64	Status of the transaction (1 for success, 0 for failure).
to	str	Address of the receiver of the transaction.
transactionHash	str	Unique identifier for the transaction.
transactionIndex	i64	Index of the transaction within the block.
type	i64	Type of transaction divided into 5 categories. Refer to Table 4 and to ZKsync documentation [9] for details.

Table 5: Description of transaction receipt attributes on ZKsync.

Attribute	Type	Description
address	str	Address of the contract that generated the log.
blockHash	str	Unique identifier of the block containing the transaction.
blockNumber	i64	Block number or height containing the transaction.
data	str	Data contained in the log. This can be used, for example, to extract the amount of tokens transferred from one user to another.
l1BatchNumber	str	L1 batch number related to the log in zkRollup systems.
logIndex	i64	Index of the log within the block.
logType	null	Type of log. Set to <i>null</i> on ZKsync.
removed	bool	Indicates whether the log was removed (<i>true</i>) or not (<i>false</i>).
transactionHash	str	Unique identifier for the transaction.
transactionIndex	i64	Index of the transaction within the block.
transactionLogIndex	str	Index of the log within the transaction. In HEX format.
topics ₀	str	First topic of the log. This is typically referred to as the name of the event encoded in hexadecimal (HEX) format.
topics ₁	str	Second topic of the log.
topics ₂	str	Third topic of the log.
topics ₃	str	Fourth topic of the log.

Table 6: Description of transaction log attributes on ZKsync.

Attribute	Type	Description
blockHash	str	Unique identifier of the block containing the log.
blockNumber	str	Block number or height containing the log.
isService	bool	Indicates whether the log is a service log (<i>true</i>) or not (<i>false</i>).
key	str	Key associated with the log that could be used to carry some data with the log.
l1BatchNumber	str	L1 batch number related to the log in zkRollup systems.
logIndex	str	Index of the log within the block.
sender	str	It is the value of <i>this</i> in the frame where the L2→L1 log was emitted.
shardId	str	It is the id of the shard the opcode was called. It is currently set to 0.
transactionHash	str	Unique identifier for the transaction associated with the log.
transactionIndex	str	Index of the transaction within the block.
transactionLogIndex	str	Index of the log within the transaction.
txIndexInL1Batch	str	Index of the transaction within the L1 batch.
value	str	Value associated with the log that could be used to carry some data with the log.

Table 7: Description of L2 to L1 log attributes on ZKsync.