EY and Huawei present

Making the network the cloud

How blockchain is pushing the cloud all the way to the edge

HUAWEI

EY
Building a better working world

contents

Executive summary

At some point, probably in the early 1990s but perhaps as early as the 1980s, the amount of computing power and data storage in individual PCs, departmental servers, smart devices and mobile phones began to exceed all the storage power and compute capacity in all the world's major data centers combined. Today, the computing power in pockets and on desks and in light bulbs or smart refrigerators exceeds all the world's data center capacity by orders of magnitude. And most of this power is idle or empty nearly all the time.

Historically it has been easier, cheaper and more reliable to build new centralized data centers than to make use of these idle systems. Numerous obstacles existed ranging from a lack of tools to manage the workloads down to a lack of security and payments infrastructure that made truly distributed computing uneconomical.

Over the next few years, the convergence of several separate streams of innovation are going to make truly distributed computing and storage a viable economic and technical proposition, and in the process, transform both the traditional data-center business and the networking business.

Game changing enablers

Decrease in edge device costs

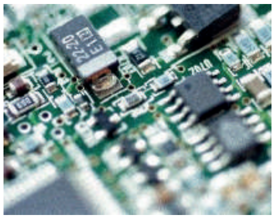

Commercial cost of fully functional computer (Raspberry PI Zero) is US$5.

Increase in computing power and mobility

Computing power of a phone exceeds total power of all NASA computers used for Apollo 11 mission to the moon.

Decrease in cost of connectivity

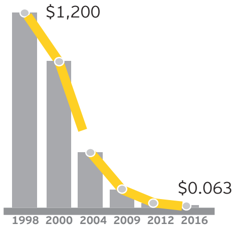

Transfer cost of 1mb of data dropped from US$1,200 (1998) to US$0.063 today.

On-demand infrastructure

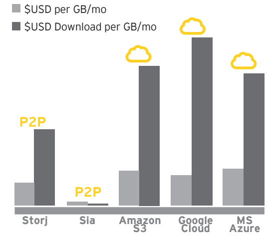

Storage cost of 1gb of data decreased from US$569 (1992) to US$0.002"

The cost of making edge devices, connecting them, and the cost of storage on edge devices is going down.
Distributed computing will not be applicable to every workload. However, the rise of truly peer-to-peer infrastructure, driven heavily by containerization and with workload allocations and payments taking place through blockchain, is likely to have two significant effects: first, it will contribute to yet another dramatic cost reduction in the cost of operating cloud computing services; and second, it will trigger another dramatic increase in bandwidth consumption.

This paper will look at the current state of distributed computing and storage, key architectural changes coming in the near future state, and the likely feasibility and economics of different business models that may emerge. While the exact shape of this market is not yet clearly visible, we believe that the organizations that lead this transformation will have the opportunity to upend both the computing and networking industries simultaneously and stake out entirely new leadership positions.

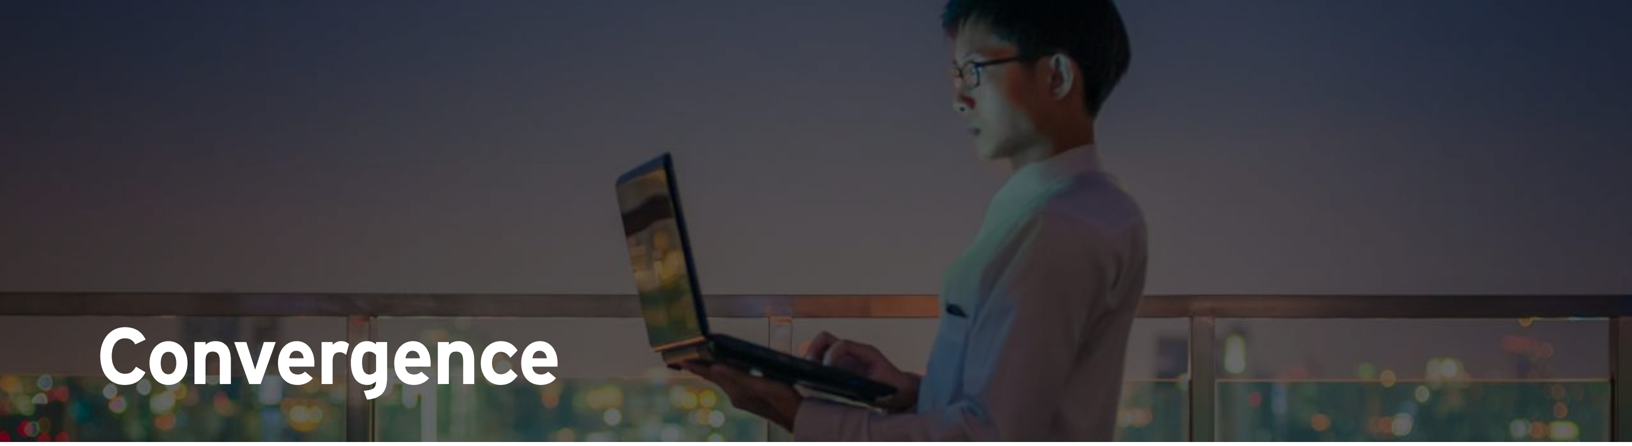

Convergence

The exponential rise in data traffic growth has been driven by high-bandwidth video streaming along with rising adoption of new applications such as augmented and virtual reality, 4K video, artificial intelligence, autonomous vehicles, etc. that will demand advanced bandwidth requirements and greater uplink / upload speed.

To meet the increasing demands of the network, newer network architectural models such as distributed or decentralized networks will be more cost-effective and efficient due to the increased proliferation of edge devices having idle computational and storage capabilities. The trust once placed in centralized systems will yield to economic necessity; as increased network demand from future users leads to unsustainable compute and storage costs, centralized infrastructure will give way to distributed networking.

Edge devices have the potential to drive toward a fully distributed architecture where every computing device is a micro data center. As the number of edge devices with significant processing and storage capabilities is increasing, it is making it easier for companies to push computing applications, data and services away from centralized cloud providers to the extremes of the network. Such a peer-to-peer model will reduce the costs associated with installing and maintaining large centralized data centers and will secure users from a single point of failure.

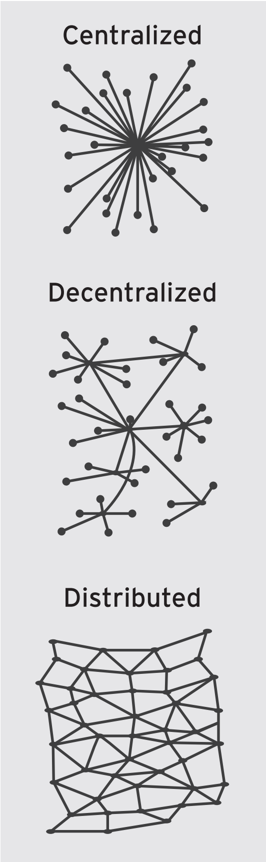

Convergence — why now?

Container technologies

Containerization, or the process of isolating tools and libraries necessary to run software from the rest of the system, has changed the IT development and deployment paradigm. Being more efficient than virtual machines, containers provide quick provisioning of a runtime environment and application portability. Among the variety of uses this can be applied are the Internet of Things (IoT) and blockchain, which require lighter, modular services to seamlessly interact with each other. Significant efficiency has been gained in application development by using the microservices architecture of breaking up applications into modular pieces and using containers for deployment. This concept of containers hosting microservices can be tweaked and applied to build decentralized applications on the blockchain, which will provide scaling and modularization benefits.

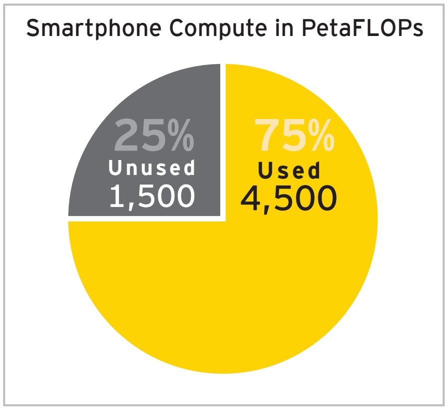

Source: IANS, January 28, 2018

Capacity at the edge

A key to successful digital transformation will be empowering digital business flows between people and data, all of which live at the edge of the business.

By 2020 there are expected to be over 200 billion connected devices in use. There are anticipated to be almost 6 billion smartphones alone in that number. Using the estimate that one million smartphones represents roughly one PetaFLOP of processing power, this would equate to almost 6,000 PetaFLOPs of potential processing power just in smart phones that sit idle for roughly 18 hours per 24 hour period.

An average individual would not need anywhere near this amount of computing power for even complex computing. At most, an average user would never demand any more than .0001 petaflops of power.

In comparison, the maximum advanced processing clusters that a large enterprise might lease could offer anywhere from 1 to 11 petaflops.

Edge devices are now powerful enough to multitask, and it is effectively cheaper to make them really smart than minimally intelligent.

Blockchain

The movement of computing to the edge, driven by the high cost of cloud computing, required a new means of establishing trust and payment across large-scale and highly distributed environments. Blockchain, as it is generically defined, has the potential to become the underlying platform on which all edge transactions are run because blockchains can create anonymous, peer-to-peer trust while providing an infrastructure to distributed computing.

Disruption impact

Each converged technology has its own promise of market growth predictions, which when taken as a whole suggest a paradigm shift in how consumers and providers will interact.

There is no big boom in disruption, but the speed of change is increasingly compressing. Telecommunications manufacturers and providers as a market will be obvious beneficiaries due to the expansion of networking bandwidth for millions of devices and users.

Divesting from centralized data centers and moving application workloads to the public cloud have been discussed for years. This is not a wholesale move of the current networking state to the P2P distributed model, but it will be a significant movement where most of the use cases will benefit from the availability and matching of specific resources to process specific workloads.

The emerging trend today is that much of data analytics is shifting its creation and processing closer to the edge. New lower latency services such as augmented reality, autonomous vehicles, remote patient services, etc., will only accelerate these trends. As a result, carriers are increasingly pursuing intelligent distributed computing platforms (e.g., centralized RAN/ cloud RAN, mobile edge computing, and data center interconnection fabric), embracing distributed functionality and peer-to-peer applications to improve performance and structurally reduce the cost curve of running their networks.

Current efforts

Current efforts to advance distributed systems utilizing blockchain are only focusing on small use cases or minimally impactful solutions. These efforts generally consist of joint ventures between firms, or small consortiums. For example, Verizon Enterprise Solutions has partnered with Guardtime to create a blockchain offering to enable their enterprise clients to better serve their consumers' security needs and concerns. This product is for a niche customer and the primary goal is solely focused on more effective security.[1]

Other blockchain initiatives utilize Hyperledger which pushes organizations into siloed blockchain ecosystems instead of allowing for one blockchain environment accommodating multiple organizations at once. For industry wide use cases, segregated blockchains will not be as competitive as scalable public blockchains. They fit into the same categories as minimum blockchain technology adoption, instead of taking the full potential of the P2P network coming from blockchain.

A more promising example would be the consortium between Softbank Corp, Sprint, Far EasTone and TBCASoft. TBCASoft is a US.-based startup developing an inclusive cross-carrier blockchain platform for the telecom industry to allow customers to utilize the services of multiple carriers with simple payments. Softbank Corp., Sprint, and Far EasTone are the initial carriers in the consortium, with more to be added in the future. Despite these efforts, current solutions are short-term and low significance because they are only tackling one use case at a time. What is needed is to utilize blockchain across the whole telecom industry to build an ecosystem that optimizes the convergence of multiple technologies.[2]

1 "Verizon Enterprise Solutions will create global blockchain platform services," Verizon website, http://www.verizon.com/about/news/verizon-enterprise-solutions-will-create-global-blockchain-platform-services, accessed 30 April 2018.

2 "SoftBank, Sprint, Far EasTone and TBCASoft launch blockchain consortium for telecom carriers," Sprint newsroom website, http://newsroom.sprint.com/softbank-sprint-far-eastone-and-tbcasoft-launch-blockchain-consortium-for-telecom-carriers.htm, accessed 30 April 2018.

Market status

All intermediary transactional businesses are building strategies around how to be relevant in an ecosystem where peers can both consume and provide technology services. For networking providers, this then means that huge business sectors and consumers expect networks to be future-ready.

While some projects are a right step forward for cloud computing, smaller visionary actions are made at this stage. **Cisco** has filed a US patent, "Block Chain Based IoT [Internet of Things] Device Identity Verification and Anomaly Detection," seeking to play a role in verifying the identity, safety and trustworthiness of connected devices operating with blockchain-based distributed ledger technology.[3] [4] [5]

Ericsson is in partnership with Guardtime to integrate Keyless Signature Infrastructure (KSI) Blockchain technology into their cloud portfolio. They have developed a combination of blockchain, microservices and open APIs to deliver "Blockchain Data Integrity" that is implemented as a microservice for the GE Predix Cloud platform.[6] [7] [8]

Verizon Enterprise Solutions launched a portfolio of blockchain platform services based on Guardtime's KSI Blockchain capabilities. Using this, enterprise and government customers will be able to integrate blockchain solutions into their existing business processes. And Verizon Ventures invested in Filament, a blockchain startup that develops a secure communications platform for devices operating in distributed environments.[9] [10]

Cloud providers are entering the market by providing blockchain services. Google is working on blockchain-related technology to support its cloud business by developing its own distributed digital ledger that third parties can use to post and verify transactions. Amazon recently released AWS Blockchain Templates to provide a fast and easy way to create and deploy secure blockchain networks. With AWS Blockchain Templates, Ethereum and Hyperledger Fabric frameworks can be deployed using AWS CloudFormation templates.

There are already numerous startups tackling distributed computing based on blockchain technology, such as Golem Network, SONM, iExec and Akash Network. Other groups are focused on uniting peers on a distributed computing platform for scientific research. For example, SETI's analysis of signals convincing millions of users to devote idle computing power to help analyze signals from space.

ServerCube is offering a peer-to-peer network of globally connected devices, designed to operate in private residences or in small clusters. The devices are small high-powered processing computers connected via the internet and centrally managed. The complete ServerCube ecosystem consists of the ServerCube Device plus ServerCube Network and ServerCube Connect (middleware fabric).[11] [12]

Caveat

The proposal in this paper will not address the many local and national interventions that could occur in a distributed ecosystem model. At this point in time, the definition of blockchain is not even fully defined enough to put proper controls in place. Furthermore, privacy regulations may be a challenge to comply with given the immutable nature of blockchains. Data storage requirements in different countries may be different and may be subjected to privacy regulations such as the EU's General Data Protection Regulation (GDPR). Tracking and keeping up with such privacy regulations of different countries may prove to be a major challenge. For instance, as the transactions on a blockchain are considered to be immutable, this is not in line with the GDPR requirements of data erasure and data deletion after the lawful use of it.

---

3 "Blockchain in enterprise: how companies are using blockchain today," Blockchain at Berkeley blog website, https://blockchainatberkeley.blog/a-snapshot-of-blockchain-in-enterprise-d140a511e5fd, accessed 30 April 2018.

4 "Cisco files patent application around blockchain, IoT integration,"ETHNnews website, https://www.ethnews.com/cisco-files-patent-application-around-blockchain-iot-integration, accessed 30 April 2018.

5 "Cisco's IoT blockchain merely scratches the surface of distributed ledger technologies,"IPWatchdog website, http://www.ipwatchdog.com/2018/01/13/ciscos-iot-blockchain-distributed-ledger/id=91936/, accessed 30 April 2018.

6 "Ericsson and Guardtime create secure cloud and big data," Ericsson website, https://www.ericsson.com/en/press-releases/2014/9/ericsson-and-guardtime-create-secure-cloud-and-big-data, accessed 30 April 2018.

7 "Implementing blockchain as a microservice for IoT platforms,"Ericsson Future Digital blog website, https://cloudblog.ericsson.com/digital-services/blockchain-microservice-data-integrity-iot-platform, accessed 30 April 2018.

8 "Ericsson launches Blockchain Data Integrity for GE's Predix platform," Ericsson website, https://www.ericsson.com/en/press-releases/2017/5/ericsson-launches-blockchain-data-integrity-for-ges-predix-platform, accessed 30 April 2018.

9 "Verizon Enterprise Solutions will create global blockchain platform services," Verizon website, http://www.verizon.com/about/news/verizon-enterprise-solutions-will-create-global-blockchain-platform-services, accessed 30 April 2018.

10 "Filament receives $15M in new funding for industrial IoT: industrial application of IoT and blockchain technology ready for high-scale production," BusinessWire website, https://www.businesswire.com/news/home/20170330005189/en/Filament-Receives-15M-New-Funding-Industrial-IoT, accessed 30 April 2018.

11 "New York startup ServerCube launches ICO to build the largest decentralized high-capacity computing network for crypto-mining, deep-learning AI, VR modeling and BaaS," TechStartups website, https://techstartups.com/2018/03/05/new-york-startup-servercube-launches-ico-to-build-the-largest-decentralized-high-capacity-computing-network-for-crypto-mining-deep-learning-ai-vr-modeling-and-baas/, accessed 30 April 2018.

12 "ServerCube: building the backbone of decentralization for the blockchain generation,"WixStatic website, https://docs.wixstatic.com/ugd/771275_5de78e10ea0645dbb962f8fcd1c1bd42.pdf, accessed 30 April 2018.

Current state

Networking

Traditionally, carriers provide telecommunications services (e.g., text messages, calls, voicemail) and the connectivity and bandwidth (pipes and transport devices) that support those services. Nowadays, consumers often rely on messaging applications whose functionality includes text and voice messaging, telephony, conference calls, video calls and more. Services that have traditionally been the realm of carrier services may now be provided by sites such as Google, Facebook or Telegram.

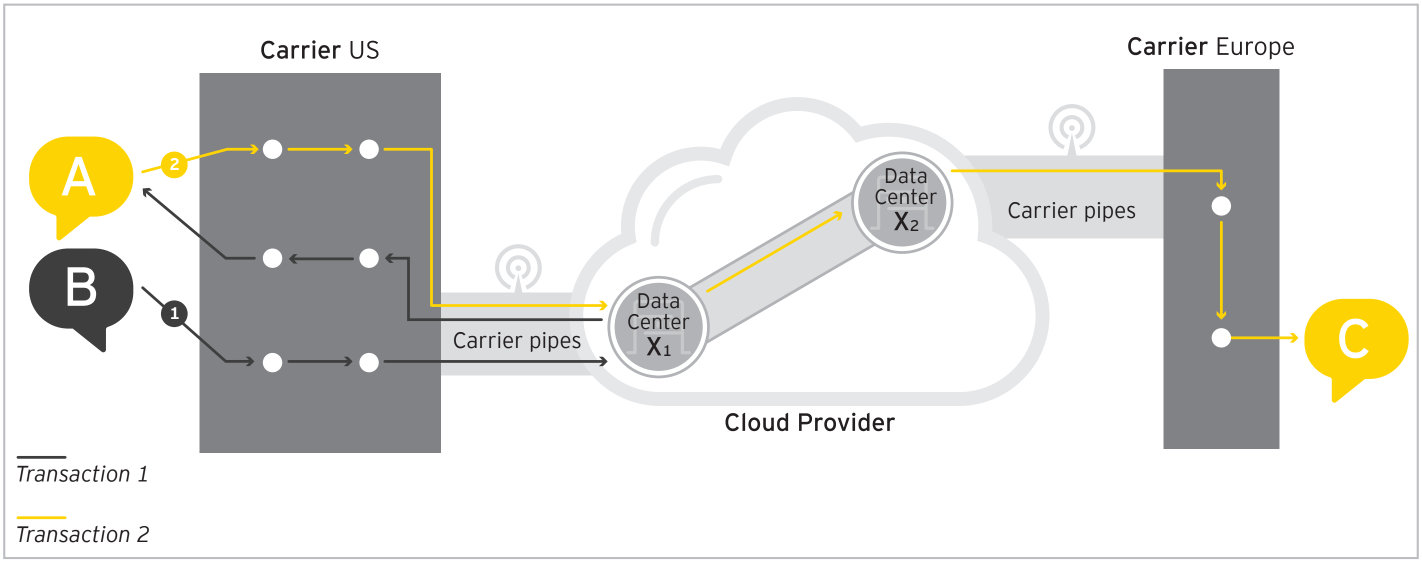

Consider this diagram: User B opens a messaging application to send a message to User A. There is no direct way for them to communicate without an intermediary. B's message will go through the pipes of the carrier, then to the data center of the application, then back through the carrier's pipes, down to User A. In this scenario the carrier only provides the pipes, while the datacenter that operates the application provides the more sophisticated service.

Not only are carriers losing part of their services business to others, they also are losing market share as data centers are now using their own pipes for inter-datacenter communications. Again referring to the diagram above, if User A is located in America and wants to have a phone call with User C in Europe, it may be more cost-effective to use one of the many messaging applications allowing calls than using a carrier's phone services. When User A initiates the call, it goes through the carrier's pipes to the application's data center; however, once in the data center, the pipes used are no longer the carrier's pipes but rather the data center's pipes.

Current State of Network Architecture Allow streaming content provider to go straight to customer, competing with the traditional offering of carrier, rendering carriers to a role of infrastructure provider (switches and routers).

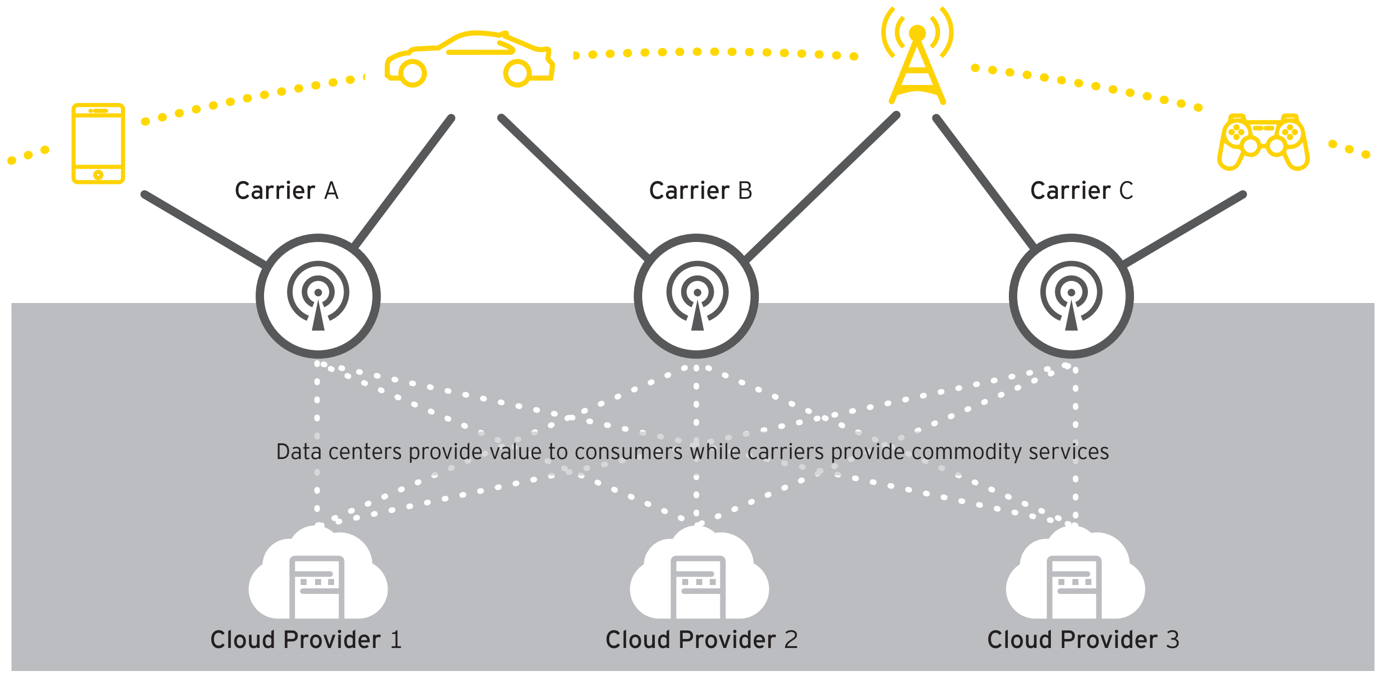

As seen in the architectural layout, data centers are becoming powerful oligopolies dislodging carriers from their traditional services.

In essence the traffic is going in the direction of the cloud providers, and once in a cloud, the services (messages, video messages, calls, etc.) tend to be operated in and by the cloud provider. What the carriers offer is the commodity hardware that allows the connectivity between users and the cloud.

Peer to peer

Peer-to-peer (P2P) historically referred to a distributed architecture that partitioned tasks between interconnected nodes or shared resources without the use of a centralized administrative system. Whereas client-server models became prevalent because of the efficiencies of a central client, social networking demands and technology advances are making P2P more economical and technically feasible. Within new social constructions, peering truly refers to the equality of all participants throughout the network.

Peering will redefine the requirements for access to capital expenditures and costly centralized infrastructure systems. As an example, the internet, as a point-to-point network, was specifically designed for participation by distributed peers without requirements of centralized hubs. However, the cost for the equipment required to participate limited users to large entities with capital funding. As device prices decrease, bandwidth becomes available, and the user base expands, the level playing field between users will enable equality between both big data centers and smarter connected devices.

Blockchain

Blockchain technology can create limitless opportunity in tapping into underutilized infrastructure and devices and provide the infrastructure to support an entirely new ecosystem.

Blockchain technology is a way to structure data without the need for a central authority. A blockchain is a distributed database that hosts a continuously growing number of records. The database stores records in blocks rather than creating them in a single file. Each block is then "chained" to the next block, in linear, chronological order, using a cryptographic signature; as a result, records cannot be revised, and any attempted changes are visible to all participants. This process allows blockchains to be used as ledgers that can be shared and corroborated by anyone with the appropriate permissions.

These distributed ledgers can be spread across multiple sites, countries or institutions. Distributed ledgers can feature smart contracts that often have logic built into code that is stored, verified and executed on a blockchain, providing a platform for self-enforcing, self-executing agreements.

Distributed ledger

Blockchain platforms rely on native marketplaces of "requesters" (clients) and "providers" (miners). These marketplaces match up the needs of requesters with appropriate providers. In some platforms, both parties utilize File Contracts to establish various terms of consideration (e.g., payment strategy, amount of data being stored) between both parties. These parties then store the contracts on the platform's blockchain for later auditing to confirm the agreements are upheld. To verify that miners are trustfully storing client data, blockchain networks utilize various consensus mechanisms to enable clients to check the security and safe handling of their data at any time. These proofs are periodically submitted to miners to verify that the fragment of data (a merkle hash) submitted by the storing miner matches the merkle root hash of the client's original file. These operations are all completed using the platform's immutable blockchain ledger.

Consensus

While there are many kinds of consensus mechanisms, some are more relevant to the distributed computing and storage use case than others. Consensus such as proof of retrievability, proof of replication and spacetime, and proof of storage are at the experimental stage, and if successful may be the enablers of a decentralized economy of storage and compute.

Proof of Retrievability (PoR), as used by Storj, allows users or verifiers to check at any time that a data provider is currently storing the data being held. In the event the data provider will not relinquish control of the data, the owner can retrieve their data without the provider's consent. This process is achieved through merkle trees and merkle proofs. The data owner can store the merkle tree and its merkle leaves and periodically initiate a challenge/response protocol with the providers to cryptographically match that both parties have corresponding merkle proofs.

Proof of Replication (PoRep) and Proof of Spacetime (PoST), as used by Filecoin, rely on the first PoRep which differs from PoR in that provers agree to store n number of replicas of the client's data in its own physical storage, and must prove that they have actually stored the n number of replicas at any given time through a challenge/response protocol. General proof-of-storage schemes only verify that the provider is holding the data at the time of challenge. It also relies on a second PoST to check that the provider has been actually storing the data for the entire agreed-upon time. These proofs are publicly available, meaning anyone at any time can check the sequence of accurate proofs to verify provider storage.

Proof of Storage, as used by Sia, is designed to regularly request proofs from providers; they are compensated for each successful proof that matches the merkle roots of the client's data. Like Filecoin, these proofs are published publicly on the blockchain for anyone to audit and verify.

At this stage it is unclear which if any of the above current consensus mechanisms will be successful at servicing the distributed computing and storage use case. However, there are good indications that some storage use cases are operational, and it is just a question of time before the compute use case is too.

Smart contracts

Smart Contracts are the mechanism that allow clients and providers to specify their terms of agreement before engaging in any network interaction. This can be the duration a provider will store a set of data, payment terms, the periodic amount of validation of storage proofs, etc. When any of these conditions are met or infringed upon, the network will automatically trigger a response either compensating the responsible party, or penalizing them with a fee for a bad action.

Incentives and tokenization

Tokenization requires a more complex infrastructure than notarization as it requires structured data. However, tokenized assets can be part of significantly more complex processes. Tokenization allows for both real and virtual assets. They can be financial or non-financial; can exist in many different locations; can be digitally represented; and can be transacted on one platform. On a blockchain, assets, currencies and other tokens can be bought, sold and exchanged but they cannot be duplicated. With tokenization, it is possible to execute full-cycle economic contracts between participants. Tokens have a value and owner(s), which allow the delivery of nearly any kind of services against that token. Exchanging tokens within contracts allows for very complex relationships to be modeled easily and reliably in one environment — the blockchain. Most people, with a very high degree of precision, have a quick access point to their online banking balance. When it comes to supply of assets and inventory, organizations have a blurry picture of what's on hand — a problem that tokenization of assets in their supply chain has a fair shot at solving.

Most blockchain platforms all have strategies designed to establish accountability among actors in their networks. This occurs through provider deposits of their own currency in escrow accounts, and for each inaccurate proof submitted, they lose a portion of their deposit. Other platforms are looking into reputation systems to create a blockchain-based history of provider trustworthiness. Tokens are also awarded to miners for validating transactions and storage proofs submitted by other miners.

Infrastructure

Emerging trends such as software defined networks, network function virtualization, and containerized application deployment using microservices architecture are all enablers to a distributed networking model relying on peer-to-peer transactions.

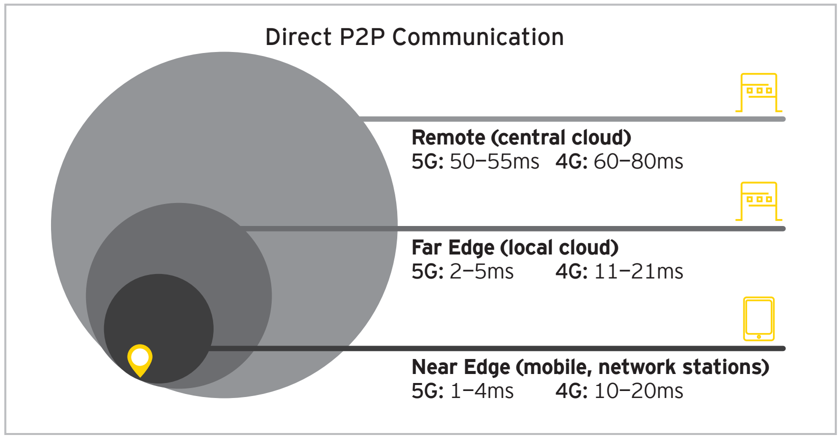

Source: Gartner

Networks

With advances in 5G technology, the increased speed capabilities will directly correlate to more data demand. This will drive the need for more sophisticated and agile networks. Moreover, with the rapid deployment of cloud infrastructure and services, traffic patterns are changing requiring the adaptation and reconfigurability of network infrastructure. By 2022, 30% of communication service providers that have 5G will also have deployed edge computing services.[13]

Telecom service providers are demanding scalable and cost-efficient network solutions to cater to escalating demand for data and higher levels of fixed-mobile convergence and to offset competition from over-the-top (OTT) competitors.

Containers

As the networking world is moving to a software-based environment for cloud-native applications, the concept of containers and microservices has only grown. This technology allows running applications and processes in software containers and enables them to operate as an isolated unit of application deployment. Containers improve overall developer experience, foster code and component reuse, and simplify operations for cloud-native applications.

Containers are highly effective in solving scalability issues. Containers can boot up in 0.20 seconds versus 10.3 seconds for a virtual machine and can be created far more quickly than virtual machines because they use the operating system kernel from the host's server.[14]

Decentralized compute

As connected devices at the edge will become more powerful and interlinked to each other, centralized data centers will be become less relevant.

Currently, around 10% of enterprise-generated data is created and processed outside a traditional centralized data center or cloud. By 2022, Gartner predicts this figure will reach 50%.

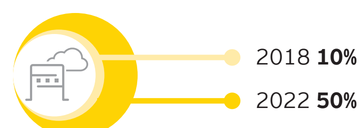

Source: Gartner, October 18, 2017

Edge devices are entry points into the a service providers network, often running specialized software stacks (for example, embedded operating systems and middleware) and performing data processing functions. Data processing at the edge can be partially or in some cases fully autonomous from core IT systems. But in a majority of current deployments, a subset or an elaborated portion of the data feeding into the edge device will make its way up to the core again.

As IoT proliferates, the number of edge devices will surge. The current centralized client-server architecture poses significant latency limitations and cannot enable the real-time insights from IoT applications. Edge-based IT environments — a distributed computing architecture in which information is processed closer to the content creators/users and speeds up collection/delivery — offer a solution. As more devices get connected, data exchange between the nodes or edge become essential. Edge servers distributed in the network reduce network traffic, lower operating costs and enhance application performance.

13 "Gartner Market Guide for CSP Edge Computing Solutions (Mar 2018)," Gartner website, https://www.gartner.com/doc/3869674, accessed 30 April 2018.

14 Kavita Agarwal, "A Study of Virtualization Overheads," thesis document, http://www.oscar.cs.stonybrook.edu/papers/files/KavitaAgarwalMSThesisSubmission.pdf, accessed 30 April 2018.

Edge devices can process data locally, can communicate with other devices directly and can share resources with other edge devices to minimize use of central cloud computing resources. This collaborative architecture is faster, more efficient and more scalable. A decentralized architecture is more private in nature, since it minimizes central trust entities, and is more cost-efficient because it leverages unused computing resources at the edge. As you get closer to the edge, the marginal cost to compute approaches zero.

Storage

There is an increased risk in storing large volumes of data on centralized databases in data centers or in the cloud. These single point of failure vulnerabilities have been exploited as evidenced by the increased frequency of large data breaches. Decentralized storage breaks up data into chunks, encrypts it and distributes it across various nodes in the network. This offers advantages by not relying on centralized storage or trust of a third party, whose systems could be vulnerable. There are additional benefits of decentralized storage, including lower operating cost, better availability and faster distribution through peer-to-peer networks. However, there are limitations that need to be solved for such as security which is proportional to the size of the decentralized network. Network technology and content delivery methods will have to improve to enable large amounts of data to reach different nodes.

Current storage solutions allow users to employ third-party service providers to store their data as well as safeguard their data's confidentiality, integrity and availability. There exist numerous policies and guidelines such as HIPAA, GDPR, etc., that organizations must adhere to while implementing storage solutions, but eventually these solutions are prone to human errors. To add to this, centralized solutions are always at the risk of their security getting compromised by unforeseen events or incidents.

It should be further noted that the centralized systems cater well to solving the issue of immense and ever-increasing data; however, they are expensive to use and consume a lot of energy. Two technologies that are solving the issue of storing increasingly large volumes of data are "swarming" and "sharding." They were designed for purposes other than peer-to-peer; however, they fit well with the technology.

Sharding and swarming

All decentralized storage networks secure client data safely through sharding, or erasure coding. Sharding protects a client's data in the event a singular provider goes offline. This is done by splitting client files into n number of shards and distributing them across n number of providers in the network. A client can specify a tolerance level of the amount of providers that could go offline and the remaining shards could be assembled together to recreate the original file. For example, a 10-of-30 erasure scheme would mean that if any 20 of the 30 providers go offline, the original files could be salvaged as 10 are still online. This means clients would only need to rely on 10 of their 30 providers to keep their data safe.

The collective storage of shards is accomplished by "Swarming." Decentralized storage uses large groups of nodes, or swarms, to store and manage data. This concept helps in reducing the latency in networks and increases speed by retrieving data in parallel from the nearest and fastest nodes. Given that a swarm comprises multiple, geographically dispersed nodes, scalability and reliability factors automatically increase.

Middleware

Middleware components will be similar to off-chain code components that are written in any language; executed within a secure, trusted container; and communicated with using secure channels. They can function as autonomous agents or bots, interacting with the world off the chain while maintaining the integrity of the blockchain. Open source technologies will serve as the building blocks for this ecosystem. Supporting open standards for protocol level implementations of P2P networking, consensus, database and virtual machines are key in establishing trust within the larger ecosystem and accelerates innovation.

A modular framework will allow consortiums to pick the best of breed components and build their distributed applications regardless of the detail underneath. Additionally, it will allow for the components to change barring any dependencies created above the core layer.

To allow for a blockchain-powered industry change, the blockchain at the platform layer should not be the only piece of distributed infrastructure. Multiple components of the middleware layer, including routing, published services, lookups, and DNS should also be distributed.

Collaborative networks

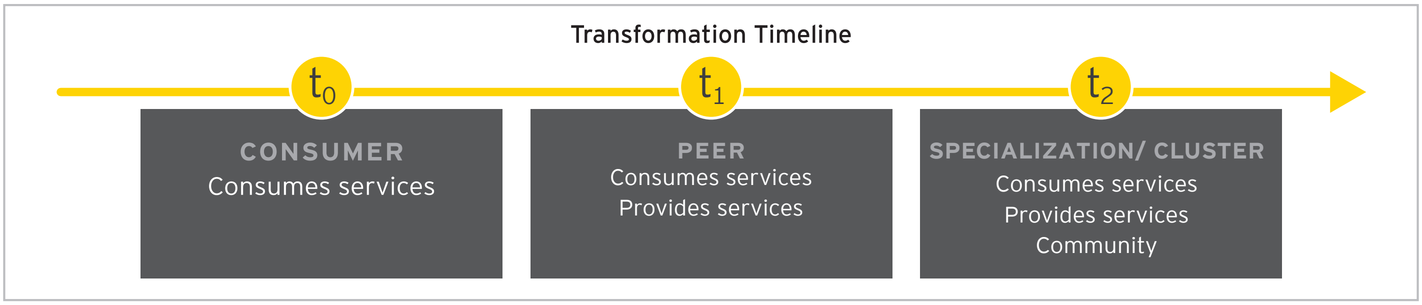

Several factors support the feasibility of fully developed, decentralized P2P collaborative networks reaching critical mass among large populations in the near future. These factors include a) growing social awareness of the dangers and vulnerabilities of keeping civilians' private data on large centralized servers, b) passive-income opportunities that will incentivize individuals to develop the necessary drive to participate, and c) quicker simplification of emerging technologies that will further enable participants to engage in collaborative networks.

In recent years, the US population has experienced breaches of user data on over 87 million people using social networks. As more and more of these cases occur and people personally feel the repercussions of data being housed on centralized servers, many will begin looking for more secure ways of storing their personal data. As the use cases of decentralized networks are realized, the mass population will see the value and the demand for collaborative networks will increase.[15]

Because these distributed networks are still in their infancy, this means that the majority of us will not have the curiosity or even the ability to participate yet, but this is normal for such early technology. The internet of the 1980s was vastly less intuitive to use than it is today. As technologies grow, the best designs and most simplistic user cases are uncovered, which plays an important role in any new product or technology reaching critical mass.

Currently there are two distinct roles within centralized networking: the consumer, who is a passive actor within the network; and providers, who include carriers and content providers. Within this ecosystem, the supply-side power created by these providers allows them to act as the gatekeepers of compute infrastructure while leaving consumers with little to no choice. However, the consumer role is evolving as emerging advancements in peer-to-peer distributed networks are enabling consumers to transact directly with one another by utilizing their own latent computing hardware.

In future P2P collaborative networks, "peers" of the network will be both consumers and providers of services operating in an efficient marketplace. To make this possible, a mentality shift is needed for consumers to realize their potential to profit from the optimization of resources and workloads. Once this occurs and a marketplace develops, the marketplace will autonomously match suitable consumers and providers based off of each party's price preference and size of computing and storage required.

From the peer model, participants will be able to decide if they want to be providers as their own agent, or form loose clusters and even participate in cluster governance to specialize in the servicing of some particular workloads.

15 "Facebook says Cambridge Analytica had data on 87 million people," c|net website, https://www.cnet.com/news/facebook-says-cambridge-analytica-had-data-on-87m-people/, accessed 30 April 2018.

Future state hypothesis

As demonstrated, peering requires a significant change to the view of the user's role. This section proposes an architectural model enabling the leasing of computing power and storage within a peering network. The industry is close to the technological tipping point where it will be possible to fractionalize compute tasks and have these tasks performed in parallel by multiple nodes of a distributed network. Not only could a distributed networking grid substantially lower the cost of computations required for complex applications, scientific calculation, and machine learning, but also high-powered computing capabilities could become substantially more accessible to all peers.

Architecture

Since the basis of a distributed network is about peers being both providers and consumers, architecting will be about aligning types of resources available to the workloads demanded. The structure will have to be smart enough to route workload requirements to resources that have the capabilities to deliver as contracted. The blockchain and other middleware components will act in the back end as the connecting fabric.

To have access to the services of the network, a consumer has to be connected to a peer node. The peer model, as discussed before, allows for both consumption of the network's resources and the servicing of the network through renting the peer node's excess capacity.

Each connected device will act as a node within the network; their node will log notable events such as payments and deposits as well as interactions with smart contracts on the network.

The workload can be either a compute or storage task. The commoditization of computing and storage capacity will require an awareness of the potential of an individual peer and of the overall network. This will allow for a marketplace with nearly complete information.

This self and network awareness includes the understanding of different devices' core competencies. Cars are outstanding tools for computing tasks. By 2016, the car industry had already provided self-driving supercomputers with the processing power of 150 MacBook Pros. Nowadays, not only is there high computing power in the vehicle's body, but also cars are connected to platforms, some of which are capable of processing up to 320 trillion operations per second. These platforms are oftentimes compatible with almost every deep learning framework. Finally, not only cars are smart, but they also generate energy. "Cars spend 95% of their time parked and only 5% of their time actually driving." It is these down times when a car can be operated as a peer node for complex computing tasks.[16] [17] [18]

It often goes unnoticed that network devices are numerous and everywhere. One carrier alone can have thousands of central offices, tens of thousands of base stations, and even several hundred thousand distribution points. Many of these physical infrastructure sites are prime candidates to perform compute tasks and operate as general purpose servers; however, they are less desirable tools for storage. They offer geographical dispersion, physical security, power plants (backup and generators) and are climate-controlled. Not only are network devices eligible candidates for compute but can also act as the hardware enablers to a distributed compute infrastructure. Network devices are great at data packet routing, which is one of their core tasks, but there is room for innovation, which would enable the better servicing of a distributed industry. Change is likely to happen either at the base station or central office layer, but would have to be further explored.

Smartphones have a use case of their own, one similar to edge computing. Smartphones could be a useful tool for social media applications, where once an application is open, its supporting device could connect with other devices to not only share direct messages, but also to store other users' messages. A smartphone could be a synchronization tool for applications. This will require the reconfiguration of content applications to utilize compute more efficiently and less bandwidth due to the decreased need for the processing to be performed somewhere else.

16 "Go, autonomous speed racer, go! NVIDIA DRIVE PX 2 to power world's first robotic motorsports competition," nvidia website, https://blogs.nvidia.com/blog/2016/04/05/roborace/, accessed 30 April 2018.

17 "Nvidia develops Pegasus supercomputer for self-driving cars," autoall website, https://autoall.com/news/nvidia-supercomputer-pegasus-self-driving-cars/, accessed 30 April 2018.

18 "Pay as you park: UCLA professor Donald Shoup inspires a passion for parking," UCLA website, http://shoup.bol.ucla.edu/PayAsYouPark.htm, accessed 30 April 2018.

How would it work?

A requestor would send a task to other peers; this task would be either compute or store. For this transaction to be executed, the requestor must first share the necessary metadata associated with the request (i.e., name, bytes, hash, and size). This information would then be used to match with the resources that would suit the workload requirements. Once one or multiple responders are selected, the query would be served, timestamps would be added to the response, adding a system of record to the requestor and responders.

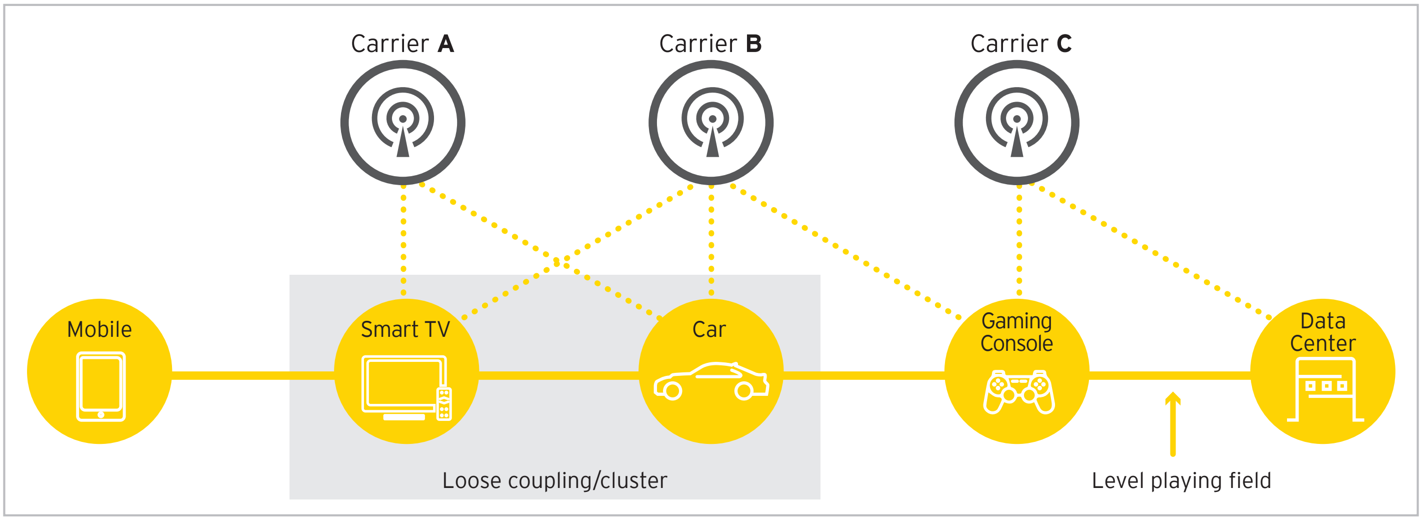

Nowadays for a phone to be useful, it has to be connected directly to a carrier. In this model, as shown in the diagram, a phone does not need to have a relationship with one carrier but can be serviced by a fellow peer node who has the connection with a carrier. Furthermore, every node is equal in voting rights and in functionalities (download, upload, storage). While there would be some difference based on operating a light node versus a full node, every full node is the same in what it is allowed to do. This means that for the first time, an industrial grade datacenter is a true peer to any connected device's computing or storage surplus and thus the enablement of a true device democracy. The power shifts from data centers being the growing oligopoly power of the network, to the peer node operated collectively by every participant of the network.

Payment

To work, this model has to be more cost-effective for the requestor than operating on a centralized model. Every indication points in that direction as current decentralized storage cost modeling is indicating a potential 37% to over 90% discount on the current cost of a gigabyte per month, with the cost of download bandwidth per gigabyte modeled to almost 99% cheaper!

While it seems that storage and compute tasks will be performed at a fraction of the current fees, it is important to consider that the network might face some busy time, and as a transparent marketplace, the suppliers will be incentivized to lease their extra computational power at a premium during peak hours. However, when the network traffic is low, the cost of compute could be free of premium and perhaps discounted. For instance, if an academic research project in artificial intelligence needs compute power for testing purposes, they may choose to optimize their consumption on downtimes; however, for an enterprise-level compute operation it may not be possible to compartmentalize consumption.

The network could be token agnostic for payment purposes and/or require a utility token native to the platform. The network could also use a tokenized representation of a nation-backed currency and do periodic settlement of total usage versus contribution balance.

An important consideration with the miniaturization of tasks into small containers or the distributions of small shards is that these shards/ containers may amount to higher total transaction fees related to distribution and payment than the value that would be created through the storage or computing task. Even though technology will allow for increasingly smaller shards and containers, there may still be tasks that will be better served when sent as relatively bigger containers or shards to a relatively smaller set of responders.

Quality

For compute activities to assure the highest quality of responses sent to requesters and to deter participants from being bad actors, peer nodes could be asked to provide collateral or to be part of a rating system. Important tasks that have to be processed swiftly could be executed at a premium by peer nodes with higher collateral (they would have more to lose), or by high-rated peer nodes. For the rating system to work, it requires that the network have a significant enough barrier to entry so that it is hard for a bad actor to set up thousands of bad nodes at a time, get bad reviews, abandon the nodes and reopen them again. The identity component will thus be important.

A diversity in the ranking of devices would probably be a sign of a healthy network, while some tasks will require a state-of-the-art high-performance computer from a college, other compute tasks would be better served on a general purpose connected device. For example, a microwave may not be the best solution device available. That does not mean it should not or cannot be used for smaller tasks. Therefore a rating can be based on making sure that the node is performing in the interest of the network, i.e., not sending to the requestor wrong answers, but can also be based on the computer's proficiency, such as GPU/CPU count.

Another way to mitigate the impact of bad actors is to use the median or majority answer from multiple peer nodes executing the same exact compute task. This can be also be done in combination of a rating and collateral system.

For storage tasks, retrievability is the main concern, combining sharding to new consensus mechanisms such as proof of replication, spacetime, retrievability, and storage, which should mitigate potential failures. However, when the use case is pushed to its extremes, some interesting dilemmas arise: in an open market, the cost of storage could fluctuate, making budgeting potentially tricky. This also means that someone could potentially be priced out from accessing their key personal information if the network becomes extremely busy.

Clusters

Clusters will form as peer nodes specialize into a workload. This means that a few cars and a network device (cell phone tower) could loosely couple for a set amount of time and act as a supercomputer that specializes in computing for a certain type of use cases — for instance, a use case that requires both the storage and compute of information with a frequent back and forth in those tasks. Once this loose coupling is no longer required, devices can either go back to acting as individual agents or act within a more organized structure; a cluster. This could then lead to a specialization of connected devices that even in their hardware and software design are modeled to perform better in synchronicity, for instance. Another option is for clusters to be organized through a central agent, who coordinates hardware supplies and upgrades for the network; in other words, if every node of a cluster is operating the same hardware and coordinates for upgrades, etc., the throughput and performance will be easily predictable. There is also room for trial and error to allow for further specialization of the hardware into lucrative tasks. Just like proof of work triggered the specialization of the hardware set to mine certain cryptocurrencies, processing pools could become what are currently called mining pools.

Sidechains

Sidechains are useful tools for transactions related to both workload and payments. They enable greater transaction throughput performance, they can be used to scale, and can create environments that are workload or application specific. In payments they lower the transaction costs and enable privacy of the payment channel.

Aside from the mainchain/sidechain relationship, it is possible and likely that multiple blockchains will operate on the network. Some may be public, some private, or mixed consensus models. A peer node could face the same dilemma that rideshare drivers face every day - "am I driving in one network or the other?" Drivers often make their decision based on profit metrics, business of the network, etc. To allow a peer node to operate in a similar way and benefit from diversity of the blockchain that could specialize in a workload, and devices, para-chains will be an important component of the interconnectedness of blockchain networks and protocols.

Node

Nodes process blocks; they take input, perform input operations, and then provide an output. A distributed network is a collection of nodes that are interlinked to one another. Nodes operate within parameters, which are the rules that they are bound by. The nodes should have architectural features that will allow for the sound distribution of cloud and compute power and storage of the ledger's transactions. First, each device should be uniquely identifiable, which is why a PKI based ID would be suggested. This will be an important aspect for rooting and allow for a competitive marketplace with easily identifiable and retrievable participants. Second, not every node will be able to operate as a full node or full copy of the ledger and may require a light node referencing a full node's copy. For instance, an average phone of 64 GB would have a hard time sustaining the full Ethereum ledger, which at the time of this paper is around 500 GB in size. Operating a light node would allow them to receive data from the network about the parts of the state that are of concern to them, and be sure that the data is correct. To do so the directed acyclic graph (DAG) will be an important part of the architecture. Finally the consensus mechanism will dictate how the node will connect with others. This will be an especially important architectural design as storage and compute protocols are experimental and combine many complex elements such as sharding, swarming, use of containers, etc.

Proposed Node Architecture

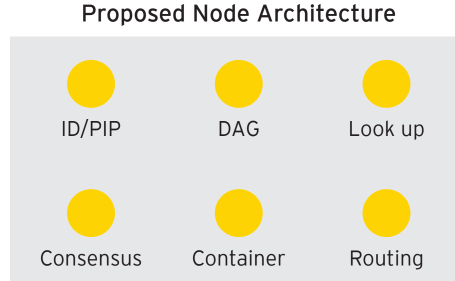

The node is the fabric on which applications run on. An example of an application that would be valuable to the network is a decentralized DNS. It would not only assure that providers' registries or records will be authenticated to be more human readable and facilitate mappings between updates, it will also avoid privacy and security issues that come with centralized DNS.

Feasibility

As with any emerging technology, challenges will continue to exist as the technology evolves. For the proposed future state, several factors will determine its viability, especially customer adoption, as we migrate to a model where consumers are also providers of compute and storage services in a sharing economy. To be part of such a network, the participants will not only need to recognize meaningful value from the distributed network,, but also the corresponding risks; can the participant be trusted and provide accurate results. Additionally, for such a network to perform efficiently, the devices will need to seamlessly negotiate and trade computational capacity with each other in real-time to achieve effective computational load distribution. Efficient matching of resources to appropriate workloads will need to occur to utilize and optimize the ecosystem effectively. Customer comfort in adopting this leap from traditional interaction models will determine the feasibility of large-scale value creation. And not all customers will have the same level of expectations as to quality of service for each workload.

Economics

The economics of distributed networking are influenced by disruptions to the traditional business model and overall marketplace. Gartner predicts that the value added by blockchain to global businesses will be $3.1T by 2030.

Traditional cloud providers offering compute and storage have been pushing the prices down in a race to zero. This race is really just a transition towards the commoditization of compute and storage services. However, due to the fact that these providers are so capital intensive, their ability to compete on price alone with distributed computing will be unlikely. Distributed storage providers, although newer to the space have aggressive pricing. Some, like Storj, have more room to reduce prices in the future by reducing the number of replicas.

As shown in the storage cost comparison example, the reduced cost impact of the movement to distributed networking is apparent. Storage services utilizing P2P and blockchain are already being offered in the market and demonstrating significant cost benefits and potential economic impacts.

Pricing would be even less with clusters which could be more efficient because the workloads could be designated to specialized resources. Incremental growth into distributed networking could occur within these specialized cluster use cases. Amazon started with S3, which was data storage, and TON, with just a messaging application. Starting with a very specialized solution that does an extremely good job and then expanding with services will be the most cost-effective approach.

Provider	$USD per GB/mo	$USD download per GB
Storj	.015	.050
Sia	.002	.001
Amazon S3	.023	.092
Google Cloud	.020	.110
MS Azure	.024	.087

The number of participants in a distributed storage network need to be high enough and generate enough volume to be lucrative for both consumers and providers. Along with scale, there are other limitations such as interconnection protocols, error handling, data consistency and security have to be overcome to create a high-quality architecture of the network. Adoption rates and the convergence of technologies in storage, security and computing will dictate the economic benefits to both customers and businesses.

Findings

As the connectivity-driven digital revolution encapsulates the world with emerging convergence trends in the networking and computing industry, communication and networking equipment providers will be pushed to adopt, adapt and transform in order to drive new growth-driven business models. With such high demand for connectivity and increased bandwidth quickly approaching, networking companies will reap the benefits of commoditization of networking devices. It becomes critical for networking companies to redesign their products to meet the new needs for smarter, faster and more specialized connectivity. For example, a traditional router that once was just a junction of basic network services would need to have additional routing features and blockchain interoperability. Routing protocols could be redesigned for dynamic evaluation of needs such as quality, bandwidth, and capacity. Their products could set the standards expected to be adopted by the rest of the market. Similar to the revolution IBM triggered 30 years ago by driving commoditization of PCs by promoting compatibility, thus making it the primary factor in microcomputer design going forward, the networking companies have the opportunity to set the standard for smart networking devices able to operate faster and with more specialized connectivity.

The networking industry also has a competitive edge due to its inherent interconnectedness. Existing cellular sites could be potentially repurposed for edge computing. Such a change would accelerate the provisioning of a large distributed network. Reconfiguring them to serve as micro data centers would not only provide additional compute and storage services, but also improve latency by bringing computing closer to those using edge devices for creating and consuming content.

And with the growth of 5G technology, which is already deploying edge computing services, the majority of enterprise-generated data will be created and processed at the edge. The advancement in 5G will only further the IoT paradigm and interaction between connected devices using peer-to-peer protocol. Efficiency in networking and edge computing will enable decentralized applications to perform at scale and propel blockchain-based ecosystems to emerge faster.

Network providers are uniquely positioned to enable an industry-wide disruption. Cloud, network and edge device market players provide a unique set of skills and experiences, as these three pillars are the foundational components to the enablement of the distributed infrastructure use case. The deep understanding of these building blocks will be a fundamental skill set for the architects of the future distributed network.

Any investments made and returns incurred are not at a level that would be prohibitive to innovations in distributed computing. Network providers have the opportunity to architect the network infrastructure of the future, which will be a driver of profit for both network and edge devices business, without impacting any potential cloud positions they may have. In fact, in the current cloud market oligopoly, even late entrants, could disrupt the oligopoly and provide a transparent marketplace to be better connected to potential demand.

Current infrastructure can support both upload and download, but the download activities on edge devices surpasses the upload ones. As the ratio of upload to download equalizes, or even shifts to more download, the current network devices will be challenged in offering the right level of bandwidth and will drive redesign to accommodate the demand.

Provider	Upstream		Downstream	
Rank	Application	Share	Application	Share
1	BitTorrent	29%	Netflix	37%
2	Netflix	7%	YouTube	18%
3	HTTP	6%	HTTP	18%
...				
6			BitTorrent	3%

Peers will not have the incentive to push their compute workload or storage tasks to centralized systems if they can solve or store themselves. On the other end, cloud providers are incentivized to drive workflow from the edge to center. This is why this innovation is unlikely to be led by the current beneficiaries of the cloud oligopoly and ought to be led by a network agent.

Blockchain innovation tends to work well when brought to market amongst competitors rather than competitors. The success of blockchain innovations relies on the network effect amongst different participant types, general users, start-ups, and larger organizations, to create an ecosystem. Startups are currently leading the way in distributed storage and computing. By working alongside the current ecosystem, there will be a lot of synergy where innovation is enabled by a larger player.

Conclusion

Jeff Bezos famously said early on in the history of Amazon.com that "your margin is my opportunity." At the time, he was talking about the profit margins enjoyed by traditional retailers and IT enterprises. In both cases, he used the internet to offer a new product at a lower cost. Amazon is not as profitable, on a unit basis, as the companies it replaced, but it has established a new benchmark in terms of both speed of growth and, increasingly, industrial scale. Though physics has laws that conserve mass in transformation, there is no guarantee in economics that profits are preserved through industry transformations.

Today, the combination of increased bandwidth, huge quantities of un-utilized and under-utilized compute and storage capacity and vastly improved IT application management tools may trigger another enormous transformation; this one swinging the pendulum back towards decentralized systems at the edges of the network and away from big data centers and their higher profit margins. Ironically, some of the key tools that will make this transformation possible and which may disrupt the new cloud leaders, would never have been invented without the need to manage large scale workloads across multiple clouds.

As with internet shopping and shared public cloud infrastructure, the disruptive new alternatives are not an apples for apples replacement, and developers and companies should not try to directly compare solutions. Efforts to copy and paste centralized solutions onto distributed architectures are likely to fail.

Distributed systems will outperform centralized ones for tasks that can be done locally, particularly in cases where connectivity is intermittent or slow. On the other hand, where work needs to be shared or integrated, the results may be quite a bit slower and the best new offerings will likely mix the two solutions to get optimal results.

All of which is to say that this revolution, though it is starting now, will take some time to bear fruit. Solution developers must learn a new way of thinking about and architecting for performance and reliability in a distributed environment. The payoff will be radically lower costs with, in many cases, better performance, but it will not come overnight.

authors

Paul Brody
Ernst & Young, LLP
Partner/Principal | Advisory, PI
paul.brody@ey.com

Suzon Crowell
Ernst & Young, LLP
Senior Manager | IT Advisory Services
suzon.e.crowell@ey.com

Melanie Marsollier
Ernst & Young, LLP
Manager | Advisory, PI
melanie.marsollier@ey.com

Jim (Jun) Xu
Futurewei Technologies | Huawei Technologies
Chief Architect
jun.xu@huawei.com

Fei Liu
Futurewei Technologies | Huawei Technologies
Director
fei.liu1@huawei.com

Huang Tao
Huawei Technologies
System Architect
huangtao@huawei.com

EYG no. 010486-18Gbl
CSG No. 1805-2679803.
ED None.