

The Three-Body Problem

How conventional IT, private ledgers, and public blockchains will work together to usher in the Web 3.0 era

Table of Contents

Table of Contents	2
Who Should Read This	3
Executive Summary	3
Introduction	4
Centralized Solutions – Classic IT Always Wins on Bandwidth and Cost	6
Private Chains – Best for Partner Data Sharing and the Environment	7
Public Chains – Best for Archiving Truly Public Data	8
Example: Building an Image Licensing Business	11
Existing (“As built”) Architecture Using a Centralized Approach	12
Next-gen Architecture Requirements	13
Are NFTs the Answer?	14
Better Together: Using all Three Approaches	15
Integrating Central IT with Private and Public Blockchains: The Real-time Data Mesh Pattern	16
Conclusions	19
Adoption Caveats	19
Going Further	20
About the Author	20

The Three-Body Problem

How conventional IT, private ledgers, and public blockchains will work together to usher in the Web 3.0 era

Who Should Read This

CIOs, CTOs, LOB leaders, cloud architects, professional developers, and investors looking to understand the relationship between public chains, private chains, cloud services, and conventional IT architectures will learn the relative strengths and weaknesses of each of these categories, and how they can be combined to provide optimal business solutions and technical outcomes.

This material is of interest to anyone contemplating or budgeting for one or more blockchain projects as well as architects faced with the challenge of building cost-effective partner data sharing solutions for businesses of any size. Readers interested in the meaning of “web3” and how cloud, blockchains, and emerging technologies will interact in the next generation of IT technology stacks will also benefit from the content.

Executive Summary

Conventional (aka centralized) application approaches, which are increasingly based on cloud services, are often portrayed as antithetical to both private and public blockchains. In reality, **all three categories are likely to emerge as key elements of modern IT solutions** because each possesses unique, durable advantages over the other two.

Emerging best-of-breed practices employ these technologies for their differentiated benefits:

- Conventional (aka “centralized”) IT solutions, increasingly built in the cloud and often on top of managed cloud services, will continue to offer the highest throughput, lowest latency, and easiest governance indefinitely. These will continue to form 80% or more of the foundational layer for SMB, mid-market, and enterprise IT.
- Private/permissioned blockchains have the advantage of innovating quickly, and can integrate seamlessly with the cloud solutions of their controlled populations, making them ideal for creating partner buses that enable businesses to share data with rigorous compliance, privacy, and governance capabilities. Conventional (“dumb pipe”) APIs will be replaced over time with distributed ledgers that offer an automatically maintained and cryptographically provable single source of truth for

cross-company and cross-org use cases such as supply chains, financial ledgers, and media/file sharing. Private chains will grow to represent 15-20% of IT stacks.

- Public blockchains offer very limited throughput due to the intrinsic requirements of Byzantine Attack resistance, planet-wide data replication, and “single box” deployment models that intentionally forgo public cloud leverage. Public chains will continue to exhibit high latency, limited data type support, and modest innovation rates (due to their open source/consortium nature). However, they will play a key role for applications that require public disclosure, non-repudiation and/or total ordering among unrelated parties, or where public data markets exist, serving key use cases such as the archival storage for NFT ownership, self-service ticket transfers among consumers, and third party auditing of (hashed, encrypted digests) of key financial and business data. Public chain integration will eventually comprise 1-2% of typical enterprise IT stacks over the long term.

Introduction

New, and especially disruptive, technologies often position themselves as antagonistic to what has come before. In the early days of the public cloud, for example, it was common to hear cloud service providers assuring would-be customers that on-premise deployments were soon to be 100% eliminated, and that only fools would consider anything but a pure-play cloud architecture. As it turned out, public cloud adoption was a more incremental adoption pattern spanning decades, and even the most “all in” enterprise customers often still pursue a hybrid strategy that includes both cloud and on-premise elements.

Blockchain advocates took a similar stance when “their” technology emerged: Not only were the “private” and “public” chain camps at war with one another for supremacy, but both positioned themselves as a one-size-fits-all replacement for the entire IT stack. Of course, with some time and perspective, it became obvious that neither public nor private chains are a true substitute for the other, and that neither is a replacement for conventional (aka “centralized”) IT architectures.

What *has* emerged is a playbook for how all three technology categories – centralized cloud and on-prem solutions, public blockchains such as Ethereum, and private/permissioned chains will interoperate to create best-of-breed IT outcomes over the next 20 years.

To understand these new design patterns better, it is helpful to begin with a clear understanding of the strengths and weaknesses of each category. Figure 1 summarizes these tradeoffs visually.




Figure 1: The continuum of capabilities. Centralized IT implementations will always offer the highest throughput, lowest cost, and easiest governance for high-bandwidth data. Conversely, public chains offer the broadest distribution model for limited amounts of data, but at the highest cost of transfer and storage. Between the two, private chains optimize for sharing large amounts of data in a reliable way with a fixed population – ideal for building partner data sharing solutions.

Blockchain advocates sometimes speak as if these distinctions will be swept away by “upcoming optimizations”. But the fact of the matter is that *this continuum is dictated by physics*: Copying data to more places, further away from the starting point, costs more time, takes more infrastructure resources, and therefore costs more money. Incenting others has a marginal cost over doing something oneself. To put it bluntly: Changing a row in 1,000 databases has to cost at least 1,000 times more than changing a row in 1 database. However those costs may be presented or apportioned, *there’s no way around the fundamental truth that, on a per-transaction basis, public chains cost dramatically more than a single, centralized resource holding the equivalent information*. This problem is exacerbated by the fact that not all information in a public chain is necessarily of interest to all of its users, while the overhead of worldwide replication is paid *regardless* of future consumption or access patterns.

Let’s take a look at each of these three categories – centralized, private chains, and public chains – and examine its unique advantages and disadvantages.

Centralized Solutions – Classic IT Always Wins on Bandwidth and Cost

Centralized solutions include virtually all existing applications and all use of existing on-premise and cloud services.¹ These solutions span the classic range of IT infrastructure: Operational databases, data warehouses for analytics, compute, file storage, and more.

Cloud-based centralized services and application designs will continue to be the majority of every company's IT portfolio for durable, structural reasons:

- **Centralized solutions will always offer the highest throughput and lowest latency and cost structure.** This is an unsurprising conclusion, because the compute, storage, and network requirements and costs are obviously minimized when an application or IT solution is “talking to itself”: With no need to replicate, achieve agreement, or transmit code or data to other geographies or parties, latency and cost are minimized, throughput and bandwidth are maximized, and costs are tightly enveloped to just what is required for the underlying storage and compute needed to perform the actual function.
- **Centralized solutions have trivial solutions for many privacy, security, and governance controls: The firewall.** A large percentage of data – and the vast majority of systems and services – within a company are never exposed to the outside world. Without intending to minimize the challenges of securing modern IT infrastructure, it is still easy to see that keeping wholesale systems inside the firewall is far easier than operating public or shared infrastructure and having to make row-by-row and API-call-by-API-call decisions regarding what to expose or hide. Centralized IT will always offer easier ways to “keep private things private” than trying to make data public while also keeping it hidden from prying eyes, as would be required when using a public chain to store data. *Complex cryptography approaches, such as zkSNARKS or homomorphic encryption, will never fully replace straightforward firewalls because they are inherently more costly and complex on a per-transaction or per-application basis.*
- **Centralized solutions offer the fastest innovation and development speed.** Again, this is obvious after a moment of reflection: If a company needs to seek permission from others – especially a worldwide consortium of protocol developers, miners, and node operators with wildly disparate needs and economic goals – to make a change, change is inherently going to be slow, incremental, and fraught with challenges. When the only person a department or company needs to consult with

¹ Existing production deployments of private chains such as Hyperledger Fabric, Ethereum-based apps, or use of cloud services explicitly designed to target blockchains, such as Amazon Managed Blockchain, excepted, of course. Today, these make up less than 0.01% of IT codebases.

is themselves, adopting a newly launched cloud service, integration with another system, or adjusting an authentication solution will always offer the most efficient decision making process possible.

- **Centralized solutions are intermediate in their environmental impacts.**

Because they typically serve only their owner and need to scale to maximum potential usage rates, centralized systems often suffer from low utilization, with 10-15% (i.e., 85-90% waste) being typical for enterprises. This makes them worse than cloud-based private chains (see below) but better than legacy public chains based on environmentally destructive Proof of Work algorithms.

Private Chains – Best for Partner Data Sharing and the Environment

Private/permissioned chains occupy an important space between centralized services and fully public chains. Because they support a *controlled* population of business partners – such as the part suppliers, logistics providers, and manufacturing partners in a conventional supply chain relationship – they make data sharing much easier than for centralized IT, where creating a “single source of truth” from the raw building blocks of cloud or on-prem infrastructure is a heavy lift. Precisely because this population *is* limited (and because private chains specialize in asymmetric sharing patterns), the challenges of trying to place “private information in a public space” that would occur with a public cloud doesn’t need to be solved, eliminating the requirement to adopt complex approaches such as zkSNARKs.

- **Private chains offer a single source of truth among business partners for operational data.** Because they can replicate large amounts of data with high throughput and low latency (unlike public chains), private chains make it easy to share operational data in real time. Smart contracts make it easy for those same partners to also share common policies, workflows, and data integrity constraints, none of which need be exposed to the public.
- **Private chains offer the most sophisticated data privacy, data protection, and data governance of all three options.** Because private chains are designed from the ground up with selective information sharing as a key feature, they generally have far more sophisticated (and well tested) data controls than “roll your own” centralized solutions built in-house *or* public chains, with their one-size-fits-all symmetric key encryption. Private chains delivered as SaaS solutions get all the conventional SaaS economic benefits of amortizing the costs of building, securing, and maintaining complex governance and data access controls over many customers, a benefit unavailable to one-off centralized solutions.

- **Private chains offer balanced innovation speed with rapid development time.** Again, it may seem odd at first blush, but private chains usually offer the fastest time to market: They require less manual buildout than a built-from-scratch centralized application, while offering more built-in capabilities (especially when offered in a SaaS deployment model) than a public chain, making them the closest to a working solution out of the box. Private chain users can innovate as fast as their joint decision making allows, without the overhead of buy-in from miners, node operators, and open source developers who could be anywhere on the planet, and who may not align with the business partnership's needs particularly well.
- **Private chains can offer the lowest infrastructure costs and environment footprint of all three categories.** It may seem odd that the "intermediate" solution is better than either of the extremes, but private chains have the unique advantage that they can be delivered as SaaS offerings. This enables them to offer fast innovation, minimal operational burdens (or infrastructure costs), and to take advantage of multi-tenanted approaches that offer significant economies of scale to users at high degrees of utilization. By contrast, centralized IT teams can rarely achieve the breadth and variety of use to achieve high utilization. Public chains, especially those that, like Ethereum, rely on Proof of Work to mint transactions, can have devastatingly large environmental impacts, and all public chains replicate data *everywhere*, regardless of consumption or access patterns, which can create a much higher carbon footprint than the highly targeted replication of a private chain.

Public Chains – Best for Archiving Truly Public Data

Not all applications or solutions require placing data into a planet-spanning public record, but when they do, *public chains* offer the broadest degree of sharing, and require the least amount of trust, to achieve this outcome.

- **Public chains are best for sharing public and permanent information.** With broad replication across geopolitical boundaries, nation states, cloud service providers, and a mix of both corporate and individual ownership, public chains such as Ethereum² effectively create a public and permanent (irreversible and irrefutable) archive. Everywhere they're stored, these facts are held in the same order, so it's also possible to establish cause and effect (i.e., whether one fact came before or after another fact), making it possible to model and audit changes in ownership, such as transfers of value or the ownership of objects, both real and virtual.

² This article focuses on the use of public chains as elements of an operational IT solution, not for purposes of currency speculation.

- **This planet-wide dispersal of information comes with both literal and figurative costs.** Ethereum, the current “gold standard” in public chains that support data and smart contracts, transaction rates are 7-8 orders of magnitude lower than similar centralized IT transactions while latencies can be 7 orders of magnitude higher, placing the price/performance ratio at up to 15 orders of magnitude worse than a typical cloud database such as Amazon DynamoDB. Storage costs for IPFS (especially when coupled to Ethereum to store hashes of files), are also orders of magnitude higher than public cloud storage, particularly when similar levels of durability and availability are required, making this a good solution only when the broadest degree of sharing is mandated by an application’s needs.
- **Public chains can offer sophisticated on-chain (“token”) currency mechanisms.** Though not intrinsically required, most public chains also feature one or more inbuilt currency mechanisms, enabling them to act as a store, transfer, and (at least in some cases) exchange of value through inter-chain / inter-token swap mechanisms. This can facilitate building cryptocurrency-based applications, though it typically does not provide useful leverage when integrating with a conventional (“fiat”) payment or billing system, such as Stripe or Plaid.
- **Public chains innovate slowly and offer minimal integration solutions.** Integration with cloud services is minimal to nonexistent because, *by design and intent, public chains eschew any form of deep integration with the public cloud.* Changes to these systems, once they are in production, are generally slow, as open source developers, miners, and operators all have to collectively agree to migrate a complex arrangement of on-chain incentives while keeping the codebase secure, available, and backwards compatible. This complex interplay of economic and technological incentive structures makes public chains slower to innovate and migrate than private ones and far slower than a centralized IT solution can adapt and evolve.
- **Support for heterogeneous data types is limited and often complex to achieve,** often requiring stitching multiple “on” and “off” chain solutions together to represent common business artifacts, such as files. Also, because they must service essentially any data model from any customer anywhere in the world, public chains are “typeless”, offering only key/value stores that require adopters to erect data models on top of them, akin to creating a complex ORM abstraction on top of an underlying database (but one with limited query and update functionality).
- **Public chains vary in their environmental footprint, but the “gold standard” (Ethereum) currently has a very poor environmental record.** Newer “Proof of

Stake” approaches improve on the cost of minting transactions, but all public chains are structurally required to spend a higher percentage of their compute (and thus of their carbon footprint) achieving Byzantine resistance and worldwide consensus than centralized or private chain approaches. They also are required to replicate data everywhere, regardless of actual need. **These multiple forms of overhead make them permanently more compute (and carbon) intensive on a per-transaction basis.**

Due to the limitations described above, most IT solutions that do employ public chains will use some form of digest (Merkle Tree) approach to store only a *synopsis* of their data, rather than treating the public chain as a real-time operational data store, allowing them to amortize the cost, environmental impact, and latency of public chain overhead.

Example: Building an Image Licensing Business

As a way of comparing (and, ultimately, integrating) the three approaches, consider an existing digital asset distribution business, such as Adobe Stock™. While business relationships can be complex, for discussion purposes this example will consider five roles:

1. An *artist* creates a digital work and sells it to a content company in exchange for royalties.
2. The *content company*, having consolidated digital media from multiple artists, now resells its catalog through retail distribution channels.
3. The *retailer*, such as Adobe Stock™, handles aggregate catalog and content integration as well as licensing and billing to end users.
4. *End users* license individual works of art for inclusion in presentations, websites, or other content that requires media elements.
5. *Viewers* of those presentations, websites, etc. consume that content through streaming services, downloaded collateral, or other means. Viewers may also play the role of *auditors* if they attempt to discern whether the material they are viewing has been properly licensed.

Figure 2 illustrates these relationships in the lifecycle of a media work from artistic creation through consumption.



Figure 2: Overview of Adobe Stock™ business example.

Existing (“As built”) Architecture Using a Centralized Approach

Since these businesses have existed for some time and obviously predate blockchain technology, they’ve been realized using conventional (“centralized”) architectures, both on prem and in the public cloud. While a detailed illustration of a typical architecture is beyond the scope of this article, the key elements of its data ownership and storage suffice.

Broadly speaking, both the reseller and the distributor maintain a full copy of media assets in some type of file system. They also maintain one or more database tables to model the key elements of their business relationships, such as:

- A media *catalog* (list of media assets and their prices)
- A list of *customers* or reselling partners
- A many-to-many *mapping* of which customers (or reselling channel partners) purchased which assets

On top of that database storage layer various web and mobile applications (along with billing and other back office services) manage the interaction with the end users, including navigating and purchasing licenses. The integration between media content companies and distributors may also be mediated through one or more public APIs.

The existing, centralized architecture has some important advantages that are worth preserving:

- **Low storage costs.** With many terabytes (possibly exabytes) of media assets, keeping storage costs low is critical. Both the content company and the distributor need storage cost structures near or at market rates for cloud file storage in order to make their business margins tractable.
- **Low transaction costs.** A license for a small image may be as little as a few dollars. That implies that the infrastructure cost of performing the licensing operation (e.g., updating a database table by adding another row containing the customer id and the asset it) must be a small percentage of those few dollars or the company won’t be able to make a profit. *Modeling a \$3.00 asset license purchase as a \$30 Ethereum transaction, e.g., would make the business model “upside down” on margins.*
- **Low latency (and access controlled) distribution.** To keep the customer experience viable, the time between purchasing a license and the end user being able to use the media asset must be single digit seconds in most cases. Low latency downloads from anywhere in the world that reflect recent purchase authorizations

in real time is a hard requirement, and only users who have licensed a particular asset should be able to download it.

- **Operational data stores and billing are easy to integrate.** The company's business data depends on being able to read and write operational data – i.e., who owns which license – quickly and at low cost. Exhausting this data into billing systems as well as analytic systems for BI (Business Intelligence) purposes and AI recommendation engines are all hard requirements, and cannot be subjected to high latencies or data retrieval costs. User billing can be handled through simple integration with SaaS solutions such as Stripe.

Next-gen Architecture Requirements

While the existing centralized approach addresses basic business needs and offers the benefits above, it also *lacks* some highly desirable capabilities:

- **The original authors of the media have no insight into how their works are being used, and are required to simply trust that compensation is fair and equitable.** This is a classic centralized application limitation: The “proof” of what was transacted is locked up in the private database tables of one or more companies, implying that no one else can see that content without their assistance and permission. This is more than just a conceptual limitation; it's also a *business model* limitation, because without this capability, offering flexible rates or pass-through income to media authors isn't readily possible.

Ideal solution: Artists can track sales and use of their works over time and validate their compensation is accurate without reference to any private company data.

- **Similarly, a third party auditor (or even just an interested party) can't discern whether use of a media image is valid or not.** For example, if one saw an image that was subject to copyright in a presentation on the web, getting the owner of the presentation to validate their license would require the intervention of the licensing service, who would need to look at their private database tables again.

Ideal solution: Viewers can independently validate appropriate image use without reference to private, centralized database content.

- **Reselling content is challenging.** For the content company and the distributor to share media and revenue, the IT departments of both companies need to build dedicated implementations. Without the benefit of a partner data bus that inherently models ownership, access and sharing controls, and versioning, both parties are required to erect complex (and thus likely, expensive) DIY sharing solutions. Furthermore, due to the increased challenges that real-time sharing would present, they likely do so with “batch” approaches that require full-scale copying of all assets and periodic (versus continuous) billing and usage updates.

Ideal solution: Partners can share media assets securely without writing code, and enjoy real-time notifications with respect to end user licensing events and co-sell revenue.

Are NFTs the Answer?

Given the media attention accorded to NFTs, it might seem at first blush that a newer, better implementation strategy exists: Simply place a hash (digital summary) of the artwork into an NFT, and use that to represent its ownership (see Figure 3).



Figure 3: Alternative implementation of digital asset licensing using NFTs.

Unfortunately, a simple, singular model of ownership doesn't map well into the reselling and licensing approach required by the actual commercial relationships in our sample problem. NFTs as implemented today also tend to suffer from structural challenges, such as high transaction costs – ignorable for a \$50,000 painting, but problematic when paying

\$0.99 to license a GIF for a slide deck. They also don't provide built-in mechanisms to handle configurable end user authentication, authorization, access controls, or CDN distribution. In short, this approach fails to capture the benefits of the existing approach identified above: low storage and transaction costs, low latency (and access controlled) distribution, and easy license management and billing. NFTs also don't *directly* accomplish the additional goals, either, because they don't model all five personas and their relationships with sufficient fidelity.

Better Together: Using all Three Approaches

Conventional (centralized) architectures provide a key set of benefits and suffer from a key set of gaps. NFTs offer the promise of a global, tamperproof record of ownership but don't support the business workflows and relationships required to commercialize this example. A well-architected "3-body" approach seeks to preserve the advantages of existing (centralized) approaches while adding key functionality through the careful incorporation of private and public chain technologies. In this approach, each category plays a key role, chosen to align with its strengths and weaknesses profile as discussed above:

- **Centralized IT** continues to manage the b2c consumer (and any partner portal) websites, billing, payment and other back office services, and business analytics, including BI and AI. These are all effectively "private" operations that optimize for low cost and development/operations efficiency rather than partner or public sharing.
- **Private chains** enable bidirectional reselling, including handling all media, CRM, and purchase/use data sharing as well as access controls, versioning, governance and security. They integrate natively with all major clouds and CDN services to offer low-margin storage and distribution of the underlying assets at cloud "line rates" of latency and throughput, keeping the customer experience intact even through multiple transitive hops of ownership. Private chains may also be used over time to integrate with backend SaaS providers, such as billing, or third-party auditors as those firms adopt blockchain technology themselves.
- **Public chains** provide a permanent, public, global record of "who's licensed to use what" that makes it possible for anyone, including the original artists, to understand how their works are being used and whether a given occurrence of a media asset is appropriate. Lower cost aggregate information is stored in a service such as IPFS, mapping public tokens (keys) for users to their list of licensed asset ids, and then hashed into a public chain such as Ethereum to create an immutable, totally ordered, irrefutable custody store. To protect the privacy of purchases, this

information is also encrypted, giving each media purchaser the “key” (quite literally) to reveal to a third party their purchase receipt without reference to the licensing company’s internal databases, and enabling the original artists to unlock purchase information (but not purchaser PII) sufficient to verify pass-through revenue accuracy.

Figure 4 illustrates how all three categories contribute to this composite architecture.

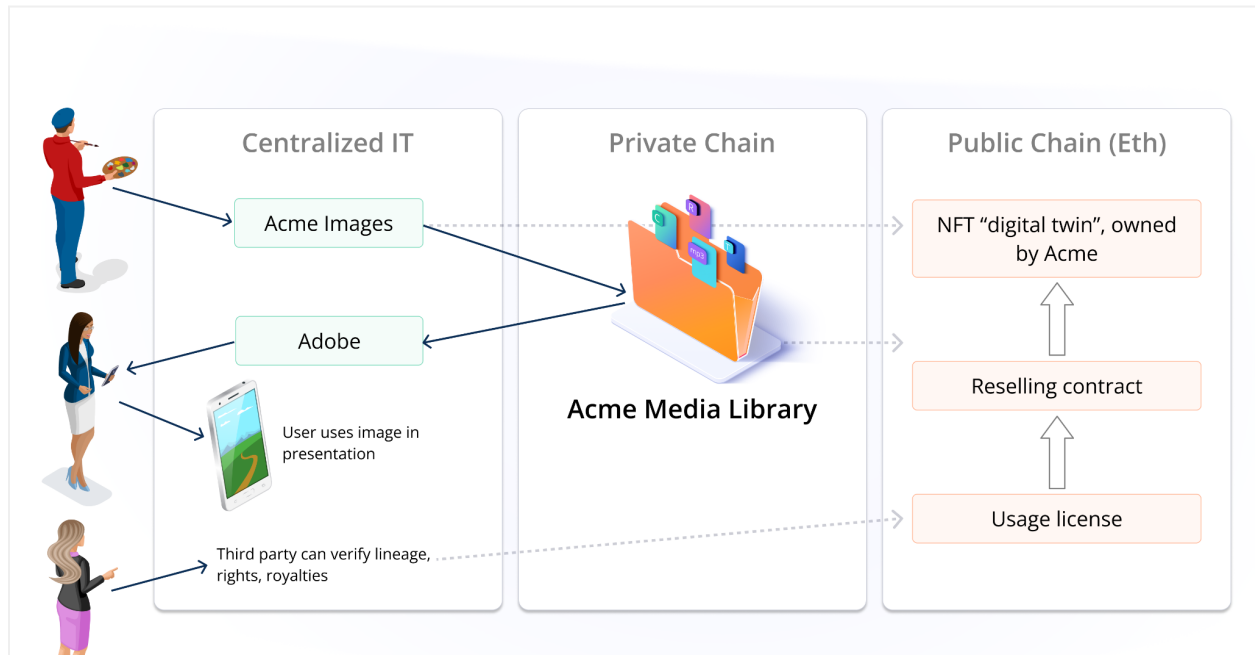


Figure 4: Combining centralized, private, and public chain approaches yields the best business outcome, with each contributing its unique benefits.

Integrating Central IT with Private and Public Blockchains: The Real-time Data Mesh Pattern

The example above illustrates that no single approach is ideal for modeling the complex, multi-party data sharing challenges inherent in real-world applications and use cases. Neither public nor private chains are a true substitute for the other, and neither is a replacement for conventional (aka “centralized”) IT architectures. What has emerged instead is a playbook for how all three technology categories – centralized cloud and on-prem solutions, public blockchains such as Ethereum, and private/permissioned chains will interoperate to create best-of-breed IT outcomes over the next 20 years.

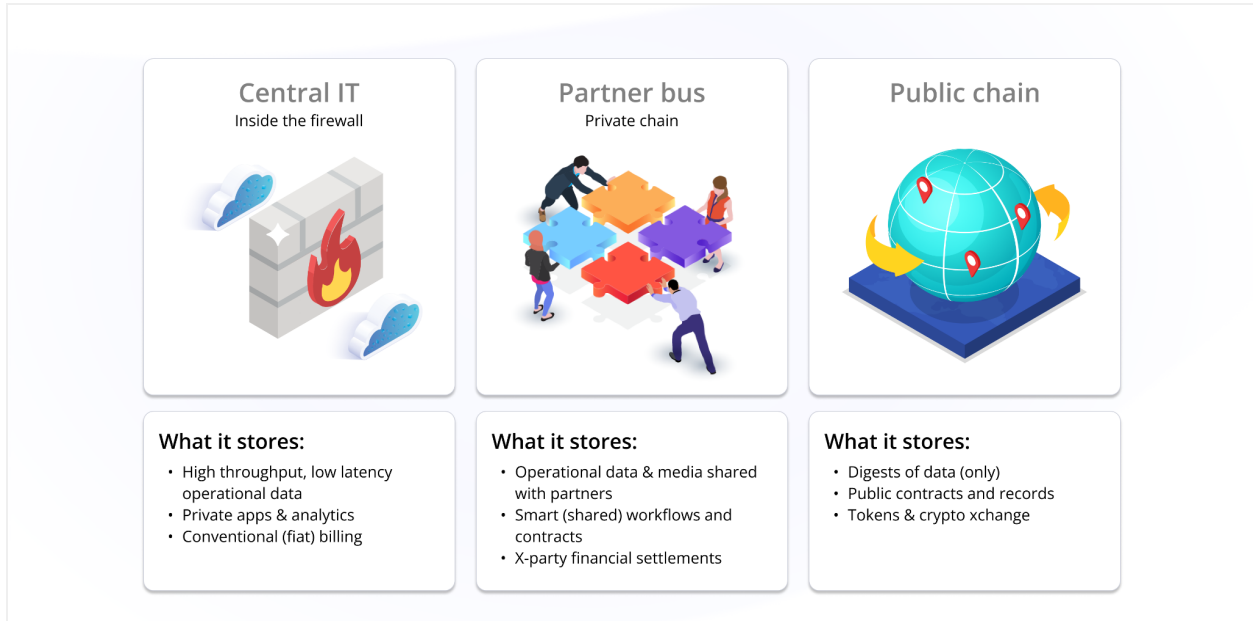


Figure 5: In a best-practice application architecture, all three categories are employed to their advantage, with private operational data held in conventional storage and compute solutions (typically in the cloud), partner data sharing accomplished through high speed, highly governable private networks, and public chains used for digests and cryptocurrencies.

Conventional (aka “centralized”) IT solutions, increasingly built in the cloud and often on top of managed cloud services, will continue to offer the highest throughput, lowest latency, and easiest governance indefinitely.

- They’re the right place to locate private applications and data, such as analytics and BI exclusive to the owning company.
- Many systems continue to be internal-only, with no need to pay for the overhead (or security concerns) of outside-the-firewall distribution or coordination.
- They’re also ideal for conventional fiat billing ... such as Stripe integration.
- They will continue to form 80% or more of the foundational layer for SMB, mid-market, and enterprise IT.

Private/permissioned blockchains have the advantage of innovating quickly, and can integrate seamlessly with the cloud solutions of their controlled populations, making them ideal for creating partner buses that enable businesses to share data with rigorous compliance, privacy, and governance capabilities.

- Partner data buses and shared applications center on private chains for their simplicity of governance & security and throughput/cost ratio advantages.
- Conventional (“dumb pipe”) APIs will be replaced over time with distributed ledgers that offer an automatically maintained and cryptographically provable single source

of truth for cross-company use cases such as supply chains, financial ledgers, and media/file sharing. This will simplify today's IT challenge of building secure, high throughput, APIs for sharing real-time data with partners.

- Private chains are ideal for cross-party financial settlements and payments.
- Over time, they will grow to represent 15-20% of IT stacks.

Public blockchains offer very limited throughput and will continue to have high latency, limited data type support, and modest innovation rates (due to their open source/multi-player ecosystem natures). Due to their high storage and transaction costs, public chains will remain a small percentage of code and data, but play a critical role in creating trusted, planet-wide stores-of-record.

- They will play a key role for applications that require public disclosure, non-repudiation and/or total ordering, or where public data markets exist, serving key use cases such as archival storage for NFT ownership, self-service ticket transfers among consumers, and third party auditing of (hashed, encrypted digests) of key financial and business data.
- Cryptocurrencies and other on-chain tokens and related capabilities will exist in, and be accessed through, public chains.
- Public chain integration will eventually comprise 1-2% of typical enterprise IT stacks over the long term, though they will be used to model critical elements, such as cryptocurrency financial rails and public records.

Figure 5 summarizes the best practices for each technology category, and Figure 6 illustrates their long-term relationship in best-practice IT portfolios.

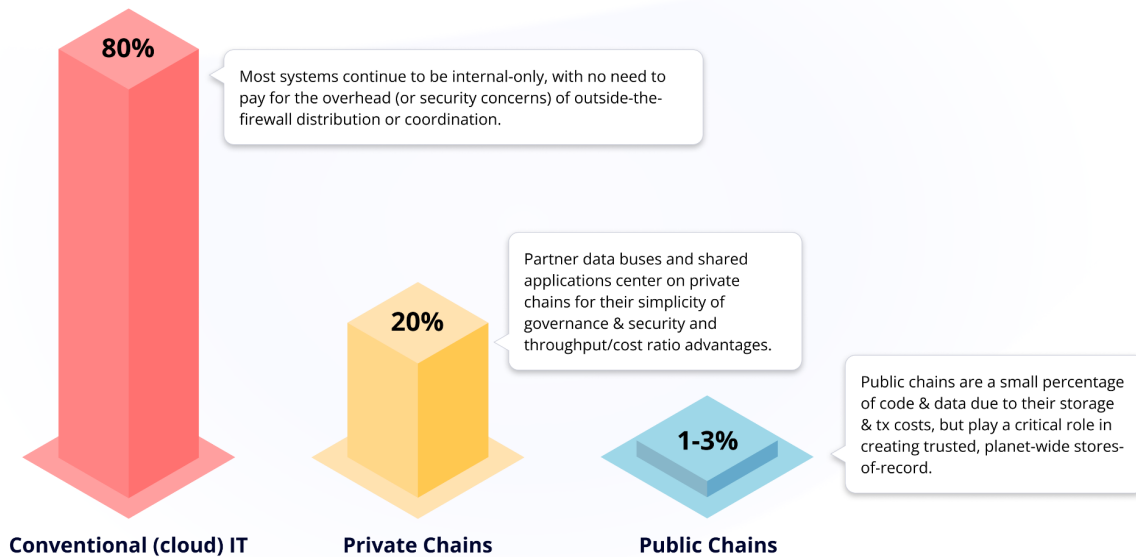


Figure 6: Long-term relative sizes of conventional IT, private chains, and public chains in a typical IT portfolio (by code size).

Conclusions

Over the next 20 years we will see all three technology categories – centralized solutions, public blockchains, and private/permissioned chains – serve as the foundation of mission-critical applications. Existing approach today, including public APIs for partners, will be upgraded through the use of private blockchain technologies, and real-time “data meshes” that span companies and clouds will become the norm, rather than a novelty. Organizations will leverage all three of these technology categories together to deliver best-of-breed IT outcomes.

Adoption Caveats

The benefit of hindsight exposes common pitfalls many organizations fall into when adopting technology. Avoid the temptation to interpret media attention and “hype curve” trends incorrectly:

- *Public cloud growth continues to be staggering, and will not be “replaced” by web3.* Rather, best of breed solutions will employ both emerging (blockchain, web3) capabilities *and* existing cloud technologies.
- *Public chains, despite modest optimizations, cannot violate the laws of physics.* Assuming they will become a replacement for operational data stores or private chains that enjoy cloud-derived economies of scale and performance is betting

against science. Carefully consider the achievable transaction rates, storage, and development costs required in using public chains, and leverage them for their unique advantages, not as a replacement for centralized or private chain data stores.

Successful companies will win by:

- Optimizing tech selection for the best outcomes, rather than pursuing a “pure play” approach (i.e., only centralized / only private / only public).
- Building as little as possible in order to focus time and energy on their unique, differentiated value add.
- Leveraging the best of cloud, central IT, and both public and private chains to minimize time to market for new and replatforming applications.


Going Further

Looking to learn more about integrating centralized and blockchain capabilities for IT use cases? The Vendia blog has a number of articles, including how these features surface in [modern applications](#) and get exposed through [data-aware APIs](#).

If you have a business need not being met by your existing blockchain provider you can visit www.vendia.net to learn more about supported use cases or [contact the team directly](#) to dive deeper.

About the Author

Dr. Tim Wagner is the inventor of AWS Lambda and a former general manager of AWS



Lambda and Amazon API Gateway services. He has also served as VP of Engineering at Coinbase. Tim co-founded Vendia to help organizations of all sizes share data more effectively across clouds and companies and serves as its CEO. He is also the driving force behind the Lean Apps Movement, holds dozens of patents in serverless, cloud compute, and database technologies, and is a frequent keynote speaker at software conferences.

[linkedin.com/in/timawagner](https://www.linkedin.com/in/timawagner)

[@timallenwagner](#)