DigitalOcean

Unlock the Power of Blockchain

An introduction to blockchain and its cloud computing needs

Unlock the Power of Blockchain

An introduction to blockchain and its cloud computing needs

Blockchain technology has been a disruptive force in the era of digital transformation. According to Statista Research, worldwide spending on blockchain solutions is expected to grow to $19 billion USD by 2024.

Blockchain works as a distributed ledger technology, meaning it uses independent computers to record and share activities. This decentralizes the record of transactions and makes data stored using blockchain more secure and transactions more transparent than those housed on a single server or controlled by a single entity.

While blockchain has gained notoriety through cryptocurrency, the technology is used by many businesses across industries. Any type of data can be stored using blockchain, and fields from finance to healthcare, and retail operations to supply chain management all benefit from the security, transparency, and efficiencies blockchain technology provides.

In a recent TechRepublic Research study, 64% of professionals said that they expect blockchain to affect their industry in some way, with most expecting the impact to be positive. Blockchain has the potential to completely change how people do business, and innovative startups are taking notice. In a recent survey, Deloitte found that 90% of emerging blockchain disruptors are actively hiring blockchain talent. They learned that these founders aren't approaching blockchain as incremental change, but are interested in fundamentally changing how business is conducted.

The opportunities are endless. Organizations are using blockchain to actualize ideas like digital verification systems that ensure creators are properly paid for their work, systems for safer and more efficient data transfer, and technologies supporting fast and secure digital identity verification. With these visions, startups and SMBs around the globe are prepared to usher in a new era of digital innovation. As more organizations seek to take advantage of the benefits that blockchain provides, and still others begin to create blockchain-based services, it's important to consider the computing and infrastructure needs of these potentially world-changing businesses. In this guide, we'll explore the fundamentals of blockchain for business and what founders should keep in mind as they build their businesses.

Blockchain words to know

Address (Wallet Address)	Used to send and receive transactions on a blockchain network. An address is an alphanumeric character string, which can also be represented as a scannable QR code.
Consensus mechanism	The process by which a group of peers is responsible for maintaining distributed ledger use. Proof of Work and Proof of Stake are examples of consensus mechanisms.
Decentralized application (dapp)	An open-source software application with backend code running on a decentralized peer-to-peer network rather than a centralized server. Also referred to as Web3.
DeFi (Decentralized Finance)	Refers to cryptocurrency and economic activity occurring on a blockchain.
Genesis block	The first block in the blockchain.
Mainnet	The primary network where actual transactions take place on a specific distributed ledger. For example, The Ethereum mainnet is the public blockchain where network validation and transactions take place. Contrast this with a sidechain (below).
Node	Any computer connected to the blockchain network is referred to as a node. A full node is a computer that can fully validate transactions and download the entire data of a specific blockchain. A lightweight or light node does not download all pieces of a blockchain's data and uses a different validation process, relying on the backplane nodes for full chain integrity. At DigitalOcean, we recommend CPU Optimized Droplets for full nodes, but Basic Droplets are typically sufficient for light nodes.

Non-Fungible Token (NFT)

Fungibility refers to an object's ability to be exchanged for another. For example, an individual dollar is considered fungible as we can trade dollars with one another. Artwork is usually deemed non-fungible because it's likely to be unequal in quality or value. A non-fungible token is a type of token that is a unique digital asset and has no equal token.

Shard

Sharding refers to splitting the entire network into multiple portions called "shards." Each shard would contain its own independent state, meaning a unique set of account balances and smart contracts. Usually, shards must be tightly coupled and sidechains must be loosely coupled. This is very similar to database sharding.

Sidechain

A chain that runs in parallel to a mainnet, but only interacts with the mainnet for validation or other activities. Sidechains have a two-way bridge to a mainnet. Sidechains are not required to use the same proof as the mainnet.

Smart contracts

Smart contracts are programs whose terms are recorded in a computer language instead of legal language. Smart contracts are automated actions that can be coded and executed once a set of conditions is met. Haskell is a common language for this.

51% attack

If more than half the computer power or mining hash rate on a network is run by a single person or a single group of people, then a 51% attack is in operation. This means that this entity has full control of the network and can negatively affect a cryptocurrency by taking over mining operations, stopping or changing transactions, and double-spending coins.

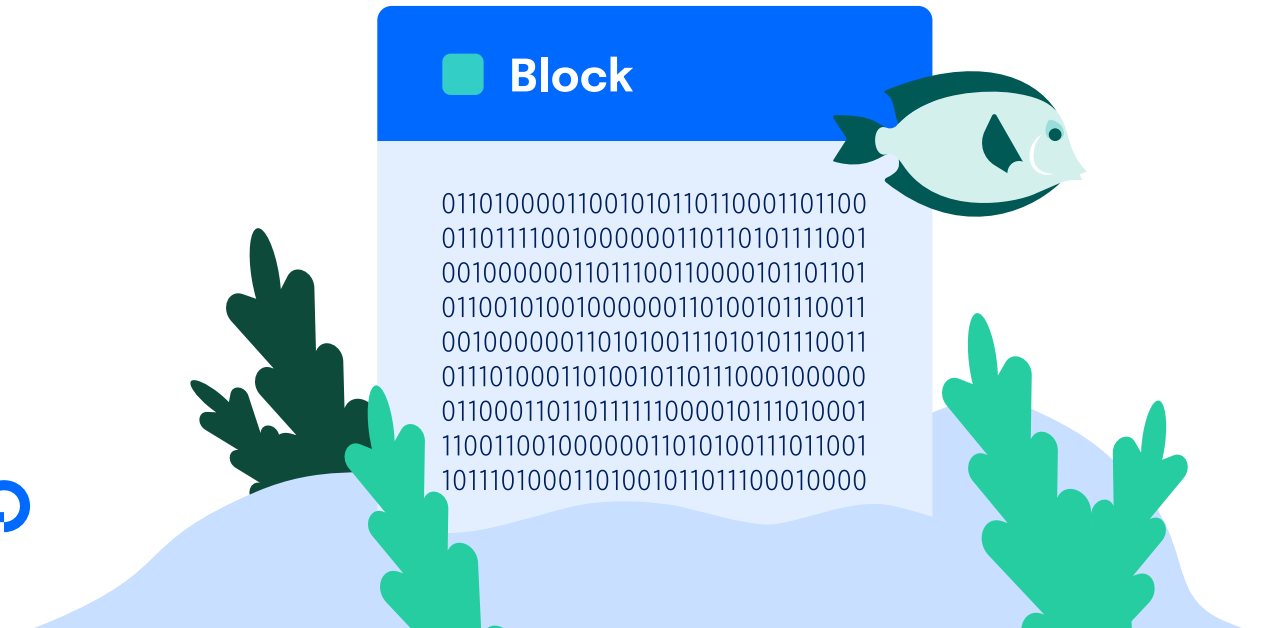

Blockchain vs. cryptocurrency

The most well-known use for blockchain technology is cryptocurrency, and people often use the terms synonymously. Blockchain and cryptocurrency, however, are not the same. Blockchain is the technology that enables cryptocurrency, but that's far from its only use. Cryptocurrencies use a distributed ledger but are based on a digital store of value, and primarily exist as a source of electronic currency. Cryptocurrencies, like Bitcoin, rely on distributed ledger technology to keep anonymity. While anyone can look at the ledger and see that a transaction took place, no one will know who or what was involved. Alternatively, businesses who use blockchain keep careful internal records allowing for the transparency of transactions through the distributed ledger while keeping the individual data secure with the block's encryption.

Understanding the basics of blockchain

At its core, blockchain enables record-keeping in a secure, immutable way. Anything can be tracked using blockchain technology, from tangible resources to intangible pieces of information. In blockchain, the data is encrypted within blocks. Each block contains the data itself, a hash that's unique to the block, and the hash of the previous block. New blocks are created by generating a new hash. These new hashes are validated through consensus mechanisms, often called proofs. When validated blocks are accepted to the network, the block is added to the blockchain.

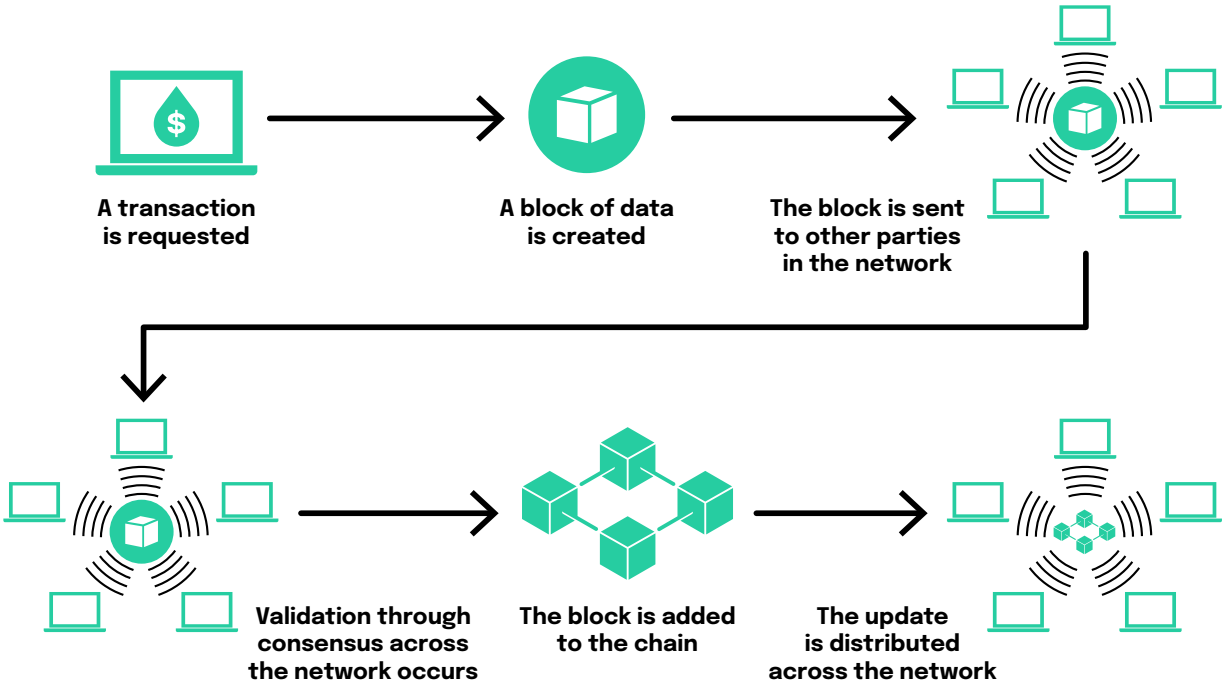

If anyone tampers with the data in one block, the hash changes. Thus, all the hashes in subsequent blocks are rendered incorrect. In order to make a change to data within a blockchain, someone would have to make the change, creating a new hash, then create new hashes and pass the proof mechanisms of the network for each subsequent block in the chain. Since the validation mechanisms are distributed across multiple computers and owned by different individuals, it's nearly impossible to tamper with the data in a blockchain.

Choosing a consensus mechanism

Consensus mechanisms allow distributed networks to stay secure by requiring a general agreement to the proposed change or addition in the system. There are a few different ways consensus is reached in order to allow a block to be added to the network. The core tenet of proofs is that they should be hard to produce, but easy to validate. The most common types of proofs are Proof of Work and Proof of Stake. There are pros and cons to each consensus mechanism, depending on the goals of the network and the resources available.

Proof of Work

Proof of Work is the consensus mechanism most people are familiar with. Many well-known cryptocurrencies, like Bitcoin, use Proof of Work as their means of validation. Proof of Work requires individuals to solve cryptographic puzzles to produce blocks. Creating new blocks (often called mining) requires powerful hardware that is able to solve the complex math problem needed to create the new hash for the block, and computers compete against each other to be the first to find the answer to the puzzle. Because of the complexity required to create each block, Proof of Work validation methods are exceptionally secure. Proof of Work is extremely energy-intensive and is the primary source of concern regarding the environmental impact of cryptocurrency. Some cloud providers like DigitalOcean do not allow Proof of Work on their platforms.

Proof of Stake

Proof of Stake is the consensus mechanism many major blockchains are moving toward. With Proof of Stake mechanisms, validators are required to “stake” coins or tokens from the blockchain network in a staking pool, locking them for a period of time. Staked coins or tokens never leave your wallet, but they can’t be transacted with until the stake is withdrawn. Exchanges will never ask you to withdraw tokens in order to stake. After coins or tokens are staked, the network uses a randomized process to choose a validator to produce the next block. Proof of Stake uses much less energy to secure the blockchain and is much faster than Proof of Work.

When choosing a consensus mechanism for blockchain, it’s important to consider the goals of your network. Consensus mechanisms should be resilient, high-performing, and efficient. Nodes should easily be able to follow the rules and participate as needed.

“Elrond is the first Proof of Stake network to implement all three types of sharding: state, network, and transactions. Maintaining this architecture in a high throughput environment of up to 15,000 transactions per second requires dependable network connectivity and dynamic provisioning capabilities, which DigitalOcean reliably provides to support us on our mission.”

Lucian Mincu, Elrond Network CIO.

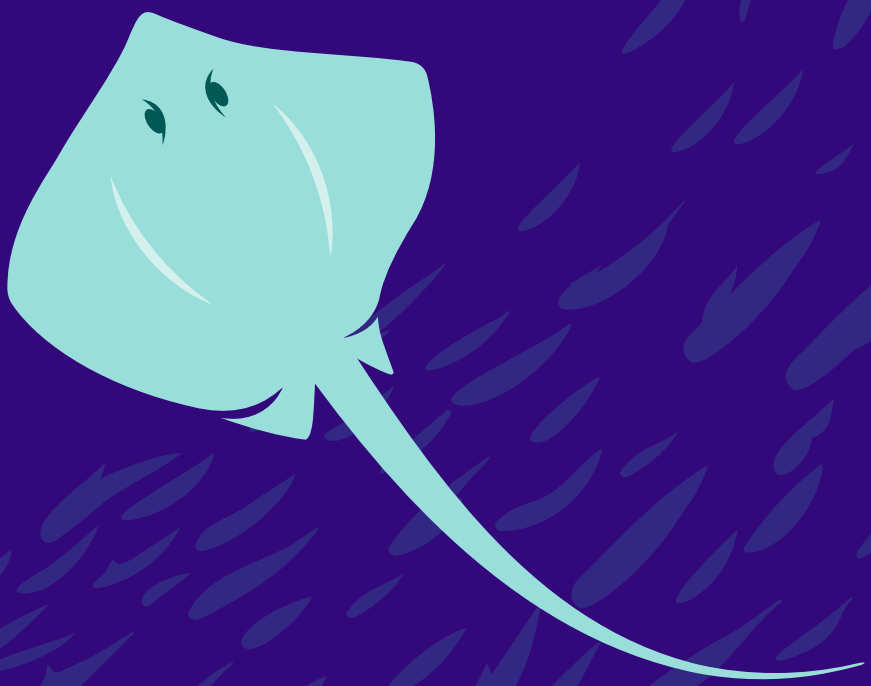

Choosing nodes for blockchain applications

Blockchain requires coordinated activity from multiple computers. Any computer connected to the blockchain network is referred to as a node. **Nodes are the framework of a blockchain, storing, spreading, and preserving the data in the chain.** There are different types of nodes, each having a different job in the validation process and each with unique computing power requirements. Two of the most common types of nodes are full nodes and light nodes.

Full nodes

- Fully validate transactions
- Download the entire history of the data
- Requires significant computing power
- Can't tolerate hypervisor steal or disk latency

VS.

Light nodes

- Rely on full nodes for full chain integrity
- Only download and validate the most recent transactions
- Faster than full nodes
- Require less computing power

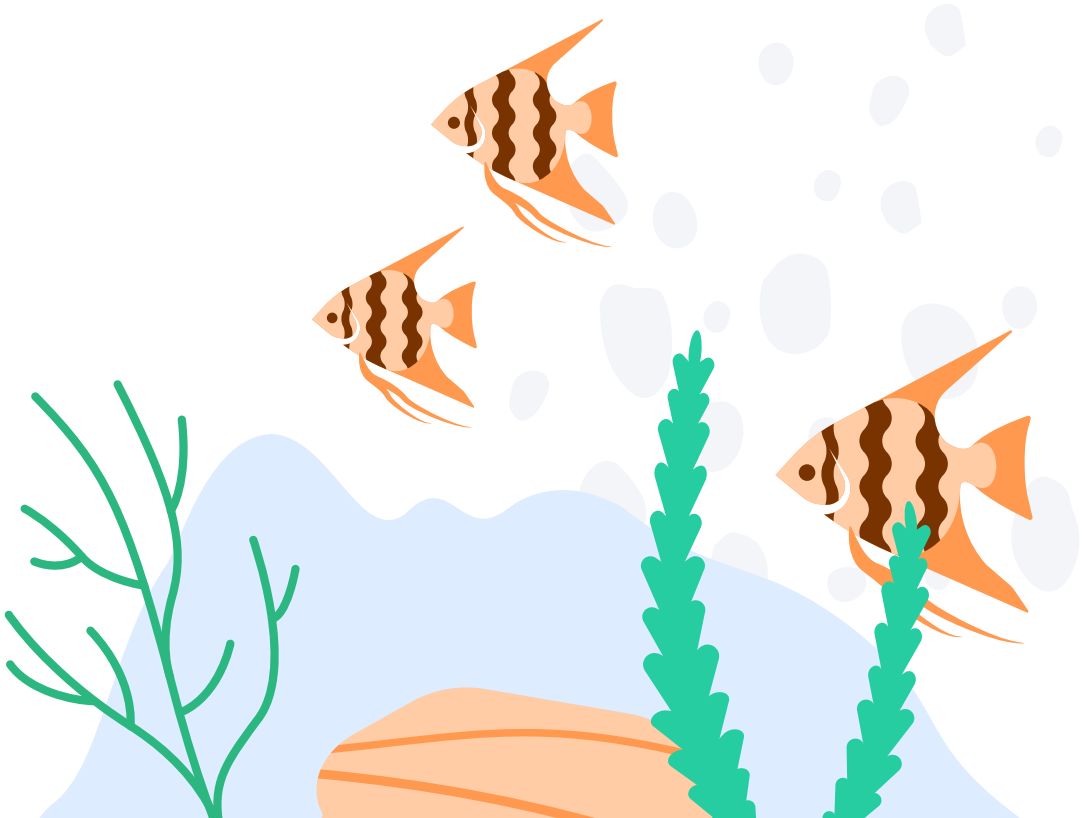

Full nodes fully validate transactions and download the data of a specific blockchain in its entirety. Full nodes support and provide security to the network. They download the entire history in order to observe and enforce the rules. Since they contain the transaction and data history in its entirety, full nodes also provide increased security to the networks, because each one is capable of carrying on the network itself. In order to lose the data, all full nodes would have to be destroyed—an unlikely scenario in a ledger distributed across multiple computers. Full nodes require a significant amount of computing power and since they are performing both validation and consensus work, they are extremely latency-sensitive both in terms of disk and network. They can't tolerate hypervisor steal or disk latency.

Full nodes can be "pruned" to save space, essentially taking out some of the older blocks to reduce the disk usage required, but then they require archival nodes—nodes that have the complete historical data of the blockchain—to support them. This typically won't reduce the other hardware requirements; they still need significant processing power and RAM to perform the work.

Lightweight or light nodes don't download the blockchain's data in its entirety. Light nodes rely on full nodes for full chain integrity and have to connect to the full node to be able to participate in the validation process. Light nodes download and validate the most recent transactions, making them faster than full nodes and able to run with less computing power, meaning that they need the full nodes to access the history of the data in its entirety. Light nodes have lighter hardware requirements and looser timing requirements and offer significant resource savings.

The core computing requirements depend on the blockchain, but in general full nodes take more CPU than light nodes. Full nodes are more latency-sensitive than light nodes due to their increased processing requirements. If users interact directly with full nodes, the interactions are often cumbersome and inefficient. At DigitalOcean, we recommend running full nodes on dedicated threads.

Full nodes and light nodes can work together, and many proof of work chains take advantage of the strengths of each type of node. In proof of stake chains, full nodes are primarily used for validation and consensus, although the chains are working toward using a full node and light node standard similar to proof of work chains.

Example infrastructure

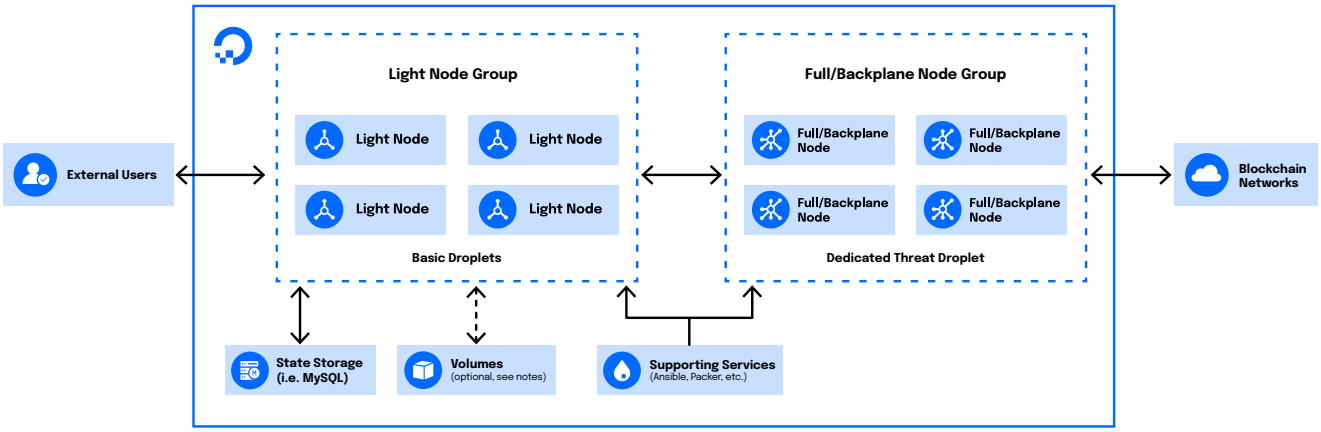

One way to set up the backend infrastructure for a blockchain network includes support from full and light nodes. In the above diagram, the users interact with light nodes and light nodes store the most recent blocks and then sync with the full nodes as transactions occur.

A few things to keep in mind:

- Light nodes typically don't have dedicated thread requirements, but full nodes do.
- Full nodes typically don't require any additional state support as they contain entire copies of the chain and have extremely sensitive disk latency requirements. If you're running a full node and run out of local disk, you must upsize to the next slug.
- Supporting services will typically be used for both groups from the same place.

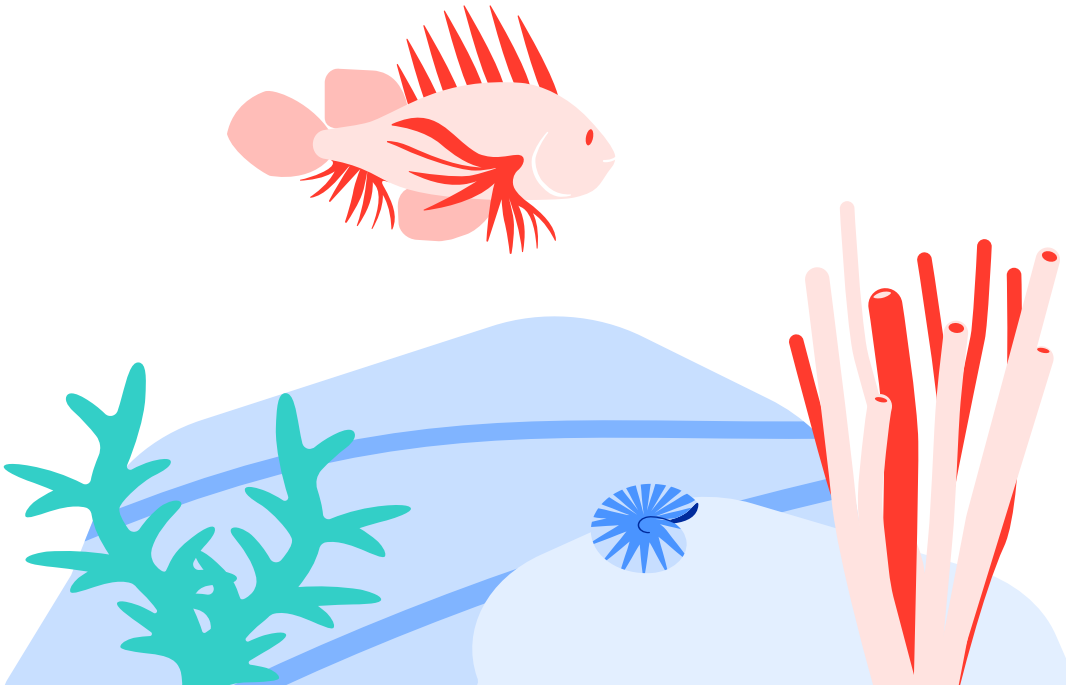

Choosing a cloud provider

Blockchain is hard. It's an extremely complicated field of tech that's constantly evolving and changing. There's no reason to add even more complexity with your cloud provider. When researching a cloud provider for your blockchain business, consider the following:

CPU considerations

Blockchain requires a significant amount of ingress due to chain download and updates and does require data to be egressed to the rest of the network. Take a close look at bandwidth models, benefits, and pricing of potential providers. For example, DigitalOcean provides free ingress and bandwidth pooling.

Uptime considerations

Blockchain networks should run well regardless of what provider they are on. Decentralized applications may run anywhere, but that doesn't mean they'll run optimally anywhere. Research the cloud agnosticism of potential providers. Find out if there's lock-in with long contracts or how well they support cloud native computing. Blockchains need extremely reliable services, so finding out about the reliability of any provider you're considering is important.

Scalability considerations

Chains also naturally grow over time. It would be significantly disruptive to destroy and create servers in order to meet chain growth. Choose a cloud provider that enables you to resize your servers. For example, DigitalOcean allows developers to spin Droplets up and down with the click of a button. Since blockchain services frequently need to scale up in response to demand, consider choosing a cloud partner that provides a transparent pricing model, enabling you to understand what the costs will be for various levels of demand.

> **"We started with DigitalOcean and we grew with DigitalOcean. We were pretty ecstatic when we were doing three and a half billion requests a month, which translates to 3,500 requests a second, and now we're doing 10X that, so DigitalOcean has helped us scale throughout our journey."** – QuickNode

Ease of use considerations

Find a cloud provider that enables you to move quickly and efficiently. Blockchain can be extraordinarily complex. A cloud provider that is equally complex can hinder administration and growth. For example, getting an environment up and running in a hyperscaler typically takes a significant amount of time due to configuration accounts, IAM, and more. To move quickly, consider a cloud platform that provides simple, easy-to-use solutions and robust documentation and support. Providers that enable automation via an API and third-party tools like Terraform make managing workloads much easier.

> **"DigitalOcean has a great workflow that enables us to deploy blockchain nodes quicker and easier, which are key factors that have allowed us to grow as large and as fast as we have to serve more customers."** – Blockspaces

Blockchain businesses building on DigitalOcean

DigitalOcean's simple, low-cost Droplet virtual machines have proven to be an excellent match for the needs of blockchain businesses. Companies including Blockspaces, QuickNode, and Elrond leverage DigitalOcean's cloud services to run their blockchain networks and serve thousands of users around the globe.

QuikNode

QuickNode enables developers to integrate with multiple large blockchain networks, including Bitcoin, Ethereum, and Solana, through elastic APIs and dedicated nodes. They provide customers with in-depth analytics and tools that enable businesses to easily create and scale blockchain applications without needing to worry about the blockchain infrastructure. QuickNode recently raised a $35 million Series A funding round and has been a DigitalOcean user since its inception.

elrond

The Elrond blockchain is the second-largest Proof of Stake network in the world after Ethereum 2.0, with 3,200 validator nodes supporting its geographically dispersed main network. When counting the main and backup nodes, as well as several public test networks and auxiliary systems, the Elrond infrastructure spans close to 10,000 servers.

BLOCKSPACES

BlockSpaces offers an integration platform that supports connectivity between business applications and blockchain networks through managed infrastructure, no/low code workflows, and robust performance analytics. The BlockSpaces platform is specially designed for companies that need to deploy blockchain solutions without dealing with the complexities involved in integrating the technology into existing systems. The company has seen 257.2% quarter over quarter growth and 1,678% year over year growth since raising a second seed round of capital in late 2021. BlockSpaces leverages DigitalOcean Droplets, our low-cost virtual machines, and our simple-to-use API and CLI to support their growth.

DigitalOcean

About DigitalOcean

DigitalOcean simplifies cloud computing so developers and businesses can spend more time building software that changes the world. With its mission-critical infrastructure and fully managed offerings, DigitalOcean helps developers, startups, and small- and medium-sized businesses (SMBs) rapidly build, deploy, and scale applications to accelerate innovation and increase productivity and agility. DigitalOcean combines the power of simplicity, community, open source, and customer support so customers can spend less time managing their infrastructure and more time building innovative applications that drive business growth.

> To get started, **sign up for an account at DigitalOcean.com.** For more information or help migrating your infrastructure to DigitalOcean, **speak to a sales representative.**