


# TRANSFORMATION IN THE 4<sup>TH</sup> INDUSTRIAL REVOLUTION



## Anthony Dutra

Presales Systems Engineer  
Dell EMC  
anthony.dutra@emc.com

## William Lupo

Technical Strategist  
Club Benchmarking LLC.  
williamdelmontlupo@gmail.com

## Varshith Hakkithimmanahalli Anilkumar

Software Engineer  
Boston University  
varshith@bu.edu

**Table of Contents**

- Introduction** ..... 3
- Shifting the Social Conscious** ..... 3
- The Fourth Industrial Revolution** ..... 6
- Digital Transformation** ..... 9
- Transformational Journey** ..... 11
  - IT Transformation.....11
  - Workforce Transformation.....12
  - Security Transformation .....12
- Democratization of Data**.....12
  - Challenge of Data Democratization.....13
  - Data Democratization Technologies.....14
  - Privacy in Trade - Cryptocurrencies .....15
- Transformations in the Cloud**.....18
  - Containers .....19
  - Micro Services .....20
- Going Server-Less**.....21
  - What is Serverless Architecture? .....21
  - What Economically Differentiates Serverless Architecture? .....23
  - What are the Limitations of Serverless Architectures? .....24
  - Cheddar: Serverless Architecture in Practice .....25
- Virtually Aware** .....28
  - Introduction .....29
  - Finite State Machines for individual components in gameplay.....30
  - Game intervals and data caching*.....32
  - NoSQL Key-Value Distributed data store* .....33
  - Synchronizer API* .....33
- On to Future Revolutions**.....34
- Appendix**.....35

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell EMC’s views, processes or methodologies.

## **Introduction**

A steadily marching force is digitally transforming the way we organize our lives. Starting from the time we wake up and continuing far past the time we go to bed, we are constantly connected through our mobile devices. Social Media and Big Data analytics enable us to better understand ourselves and each other, while Cloud technologies provide a platform for us to create almost anything as a service. As ostentatious as it may be to use such buzzwords to paint a grandiose portrait of a utopian machine-driven future, they are already cemented into the daily rhythm of our social fiber.

Then what's next? What technologies will then sow the seeds of our future society?

This article explores the new direction that technology is taking humanity. It starts by peering into the history of technologies' influence on humanity. This will include touching briefly on the first Industrial Revolution to how e-commerce has reshaped the way we do business.

Next, this article will delve into specific technologies and how they are shifting the social consciousness. It begins by exploring the notion of democratization of data through investigating the influence of technologies such as virtualization, Cloud, Big Data, the Internet of Things (IoT), and cryptocurrencies. Then, the article explores how Cloud technologies are driving the new ways we operate our businesses, both on the front and back end. This idea traverses through the emerging technology that is serverless architectures. The article discusses current business-use cases of serverless technologies and details an application that leverages a serverless architecture to better comprehend the ecosystem.

Finally, the article examines the world of virtual and augmented reality. Virtual and augmented reality, as we know it, is one of the most intense ways to deliver an immersive environment to users. Individuals experiencing these immersive environments exhibit psychological responses which mimic those in the real world. Information generated from these responses is extremely relevant and useful in creating awareness amongst many complex societal topics.

In total, the breadth of this article will offer an industry and social level view of a few technologies that are heavily influencing the direction of our society.

## **Shifting the Social Conscious**

This section will showcase the ability of technology to influence society abstractly. From history to culture to government and more, all play a significant role in the development of a nation or a group of people (such as a business), but it is technology that always seems to be a constant topic that tends to have implications that redefine social norms. It starts with a brief glimpse into the Industrial Revolution and how this movement created a familiar pattern between our relationship with each other, our relationship with technology, and the synergy between how we utilize technology in relation to the development of our society. This is followed by a brief analysis of the history of eCommerce and its technology-driven events that have spawned the vast market adoption of the digital web, mobile devices, and social media. Without these two

events, the interactions within our society – both for profit and not-for-profit – would be disparate from current day social norms.

Decades ago, it was common for people to live on or close to the land that provided their food. About 80% of the world population were farmers<sup>(1)</sup>, education was a privilege not a right, and no one ever used a single item outside of their own communities. This all changed around the 1700s, with many of the inventions starting in Britain. It began with the invention of the Flying Shuttle by John Kay in 1733 which dramatically increased the speed of weaving, which in turn created demand for yarn, which led to inventions like the Spinning Jenny and the water frame<sup>(1)</sup>. These processes were once mechanized using water power until the steam engine became cost efficient enough to widely adopt and improve the performance of these machines. It was the convenience of the technology and whatever energy source that was most cost effective and convenient at the time to utilize.



Flying Shuttle Advertisement<sup>(2)</sup>

It is increases in production such as this that brought about the mass adoption of machines and characterized by the user of new energy sources resulting in what is known today as The Industrial Revolution.

The Industrial Revolution was all about leveraging new forms of technology and energy to automate production. In England, the fact that large supplies of coal that were near the surface meant that it was cheap to mine, so it quickly replaced wood for heating purpose. With all this incentive to get more coal out of the ground, steam engines were invented to pump water out of the mines. Since those early steam engines were very inefficient, a cheap and abundant source of fuel in order to work was needed. As coal was abundant and readily available, steam engines used cheap British coal which kept the price of coal low. This, in turn, generated the opportunity for everything from railroads to steel, which like so much else in the Industrial Revolution,

created a positive feedback loop. Mainly, since people now had the wealth (both in knowledge and financially) to experiment with technology combined with cheap fuel cost, it became economically efficient for manufactures to look to machines as a way of lowering their production costs. These prices led directly to the Industrial Revolution by giving firms strong incentives to invent technologies that substituted capital and coal for labor.

The recipe of cheap energy and incentivized individuals and organizations blended with the technology of the time produces innovation that revolutionizes the human experience. Nothing exemplifies this growth more than our exponential population increase from that time up until more recent years. In 1700, the world population stood at 670 million<sup>(3)</sup>, by 2011 the world's population had reached 6.7 billion, a 10-fold increase in just 300 years. The world economy now is flourishing about 14-fold, the per capita income grew almost fourfold, and the use of energy expanded at least 13-fold. This kind of growth has never before occurred in human history<sup>(3)</sup>.

Decades later, a true blend of technology and human connectivity shifted the dynamics of our social fiber. As did the Industrial Revolution, this movement took the physical limitations of the way we communicate and acquire our needs. It reformed how we interact with each other and how we present ourselves.

eCommerce, starting as the B2B community sharing documents over the internet in the 1960s has grown to a \$279 billion online retail industry<sup>(4)</sup>. This movement of man and machine was born from technological developments that made the exchange of data electronically possible for the first time. The development of the Electronic Data Interchange (EDI) replaced traditional mailing and faxing of documents with a digital transfer of data from one computer to another<sup>(5)</sup>. This made transferring orders, creating invoices, and performing other daily business transactions easier and streamlined a data format that met the ANSI ASC X12, the predominant set of standards in North America<sup>(5)</sup>. Data could flow seamlessly through EDI, without any human intervention.

As improvements of the EDI system began to take shape in the late 1970s, a collaboration between manufacturers, retailers, servicers, and numerous other organizations began to take shape, developing Interorganizational systems (IOS). IOS allows the flow of information to be automated between organizations to reach a desired supply-chain management system, which enabled the development of competitive organizations<sup>(6)</sup>. This organization of management systems quickly bred competition between both business and consumers which was further fueled by the widespread adoption of the World Wide Web, also known as the Internet.

From the 1990s onward, eCommerce spawned technologies that would begin to mold our social fiber. The mass market adoption of broadband sped the world of the internet through Digital Subscriber Lines (DSL), cable, satellite, and fiber optics. The popular eCommerce blog *Statement* best describes the benefit of this adoption of broadband in the market: "...the introduction of broadband to people's homes didn't only mean that more people were buying online, it also meant that people were starting to conduct more research before placing an order. With easy access to the internet increasing, the average time spent researching products

and searching for more competitive prices also increased<sup>(4)</sup>. Broadband abetted the internet at a speed that best suited our demands, but there was still a need to feel secure online and soon this was possible through Secure Socket Layers (SSL).

SSL is an encryption certificate developed by Netscape in 1994 that provided a safe means to transmit data over the Internet<sup>(5)</sup>. A web browser could now identify if a website could be trusted with an authenticated SSL certificate, creating trust in a once seemingly complex and distrusted digital space. Secure transactions drove up the volume of online purchases to astronomical numbers. For example by the end of 1997, Dell became the first company to announce a single day sales record of \$1 million online<sup>(4)</sup>. Other key players emerging in this market were eCommerce giants such as Amazon, Google, and PayPal. These companies ingrained themselves into our society. Today we “Google” an answer to a question or expect to be able to purchase anything we desire through Amazon. Paypal changed the way we think of online payments by streamlining the shopping experience for its users. Instead of inputting customer and card details every time a consumer wants to make a purchase, all they need to do is checkout using Paypal and log in with their account. Plus, with the ease of using Paypal on a mobile device, the platform has lent a hand to the rise in mobile commerce<sup>(4)</sup>.

The increase in mobile phone users shifted the dynamics of shopping from behind the desk to on the smartphone. Couple that with the growing number of users on social media sites such as Facebook and Twitter, and the cycle of innovation and technological adoption continues.

Our once physical world has become a digital one, all started with a simple electronic interchanging of data. One technology gave rise to another and that to another and so on. Through this, businesses have transformed their service models to thrive in the digital age, some in response to societal digital adoption and others as leaders. However, it was those organizations who recognized eCommerce beginning to spill into the space of enterprise resource planning systems (ERP), data mining and data warehousing<sup>(5)</sup>, that are better positioned to foster success in the oncoming Fourth Industrial Revolution.

## **The Fourth Industrial Revolution**

Technology has evolved us. The way we live, work, and relate to one another has diverged exponentially from the norms of yesterday. The force leading the charge and changing the face of humanity is enormous and has been coined “The Fourth Industrial Revolution”. As Klaus Schwab, Founder and Executive Chairman, World Economic Forum Geneva warns of this immense movement in his blog post on the topic:

“In its scale, scope, and complexity, the transformation will be unlike anything humankind has experienced before. We do not yet know just how it will unfold, but one thing is clear: the response to it must be integrated and comprehensive, involving all stakeholders of the global polity, from the public and private sectors to academia and civil society<sup>(7)</sup>”.

He continues on to describe how this revolution differs from the previous:

“The First Industrial Revolution used water and steam power to mechanize production. The Second used electric power to create mass production. The Third used electronics and information technology to automate production. Now a Fourth Industrial Revolution is building on the Third, the digital revolution that has been occurring since the middle of the last century. It is characterized by a fusion of technologies that is blurring the lines between the physical, digital, and biological spheres<sup>(7)</sup>”.

What makes The Fourth Industrial Revolution significantly more diverse than its predecessors is the mere velocity of current breakthroughs in technology and how it just as quickly is having a direct and almost immediate impact on society. Klaus continues to write: “When compared with previous industrial revolutions, the Fourth is evolving at an exponential rather than a linear pace. Moreover, it is disrupting almost every industry in every country. And the breadth and depth of these changes herald the transformation of entire systems of production, management, and governance<sup>(7)</sup>”.

Understanding the challenges and opportunities of the Fourth Industrial Revolution is important to discovering how technology and which technology will propel humanity forward. Much of what we have to thank the Fourth Industrial Revolution for up to this point comes in the form of mobility: calling for a ride, ordering a pizza, etc. all can be done remotely through our mobile devices. These are mainly consumer-focused, but technological innovation such as virtual reality, cryptocurrencies, and Big Data could have explicit social impacts. Schwab writes about this transformation by stating “transportation and communication costs will drop, logistics and global supply chains will become more effective, and the cost of trade will diminish, all of which will open new markets and drive economic growth<sup>(7)</sup>”.

He continues by unearthing the difficulties in the current socio-economic environment as a result of the mobile and social media explosion, a concerning issue of the Fourth Industrial Revolution:

“This helps explain why so many workers are disillusioned and fearful that their own real incomes and those of their children will continue to stagnate. It also helps explain why middle classes around the world are increasingly experiencing a pervasive sense of dissatisfaction and unfairness...”

Discontent can also be fueled by the pervasiveness of digital technologies and the dynamics of information sharing typified by social media. More than 30 percent of the global population now uses social media platforms to connect, learn, and share information. In an ideal world, these interactions would provide an opportunity for cross-cultural understanding and cohesion. However, they can also create and propagate unrealistic expectations as to what constitutes success for an individual or a group, as well as offer opportunities for extreme ideas and ideologies to spread<sup>(7)</sup>”.

The final challenge of the Fourth Industrial Revolution is one that has plagued technology for decades, the topic of privacy. We instinctively understand why it is so essential, yet the tracking

and sharing of information about us is a crucial part of the new connectivity<sup>(7)</sup>. Debates about fundamental issues such as the impact on our inner lives of the loss of control over our data will only intensify in the years ahead<sup>(7)</sup>. This matter may find resolution in cryptocurrencies, an idea explored in detail further in a later module.

Challenges aside, opportunities of the Fourth Industrial Revolution are most favorable for businesses engaged heavily in technology.

Businesses recognize a relationship between society and technology. Innovators such as Uber, Fresco, and Mint have developed completely new business models through connectivity rather than market fundamentals. These million dollar moguls started with their IT budgets in the Cloud and came down to create new and interesting ways to provide consumers services.

We get a lot of information about industry transformers and innovators, but how does that relate to what an existing business does? Most organizations, like financial services and healthcare, are so used to doing what they already know that works. However, if they incrementally incorporate technologies into their business units, then they too can begin to embrace transforming in the digital age.

Schwab supports this observation by stating: “An underlying theme in my conversations with global CEOs and senior business executives is that the acceleration of innovation and the velocity of disruption are hard to comprehend or anticipate and that these drivers constitute a source of constant surprise, even for the best connected and most well informed. Indeed, across all industries, there is clear evidence that the technologies that underpin the Fourth Industrial Revolution are having a major impact on businesses<sup>(7)</sup>”.

Many trades are introducing new technologies that create entirely novel ways of serving existing needs and significantly shake the value chains of industry. Flexible, innovative competitors whose guise, access to global resources and platforms, marketing, sales, and distribution, can overthrow incumbents more rapidly than ever by improving the quality, speed, or price at which value is distributed. Demand of products and services has also shifted. Transparency, consumer engagement, and new patterns of consumer behavior have become staples of a trusted business. Technology-enabled platforms are an amalgam of these new demand and supply criteria to disrupt existing industry and economic structures, such as the “sharing” economy.

Schwab summarizes the effect of the Fourth Industrial Revolution on business in four main points:

“On the whole, there are four main effects that the Fourth Industrial Revolution has on business—on customer expectations, on product enhancement, on collaborative innovation, and on organizational forms. Whether consumers or businesses, customers are increasingly at the epicenter of the economy, which is all about improving how customers are served. Physical products and services, moreover, can now be enhanced with digital capabilities that increase



their value. New technologies make assets more durable and resilient, while data and analytics are transforming how they are maintained. A world of customer experiences, data-based services, and asset performance through analytics, meanwhile, requires new forms of collaboration, particularly given the speed at which innovation and disruption are taking place.

Overall, the inexorable shift from simple digitization (the Third Industrial Revolution) to innovation based on combinations of technologies (the Fourth Industrial Revolution) is forcing companies to reexamine the way they do business. The bottom line, however, is the same: business leaders and senior executives need to understand their changing environment, challenge the assumptions of their operating teams, and relentlessly and continuously innovate<sup>(7)</sup>. This transformation of businesses, in general labeled as the “Digital Transformation” is explored further in the succeeding section.

In summation, technologies of the Fourth Industrial Revolution influences every aspect of our lives. As time progresses the technologies of the Fourth Industrial Revolution will have its challenges but the opportunities it brings to business; a network of organizations who produce value that deeply influence society through their products and services, which in turn influences our society. This changes not only what we do but also who we are. The Fourth Industrial Revolution affects our identity and all the issues associated with it: our sense of privacy, our notions of ownership, our consumption patterns, the time we devote to work and leisure, and how we develop our careers, cultivate our skills, meet people, and nurture relationships<sup>(7)</sup>.

## **Digital Transformation**

How do businesses benefit from the changes being brought forth by the Fourth Industrial Revolution? The answer is through “Digital Transformation” – a profound and accelerating transformation of business activities, processes, competencies and models to fully leverage the changes and opportunities of digital technologies and their impact across society in a strategic and prioritized way<sup>(8)</sup>. The Digital Transformation, simply defined, is about people leveraging technology to make better products and services for one another. It’s about filtering our business fundamentals through the lens of technology to sell our consumers the most efficient products or services we can.

The world is dramatically shifting as a result of the Fourth Industrial Revolution. Many organizations recognize the change in the conversations of market dynamics and the prevalent trends. These can be summarized as:

- 1) Rapid growth in data and data services
- 2) Security is baked into every aspect of every layer of the business and its applications
- 3) Best in class user experiences whether it be from performance and responsiveness to ease-of-use, there is a focus on using technology as a tool that enables us to achieve goals better and faster
- 4) Consumers of this technology expect this data to be available, correlated, and provides added value to them.

In order to capitalize on these trends, organizations must digitally transform all activities of their business down to the business model. In the digital transformation, everything from functions to processes are as interconnected as possible. This means a complete removal of IT silos and agile integration of applications that have input across all business units. According to McKinsey, digital transformation can reshape an enterprise through interdependency of attributes ranging from customer experience to risk optimization because technology drives value in businesses in four ways: enhanced connectivity, automation of manual tasks, improved decision making, and product or service innovation<sup>(9)</sup>:



Expert interviews; McKinsey analysis<sup>(9)</sup>

Different technology-based organizations such as Dell Technologies, have developed their own interpretations and paths for organizations to embrace Digital Transformation. In order to obtain the synergy diagrammed above, businesses should begin to consider investing in three key competencies:

- **Derive real-time insights using Big Data & Analytics**

In the digital era, businesses need to focus on getting data-driven insights in real time. For this, they need the ability to store, manipulate, and analyze the data from a variety of existing and new, structured and unstructured data sources. Only then can businesses make real-time decisions that help tailor to the specific needs of their customers.

- **Enable software-centricity across your organization and products**  
In the digital world, staying ahead of the competition comes down to who has the best product/service. If your product's smarts come from software, you can continue to improve and iterate on it even after it's in the customers' hands. This means that in the constant race to lead the market, your advantage comes down to who has the best capability to develop and deploy software. This capability spans every industry – from manufactured goods like cell phones to automobiles, to online retailers and service providers.
- **Embrace IoT and instrumentation**  
Digital transformation relies on the ability of businesses to measure and quantify how their products interact with their customers and the environment. Getting information about product status, health and its usage can help organizations improve their products over time, identify problems sooner, and even deliver innovative business models on how products are paid for.

These competencies require time to develop. They are both labor- and cost-intensive to implement particularly at an enterprise level, but they are all part of a journey. Discussed briefly in the next section, the three areas of this Digital Transformation provide a baseline for an organization looking to adopt digital into their business. These three areas will also relate back to preceding technology sections to add perspective to their use cases.

## Transformational Journey

### IT Transformation

Many organizations that were not born in the digital age suffer with legacy systems and operating models that hinder achievement. Technology companies have developed a comprehensive and collaborative approach to IT Transformation: organizations need to transform and optimize IT by modernizing infrastructure, automating delivery of IT services and transforming the people and processes they have in place today. They can then begin to move from traditional IT models to IT as a Service (ITaaS) Cloud models, built on flexible, scale-out platforms that can deliver transformative value to their business.

Modernization is best achieved with flexible converged systems and hyper-converged infrastructures for a rock-solid, easily scalable foundation. This foundation allows easy migration to a cloud consumption model and the automation of IT services. Use next generation datacenter technologies to standardize, simplify and drive out cost of your legacy infrastructure. One example is to consider All Flash (i.e. solid state disks) options within data center technologies such as storage arrays to take advantage of technologies such as dynamic pools and compression which can increase the scalability and well as provide cost flexibility in the data center. Next is to automate, which means to create a software defined datacenter that is highly instrumented, automated and predictive. Ultimately, this enables an organization to completely transform their IT operations. They can now refocus their efforts to incorporating people and processes into their IT infrastructure planning. This enables the IT department to

operate much like a hybrid cloud and begin to develop that “as a service” feel. From there, the organization can move towards a multi-cloud model to deliver the most optimized cost models, as well as fully leveraging both on-premises and off-premises IT assets.

## **Workforce Transformation**

The way business is done has changed. As mentioned in previous sections, mobility has made it possible to work from anywhere, at any time. A 2017 New York Times article by Niraj Chokshi on the topic of working remotely reports that in 2016, 43% of employed Americans said they spent at least some time working remotely, according to a survey of 15,000 working adults<sup>(10)</sup>. Best practice in transforming an organization’s workforce dictates that employees be provided new ways of collaborating, accessing data, and meeting customer demands from anywhere with innovative devices, digital workspaces, and seamless experiences. Businesses should equip their technology users with more intuitive experiences, and anytime/anywhere access to applications and data by connecting users with innovative methods that enable more effective collaboration and provide access to critical customer and business insights wherever and whenever needed.

## **Security Transformation**

As the fundamentals of business change and the workforce becomes more remote, it stands to reason that the security posture must change. To collaborate, communicate, and innovate securely, cybersecurity needs to be transformed by implementing effective and efficient security strategies to stay safe and manage risk in the face of advanced cyber threats, and future digital challenges. Best Practices in this space include securing the business and its data so new technology is embraced without incremental risk to the business.

No matter what journey an organization takes on its path towards digital transformation, they all seek to take advantage of the technologies which are responsible for that shift. Throughout most of this document, technologies such as Big Data, IoT, and Cloud have been mentioned. The remainder of this document will explore these topics as they relate to societal change during the Fourth Industrial Revolution. However, obscure technologies, such as cryptocurrencies, “serverless” architectures, and virtual reality are introduced and presented as arguably the coming technologies that will propel an organization attempting to become a leader in the Fourth Industrial Revolution.

## **Democratization of Data**

Knowledge is power and data is knowledge in its rawest form. Organizations, more specifically analysts within an organization, have access to an abundance of data and the information created from that data constrains time and resources to analyze and visualize into knowledge. The Internet of Things (IoT) is one of the reasons behind this massive data growth. IoT is a broad concept that refers to the immense collection of data from connected devices (i.e. mobile phones, sensors, etc.). The amount of data being created is massive and the ability to analyze it properly has become cumbersome. A study by IBM found that key executives spend 70% of their time finding data and only 30% analyzing it<sup>(11)</sup>. This alarming statistic means that the

knowledge derived, if not properly evaluated through a proven analytical process (i.e. scientific methodologies followed by data scientists), could be utterly false.

Now that companies have access to more data, they have come to realize that the key to unlocking its value is in making it available to the people who need it most. At the same time, employees are increasingly asking for access to information which goes directly against how data governance within an organization usually works. As explained in the white paper “The democratization of data, how information can give power to your people”: In most companies, the IT department has long been the gatekeeper of BI and analytical tools, not because of a desire to control information but out of necessity<sup>(11)</sup>. The tools available were simply too complex for the average employee to use effectively, and analyzing data was a cumbersome and complicated process<sup>(11)</sup>. As a result, reporting was most often done to satisfy the needs of only a handful of top executives<sup>(11)</sup>. Generally, if a business user wants to go outside the scope of their data needs to work on an idea, they must justify use for the data and tend to fall victim to bureaucratic scrutiny. The data is held hostage.

One initiative that seeks to equalize the distribution of data (within an organization) without derogating the integrity of data that is brought forth by proper analysis methodologies is “data democratization”. The democratization of data is the process of expanding business information and the tools to analyze it out to a much broader audience than has traditionally had access<sup>(11)</sup>. This means putting data directly into the hands of business users and trusting that the analysis performed on the data produces knowledge that is correct. The ideal is to ensure every stakeholder will have access to all of the data with limits on sensitivity and privacy. The only thing they need analysts for is that which is beyond their capabilities.

### **Challenge of Data Democratization**

In an idealized state, the democratization of data within an organization seems to be a fantastic idea. However, this notion has its flaws. One inhibition to data democracy is the lack of checks and balances on the data. Given that anyone within an organization may have access to the data, there may be rogue actors who are motivated by political aspirations or personal gains within an organization to promote an agenda that doesn't help make the optimal decision for the betterment of the business.

Currently, most organizations have tasked their IT organizations to control their business critical data. Analysts within these organizations provide data integrity and confidence as this is what they are trained to do. The practice of data democratization is left to the data governance policies of a business seeking to adopt it. As long as strong ethical boundaries, training, and data governance are in place, a manageable form of data democratization can be implemented within any organization.

While it's still too soon to know the full impact of data democratization across enterprises, there's widespread belief that it will revolutionize our business dynamics by enabling employees at all levels to gain access to and insights from the data their organizations collect.

## Data Democratization Technologies

Businesses are actively seeking to make data democratization a reality, and this is why it is important to understand which technologies could assist in making this possible. It starts by creating a structure around the presentation and visualization of data so that so that it is not only open to all (or at least most) but simple enough for any contributor to perform an analysis through simple General User Interfaces (GUIs), dashboards, and so on. Bernard Marr, a best-selling author and keynote speaker on business, technology and Big Data, describes a few other technologies that make data democratization possible<sup>(12)</sup>:

- **Virtualization:** Data virtualization software makes it possible for an application to retrieve and manipulate data without knowing the technical details about it. This eliminates the need for labor-intensive processes.
- **Data federation software:** This software compiles metadata from a variety of data sources into a virtual database so it can be analyzed.
- **Cloud storage:** The adoption of cloud storage has been instrumental in breaking down data silos to create a central repository for data.
- **Self-service BI applications:** These provide non-technical users with tools that make data analysis easier to understand.

One technology missing from the list-which may surprise those who are familiar with Big Data technologies are data lakes, a technology that is synonymous with anything related to Big Data. A data lake is a storage repository that holds a vast amount of raw data in its native format, including structured, semi-structured, and unstructured data. The data structure and requirements are not defined until the data is needed<sup>(13)</sup>. This raw collection of data may seem like a simple way to provide access to individuals and promote the idea of data democratization. However, as Brian McKenna, a Business Applications Editor for *Computer Weekly* explains, this proves to be more of a hindrance to the idea than a benefit:

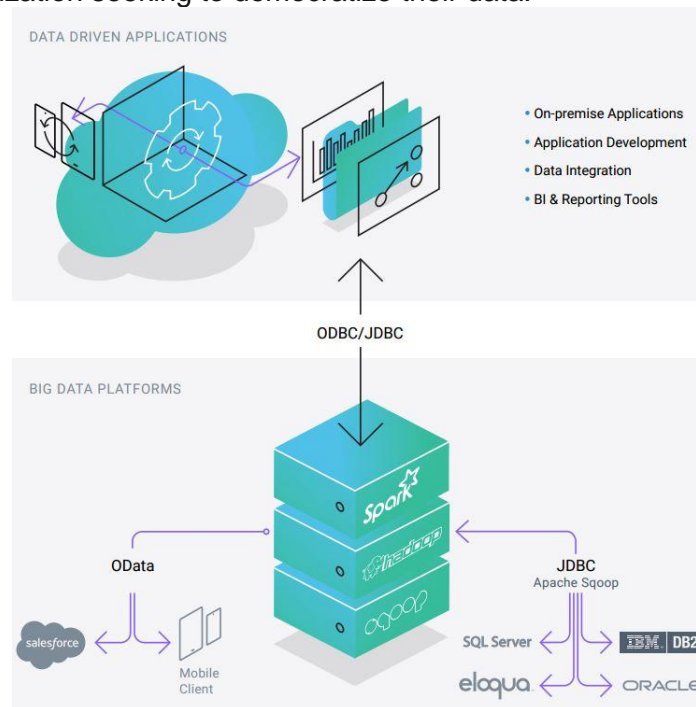
“Data is always structured by the applications that access it. So businesses would still have to develop ad-hoc pathways to the data, for each application... Unless you replace all of your separate applications by a single one that handles all functions, the data can’t reside in one place. Data lakes alone won’t democratize the data<sup>(14)</sup>.”

He continues this justification further by stating an important point:

“The only way to liberate the data is to leave it exactly where it is; in its existing, separate applications, but provide a layer of intelligence above the disparate sources that can integrate the data without replicating it. To set the data free, this intelligent layer would need to provide business users with an up-to-date, unified view of all the data in an organization, no matter the source<sup>(14)</sup>”.

In summary, organizations may look to other technologies or find a workaround to consolidate the data without necessarily needing to replicate all of it. Hybrid environments, for example, may prove to be a solution. Suzanne Rose, a contributor of the website *DZone* summarizes a study

from the IT analysts *451 Research* who state: “Hybrid IT embracing on-premises and hosted private cloud, along with public cloud, SaaS and existing client-server applications sitting on Cloud infrastructure, is the future of enterprise IT...The optimal way to approach democratization in hybrid environments is through industry-standard interfaces such as Open Database Connectivity (ODBC), Java Database Connectivity (JDBC)...<sup>(15)</sup>”. This strategy best blends the use of Big Data and Cloud technologies to produce a unified IT environment that can best serve an organization seeking to democratize their data:



**Hybrid Environment for Data Democratization<sup>(15)</sup>**

In this hybrid configuration, there is no need to maintain multiple application programming interfaces (APIs) or manage different versions of code. Everything in the environment is treated as a relational database and works with relational database-driven tools to deliver full access across numerous data sources in real time.

### Privacy in Trade - Cryptocurrencies

Privacy is key in the development of data democratization. Technologies such as Cloud, Big Data, and IoT provide the platforms to host these initiatives but are lacking when it comes to maintaining a sufficient level of security. One major stream of data that is most important to consumers and is most sought after by companies is transactional data. All of us have purchased something at one point, and that purchase – the item, the price, and the means of the exchange – has been recorded. This data is important to an individual as it can provide insight into their personal lives and could be used to make predictions as to what future purchases they may make.

For example, effective January 1<sup>st</sup> 2018, PayPal shared within its legal agreements a list of third parties with whom personal information may be shared to. This document lists organizations and reasons to share such information to a listed company such as Payment Processors like JPMorgan Chase Bank (UK, USA) to allow payment processing settlement services, and fraud checking, to Operational Service Providers such as Rackspace US, Inc. (USA) to provide hosting and storage services to assist and/or enable PayPal to provide services to customers<sup>(15)</sup>.

Individuals who utilize PayPal for payment services may find it unnerving to learn that the company they trust to handle their monetary transactions is willing to share their data with other financial organizations that may try to leverage it for their own gain. This is where the obscure technology of cryptocurrency to handle such transactions will play a beneficial role. As summarized from *What is Cryptocurrency: Everything You Need to Know*, cryptocurrency is a form of digital currency in which encryption techniques are used to regulate the generation of units of currency and verify the transfer of funds, operating freely from a central bank<sup>(16)</sup>. As is the case with many forms of currency worldwide (such as the Euro), cryptocurrency has no intrinsic value in that it is not redeemable nor is it backed by another commodity, such as any precious metal. The only apparent value generated is the one in which individuals are willing to trade real goods and services, and believe that others will do the same. It is not legal tender, and is not currently backed by any government or permissible entity.

Cryptocurrency has no physical form and the supply is generally not determined by a centralized entity –whether or not is based on the rules initially set up by the coin creator. The network structure of a cryptocurrency transaction is completely decentralized, as all transactions are performed by the users of its system utilizing a peer-to-peer architecture. Consumers decide for themselves what each unit of a certain cryptocurrency represents – whether it is a portion of a car or a kilowatt of electricity. They also decide the cryptocurrencies properties, as they have the ability to divide it into 100 million units, all of which are both independently classifiable and programmable. The term cryptocurrency is used because the technology is based on public-key cryptography, meaning that the communication and transactions are removed from the view of “Trusted Third Parties”. These third-parties are entities that facilitate, regulate, and approve financial transactions, such as governments, banks, accountants, and/or notaries. Third parties, such as banks or certification authorities (CA), are used in the electronic transfer of secure data. The third party uses cryptography and other security measure to authenticate the identity of the sender, the security of the data during transmission and to verify delivery to the intended recipient<sup>(17)</sup>.

Arriving full circle, herein lies the security and trust concerns previously mentioned as per the previous PayPal example. If companies such as J.P. Morgan Chase have the ability to mine consumers’ transactional data through Trusted Third Parties, then how can consumers be sure that the third party has the consumers’ best interest in mind?

Monero, a private cryptocurrency whose mission is to propel forward the idea of money as an entity completely decentralized from the fallacies of Trusted Third Parties is the perfect example of such a cryptocurrency.



Released in 2014, transactions with the Monero cryptocurrency allows users to do so anonymously. Unlike most of the digital coins, Monero uses a special technology called “ring signatures” which shuffles users’ public keys to eliminate the possibility to identify a particular user<sup>(18)</sup>. These ring signatures are a type of digital signature that can only be performed by individuals who have the keys. This ring signature makes it computationally infeasible to determine which of the group members’ keys was used to process the signature.

Compared to an ordinary digital signature where users would both need to sign and verify with public and private keys, a ring signature makes use of your account keys and a number of public keys (also known as outputs) pulled from the blockchain (a network of compute nodes all maintaining a single version of transactional “truths” on a ledger) using a triangular distribution method<sup>(19)</sup>. Over time, past outputs could be used multiple times to form possible signer participants. In a “ring” of possible signers, all ring members are equal and valid. There is no way an outside observer can tell which of the possible signers in a signature group belongs to your account. So, ring signatures ensure that transaction outputs are untraceable<sup>(19)</sup>. Moreover, there are no fungibility issues with Monero given that every transaction output has plausible deniability (e.g. the network cannot tell which outputs are spent or unspent) <sup>(19)</sup>.

## Ordinary signature



## Ring signature



Image adapted from [Bitcoins anonymity by Mai-Hsuan Chai](#)<sup>(20)</sup>

An example from the Monero website best showcases this functionality: “For instance, a ring signature could be used to provide an anonymous signature from “a high-ranking White House official”, without revealing which official signed the message. Ring signatures are right for this application because the anonymity of a ring signature cannot be revoked, and because the group for a ring signature can be improvised (requires no prior setup)<sup>(19)</sup>”. However, this inability

to trace the transaction isn't enough to ensure anonymity. It doesn't protect a receiver from defining a user's balance through examining inbound messages to the user's public address.

Therefore, Monero employs a specific protocol which generates multiple unique one-time addresses that can only be linked by the payment receiver and are unfeasible to be revealed through blockchain analysis<sup>(19)</sup>. This protocol is called the "stealth address", a requirement that the sender must employ to produce a random, one-time use address for every transaction on behalf of the recipient. As summarized by the Monero website: When a Monero account is created, the user is given a private view key, a private spend key, and a Public Address<sup>(21)</sup>. The receiver uses the one address for this transaction, but has the remainder of the payments routed to their unique address on the blockchain – only two users transacting can govern where a payment is to be directed<sup>(21)</sup>. The spend key is used to send payments, the view key is used to display incoming transactions destined for a user's account, and the Public Address is for receiving payments<sup>(21)</sup>. Both the spend key and view key are used to build a Monero address. Users can have a "watch only" wallet that only uses the view key. This feature can be used for accounting or auditing purposes but is currently unreliable due to the inability to track outgoing transactions. A user can decide who can see their Monero balance by sharing a view key<sup>(21)</sup>.

Finally, the Monero cryptocurrency makes certain security and privacy between two users in a transaction by being analysis resistant. No outsider can analyze the transaction as a modified version of the Diffie-Hellman exchange protocol – a protocol that generates multiple one-time public addresses – can only be simply gathered by the message receiver<sup>(19)</sup>.

Privacy in the digital age is a necessity for the advancement of society. The security of our transactional data is a step toward data democratization. Consider this technology being implemented at the National level – each individual citizen can make a purchase just as they do today (whether it'd be online or in a store), but the data they generate is completely their rightful property. If they so choose to sell this data, they may generate a basic form of income.

In summary, The Fourth Industrial Revolution is living up to its name. Ideas such as data democratization and basic income through the privatization of transactional data sparks movements throughout society. The uses of technology have moved far beyond making our lives easier; they are changing the very thread by which our society is woven to create a dynamically democratic society.

## **Transformations in the Cloud**

Aside from its role in the democratization of data, Cloud creates other interesting opportunities that change the dynamics of IT and business. Revert back to the lesson learned from the Industrial Revolution, a cheap, abundant resource mixed with a strong purchasing power breeds technological innovation which fulfills a need of ours.

This is exactly what makes virtualization the bridge into the Cloud model of IT. Virtualization is an abstraction and consolidation of compute, network, and/or storage resources. The

abstraction and management of these resources meant that no longer were IT organizations bound to the capabilities physical infrastructure. As long as the cluster of commodity servers could do their job effectively, the underlying infrastructure really didn't matter much. We were given the ability to consolidate and manage our once complex, monolithic IT environments into just a few clicks on the screen.

This is the essence of Cloud computing – a mixture of business and architectural models which makes this way of computing the most sought after destination for so many C-level executives today.

To a business, moving to the Cloud may be considered similar to provisioning an IT environment built on an economic model. These Cloud services fall under three Service models, explained as:

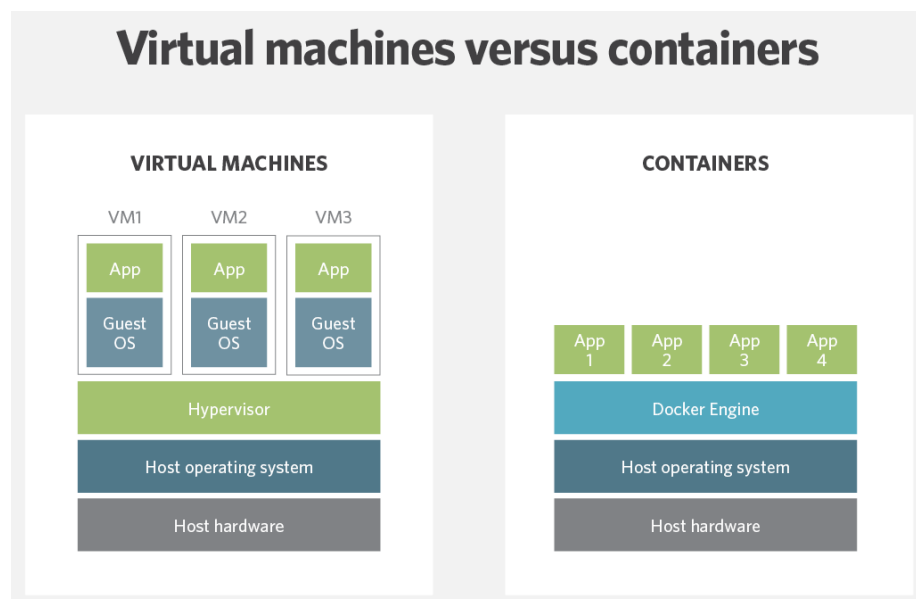
- 1) **Infrastructure as a Service (IaaS)**: The end user is purchasing the ability to provision processing, storage, networks, and other essential computing resources in order to deploy their software, operating systems (OS), some networking capabilities, and applications of which they have full control.
- 2) **Platform as a Service (PaaS)**: The end user is obtaining the space to deploy their created or procured applications utilizing the programming languages, libraries, services, and tools that are supported by the Cloud provider only. The end user has no management insight or control to the underlying cloud infrastructure yet has free rein over the deployed applications.
- 3) **Software as a Service (SaaS)**: The end user has no line of sight into any of the Cloud infrastructure (including network, servers, operating systems, storage) and has limited access to the application capabilities itself. In this service model, the end user is given the right to run the Cloud provider's application(s) on the Cloud infrastructure. The application(s) are easily accessible from multiple end user devices.

These on-demand, scalable, usage-based pricing models and optimized capital costs have impacted IT environments within businesses across the world. However, the benefits from moving to the Cloud isn't complete once a business has lifted and shifted all of its infrastructure to the Cloud. It requires experimentation with Cloud-driven technologies such as containers, micro services, and a new technology that will change the very nature of businesses – serverless technology.

## Containers

Much of application development has been based on monolithic lines of code that were deployed to a single or cluster of servers. Today, applications are constantly being developed and are required to be deployed often on a multitude of servers. This capability is possible through containerization. Containerization is an operating system- (OS) level virtualization method for deploying and running distributed applications without spinning up an entire virtual machine (VM) for each application<sup>(22)</sup>.

Instead, multiple isolated systems are run on a single control host and access a single kernel – the central component for managing resources between the hardware and the software. The containers hold the components such as files, environment variables, and libraries necessary to run the desired software and run only the required resources<sup>(22)</sup>. In layman’s terms, a container is pieces of an application in a box that include the runtime components – such as files, environment variables and libraries – necessary to run the desired software<sup>(23)</sup> consuming fewer resources. The complete set of information to execute in a container is the image and the container engine deploys these images on hosts<sup>(23)</sup>. This container consumes less memory, CPU, and storage as they do not have the same overhead as virtual machines.



Adapted from Tech Target<sup>(23)</sup>

As long as the container runs on identical servers across systems, a container can run on any system and in any Cloud without requiring code changes and without any guest OS environment variables or library dependencies to manage<sup>(23)</sup>.

### Micro Services

Multiple sets of containers can be compiled into a modular packages to create micro services. Each module supports a specific business goal and uses a simple, well-defined interface to communicate with other sets of services, usually through mechanisms such as APIs or HTTPS<sup>(24)</sup>. Once an application is split into multiple components, or micro services, it enables flexibility to deploy each one separately on completely different infrastructures, if that’s what’s best for the needs of the business.

Recently, containers and micro services are becoming highly adopted within Cloud-based developer environments to leverage that elasticity of these loosely coupled components. According to *451 Research*, container technology such as Docker and Kubernetes will see the fastest growth compared to other cloud-enabling technologies, with an estimated compounded

annual growth rate of 40% through 2020<sup>(25)</sup>. The speed and safety of deploying applications as micro services and containers helps lessen the stress of the overall developer process.

There is a technology, however, that abstracts application development even further and it is completely removing the need for a server.

## Going Server-Less

The availability of cloud resources and an industry adoption of the cloud has had a transformative effect upon the process in which business applications are designed, implemented and maintained. Modern cloud business applications are designed and optimized for decentralized services, elastic scaling, multi-storage technologies, parallel and asynchronous compute processing, as well as automated server infrastructure management<sup>(26)</sup>. The continued evolution of cloud infrastructure development has produced a new category of cloud architecture that further exemplifies these characteristics – serverless computing.

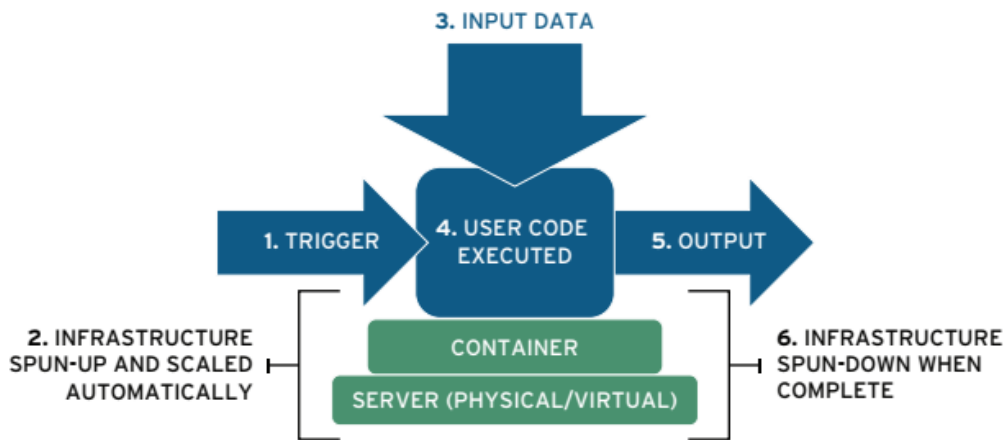
### What is Serverless Architecture?

Unlike what the title implies, “serverless” computing does in fact utilize a server to execute and run services and application code. Still, serverless architecture eliminates the complexity and maintenance burdens of managing Cloud server resources, such as Virtual Machines and Containerized Services. Serverless computing is defined as the “abstraction of servers, infrastructure and operating systems from application development<sup>(27)</sup>”. Serverless applications do not require the provision or direct management of server technology, but rather abstract and automate those services in order to allow businesses to allocate additional resource and capital towards application design and development. In this regard, serverless computing introduces a deeper level of abstraction in cloud infrastructure design from a Platform as a Service (PaaS) to a Function as a Service (FaaS).

As described in Serverless Computing: Current Trends and Open Problems, serverless applications are characterized by several distinct properties<sup>(28)</sup>:

- **Cost:** Usage is metered, and business need only pay for the time and resources used when serverless functions are running. This ability to scale to zero running instances is one of the key differentiators of a serverless platform.
- **Performance and computational limits:** Cloud service providers enforce processing limitations upon FaaS applications, including the number of concurrent requests, and the maximum memory and CPU resources available to a function invocation. Some limits may be increased when users’ need grow, such as the concurrent request threshold, while others are inherent to the platforms, such as the maximum memory size.
- **Programming languages:** Serverless services support a wide variety of programming languages including JavaScript, Java, Python, Go, C#, and Swift. Most platforms support more than one programming language.
- **Programming model:** Currently, serverless platforms typically execute a single main function that takes a dictionary (such as a JSON object) as input and produces a dictionary as output.

- **Composability:** The platforms generally offer some way to invoke one serverless function from another, but some platforms provide higher level mechanisms for composing these functions and may make it easier to construct more complex serverless apps.
- **Deployment:** Platforms strive to make deployment as simple as possible. Typically, developers will write function code within a browser IDE or provide a file with the function source code. Consequently, tools to version or group functions are generally not available.
- **Monitoring and debugging:** Each platform provides basic debugging by using print statements that are recorded in the execution logs. This is generally a very basic debugging feature in comparison to full-featured debugging environments.



Adapted from Economics of Serverless Computing<sup>(29)</sup>

At a high level, a serverless application is a collection of loosely coupled functions that reside within a cloud service. Each function independently executes in response to one or more of triggers or events. Example events could be scheduled triggers, webhook triggers, IoT device events, database-triggered events, or HTTP requests initialized from any source (including other serverless functions). Because of this functional independence and loose coupling, serverless infrastructure possesses two unique benefits: greater resource flexibility and significant programming language collaboration.

This programming language flexibility is exemplified in that each function is not restricted to a single programming language. This can potentially be very powerful in that each function is capable of being written in whichever language suits the development team working on a business function. It also allows a development team the capability to utilize whichever language is best suited for that particular functional task. As an example, the backend logic API for a mobile application may be written in C# due to that being the specialization of the mobile development team, but the functions that consume and analyze customer usage data for marketing purposes are able to be written in a language better suited for data analysis (such as Python or R). The figure below illustrates a serverless function responding to a database triggered event to persist database generated errors in a secondary Datastore.

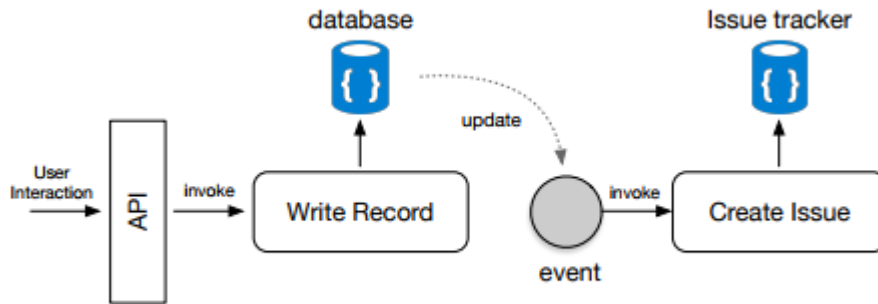


Fig. 8 Batched invocation for issue tracking

Adapted from *Serverless Computing: Current Trends and Open Problems*<sup>(30)</sup>

The second significant benefit of this functional independence, is the segregation of computation and memory usage between functions. When an event triggers a serverless function, the FaaS will allocate the required memory and provision the required underlying infrastructure (most typically a container on a virtual machine) before executing the application code. The figure above illustrates this process flow. The benefit of these independent computational “nodes” allows for asynchronous operations between functions and a very granular decentralization of computation. For example, if a single function is undergoing heavy usage, there is no performance impact upon the other FaaS functions within that application as each node independently manages its own local memory and processing.

### What Economically Differentiates Serverless Architecture?

From an economic standpoint, there are two key characteristics that distinguish serverless-backed applications from other cloud services: the exceptionally granular pricing model and the ability to develop backend application logic at a functional level<sup>(31)</sup>.

From a theoretical point of view, the cost of implementing a serverless application is derived from the runtime usage of the application code. Put simply, a company will only incur charges for the time in which their applications are actively in use. This model promotes an increase in compute resource utilization, and greatly reduces the “sunk costs” of maintaining idle applications when not in active use or are not performing meaningful work. This resource- and cost-optimized pricing model provides a clear economic benefit for infrequent/sporadic process execution (such as webhooks or IoT event handlers), interval processes that run on a timed schedule or other processes that simply do not require uninterrupted availability.

More precisely, FaaS pricing is derived from three metrics: the duration for which the function is executing, the memory resources assigned and utilized during execution, and how often the function is executed. The duration of the function runtime is calculated as “the aggregate sum of the time the code is executed over the course of a month. Each execution duration is rounded up, often to the nearest 100ms<sup>(32)</sup>”. In terms of memory resource usage, each cloud provider utilizes a different pricing approach, but the generally accepted standard is a flat charge based on the amount of preset memory assigned to a function. Microsoft has taken this model into a more

dynamic approach in that their memory pricing model calculates cost based upon the actual amount of memory used per each function in incremental blocks of 128MBs. Lastly, each cloud provider factors in the number of executions of each function over the course of a month, in a range of two hundred-thousandths to four hundred-thousandths of a US cent per request. An example pricing model for Amazon, Microsoft, Google and IBM is displayed in the table below.

	AWS	AZURE	GOOGLE	IBM
<b>FREE REQUESTS PER MONTH</b>	1,000,000	1,000,000	2,000,000	
<b>FREE RAM PER MONTH</b>	400,000	400,000	400,000	400,000
<b>FREE CPU PER MONTH</b>			200,000	
<b>PRICE PER GB-SECOND</b>	\$0.0000167	\$0.0000160	\$0.0000025	\$0.0000170
<b>PRICE PER REQUEST</b>	\$0.0000002	\$0.0000002	\$0.0000004	
<b>PRICE PER GHZ-SECOND</b>			\$0.0000100	

Adapted from Economics of Serverless Computing<sup>(33)</sup>

### What are the Limitations of Serverless Architectures?

While serverless architecture provides benefits such as flexible pricing, function isolation and dynamic resource scaling, there are three distinct consequences of implementing FaaS architecture: complex software design and maintenance, higher computational latency as compared to traditional cloud architecture, and a significant switching cost when changing cloud service providers once a serverless application has been established.

Due to flexible language constraints and function isolation, serverless applications require meticulous and time-consuming design. The process of breaking an enterprise-scale application into micro services requires disciplined coordination between software development teams due to the total independence of each function. Development teams must communicate effectively to collectively build and test functions as they interact with one another as well as the other processes within the application environment. This added integration complexity further complicates the debugging process of a distributed system once the application has reached its maintenance stage. Lastly, because FaaS is a relatively new cloud technology, there are few tools that can assist development teams with monitoring and debugging production serverless environments<sup>(34)</sup>.

Regarding compute performance, serverless architecture tends to produce higher levels of latency. This latency is the result of the additional overhead that is required to provision server resources (such as a virtual machine and memory) before code execution. Thus, for FaaS applications to benefit from automatic resource scaling and “on demand” resource generation, FaaS functions generally require more time to execute as opposed to predesignated, “always on” compute resources. Thus, to offset the initial added latency, FaaS resources remain temporarily active for any subsequent events. For example, resources generated for an AWS Lambda function will remain active for 10 minutes before needing to be recommissioned.

The final limitation is the potential cost of an application migration between cloud service providers. Due to the specific nature of how each cloud service provider offers FaaS application



development, migrating to a new service provider would most likely entail program code changes, architecture redesign as well as the adoption of new development and operational tools. While this is a potential cost, as opposed to a direct cost, it is certainly a factor that should be considered when a business evaluates serverless adoption.

In this section, serverless technology has been defined and examined in both application design and in terms of business economics. In summation, serverless architecture is the product of a cloud evolution that strives for higher levels of abstraction between application development processes and server oversight. While FaaS provide unparalleled flexibility and an enticing cost structure for businesses, the limitations and added software complexity of this cloud architecture better suit certain use cases. Such use cases include chaining together established applications and data sources in new ways (such as the example of storing database events in a separate data store), or new applications that can benefit from outsourcing sporadic or infrequently used compute resources away from the primary compute resource. One such example that will be examined in depth is Cheddar, a personal finance application that utilizes serverless architecture as a core business model.

### **Cheddar: Serverless Architecture in Practice**

Cheddar, an intuitive personal financial planning web application, allows users to aggregate financial account data while providing near real-time financial insights coupled with machine learning-assisted budgeting and goal setting recommendations. As a newly developed software application, Cheddar benefits from the plethora of cloud services available to deliver a product that utilizes cutting edge cloud architecture, such as serverless architecture. In conjunction with the financial API Plaid, Cheddar implements FaaS architecture as the core resource for handling several key application services: handling client account authentication and onboarding, accessing and storing historical and real-time transactional data as well as several other operational application routines.

The process flow for onboarding a new Cheddar user (aside from email, SMS and reCaptcha authentication measures) is the process of selecting and authenticating user financial accounts. This process is initiated and maintained by the Plaid API Link Integration, and is initiated by way of a Plaid maintained JQuery function. This function accepts and sends several parameters when establishing a connection to the Link portal – namely the application key, development environment and most notably, the Cheddar webhook URL.

Once the Link portal has been established, the user provides Plaid with a specified institution and the username and password associated with that institution. Plaid (utilizing ACH measures) then creates a corresponding “Item” on their servers. This Item acts as the link between the specified Institution and Plaid, acting on behalf of the user. Once successfully authenticated, Plaid provides Cheddar with the name of the authenticated institution as well as a short-lived public token. For added authentication, Plaid then requires the temporary public token to be exchanged for a permanent access token and an item identifier token. Both tokens are securely stored by Cheddar and the stored tokens are then utilized by Cheddar when submitting future data requests to Plaid.

The webhook that is specified during item creation acts as the access point in which all future webhook requests are sent from Plaid. Once this authentication process has been completed, Plaid will send a POST HTTP request to the webhook URL with a JSON blob containing the webhook type, the webhook code, the item identifier and the number of available transactions for the given user account.

```
{
  "webhook_type": "TRANSACTIONS",
  "webhook_code": "INITIAL_UPDATE",
  "item_id": "wz666MBjYWTp2PDzzggYhM6oWmBb",
  "error": null,
  "new_transactions": 19
}
```

Sourced from Plaid API Documentation

Cheddar specifies this webhook URL as the unique URL to a serverless function that decodes and disseminates the request to other subsequent serverless functions based upon the type of webhook that is received. In the ongoing example, there is a separate function that handles initial

```
{
  "accounts": [{object}],
  "transactions": [{
    "account_id": "vokyE5Rn6vHKqDLRXEn5fne7LwbKPLIXGK98d",
    "amount": 2307.21,
    "category": [
      "Shops",
      "Computers and Electronics"
    ],
    "category_id": "19013000",
    "date": "2017-01-29",
    "location": {
      "address": "300 Post St",
      "city": "San Francisco",
      "state": "CA",
      "zip": "94108",
      "lat": null,
      "lon": null
    },
    "name": "Apple Store",
    "payment_meta": Object,
    "pending": false,
    "pending_transaction_id": null,
    "account_owner": null,
    "transaction_id": "1PNjeWlnR6CDn5okmGQ6hEpMo41LLNoSrzdDje",
    "transaction_type": "place"
  },
  ...
],
  "item": {Object},
  "total_transactions": Number,
  "request_id": "45QSn"
}
```

account onboarding by swapping the item identifier with the corresponding access token, and initiating a new Plaid API request for the specified number of transactions.

#### Example JSON snippet of a Plaid Transaction Payload

Simultaneously, once a new item has been created on the Plaid servers, Plaid initiates a historical transaction request with the institution. In most cases, this entails two years' worth of transactional data. Conveniently, Plaid will initiate this request as a webhook (in a similar fashion to an initial request, exemplified in the code snippet below) once their server has received the transactional data. To manage the possible compute strain of requesting, parsing and saving this potentially massive number of transactions, Cheddar offloads this process into a separate collection of serverless functions.

```
{
  "webhook_type": "TRANSACTIONS",
  "webhook_code": "HISTORICAL_UPDATE",
  "item_id": "wz666MBjYWTp2PDzzggYhM6cWmBb",
  "error": null,
  "new_transactions": 231
}
```

Sourced from Plaid API Documentation

The final two serverless examples to be examined are the processes that handle new, real-time transactions and the operations for handling pending transactions. Plaid handles new transactions in a manner that supports real-time processing – once an institution posts a new transaction (pending or otherwise), Plaid will save the corresponding data to its servers and will notify the appropriate webhook. Once the transactions have 'settled' (in that they are verified by the institution as no longer pending) Plaid will send delete webhook notifying Cheddar that the transaction has been removed/replaced. Plaid will then resend a transaction request specifying the finalized transaction data. Though I/O intensive, this process enables very accurate transaction recording and eliminates the need for duplicate transaction deletion database tasks.

```
{
  "webhook_type": "TRANSACTIONS",
  "webhook_code": "TRANSACTIONS_REMOVED",
  "item_id": "wz666MBjYWTp2PDzzggYhM6cWWmBb",
  "removed_transactions": [
    "yBVBEwrPyJs8GvR77N7QTxnGg6wG74H7dEDN6",
    "kgygNvAVPzSX9KkddNdWHaVGRVex1MFm3k9no"
  ],
  "error": null
}
```

Sourced from Plaid API Documentation

Cheddar utilizes several other serverless functions as a means of interacting with Plaid and the application database in operational functions (handling and storing errors, initiating the authentication of accounts, and token swapping for example). This process flow of handling webhooks and HTTP events is ideally designed for serverless architecture. These infrequent, compute-intensive operations are offloaded from the main Cheddar API compute source and provides a means to abstract and isolate database I/O operations away from the client application. This architecture also greatly reduces costs that would normally be required to provision resources that would otherwise be idle while not in use.

The Cloud has offered more than a shift in an operating model for provisioning and utilizing IT resources. It has become a movement within both business and developer circles as a way to decouple portions of an application in order to increase performance or reduce application overhead. Serverless technology is further shifting application development and its relationship to business by disrupting the once rigid synergy between the two.

Abstraction through technology is taken one step further as the final piece in transforming to the Fourth Industrial Revolution is virtual reality. It explores a system architecture currently being developed by an author that could help shape our psychological responses relevant to creating awareness amongst individuals on complex societal topics.

## Virtually Aware

Virtual reality (VR) and augmented reality (AR) as we know it is one of the most intense ways of delivering immersive environment to users. An individual experiencing these immersive environments displays psychological responses which mimic those in a real world. The stimuli in an immersive environment created by VR/AR has the potential to trigger these psychological responses. We are focusing on the immersive environment created by VR in the scope of this paper. We propose to develop a system which can create custom psychological responses through specifically developed components inside of an immersive environment. These components are designed to provide stimuli which can trigger the required psychological

responses. Furthermore, the proposed system is extensible to generate desired cognitive responses coupled with psychological responses. This is extremely relevant and useful in creating awareness among individuals on abstract and complex topics such as gender equality in a workplace and organizational behavior, or issues such as causes and implications of global warming.

The components in the immersive environment will be developed based on proven research conclusions in learning methodologies, cognitive psychology and quantitative psychology. The system to be developed is then proposed to have the analytical component which takes in the user input which will be mapped to the cognitive and psychological responses to learn and evolve the existing components in the VR environment. The analytical component will be implemented through machine learning and deep learning models which will be built to accommodate a feedback loop generated through user inputs mapped to cognitive and psychological responses.

## **Introduction**

This section will explore a proposed system architecture with three distinct layers, an immersive environment game design layer, a SQL cache layer for dynamic storage and a NoSQL database layer on the cloud which represents a key value data store for each user in an active game. The three layers in the architecture create an intelligent immersive environment in VR which leverages machine learning strategies to evolve a VR environment at the user's end by learning his/her cognitive responses based on stimuli created by the environment<sup>(35)</sup>.

The VR environment provides the game interface and the gameplay mechanism through different assets or components with specific behaviors scripted for each component. Features are extracted through behaviors associated with each component. Each behavior associated with the component is cached during gameplay at certain intervals using a SQLite database on the user end. Before the start of each set of intervals, the cached data is transmitted to a key value data store on the cloud where it is fit to a relevant model to understand user behavior and the predictions are transmitted back to the local user which triggers specific behavioral scripts on the client side which alters the next interval of gameplay.

## Architectural Overview

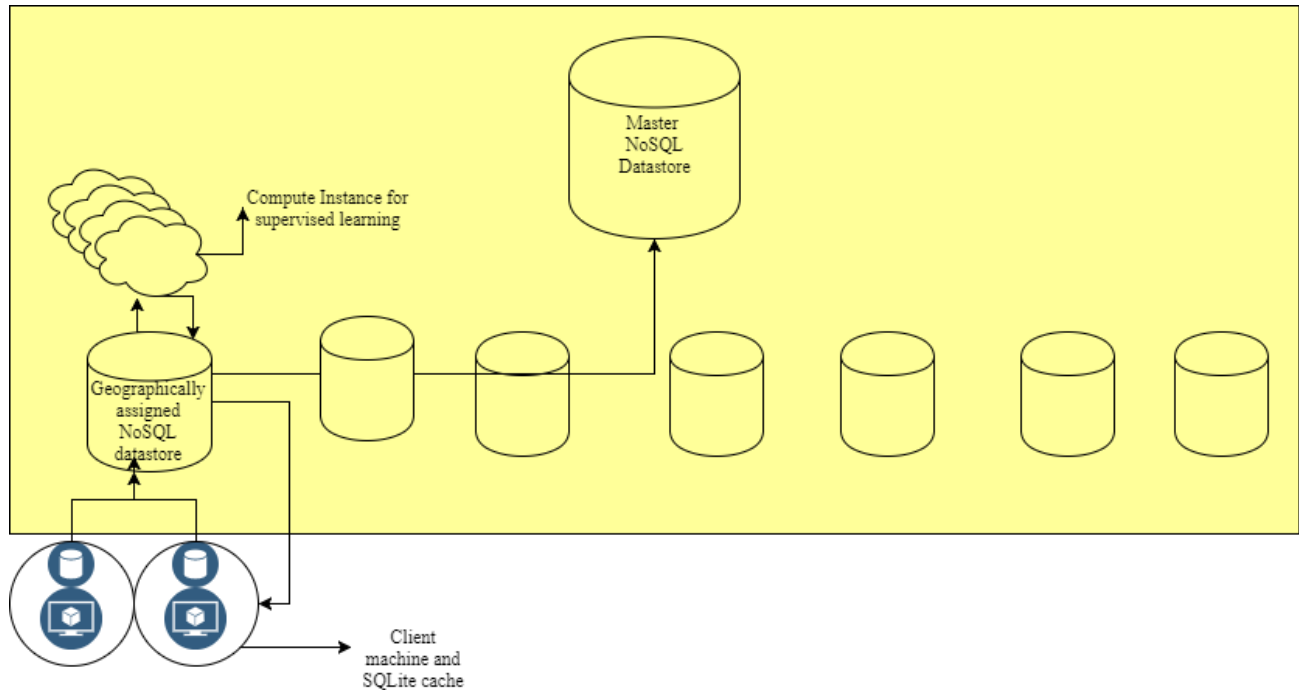



Figure 1.0 A high level overview of the technical system architecture

### **Finite State Machines for individual components in gameplay**

Each element or component in gameplay has a specific set of behaviors and are defined in terms of finite state machines or finite automation.

A finite state machine is a computational model which abstracts an object or a machine such that it can only be in a unique state at a time from a given set of states. In the immersive environment we developed, each asset or component has a set of states defined by behavioral scripts which are triggered for certain user actions as well as events in gameplay. A finite state can be considered analogous with that of a brain or intelligence associated with a game object or asset. We will evolve these finite state machines which in turn makes these game objects more intelligent.

## Reinforcement learning for game objects



**Figure 1.1** A representation of a finite state machine defining the behaviors of a game object. **S** represents a state, **a** represents an action from a state and **e** an event of trigger and transition.

As previously mentioned, behaviors for each game object is defined through a computational model known as a finite state machine. A game object can be at only one state at a given instance of time or gameplay. We can think of each state as a node and the event which triggers that state as an edge in a graph as shown in the above figure (1.1). Game objects have their finite state machines modelled in a probabilistic way using Markov Decision processes. Each edge or event in the finite state automata has a probability or a weight associated with it. These weights provide a heuristic to the node it is directed to, which in turn decides the final action that a game object will be performing. The actions are triggered based on a reinforcement learning mechanism as the finite states evolve. The evolution of finite states is based on the positive or negative reinforcement that a game object receives during gameplay.

### Markov Decision process modelling for reinforcement learning

Each state represented by a node has an action set  $A$  which defines the behavioral scripts that will be triggered for specific actions in the action set. The specific action triggered at the node is decided based on the weight of the incoming edge which represents the event. Each action has a specific range of weights associated with it.

Initially there is no heuristic from the events or edges to trigger specific actions at different states and is randomly triggered at each state. As the gameplay progresses, due to the chain of events triggered by mutual game objects and user actions, the finite states associated with each game object evolves through positive and negative reinforcement<sup>(36)</sup>.

The positive and negative reinforcements are decided based on the action triggered at each state for a finite state machine. An action which is rewarded receives a positive reinforcement

and hence the resulting event from that action receives a higher weight when transitioning to the next state. An action which is taxed or hit receives a negative reinforcement and the resulting event from that action receives a lower weight when transitioning to the next state. Because of this process, the state machine learns through positive and negative reinforcement as it begins to favor actions at each state which result in positive reinforcement and the weights of the edges map the probabilistic modelling of the learning process.

### **Game intervals and data caching**

The gameplay developed will consist of specific intervals which collect data relevant to gameplay status, actions and events triggered, metrics surrounding game objects and user actions. This data is used to evolve the next interval of gameplay as well as to collect insights on specific cognitive responses associated with user actions for the gameplay interval.

### **Vectoring the data**

The data required to model user behavior and cognition must be tokenized and converted into a numeric array. A vectorizer is a function which deals with data such as gameplay status, actions and events triggered, metrics surrounding game objects and user actions and converts them by tokenizing them or providing a normalized numeric value which allows the data to be modelled later. The tokenized data is then cached in a client side Datastore.

In addition to tokenizing the data, the vectorizer also generates a key based on the type of modelling that the data has to undergo. This is specific to supervised learning algorithms such as support vector machine, random forest classifier or a polynomial regression.

### **Client-side SQLite Cache**

A vectorized form of data is cached throughout a gameplay interval. This dataset is also known as the set of feature vectors which will be used for modelling user behavior later.

A relational SQLite database is used to store the feature vectors generated for each gameplay interval. The SQLite cache is a representation of labelled data needed for supervised machine learning – an example would be analyzing the correlation between labelled features corresponding to actions triggered and the performance metrics surrounding gameplay objects. This is a potential use case when user behavior and decision-making processes can be tied to certain metrics surrounding gameplay objects.

The SQLite cache has a unique identifier tied to the client or user and at each gameplay interval, a unique interval ID is generated as well as the key generated by the vectorizer for supervised learning<sup>(37)</sup>. Before the start of each interval, the data from the cache is uploaded to a Datastore on the cloud and the local cache is cleared to make way for new features for the next interval.



## **NoSQL Key-Value Distributed data store**

The data collected at each client cache is uploaded to a key value store on the cloud. Each client or user is assigned a data store based on geographical IP. Each of these data stores generate a document as a value for each unique identifier assigned as a key based on the SQLite cache and client ID. The generated document has the labelled feature vectors which will be used for supervised machine learning through a provisioned instance for the client ID.

## **Supervised Machine learning on the cloud**

A CPU instance is assigned to each generated document and client ID. Each document containing the feature set will be fed to the CPU instance which runs the appropriate training model based on the key associated in the client ID formerly generated by the vectorizer on the client side.

## **Master NoSQL Datastores**

Individual NoSQL Datastores upload all of their generated documents to a primary master database which is a highly scalable distributed database in the cloud. This database has all the data associated with the users and they are segregated inside the master store based on the client ID key which also contains the specific model which was used on the data.

The data is smartly organized and segregated inside the master Datastore based on the model which was used on the data. It is analogous to containing partitions of documents trained using the same supervised learning algorithm in the former step.

The segregated documents are updated with prediction metrics generated by individual geographic compute instances which perform supervised learning based on the data. The update is sent to the geographic NoSQL Datastores as well as the Master Datastore.

## **Synchronizer API**

Synchronizer is an API service proposed for the application. The API is responsible for communicating with the Vectorizer to update the SQLite cache and manage relations between the features inside the SQLite database through an ORM.

The API is also responsible for updating the geographic NoSQL Datastore by sending the entire copy of the relational table from the SQLite cache as a new document which replaces the most recent document on the geographic NoSQL Datastore. As the name suggests the synchronizer synchronizes the data in the local cache for a client with his/her client ID key for the assigned geographic NoSQL Datastore.

## **Synchronizer for fetching predictions**

When the geographic compute instances on the cloud make a prediction, the document in the corresponding geographic NoSQL Datastore is updated through a cloud-based trigger functionality<sup>(38)</sup> which triggers the update event for the client ID when the corresponding compute instance makes a prediction. The Synchronizer API queries the updated NoSQL

Datastore document which corresponds to the previous gameplay interval and specifies the user behavior associated with the previous interval.

### **Gameplay versions and triggers**

Heuristics and metrics for different versions of gameplay are required to alter the gameplay interval based on a predicted cognitive response. References to these heuristics and metrics are maintained by the Synchronizer API which triggers certain events and actions by activating the specific behavioral scripts associated.

The concept of different game versions is emulated by having a variability in terms of state space for the finite state machines associated with game objects, heuristics for event triggers and gameplay variables.

The prediction is potentially correlated with the vector. This vector is updated by the synchronizer based on the received prediction which updates the vector representation of the prediction at every gameplay environment and object definition. This evolves the virtual environment to provide a relevant game experience based on the user behavior and cognitive responses in the previous gameplay.

In conclusion, the proposed system has major potential applications in domains including but not limited to mental health, simulated learning, digital learning experiences, social awareness programs and organizational training.

### **On to Future Revolutions**

If the First Industrial Revolution caused a boom in our means of production and the overall exponential growth of humanity, then the Fourth Industrial Revolution will cause a wave of nuanced thinking through technology. No matter how big or small the idea may seem, what will pump life into that idea is the ability for its creators to intuitively utilize technology to make it real.

Whatever is in store for the Fourth Industrial Revolution is yet to be discovered, but it is the developers and engineers – such as the authors of this paper – who hope to use technology as the innovators who preceded them did, with the best intentions and the greatest impact to society. Technology is in every aspect of our lives and it's here to stay. Why not make the best of it.

## Appendix

- 1) Green, John (2013), "Coal, Steam, and The Industrial Revolution: Crash Course World History #32", Crash Course, <https://www.youtube.com/watch?v=zhL5DCizj5c>, Retrieved January 2018
- 2) Ray, Charles (2010), "Illustration for the Romance of the Nation", Look and Learn, <https://www.lookandlearn.com/history-images/M816028>, Retrieved January 2018
- 3) Brown, Cynthia Stokes (2018), "Fossil Fuels, Steam Power, and the Rise of Manufacturing", <https://www.khanacademy.org/partner-content/big-history-project/acceleration/bhp-acceleration/a/the-industrial-revolution>, Retrieved January 2018
- 4) Adamson, Lucy (2016), "The History of eCommerce", Statement, <https://www.statementagency.com/blog/2016/03/the-history-of-ecommerce>, Retrieved January 2018
- 5) Miva (2011), "The History of Ecommerce: How Did It All Begin?", <https://www.miva.com/blog/the-history-of-ecommerce-how-did-it-all-begin/>, Retrieved January 2018
- 6) Lim Xin Ying (2008), "The History and Evolution of e-Commerce", WordPress Blog, <https://ecommerze.wordpress.com/2008/06/12/the-history-and-evoltion-of-e-commerce/>, Retrieved January 2018
- 7) Schwab, Klaus (2016), "The Fourth Industrial Revolution: what it means, how to respond", World Economic Forum, <https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/>, Retrieved January 2018
- 8) I-Scoop (2018), "Digital transformation: online guide to digital business transformation", <https://www.i-scoop.eu/digital-transformation/>, Retrieved January 2018
- 9) Olanrewaju, Tunde & Willmott, Paul (2013), "Finding Your Digital Sweet Spot", McKinsey, <https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/finding-your-digital-sweet-spot>, Retrieved January 2018
- 10) Chokshi, Niraj (2017), "Out of the Office: More People Are Working Remotely", New York Times, <https://www.nytimes.com/2017/02/15/us/remote-workers-work-from-home.html>, Retrieved January 2018
- 11) Infor, "White Paper: The democratization of data, how information can give power to your people", <https://www.infor.com/content/whitepapers/democratization-of-data.pdf/?&isGated=no>, Retrieved January 2018
- 12) Marr, Bernard (2017), "Why Data Democratization Is Such a Game-Changer In Our Big Data World", Data Informed, <http://data-informed.com/why-data-democratization-is-such-a-game-changer-in-our-big-data-world/>, Retrieved January 2018
- 13) Buff, Anne (2018), "Data Lake vs Data Warehouse: Key Differences", KD Nuggets, <https://www.kdnuggets.com/2015/09/data-lake-vs-data-warehouse-key-differences.html>, Retrieved January 2018
- 14) McKenna, Brian (2016), "Data democratization in the age of big data: why data lakes won't work", Computer Weekly, <http://www.computerweekly.com/blog/Data-Matters/Data-democratization-in-the-age-of-big-data-why-data-lakes-wont-work>, Retrieved January 2018

- 15) PayPal UK (2018), "Legal Agreements: List of Third Parties (other than PayPal Customers) with Whom Personal Information May be Shared", <https://www.paypal.com/uk/webapps/mpp/ua/third-parties-list>, Retrieved January 2018
- 16) Rosic, Ameer (2017), What is Cryptocurrency: Everything You Need To Know [Ultimate Guide]", Blockgeeks, <https://blockgeeks.com/guides/what-is-cryptocurrency/>, Retrieved January 2018
- 17) Business Dictionary (2018), "Definition: Trusted Third Party Servers (TPP Services)", <http://www.businessdictionary.com/definition/Trusted-Third-Party-Services-TTP-Services.html>, Retrieved January 2018
- 18) Monero (2018), "Why Monero is Safe", <http://monero.org/>, Retrieved January 2018
- 19) Monero (2018), "Ring Signatures", <https://getmonero.org/resources/moneropedia/ringsignatures.html>, Retrieved January 2018
- 20) Chia, Mai-Hsuan (2016), "Bitcoins' anonymity", LinkedIn, <https://www.slideshare.net/MaiHsuanChia/bitcoins-anonymity-60985284>, Retrieved January 2018
- 21) Monero (2018), "Stealth Address", <https://getmonero.org/resources/moneropedia/stealthaddress.html>, Retrieved January 2018
- 22) McKenzie, Cameron (2015), "What is containerization?", The Server Side, <http://www.theserverside.com/discussions/thread/80994.html>, Retrieved January 2018
- 23) Rouse, Margret (2018), "application containerization (app containerization)", Tech Target, <http://searchitoperations.techtarget.com/definition/application-containerization-app-containerization>, Retrieved January 2018
- 24) Rouse, Margret (2018), "micro services ", Tech Target, <http://searchmicroservices.techtarget.com/definition/microservices>, Retrieved January 2018
- 25) CloudOps (2017), "Docker and Kubernetes: What is the Value of Containerization?", <https://www.cloudops.com/2017/07/docker-and-kubernetes-what-is-the-value-of-containerization/>, Retrieved January 2018
- 26) Microsoft (2018), "Azure Application Architecture Guide", <https://docs.microsoft.com/en-us/azure/architecture/guide/> Retrieved January 2018
- 27) Owen Rogers, 451 Research (2017), "Economics of Serverless Computing", Page IV, Retrieved January 2018
- 28) Ioana Baldini, Paul Castro, Kerry Chang, Perry Cheng, Stephen Fink, Vatche Ishakian, Nick Mitchell, Vinod Muthusamy, Rodric Rabbah, Aleksander Slominski, Philippe Suter (2017), "Serverless Computing: Current Trends and Open Problems", Pages 5-6, Retrieved January 2018
- 29) Owen Rogers, 451 Research (2017), "Economics of Serverless Computing", Page 1, Retrieved January 2018
- 30) Ioana Baldini, Paul Castro, Kerry Chang, Perry Cheng, Stephen Fink, Vatche Ishakian, Nick Mitchell, Vinod Muthusamy, Rodric Rabbah, Aleksander Slominski, Philippe Suter (2017), "Serverless Computing: Current Trends and Open Problems", Pages 13, Retrieved January 2018
- 31) Owen Rogers, 451 Research (2017), "Economics of Serverless Computing", Page 5, Retrieved January 2018

- 32) Owen Rogers, 451 Research (2017), "Economics of Serverless Computing", Page 6, Retrieved January 2018
- 33) Owen Rogers, 451 Research (2017), "Economics of Serverless Computing", Page 6, Retrieved January 2018
- 34) Aditi Chaudhry (2017), "What is Serverless", <https://dev.to/aditichaudhry92/what-is-serverless-1bf>, Retrieved January 2018
- 35) Gregg, L. & Tarrier, N. Soc Psychiat Epidemiol (2007) 42: 343. <https://doi.org/10.1007/s00127-007-0173-4>
- 36) S. Benson and N. Nilsson. Reacting, planning and learning in an autonomous agent. In K. Furukawa, D. Michie, and S. Muggleton, editors, Machine Intelligence 14. Oxford University Press, Oxford, 1995.
- 37) SQLite <https://www.sqlite.org/whentouse.html> . "Appropriate uses for SQLite". 2018.
- 38) Google. <https://firebase.google.com/docs/functions/database-events>. "Realtime Database Triggers". 2018.

Dell EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL EMC MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell EMC software described in this publication requires an applicable software license.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.