CBINSIGHTS

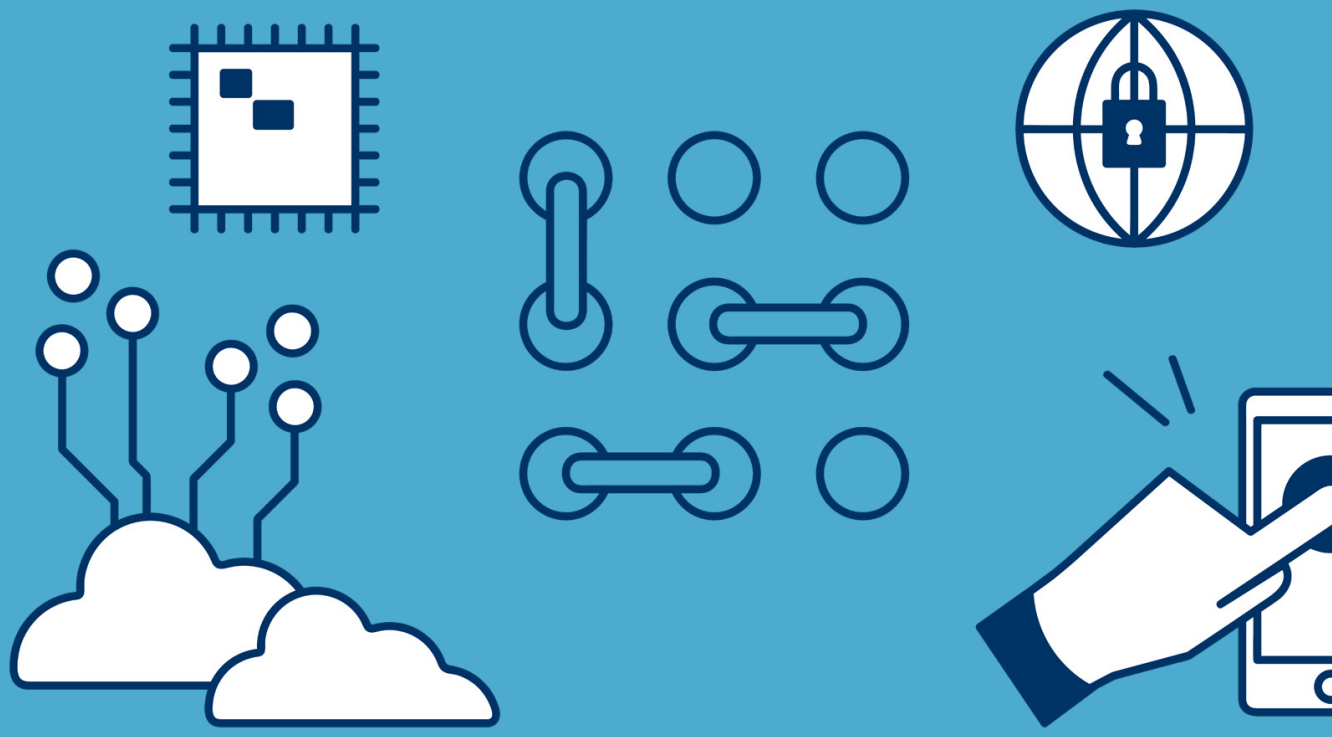

EMERGING TRENDS

What's Next In Enterprise IT

2019

Table of Contents

NExTT FRAMEWORK

Enterprise IT Trends to Watch in 2019

- Infrastructure
- Platforms
- Software

NExTT Trends

INDUSTRY ADOPTION: High / Low

TRANSITORY	NECESSARY
Trends seeing adoption but where there is uncertainty about market opportunity. As Transitory trends become more broadly understood, they may reveal additional opportunities and markets.	Trends which are seeing wide-spread industry and customer implementation / adoption and where market and applications are understood. For these trends, incumbents should have a clear, articulated strategy and initiatives.
EXPERIMENTAL	**THREATENING**
Conceptual or early-stage trends with few functional products and which have not seen widespread adoption. Experimental trends are already spurring early media interest and proof-of-concepts.	Large addressable market forecasts and notable investment activity. The trend has been embraced by early adopters and may be on the precipice of gaining widespread industry or customer adoption.

Low — MARKET STRENGTH — High

The NExTT framework's 2 dimensions:

INDUSTRY ADOPTION (y-axis)
Signals include:


momentum of startups in the space


media attention


customer adoption (partnerships, customer licensing deals)

MARKET STRENGTH (x-axis)
Signals include:


market sizing forecasts


quality and number of investors & capital


investments in R&D


earnings transcript commentary


competitive intensity


incumbent deal making

Necessary

SD-WAN

The widespread adoption of software-defined wide-area networking will support the proliferation of internet of things sensors and the deployment of the next generation of wireless systems.

Software-Defined Networking (SDN) is a modern network architecture that provides advantages over traditional architectures, such as reduced costs, increased bandwidth, greater security, and improved application performance. It is the ideal deployment for today's data-hungry and dynamic software deployments.

While the concept of SDN has been around for quite some time, adoption didn't take off until Facebook introduced the Open Compute Project (OCP) in 2011. The OCP promotes the redesign of hardware technologies to efficiently support the growing demands on today's compute infrastructure.

What's unique about SDN is the separation of the network planes — the application plane and the data plane. This separation makes the network programmable and therefore more flexible and customizable to an enterprise customer.

Software-Defined Wide-Area Networking (SD-WAN) soon followed the introduction of SDN. SD-WAN is simply an extension of SDN, and brings the programmability and flexibility of software-defined networking to wide-area networks.

In addition to flexibility, SD-WAN also provides greater network bandwidth, better network performance, and extended network reach (for remote, underserved locations).

It does so by by using the public internet to connect disparate locations of the network together (as opposed to a VPN). As a result, SD-WAN can use a variety of different internet service providers (ISPs) throughout the network, allowing it to achieve each of these benefits at low cost.

While there are certainly advantages to traditional SDNs, the benefits of SD-WAN are more suitable for today's evolving businesses.

For example, a multi-national bank may leverage SD-WAN to connect its disparate branches together, without having to worry about which ISP is used by each respective branch.

And although SD-WAN uses the public internet instead of a private network, all data is encrypted, end-to-end, which provides even greater security over traditional private networks.

SD-WAN has already seen significant adoption by enterprises over the past few years. This adoption is expected to accelerate in 2019 and beyond, pushed along especially by expanded IoT use and the introduction of 5G networks.

These networks will ensure security, reliability, and performance all while keeping costs low despite growing data consumption.

In November 2018, Oracle acquired SD-WAN company Talari Networks in a effort to ensure these advantages for its enterprise software customers.

The SD-WAN market is expected to reach over $8B by 2021.

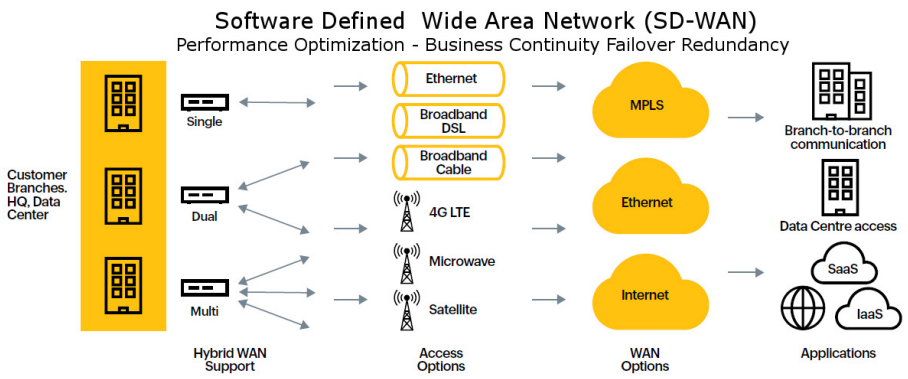

Source: Jules Bartow

TIME SERIES DATABASES

The growing amount of time-stamped data originating from internet of things devices is driving the revitalized interest in this decades old database system.

Since the digitization of financial data, there have been time series databases (TSDB). These databases provide the ideal format for storing time-stamped information and measuring changes in the data over time.

And while Wall Street continues to be a primary user of time series databases, many other industries are now adopting it as well.

Today, data is being collected continuously by every technological device. This has created a growing demand for time series databases that can store time-stamped information, which become a source of meaningful insight over time.

Enterprise technology, especially, benefits from the growing adoption of TSDBs.

According to estimates, the total data produced is expected to grow nearly fivefold to 175 zettabytes by 2025. Whether its data from sensors monitoring energy infrastructure in a remote environment or data on the performance of a software application running in a public cloud, a majority of it will be the time-stamped data types that work especially well with time series databases.

Though nearly fivefold growth in total data production won't necessarily lead to fivefold growth of time series database adoption, the growth of TSDBs is directly correlated to the growth in data.

Source: influxdata

SERVICE MESH

While microservices have improved software development, they have also introduced complexity. Service mesh technology can help manage and orchestrate these complicated deployments.

Microservices have become one of the most popular application development systems of choice in recent years.

Microservices operate as smaller, individual services (blocks of code) that connect together to form a comprehensive application. For example, a retailer's e-commerce app may feature a variety of microservices — one for the login authentication, another for the store locator service, etc. — working together.

These microservices live within containers, which are virtual receptacles used to store and run these blocks of code.

When an application needs to be updated, developers only need to edit the individual microservice within the individual container. There is no longer the need to update an entire, monolithic application.

But while the introduction of microservices and containers have improved software development, deployment, and run time, they have also introduced additional complexity — securing, monitoring, and connecting numerous microservices together can be difficult.

To reduce this complexity, a service mesh can be used to automate various aspects of container management and orchestration.

A service mesh acts as the programmable infrastructure layer used to support the microservices. It allows them to communicate in a flexible, reliable, and fast way.

A service mesh does so by separating the development processes from operational processes. Developers can deploy or modify containers individually, while system admins can monitor and secure containers, collectively — without having to redeploy or modify the individual contents of each container.

As a result, developers benefit from the flexibility of using microservices, while system admins are able to effectively manage the increasing complexity they bring. And their contributions work in harmony, rather than interfering with one another.

With the microservices market expected to reach $33B by 2023, adoption of the service mesh will grow alongside.

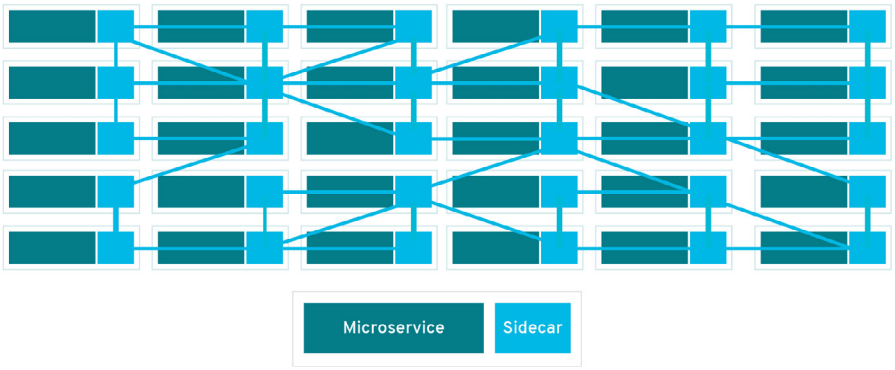

Source: Red Hat

Experimental

MULTI-CLOUD MANAGEMENT

As organizations transition to the cloud, many are using multiple providers. The growing adoption of multi-cloud strategies can provide businesses with greater flexibility, improved performance, and lower costs.

Businesses around the world are looking to adopt cloud services as a way to lower infrastructure costs, accelerate software deployment, and increase operational flexibility.

Many are already leveraging the products and services from market leaders like Amazon Web Services, Microsoft Azure, or Google Cloud Platform — which all offer comprehensive and reliable cloud services.

But with a growing number of trustworthy, third-party tools, it's becoming easier for businesses to migrate, manage, and monitor data across multiple cloud providers — not just one service.

The adoption and utilization of services from multiple providers is called a multi-cloud strategy. These strategies benefit organizations looking to avoid vendor lock-in (dependency on one provider), increase application reliability, reduce costs, and/or leverage the best services that each cloud provider has to offer.

And in addition to the use of multi-cloud management tools, the growing use of microservices and containers has also helped to facilitate multi-cloud adoption.

In June 2018, container management platform Docker announced its initial support of distributed applications across multi-cloud environments. Chief Product Officer at Docker Scott Johnston highlighted the need for a single platform that can manage distributed containers in varying environments:

"With an estimated 85% of today's enterprise IT organizations employing a multi-cloud strategy, it has become more critical that customers have a 'single pane of glass' for managing their entire application portfolio."

And companies like Snap have discussed their multi-cloud strategy on earnings calls. According to former CFO Drew Vollero, Snap's multi-cloud strategy has saved the company money:

"We've been able to moderate user cost growth through the successful execution of our multi-cloud strategy. Specifically, hosting costs per user dropped from $0.72 a year ago to $0.70 in the quarter."

And Snap isn't the only company adopting and discussing these types of strategies. Multi-cloud strategies have been a popular topic on recent earnings calls, according to CB Insights Earnings Transcripts tool. In Q3 2018, mentions of multi-cloud reached an all-time high of 99 mentions.

As adoption of microservices and containers continues to grow, so too will the adoption of multi-cloud strategies.

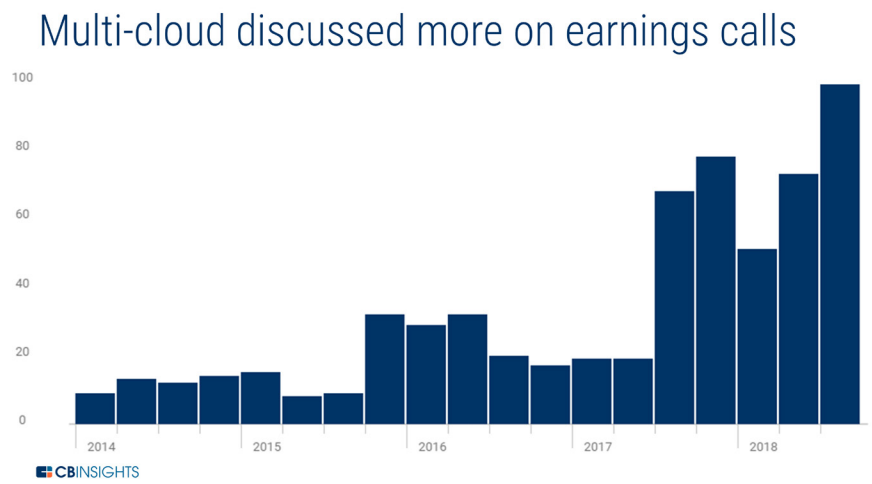

STATELESS HYPERCONVERGED INFRASTRUCTURE

This new type of computing infrastructure architecture makes hyperconverged infrastructure cheaper, while maintaining reliability.

Hyperconverged Infrastructure (HCI) is an IT architecture that combines the compute functionality and storage functionality of a server in a way that improves overall efficiency and limits many bottlenecks present in earlier forms of converged infrastructure (CI).

However, HCI also comes with its own set of limitations. By combining the compute and storage functionalities, HCI infrastructure becomes "stateful," meaning that each HCI component (or node) stores its own unique set of data.

For example certain files (like a Comma Separated Values, or .CSV, file) may be stored on one HCI node, but not another.

With HCI, if an individual node were to fail, or become inactive, so too would its unique set of data (for as long as that node remains inactive). As a result, that data would be inaccessible to other, active nodes in the system — even if that data was critical to the overall functioning of the system or the organization itself.

To limit this possibility, HCI nodes are duplicated for backup and recovery. In that same scenario, if an HCI node were to fail, the duplicate node would "spin-up," or start, to take the place of the failed node. The unique set of data would therefore remain available to other nodes in the system.

But, unsurprisingly, this infrastructure model can become expensive — especially for large enterprises running tens of thousands of primary and duplicate nodes. And every time an active HCI node makes changes to its data, so too does its duplicate.

As a result, a new form of hyperconverged infrastructure is emerging. While this architecture doesn't yet have an industry-accepted name, it can be thought of as "stateless," hyperconverged infrastructure.

By using a shared data server, there is no need for the nodes to store their own unique sets of data. Rather, they operate as "stateless" nodes.

These stateless nodes are no longer unique. If a node were to fail, a generic node would be able to take its place since there is now a shared source of data.

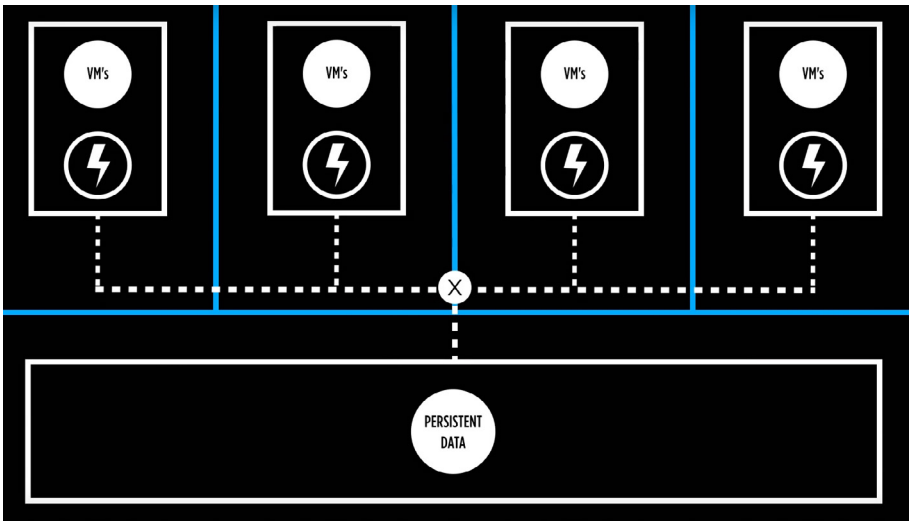

Source: Datrium via YouTube

This new form of HCI architecture still has the benefits of traditional hyperconverged infrastructure, while limiting the potential loss of unique data sets. It also allows for more predictable infrastructure performance and is ideal for hybrid-cloud environments.

Today, NetApp and Datrium, are two of the primary companies offering this form of stateless, hyperconverged architecture.

As flexible, hybrid-cloud computing remains a focus for enterprises, this new architecture is likely to become a primary offering of every other HCI provider in the market. The market for hyperconvereged infrastructure is expected to reach $7.5B by 2021.

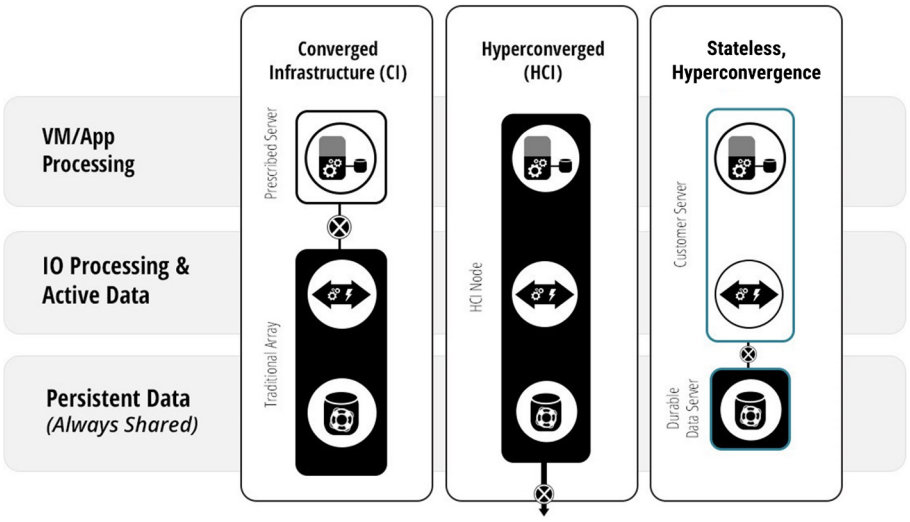

Source: My Virtual Cloud

ASICS & ASSPS

With over 20B internet of things devices expected to be deployed by 2020, new application-specific chips are providing a way to improve device performance and limit energy consumption.

Most computer chips and components are intended to do a variety of different tasks, in different ways, all at the same time. Both smartphones and personal computers are examples of general purpose devices that are designed to cater to a broad demographic of users with various technological needs.

But with all this flexibility and multi-management comes limitations and inefficiencies. This is where Application-Specific Integrated Circuits (ASICs) and Application-Specific Standard Products (ASSPs) are most valuable.

ASSPs are an iteration of ASICs that offer slightly more flexibility in terms of what they can do and where they can be used. However, both ASICs and ASSPs are ideal for consistent and continuous computing needs. Instead of managing many different tasks, ASICs and ASSPs complete one type of task, repeatedly, maximizing power and efficiency.

For example, ASICs or ASSPs are now used for cryptocurrency mining. Instead of using generic technologies, organizations are building custom ASICs and ASSPs to maximize mining output and limit energy consumption.

While the first deployments of ASICs date back to the 1980s, their popularity is rising again with the emergence of technologies like blockchain and the proliferation of IoT sensors. In both of these examples, ASICs and/or ASSPs may be more ideal computing components because they are optimized for completing the same processes again and again.

The ASICs market is expected to reach $32B by 2023.

In addition, using CB Insights Trends tool we found that Application-Specific Integrated Circuits received an all-time high number of press mentions in October 2018.

While current chip manufacturers may cater to the diverse needs of many different enterprise organizations, in the future there will likely be a number of emerging chip manufacturers that specialize in developing powerful and energy efficient chips that serve a very specific need.

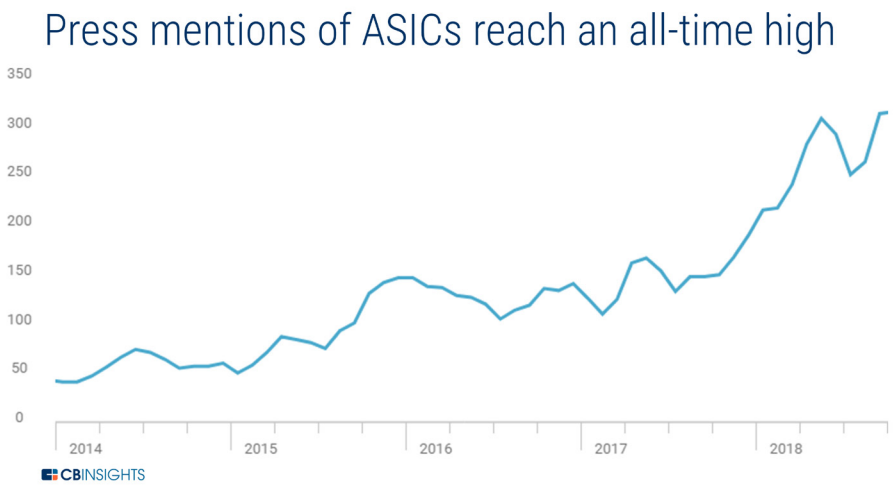

QUANTUM CLOUD COMPUTING

The distribution and reach of cloud computing will play an important role in bringing quantum computing to the masses.

Quantum computers leverage the basics of quantum mechanics to solve complex problems — the types of computational problems that incorporate a seemingly unlimited number of variables and that stump even the supercomputers of today.

The power of quantum computing offers potential solutions to a variety of optimization problems faced by many industries ranging from finance to genomics. These computers have already reinvented aspects of the cybersecurity industry with their ability to encrypt and decrypt nearly any form of electronic communications.

And some of the biggest players in tech — including Google, Microsoft, Intel, IBM, and Alibaba — are already exploring quantum computing, a sign that the next big computing race is already underway.

While a prominent leader has yet to emerge in the space, certain organizations are already offering commercially available quantum computing solutions. Some manufacture, deliver, and install these computers on-site for businesses, while others provide access to quantum computers with the cloud.

Amazon joked about the concept back on April Fools Day 2010, but IBM and Rigetti Computing already offer cloud-based quantum computing today. Others like Google, Alibaba, and Fujitsu are starting to experiment with cloud-based quantum computing as well.

aws

Blog Home | Category | Edition | Follow

AWS News Blog

Introducing QC2 – the Quantum Compute Cloud

by Jeff Barr | on 01 APR 2010 | in Amazon EC2, Compute | Permalink | Share

We've had more than our fair share of technical challenges along the way, but the time is right for me to talk about our newest product, the Quantum Compute Cloud, or QC2 for short.

This is the first production-ready quantum computer. You can use it to solve certain types of math and logic problems with breathtaking speed.

Ordinary computers use collections of bits to represent their state. Each bit is definitively 0 or 1, and the number of possible states is 2^n. 1 bit can be in either of 2 states, 2 bits can be in any one of 4 states, and so forth.

Quantum computers such as the QC2 use a more sophisticated data representation known as a qubit or quantum bit. Each qubit exists in all of it's possible states simultaneously, but the probability that a qubit can be in any of the states can change. Quantum computers work by manipulating the probability distribution of each state.

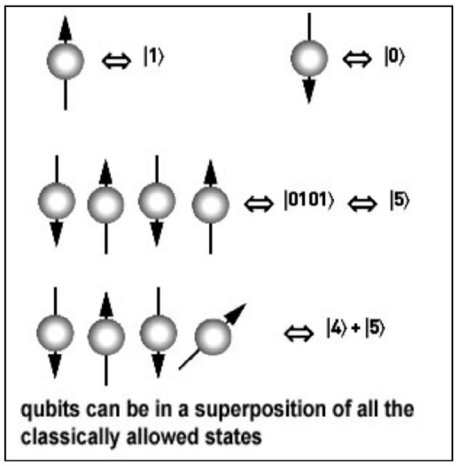

qubits can be in a superposition of all the classically allowed states

Quantum cloud computing seems to be emerging as one of the most promising distribution methods as it can provide access to nearly anyone, anywhere, and at low cost.

The quantum computing market could be valued at $493M by 2023, assuming a compound annual growth rate of 29%+ between 2017 and 2023.

Based on the current adoption trends between physical, on-premise infrastructure, and cloud computing, it's likely that a majority of this quantum computing will be delivered via the cloud.

But even if quantum cloud computing is just a fraction of the overall quantum industry, it will be big business for tech giants and startups alike.

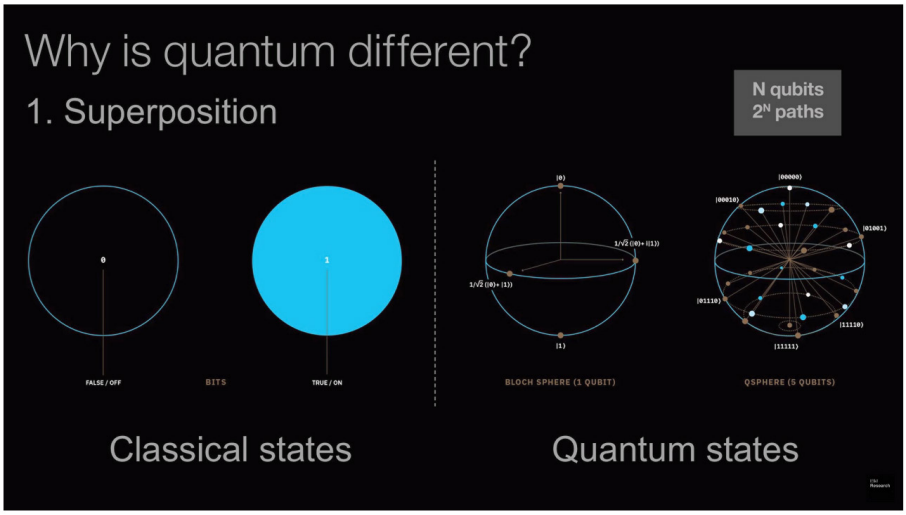

Source: Towards Data Science

Threatening

HTTP/3

This new internet protocol will help to improve speed and security for today's rapidly growing data streams.

The companies that control the majority of the world's web traffic are also in control of the development of future internet protocols.

This is the case for Google, which has the two most trafficked websites in the world, YouTube and Google, in addition to a number of the world's most popular mobile apps — including Gmail and Google Maps — as well as the world's most popular web browser, Google Chrome.

Google's most recent upgrade, due in 2019, is the third version of HTTP, known as HTTP/3.

For context, HTTP (Hypertext Transfer Protocol) is the foundation of data communication for the World Wide Web. It allows a web client, or web browser (Chrome, Safari, etc.), to communicate with a web server, where a website and its respective content is hosted.

HTTP/3, formerly known as HTTP-over-QUIC, provides a number of advantages over the current HTTP/2 network protocol.

With HTTP/3, connection establishment times will be reduced, as will bandwidth congestion. Latency will also be improved and connections can be migrated even as IP addresses dynamically change.

In addition, HTTP/3 uses what is referred to as multiplexing (or muxing, for short), which helps better utilize shared resources.

Traditionally, multiple connections were established to transport multiple packets of information between a user's internet browser and the hosting server. However, this has drastically negative impacts to bandwidth. With multiplexing, multiple packets are transported over a single connection, which drastically improves bandwidth.

In addition, Google designed Quick UDP Internet Connections (QUIC), which allows HTTP-over-QUIC (HTTP/3) to connect to a server, and encrypt that request, in a single transmission.

Again, this ensures a safe, reliable internet experience while also introducing drastic improvements to speed.

The release of HTTP/3 in 2019 is timed well with advancements in networking and telecommunication hardware (i.e. 5G). Together, a faster, more secure internet experience should soon be the norm for everyday users.

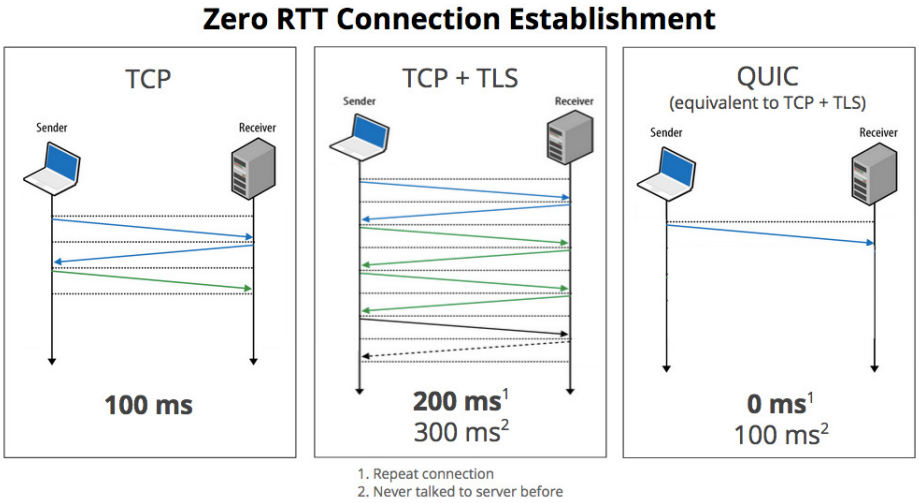

Source: TechSpot

AUGMENTED BUSINESS INTELLIGENCE

Businesses will soon be able to leverage machine learning solutions to autonomously synthesize relevant data, package it, and present it.

Artificial intelligence is infusing itself into all areas of the modern day enterprise. From the cloud to internal messaging services, AI has become a common way to improve operations and enhance a user's product experience.

One particularly interesting use case for artificial intelligence and machine learning technology is business intelligence. Despite the growing amount of unstructured data, artificial intelligence algorithms are now trained to find, visualize, and present relevant data to drive greater business insight.

With augmented business intelligence, employees will be able to ask questions related to high-level metrics or team-level performance. Then, AI will be used to find the relevant data, package it together, and present it in the most meaningful way.

And since AI models learn over time, they can also provide unprompted, but relevant, updates that may help to inform further business decisions. In short, AI will be able to figure out what type of information is most valuable to the business, how it should be presented, and when.

Microsoft is already experimenting with AI in a number of its business intelligence products. The company has implemented machine learning technology into its CRM and ERP solution, Dynamics 365, in an effort to help customers predict future business outcomes.

In addition, Microsoft has used artificial intelligence to analyze tens of millions of data points each day to detect security anomalies. These systems are far better at discovering such anomalies than human analysts. And by ingesting continuous streams of data, these AI security systems become more accurate over time.

The business intelligence market is expected to reach $136B by 2020. And while artificial intelligence within the enterprise is still in its earliest stages, it has the potential to infuse itself into the entire business intelligence industry.

It may not be long before real-time, relevant insights are autonomously delivered to us and in a format that effectively visualizes the data.

Enterprise AI startups raise over $1.6B in 2018

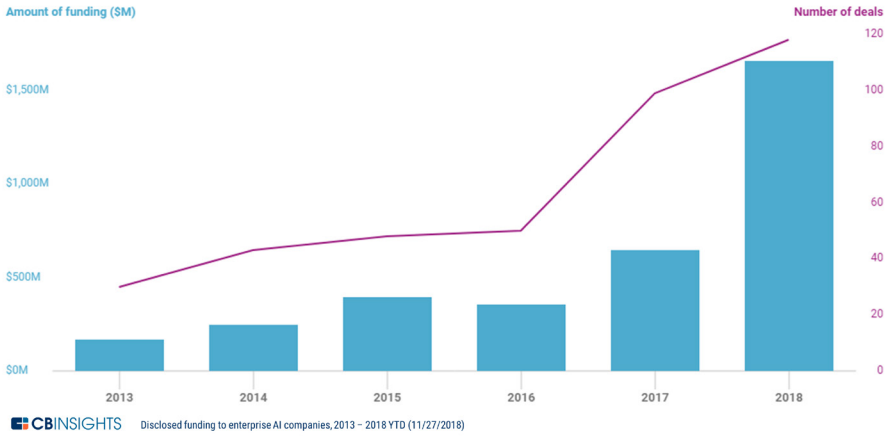

NANOSEGMENTION

This emerging security solution provides the distributed security benefits of microsegmentation, while offering the communication flexibility required by distributed microservices.

Microsegmentation is a method of distributed security for individual workloads within data centers and cloud deployments. While network firewalls are still used to monitor inbound and outbound traffic from a network, microsegmentation goes a step further by monitoring traffic as it moves throughout the network.

It does so by protecting individual virtual machines (individual computer operating systems), rather than blanketing the entire network or virtual local network (VLAN).

Microsegmentation helps to drastically reduce vulnerabilities and provides a far more more comprehensive approach to security than traditional methods.

However, the growing adoption of containers for hosting microservices is highlighting the limitations of the microsegmentation model. Unlike monolithic applications running on virtual machines, microservices running in containers need to communicate with each other.

With microsegmentation, that communication is limited or, at the very least, heavily scrutinized.

Fortunately, a new model, called nanosegmentation, is emerging to overcome these limitations. Nanosegmentation provides the distributed security benefits of microsegmentation, while offering the communication flexibility required to run microservices.

It does so by mapping the internal communications of individual microservices and whitelisting those relationships. If un-approved services attempt to communicate with the specific microservice, the connection is rejected.

The market for microsegmentation is projected to reach $2B+ by 2022. While there are no current market sizing estimates for the nanosegmentation market, the microservices market is expected to reach $33B by 2023.

If nanosegmentation provides better security solutions to just a fraction of the overall microservices market, it may soon outpace the popularity of microsegmentation.

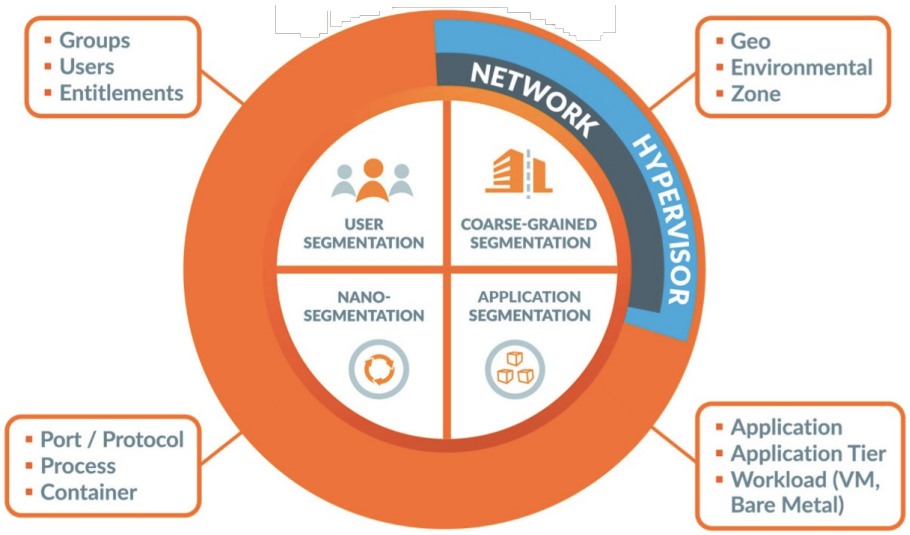

Source: Illumio

WORKFLOW AUTOMATION

As automation enters the white collar workplace, it has the potential to improve productivity, reduce operating costs, and displace employees.

Workflow Automation is playing a central role in streamlining workplace functions by automating manual, repetitive tasks such as data entry, document review, transaction processing, customer service, and more.

The category, also referred to as robotic process automation (RPA), has gained attention as it increasingly becomes the "gateway drug" for the automation of human cognition at large companies. And while not all workflow automation tools leverage artificial intelligence, many are beginning to do so.

It has also been a topic of conversation on recent earnings calls. Some businesses have discussed the opportunity in providing workflow automation services, while others have discussed its potential to reduce operating costs.

According to CB Insights Market Sizing tool, the Workflow Automation market is expected to reach $18.5B by 2023.

This could mean further disruption for white-collar jobs, including those in procurement, insurance underwriting, and the legal field. For businesses, eliminating human error is just as attractive as reducing costs.

And these systems are cheaper and far less prone to making mistakes. So while many roles may benefit from workflow automation technology, others may inevitably be replaced entirely.

Deloitte successfully implemented workflow automation technology at a major bank, by deploying 150 workflow automation bots to run 20 different processes. These systems were capable of reviewing 90K weekly requests and approximately 4M requests over an 18 month period.

In November 2018, enterprise tech giant SAP acquired workflow and robotic process automation company, Contextor. (*Contextor was featured in our research brief, The End Of White Collar Jobs? 9 Robotic Process Automation Companies Ripe For Acquisition.*)

With huge potential to transform the enterprise, we should expect to see more investment in and acquisition of workflow automation companies in 2019 and beyond.

Executives talk more about enterprise automation

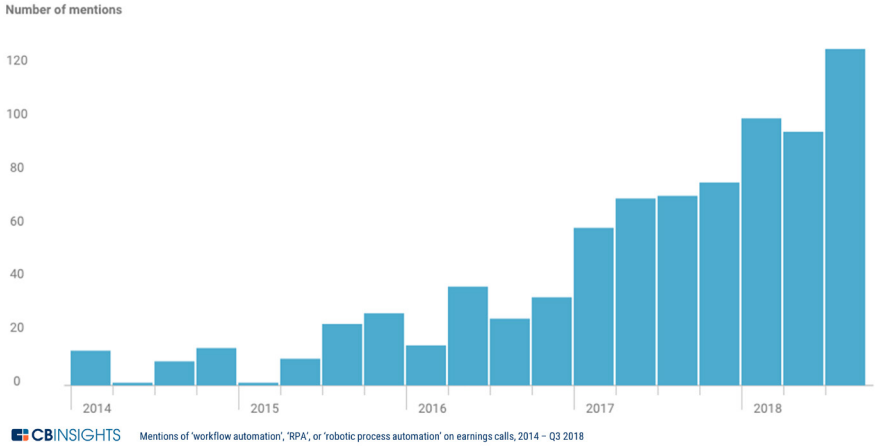

MODULAR EDGE DATA CENTERS

Deploying data centers at the far reaches of a network can help to improve processing speed and reliability for the world's distributed devices.

With the global proliferation of internet-connected devices, efficiency in data transmission and processing is becoming increasingly crucial.

While cloud computing has traditionally served as a reliable and cost effective means for connecting many of these devices to the internet, the continuous rise of IoT and mobile computing has put a strain on networking bandwidth.

Edge computing technology is now emerging to offer an alternative solution. It involves placing computing resources closer to where data originates (i.e. motors, pumps, generators, or other sensors) — or the "edge." This reduces the need to transfer data back and forth between centralized computing locations, such as the cloud.

Although nascent, the technology already provides a more efficient method of data processing for numerous use cases — including autonomous vehicles.

Tesla cars, for example, have powerful onboard computers which allow for low-latency data processing (in near real-time) collected by the vehicle's dozens of peripheral sensors. This provides the vehicle with the ability to make timely, autonomous driving decisions.

However, other edge technologies, such as wireless medical devices and sensors, lack the necessary compute capacity to process large streams of complex data directly. As a result, smaller, modular data centers are being deployed to provide hyper-local storage and processing capacity at the edge.

They are typically the size of a shipping container and are placed at the base of cell towers, or close to industrial facilities.

In addition to transportation and healthcare, these modular edge data centers are being used in industries such as manufacturing, agriculture, as well as energy & utilities. They are also helping mobile network operators (MNOs) deliver content faster to mobile subscribers. And many tech companies leverage these systems to store (or cache) content closer to their end users.

The global edge computing market is estimated to reach $34B by 2023.

From edge sensors to the centralized cloud

The edge computing ecosystem is comprised of four primary areas

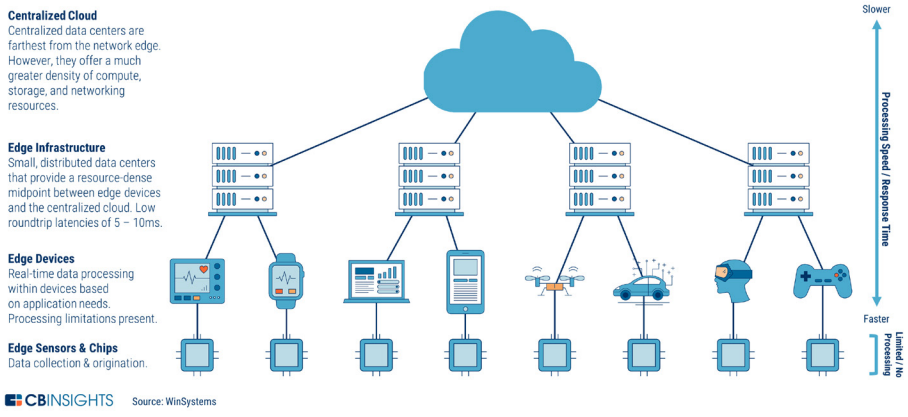

DATA LINEAGE

As datasets become larger and more complex, understanding their origins is a necessary first step in achieving compliant data governance.

Data lineage is the record of origin, movement, and modification of data over time. Mapping data lineage is one of the first, and most important steps, for organizations adopting broader data governance policies. It helps to ensure data integrity, while also providing organizations with greater understanding and control of the information they collect, store, and use.

Organizations around the world are mapping data lineage in response to specific laws like the European Union's General Data Protection Regulation. Even organizations operating outside the EU must comply with these regulations when using data that resides in or originates from the region. And as data regulation becomes more common, widespread mapping of data lineage is likely to follow.

Adopting these practices also helps to improve data quality, which, in turn, improves the insights delivered through data analysis. After all, the quality of business insights are only as good as the quality of the underlying data.

To establish a data mapping system, organizations must first formalize a process, then define the types of events that warrant a record, the specific information that will be recorded, and implement the appropriate sensors to scan an organization's system for these types of activities.

While data lineage can be useful for any organization, in any industry, the earliest adopters have been in industries that deal with highly sensitive data, such as the financial and healthcare industry.

For example, many healthcare organizations are using data lineage and governance solutions to ensure that they remain HIPAA compliant.

In addition, data lineage proves to be particularly valuable to organizations operating critical infrastructure, such as energy and utilities. With the growing number of sensors collecting data, it's particularly important for these organizations to know where their data is coming from.

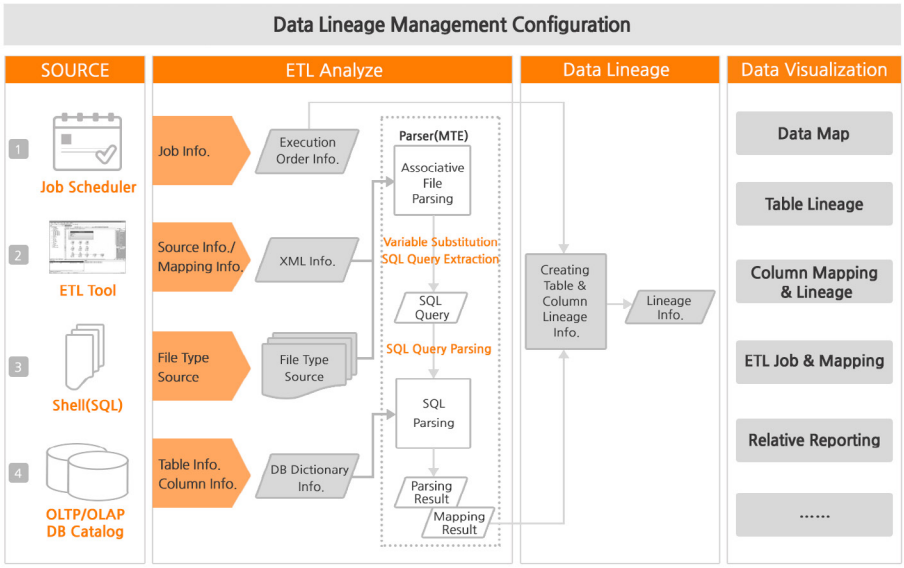

Source: Data Streams

Transitory

GRAPH DATABASES

Large datasets and complex data connections have brought about this new type of relational, NoSQL database.

The introduction of Web 2.0 companies such as Facebook, Amazon, and Google has brought with it large amounts of data and complex relationship structures.

While traditional, relational databases (structured as data tables) were able to store and manage large datasets, they failed to highlight many of the peripheral attributes surrounding the data itself — namely the relationships and connections between data.

Fortunately, graph databases have overcome certain shortcomings of past SQL database systems. For reference, SQL databases leverage the structured query language (SQL) for defining and manipulating data.

A graph database, on the other hand, does not leverage SQL and is referred to as a NoSQL database. The main advantage of NoSQL databases is that it provides a level of flexibility that is ideal for unstructured data.

In addition, NoSQL databases scale horizontally as opposed to vertically. This means that additional servers are added to the system to scale NoSQL databases, whereas compute power (CPU,RAM) must be added to scale a SQL database.

By scaling horizontally, on the other hand, the potential power and size of the system is unlimited as it is not constrained by a single machine.

This provides NoSQL databases with the ability to become larger and more powerful over time compared to their SQL database counterparts.

While the graph database is just one of many NoSQL database types, it has seen significant adoption in recent months. Since the release of Amazon's graph database in May 2018, organizations from all industries have been experimenting with this emerging database technology.

These databases are ideal for organizations building social networks, recommendation engines, or knowledge graphs. They also benefit any organization that values the relationships between data as much as the data itself. Use cases for these systems can be found in industries like life science, banking, and networking IT.

The NoSQL market is expected to reach $4.2B by 2020.

With the introduction of graph databases from a variety of tech companies like Amazon, Google, Microsoft, IBM, and Oracle, the graph database may soon become one of the primary NoSQL databases of choice.

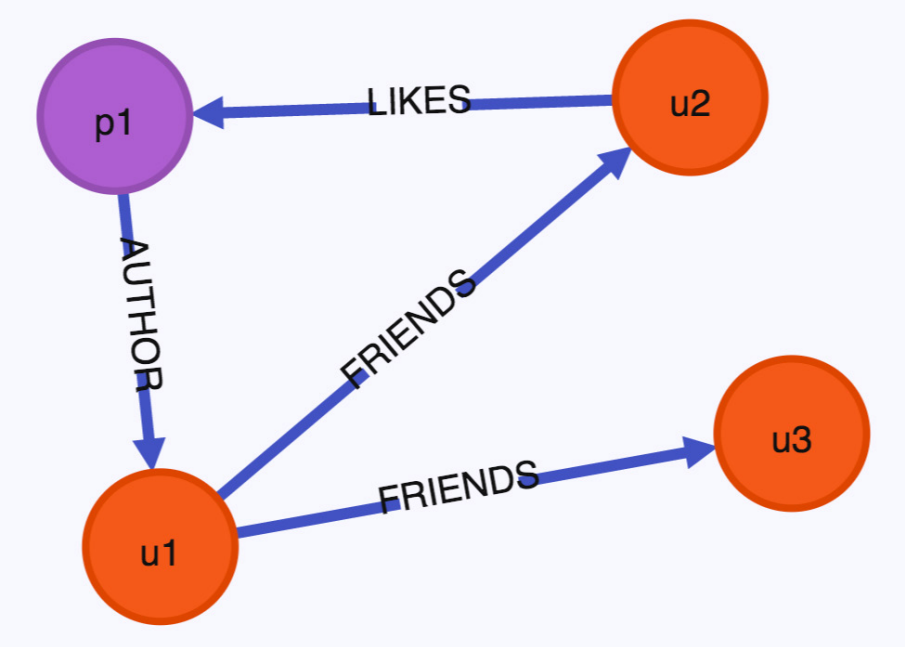

Source: Compose

SERVERLESS COMPUTING

This new cloud computing service allows developers to run small pieces of code to complete specific and repetitive requests.

Serverless is a new type of computing model that allows organizations to manage only the computing resources they need for each distinct action.

The new computing model removes the need for developers or system admins to manage the infrastructure resources offered by cloud providers. While serverless computing does, in fact, use servers to run applications, it removes the server management and capacity planning aspect of cloud computing. This gives serverless computing its name.

Serverless computing is already the fastest-growing public cloud service, growing from 12% adoption in early 2017 to 21% adoption in early 2018, according to a survey conducted by cloud management company RightScale.

While a number of other public cloud services have greater total adoption (such as database-as-a-service, or DBaaS), serverless is more popular than a variety of much-hyped emerging services including machine learning and IoT.

This is likely due to the fact that serverless can benefit all types of applications, no matter the industry or use case. It also improves the development process overall, not just the functionality experienced by the end user.

With serverless, developers and system admins only pay for the resources consumed by their application. This removes the need to pay for pre-defined computing capacity and storage. By adopting this new offering, businesses can further reduce costs and expedite application development in many instances.

A key aspect of serverless computing is what are called functions. Functions are small, independent pieces of code that complete specific requests.

For example, a business may develop a function to resize and reformat images as they are uploaded by users. As users upload high-resolution photos in unsupported formats, the function converts the images to a uniform format and size based on pre-determined specifications.

This creates a consistent experience and interface for the user on the front-end while potentially limiting storage costs for the business on the backend (as images are reformatted and resized by the function before they are stored).

Unlike virtual machines (VMs) and containers, functions in serverless computing do not operate indefinitely. Rather, they "spin up" to complete user requests or in response to certain events (like a user uploading an image), and spin down as soon as the task is complete — typically running for just a few seconds (or milliseconds) at a time.

For example, AWS' serverless product, Lambda, charges based on number and duration of requests (executed with functions). And this has introduced further modularity.

As we mentioned earlier, these functions run small pieces of code that complete specific requests. While functions don't replace monolithic applications or microservices, they do provide developers with more operational flexibility and cost savings when building, deploying, and running applications.

The management and deployment of functions is referred to as function-as-a-service (FaaS) and is synonymous with serverless computing. Cloud providers that offer serverless functionality are, thus, providing function-as-a-service.

The introduction of these services has promoted a broader interest in the technology over the last two years. In October 2018, there were a record high 100 media mentions of serverless or function-as-a-service. Growth in media interest is a direct result of the current adoption and expected growth of the industry.

The serverless computing industry is estimated to grow from approximately $1.88B in 2016 to $7.72B by 2021.

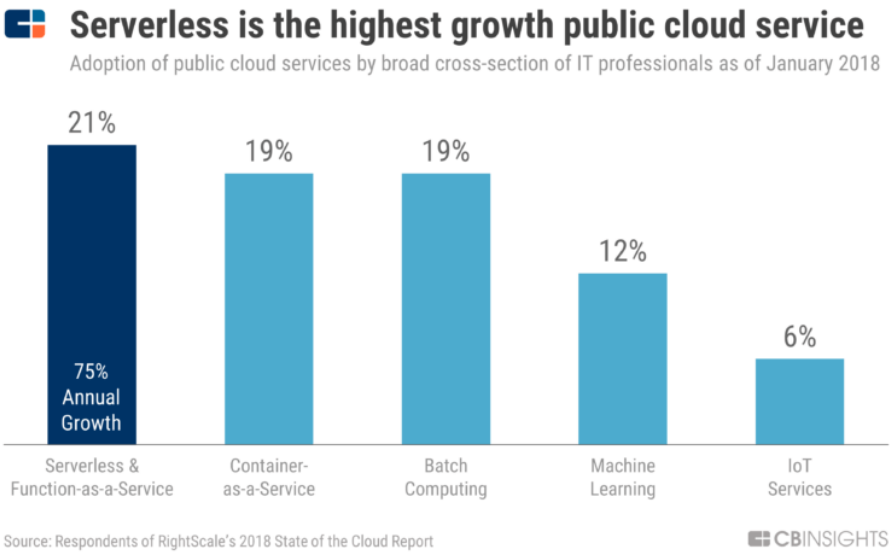

COLD STORAGE

Old forms of data storage, such as CD-ROMs and magnetic tapes, are gaining new life thanks to this method of long-term data preservation.

No business wants to find itself behind the curve in data collection. Most organizations would argue that it's better to collect as much data as possible today, even if they've yet to determine how it will be used tomorrow.

This type of unused data is referred to as dark data. This is data that is collected, processed, and stored, but not typically used for a specific purpose. According to estimates from IBM, approximately 90% of sensor data collected from internet of things devices is never used.

While there may be ways to reduce the overall quantity of dark data collected and stored, we have yet to find a solution that has had any meaningful impact. Even with advances in artificial intelligence and machine learning, businesses are still incentivized to collect and store as much data as possible to capitalize on future data-driven opportunities or avoid future pitfalls.

So for the foreseeable future, the best way to store this information is at the lowest possible cost. With cold storage, businesses are able to achieve just that.

Cold storage leverages many of the technologies of decades past, such as CD-Rs and magnetic tapes, to store data using as little energy as possible. Hot storage, on the other hand, leverages modern-day technologies, such as spinning hard disk drives and flash drives, to store data. However, these technologies consume far more energy as the data must remain readily available.

With cold storage, accessing data takes much longer than with hot storage. Only infrequently used data should be stored in these types of systems. With that said, the world's biggest tech companies like Facebook, Google, and Amazon all use cold storage as a way to store even the most granular user data points.

The cold storage market is expect to reach nearly $213B by 2025. And there are also a growing number of press mentions related to cold storage, according to CB Insights Trends tool.

In addition to major tech companies using cold storage to store infrequently used data, many cryptocurrency investors are using cold storage to store their digital assets.

Hot, warm and cold data

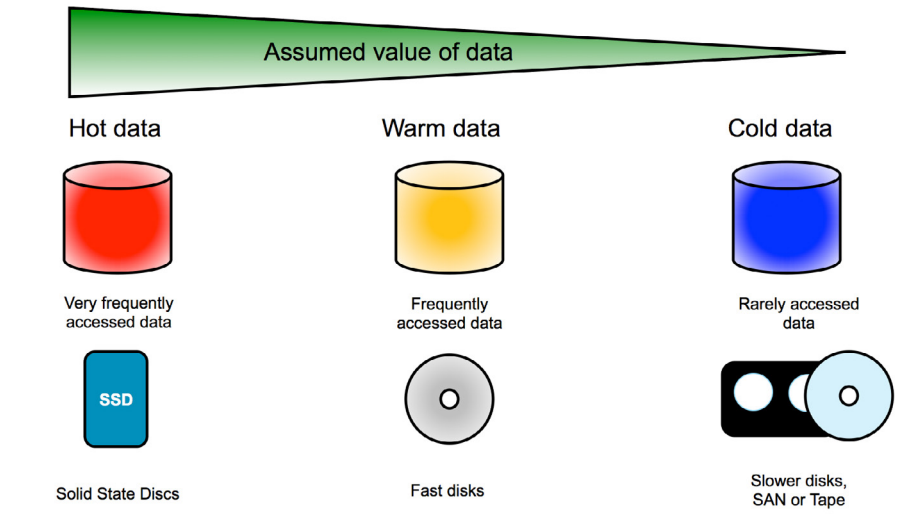

Source: IBM

WHERE IS ALL THIS DATA FROM?

The CB Insights platform has the underlying data included in this report

CLICK HERE TO SIGN UP FOR FREE