How to Read a Blockchain Explorer: OSINT Guide to Tracing Crypto Transactions

Maria Cattini | 18/05/2026 | OSINT

A lawyer sends you two wallet addresses — one Ethereum, one Bitcoin — and a single sentence: "I need to know where the money went." No names, no context, no additional data. Just two strings of alphanumeric characters and the expectation that you can reconstruct what happened.

This is not a hypothetical. It's the standard entry point for cryptocurrency fraud investigations. And the first decision — which tool to open, which field to read, which pattern signals a problem — determines whether the next two hours produce usable intelligence or a dead end.

The Public Ledger Problem

Blockchain data is radically public. Every transaction on Bitcoin, Ethereum, and most major cryptocurrencies is permanently recorded and accessible to anyone with an internet connection, without registration, without authorization, and without cost. No journalist, investigator, or analyst needs a court order to query a wallet's complete transaction history. That transparency is the foundational property that makes blockchain OSINT both powerful and frequently misread.

The misunderstanding cuts both ways. Investigators who don't work with crypto often assume pseudonymity means anonymity — that blockchain data is opaque. The opposite error is equally common: treating a wallet address as a confirmed identity when it's still a pseudonym. A blockchain explorer shows every movement, every timestamp, every counterparty address. It does not show who controls those addresses. That gap — between transaction data and attributed identity — is where most blockchain investigations stall, and where most methodological errors accumulate.

Block explorers are the interface layer between raw blockchain data and human-readable investigation. They function like search engines pointed at the ledger: input a wallet address, a transaction hash, or a block number, and the tool returns structured data about that entry. Three explorers cover the operational core: **Etherscan** (etherscan.io) for Ethereum and compatible chains, **Blockchain.com** (blockchain.com/explorer) for Bitcoin, and **mempool.space** for Bitcoin with real-time mempool visibility. These are not interchangeable. Each is chain-specific. Querying a Bitcoin address in Etherscan returns nothing — the ledgers are separate infrastructures.

Identifying which explorer to use requires only one check: if the address begins with 0x, it's Ethereum — use Etherscan. If it begins with 1, 3, or bc1, it's Bitcoin — use Blockchain.com or mempool.space.

System Map: Where the Data Lives and What It Cannot Tell You

Blockchain data is distributed across nodes, but practically accessed through a small number of high-reliability explorers. The information architecture has distinct layers.

What is **fully visible**: complete transaction history of any public address, timestamps in UTC, amounts transferred, transaction fees paid (gas, in Ethereum's terminology), the status of each transaction (confirmed, pending, failed), and — on Ethereum — interactions with smart contracts,

internal transactions generated by those contracts, and token transfers (ERC-20). On Bitcoin, the UTXO model makes inputs and outputs explicit: every transaction lists the source addresses and destination addresses with exact amounts.

What is **partially visible**: address labels. Platforms like Etherscan crowd-source and maintain tags for known entities — exchanges, DeFi protocols, mixers, flagged fraud addresses. These labels are community-generated and not authoritative, but they provide investigative leads. A destination address tagged as belonging to a centralized exchange is structurally different from an untagged address.

What is **missing**: the identity of wallet controllers. Blockchain addresses are pseudonymous. The ledger records the address, not the person. Linking an address to a real identity requires cross-referencing with external sources — exchange registration data, court records, data leaks, or behavioral patterns that connect an address to other identifiable activity. No explorer provides that link directly.

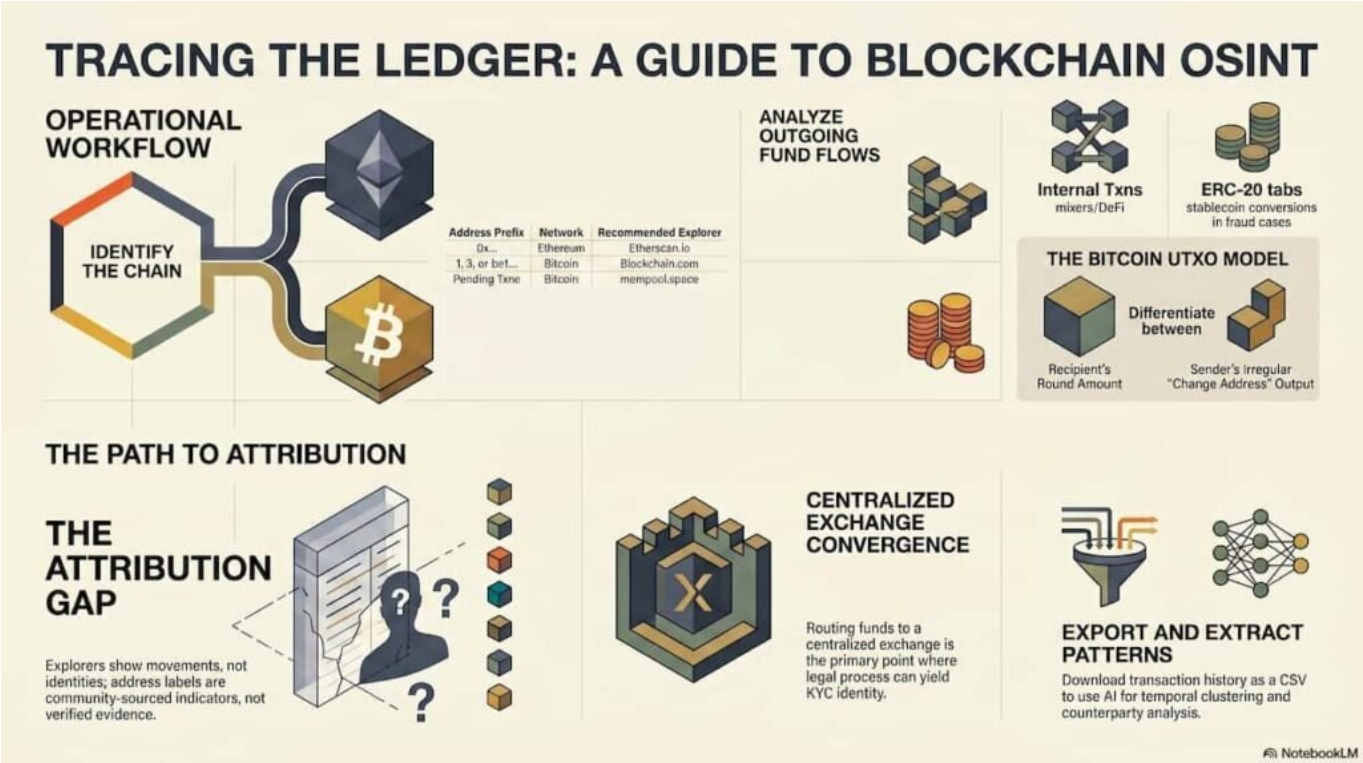

Operational Method

Step 1 — Identify the chain and open the correct explorer. Determine whether the address is Ethereum (0x…) or Bitcoin (1, 3, or bc1). Open etherscan.io or blockchain.com/explorer accordingly. No account required for either.

Step 2 — Read the wallet summary. On Etherscan, the overview shows current ETH balance, total transaction count, and token holdings. On Blockchain.com, the summary displays Total Received, Total Sent, and Final Balance. A Final Balance of zero means funds have already moved. A high transaction count — thousands of operations — is inconsistent with a personal wallet and suggests an exchange, service, or aggregation address.

Step 3 — Analyze outgoing transactions to map fund flow. On Etherscan, the Transactions tab shows ETH movements. The Internal Txns tab shows movements generated by smart contract interactions — this is where mixer activity and DeFi protocol usage appear, not in the main tab. The ERC-20 Token Txns tab captures stablecoin transfers (USDT, USDC, DAI). Stablecoins are operationally significant in fraud cases: their dollar-pegged value makes them the preferred vehicle for consolidating proceeds before moving to an exchange.

Step 4 — On Ethereum, read the transaction-level fields that carry investigative weight. Status (Success vs. Failed) matters: failed transactions remain permanently recorded. A pattern of failed transactions may indicate probing behavior. Transaction Fee (gas) signals urgency — an abnormally high fee on an outgoing transaction suggests the sender prioritized speed. The To field shows whether the destination is a standard address or a smart contract (Contract). The Input Data field, decoded via Etherscan's built-in decoder, reveals the function called — token swap, spend approval, NFT transfer. This distinguishes a direct transfer from obfuscated layering through DeFi.

Step 5 — On Bitcoin, account for the UTXO model before following outputs. Bitcoin transactions do not work like bank transfers. Each transaction has inputs (source UTXOs) and outputs (destination addresses). A single transaction can consolidate funds from multiple source addresses and distribute to multiple destinations simultaneously. One of those output addresses is typically the **change address** — the sender's own address receiving the unspent remainder. Misidentifying the change address as a third party is the most common error in Bitcoin tracing. Indicators: the change address is usually newly created (no prior transaction history), and its amount is typically an irregular number. The intended recipient's amount tends to be round.

Step 6 — Use mempool.space to assess pending transactions. Transactions broadcast to the Bitcoin network wait in the mempool before confirmation. mempool.space shows unconfirmed transactions in real time. A transaction with an abnormally low fee may stall for hours — this is detectable before it confirms. Conversely, a fee far above market rate signals urgency; in fraud contexts, that's a behavioral indicator worth documenting.

Step 7 — Export and structure the data for analysis. Etherscan provides CSV export of full transaction history from the "Download CSV Export" link below the transaction table. This file can be analyzed directly or passed to an AI system (Claude, ChatGPT) for pattern extraction: unique counterparties, recurring amounts, temporal clustering, fund concentration or dispersion across destinations. The AI processes what you give it — it does not query the explorer independently. Feed it the exported data, not the address.

Critical Issues

Pseudonymity is not a minor limitation — it is the ceiling of unaided blockchain analysis. A complete transaction trail, fully documented, still does not name the person behind the wallet. Attribution requires convergence with external data: KYC records from centralized exchanges (accessible via legal process), open-source correlations between addresses and platform activity, or behavioral cross-referencing across platforms.

Address labels in Etherscan are community-contributed, not verified. A label reading "Binance Hot Wallet" or "Known Scammer" reflects collective tagging, not official classification. It is an investigative signal, not evidence.

The change address error is structurally predictable. Any investigator new to Bitcoin tracing will follow the wrong output at least once without understanding the UTXO model. Document the reasoning behind each output selection explicitly.

AI-assisted pattern analysis amplifies speed but not certainty. Giving a transaction CSV to an AI and asking it to identify mixer behavior or fraud patterns produces hypotheses, not conclusions. The AI model does not have access to current blockchain data; it analyzes the data you provide. Its output belongs in the analytical layer of a report, not in the findings.

Privacy-focused cryptocurrencies break this method entirely. Monero and Zcash in shielded mode obscure senders, recipients, and amounts by design. Explorers exist but provide only partial information. If a case involves these assets, the tracing methodology above does not apply.

Analytical Layer

The structural pattern in crypto fraud — visible repeatedly across the transaction architecture described here — is consolidation followed by rapid conversion. Victim funds arrive at a collection

address in variable amounts over days or weeks. The collection address shows no outgoing activity during that period. Then, in a single transaction or compressed sequence, the balance moves to a second address where it converts to a stablecoin via a DEX interaction (visible in Etherscan's Internal Txns or ERC-20 tabs). The stablecoin then routes to a centralized exchange address. That routing to a centralized exchange is the only structural point in the chain where legal process can produce identity — because centralized exchanges hold KYC data on their account holders.

The dependence on that single convergence point — centralized exchange attribution — means that any fraud operation routing funds exclusively through DEXs and self-custodied wallets remains pseudonymous. The transaction history is complete; the identity remains absent. That gap is not a failure of the explorer. It reflects the architecture of the system.

The investigative value of blockchain tracing lies not in identifying perpetrators directly, but in two more limited and more reliable outputs: establishing that a specific address received funds from multiple victims (pattern documentation), and identifying whether those funds reached a regulated entity capable of producing identity data under legal compulsion. The explorer provides both. The rest is legal process.
A lawyer sends you two wallet addresses — one Ethereum, one Bitcoin — and a single sentence: "I need to know where the money went." No names, no context, no additional data. Just two strings of alphanumeric characters and the expectation that you can reconstruct what happened.

This is not a hypothetical. It's the standard entry point for cryptocurrency fraud investigations. And the first decision — which tool to open, which field to read, which pattern signals a problem — determines whether the next two hours produce usable intelligence or a dead end.

The Public Ledger Problem

Blockchain data is radically public. Every transaction on Bitcoin, Ethereum, and most major cryptocurrencies is permanently recorded and accessible to anyone with an internet connection, without registration, without authorization, and without cost. No journalist, investigator, or analyst needs a court order to query a wallet's complete transaction history. That transparency is the foundational property that makes blockchain OSINT both powerful and frequently misread.

The misunderstanding cuts both ways. Investigators who don't work with crypto often assume pseudonymity means anonymity — that blockchain data is opaque. The opposite error is equally common: treating a wallet address as a confirmed identity when it's still a pseudonym. A blockchain explorer shows every movement, every timestamp, every counterparty address. It does not show who controls those addresses. That gap — between transaction data and attributed identity — is where most blockchain investigations stall, and where most methodological errors accumulate.

Block explorers are the interface layer between raw blockchain data and human-readable investigation. They function like search engines pointed at the ledger: input a wallet address, a transaction hash, or a block number, and the tool returns structured data about that entry. Three explorers cover the operational core: **Etherscan** (etherscan.io) for Ethereum and compatible chains, **Blockchain.com** (blockchain.com/explorer) for Bitcoin, and **mempool.space** for Bitcoin with real-time mempool visibility. These are not interchangeable. Each is chain-specific. Querying a Bitcoin address in Etherscan returns nothing — the ledgers are separate infrastructures.

Identifying which explorer to use requires only one check: if the address begins with 0x, it's Ethereum — use Etherscan. If it begins with 1, 3, or bc1, it's Bitcoin — use Blockchain.com or mempool.space.

System Map: Where the Data Lives and What It Cannot Tell You

Blockchain data is distributed across nodes, but practically accessed through a small number of high-reliability explorers. The information architecture has distinct layers.

What is **fully visible**: complete transaction history of any public address, timestamps in UTC, amounts transferred, transaction fees paid (gas, in Ethereum's terminology), the status of each

transaction (confirmed, pending, failed), and — on Ethereum — interactions with smart contracts, internal transactions generated by those contracts, and token transfers (ERC-20). On Bitcoin, the UTXO model makes inputs and outputs explicit: every transaction lists the source addresses and destination addresses with exact amounts.

What is **partially visible**: address labels. Platforms like Etherscan crowd-source and maintain tags for known entities — exchanges, DeFi protocols, mixers, flagged fraud addresses. These labels are community-generated and not authoritative, but they provide investigative leads. A destination address tagged as belonging to a centralized exchange is structurally different from an untagged address.

What is **missing**: the identity of wallet controllers. Blockchain addresses are pseudonymous. The ledger records the address, not the person. Linking an address to a real identity requires cross-referencing with external sources — exchange registration data, court records, data leaks, or behavioral patterns that connect an address to other identifiable activity. No explorer provides that link directly.

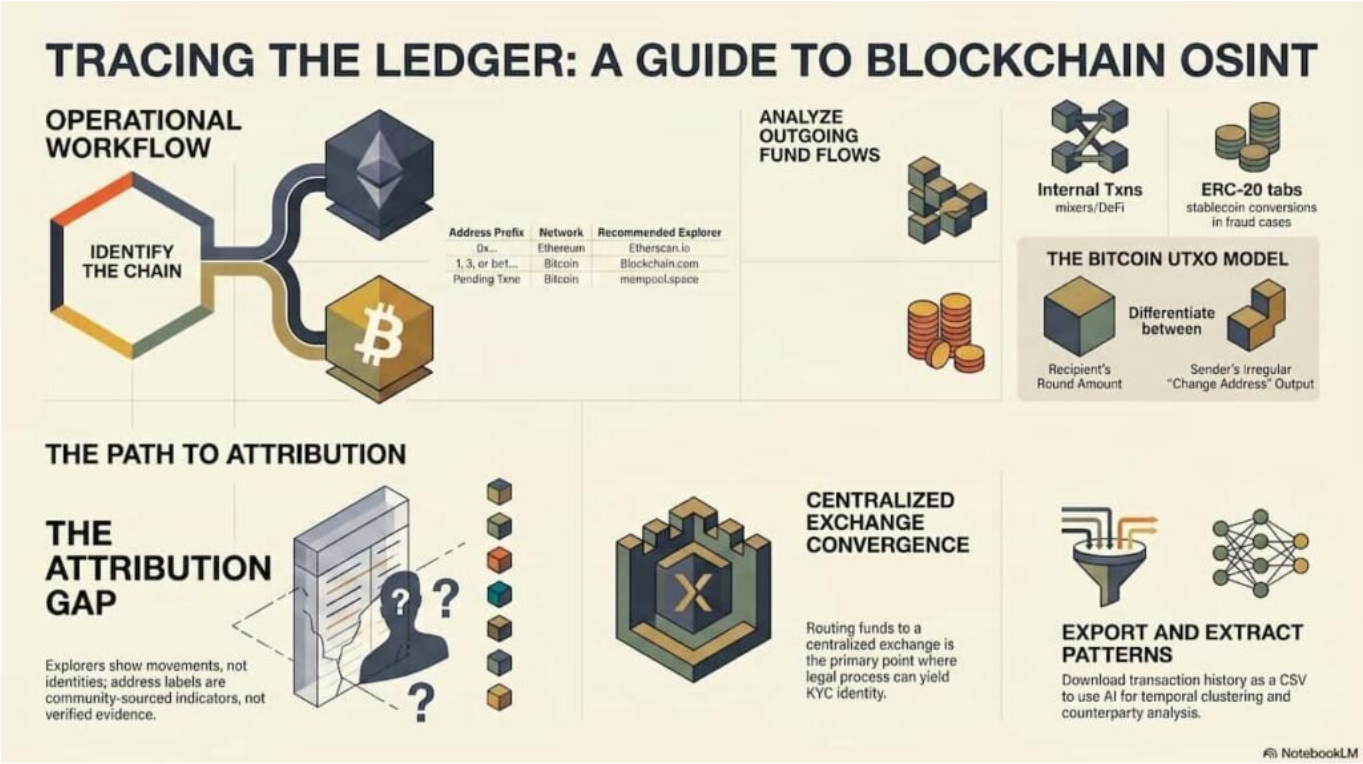

Operational Method

Step 1 — Identify the chain and open the correct explorer. Determine whether the address is Ethereum (0x…) or Bitcoin (1, 3, or bc1). Open etherscan.io or blockchain.com/explorer accordingly. No account required for either.

Step 2 — Read the wallet summary. On Etherscan, the overview shows current ETH balance, total transaction count, and token holdings. On Blockchain.com, the summary displays Total Received, Total Sent, and Final Balance. A Final Balance of zero means funds have already moved. A high transaction count — thousands of operations — is inconsistent with a personal wallet and suggests an exchange, service, or aggregation address.

Step 3 — Analyze outgoing transactions to map fund flow. On Etherscan, the Transactions tab shows ETH movements. The Internal Txns tab shows movements generated by smart contract interactions — this is where mixer activity and DeFi protocol usage appear, not in the main tab. The ERC-20 Token Txns tab captures stablecoin transfers (USDT, USDC, DAI). Stablecoins are operationally significant in fraud cases: their dollar-pegged value makes them the preferred vehicle for consolidating proceeds before moving to an exchange.

Step 4 — On Ethereum, read the transaction-level fields that carry investigative weight. Status (Success vs. Failed) matters: failed transactions remain permanently recorded. A pattern of failed transactions may indicate probing behavior. Transaction Fee (gas) signals urgency — an abnormally high fee on an outgoing transaction suggests the sender prioritized speed. The To field shows whether the destination is a standard address or a smart contract (Contract). The Input Data field, decoded via Etherscan's built-in decoder, reveals the function called — token swap, spend approval, NFT transfer. This distinguishes a direct transfer from obfuscated layering through DeFi.

Step 5 — On Bitcoin, account for the UTXO model before following outputs. Bitcoin transactions do not work like bank transfers. Each transaction has inputs (source UTXOs) and outputs (destination addresses). A single transaction can consolidate funds from multiple source addresses and distribute to multiple destinations simultaneously. One of those output addresses is typically the **change address** — the sender's own address receiving the unspent remainder. Misidentifying the change address as a third party is the most common error in Bitcoin tracing. Indicators: the change address is usually newly created (no prior transaction history), and its amount is typically an irregular number. The intended recipient's amount tends to be round.

Step 6 — Use mempool.space to assess pending transactions. Transactions broadcast to the Bitcoin network wait in the mempool before confirmation. mempool.space shows unconfirmed transactions in real time. A transaction with an abnormally low fee may stall for hours — this is detectable before it confirms. Conversely, a fee far above market rate signals urgency; in fraud contexts, that's a behavioral indicator worth documenting.

Step 7 — Export and structure the data for analysis. Etherscan provides CSV export of full transaction history from the "Download CSV Export" link below the transaction table. This file can be analyzed directly or passed to an AI system (Claude, ChatGPT) for pattern extraction: unique counterparties, recurring amounts, temporal clustering, fund concentration or dispersion across destinations. The AI processes what you give it — it does not query the explorer independently. Feed it the exported data, not the address.

Critical Issues

Pseudonymity is not a minor limitation — it is the ceiling of unaided blockchain analysis. A complete transaction trail, fully documented, still does not name the person behind the wallet. Attribution requires convergence with external data: KYC records from centralized exchanges (accessible via legal process), open-source correlations between addresses and platform activity, or behavioral cross-referencing across platforms.

Address labels in Etherscan are community-contributed, not verified. A label reading "Binance Hot Wallet" or "Known Scammer" reflects collective tagging, not official classification. It is an investigative signal, not evidence.

The change address error is structurally predictable. Any investigator new to Bitcoin tracing will follow the wrong output at least once without understanding the UTXO model. Document the reasoning behind each output selection explicitly.

AI-assisted pattern analysis amplifies speed but not certainty. Giving a transaction CSV to an AI and asking it to identify mixer behavior or fraud patterns produces hypotheses, not conclusions. The AI model does not have access to current blockchain data; it analyzes the data you provide. Its output belongs in the analytical layer of a report, not in the findings.

Privacy-focused cryptocurrencies break this method entirely. Monero and Zcash in shielded mode obscure senders, recipients, and amounts by design. Explorers exist but provide only partial information. If a case involves these assets, the tracing methodology above does not apply.

Analytical Layer

The structural pattern in crypto fraud — visible repeatedly across the transaction architecture described here — is consolidation followed by rapid conversion. Victim funds arrive at a collection

address in variable amounts over days or weeks. The collection address shows no outgoing activity during that period. Then, in a single transaction or compressed sequence, the balance moves to a second address where it converts to a stablecoin via a DEX interaction (visible in Etherscan's Internal Txns or ERC-20 tabs). The stablecoin then routes to a centralized exchange address. That routing to a centralized exchange is the only structural point in the chain where legal process can produce identity — because centralized exchanges hold KYC data on their account holders.

The dependence on that single convergence point — centralized exchange attribution — means that any fraud operation routing funds exclusively through DEXs and self-custodied wallets remains pseudonymous. The transaction history is complete; the identity remains absent. That gap is not a failure of the explorer. It reflects the architecture of the system.

The investigative value of blockchain tracing lies not in identifying perpetrators directly, but in two more limited and more reliable outputs: establishing that a specific address received funds from multiple victims (pattern documentation), and identifying whether those funds reached a regulated entity capable of producing identity data under legal compulsion. The explorer provides both. The rest is legal process.