Report
DYNAMO 106
2/4 2020

Thoughts on Digital Transformation[1]

A few years ago, we shared our thoughts on how the moats of the incumbent consumer goods companies are being eroded. We explained how the Internet fundamentally changed the dynamics of supply and demand for consumers and thereby altered the competitive landscape. Specifically, digital marketing enabled unprecedented granularity in targeted advertising, social media shifted brand advocacy from the brand owners to peers and influencers, and e-commerce created endless and dynamic shelf-space.

These structural shifts allowed for the emergence and success of a new breed of competitor. These were businesses that developed innovative go-to-market strategies which took full advantage of the new paradigms and cheap capital. Many incumbents were caught off-guard, but quickly realized that they too needed to adapt to this new landscape. The strategies to succeed in a world where consumer behavior has shifted online are now well understood, and one could argue that well-managed and well-funded incumbents might be in a position to claw back some of their lost allure. However, we think this view ignores another layer of complexity that arises when companies try to provide customers the solutions they demand.

As we spent more time with these new entrants, we learned that the Internet had not only transformed the consumer facing side of these businesses, but also their mode of operation. Companies born in the last 10 years, across all industries, are built differently from the bottom up, particularly with regards to their IT infrastructure. Consumer goods, payment providers, industrials, banks, healthcare incumbents and many others are now competing with entities that were built in an environment that enables previously inconceivable flexibility, modularity, and agility. This step-change was triggered by the emergence of cloud computing and it is forcing businesses to completely rethink their legacy IT infrastructure. In this report, we analyze the moving parts of this fundamental shift and how it can impact the way companies are organized at their core.

We believe that there will be a fundamental competitive divide between businesses that accept the need to change and those that do not. This bifurcation impacted our research process in two ways. First, we have been spending considerable time and efforts on technology companies that enable enterprises to transform. Second, we are scrutinizing the current state of digital transformation in our portfolio companies to better assess the potential risks they could be facing or benefits they could be reaping. To be able to differentiate between companies that prosper and those that falter, we think it is imperative to fully comprehend the mechanics of the digital transformation. We believe the companies that prosper will use their technological advantage to improve their customer offering, increase their productivity and efficiency, and attract a more capable workforce, while the organizations that falter will experience the opposite. Consequently, as long-term investors we should no longer assess the quality of an investment without taking the state of digital transformation into consideration.

At its core, digital transformation is the transformation of the IT stack within organizations. More specifically, it is the shift from on-premise, self-managed IT infrastructure to cloud computing, from which all new digital capabilities follow. Cloud computing is not a new concept. AWS was launched in 2006 and companies like Slack operate entirely on the service. However, most organizations have realized that they can no longer effectively compete in the future if they remain bound by their legacy IT. We have spent the last two years examining the various components of this transformation on a technical and organizational level and aim to summarize our findings in this report. We begin by explaining the general transformation of the IT stack. Next, we will show that this requires a new way of architecting applications, which in turn necessitates an organizational transformation. This organizational change requires a restructuring of the IT budget that allows an organization to take full advantage of the cloud, and results in significant operational improvements.

---

1 *This Report reproduces the* ***Dynamo Fund Report September 2020****, written by our team in London.*

The Legacy IT Stack

The IT stack can be thought of as a layer cake, where the bottom layer enables the layer on top of it. At the very bottom of this structure is hardware, which are servers and switches that handle compute, storage, and networking. This layer is referred to as the infrastructure layer. It is largely commoditized, since most computing tasks can be handled by general purpose hardware. The layer on top of the hardware is the virtualization layer. VMware, founded in 1998, was the driving force behind the adoption of virtualization software. The concept behind virtualization is that you can run multiple operating systems on the same hardware. Prior to virtualization, only one operating system could be configured to each piece of hardware. Virtualization software creates a virtual machine with software that mimics the hardware characteristics required for the operating system to run. We are all familiar with operating systems being dependent on hardware. A Mac computer runs the Mac operating system and a Dell desktop runs Windows. However, once an organization has multiple servers, it becomes very difficult to ensure that the servers are configured for the required operating systems. Virtualization dramatically increased the flexibility and utilization of hardware and therefore became a fundamental layer in the enterprise IT stack.

The next layer are the operating systems. Operating systems translate the outputs of the applications into hardware commands. There are a number of operating systems, but the two most common in the enterprise are versions of Linux and Windows. The layer on top of the operating system is the platform layer. The boundaries of the platform layer are less clearly defined, but we consider it as all the software that enables developers to build applications. This includes database software, machine learning tools, monitoring and analytics tools, app development engines, and more. For example, imagine a company wants to build an app that allows their warehouse workers to pack orders more efficiently. The developers will need a database that they can store the data in, they will need an app development engine to build an iOS or Android compliant application, a monitoring tool that allows them to track the performance of the application, and potentially some pre-built machine learning algorithms to help predict order volume. All of these tools are part of the platform layer. Finally, the top layer of the IT stack is the application layer. In the aforementioned example that would include the warehouse application, as well as any software applications that are purchased. This can include anything from Office 365 to cyber security and communication applications.

The composition of the legacy IT stack does not differ from the modern IT stack. The difference between the two is who manages the individual layers and what that enables an organization to do. In the traditional IT stack, everything is managed in-house. This means that the hardware consists of servers owned by the company. Depending on the size of the organization these are either stored in a room within their premises, collocated at third-party data centers, or located within the company's own data centers. The organizations purchase servers and networking equipment from companies like Dell and IBM, they pay for the space, electricity, and cooling of the facility housing the servers, and have their own staff to manage, fix, and replace the hardware. This is often referred to as on-premise infrastructure. Virtualization services are purchased as license agreements from companies like VMWare and operating systems are also purchased as licenses. Both are managed by the company's own IT staff. The platform layer consists of software licenses and in some cases homegrown solutions. Organizations purchase database software from Oracle or Microsoft, and their IT staff install and manage these licenses. Finally, the application layer consists mainly of homegrown applications and licensed software. The developers are in charge of writing the code for the applications, but deployment of the code and any purchase of software licenses is managed by IT. It is important to note that in the legacy IT stack, there is a very clear distinction between software developers and IT operations, with the latter far outnumbering the former.

Although, the modern IT stack has the same fundamental layers as the legacy stack, many of these layers no longer need to be managed internally. Instead, they can be managed by the leading cloud providers: Amazon, Microsoft, and Google. This shift from self-managed to cloud-managed is the core of the cloud transition. There are three approaches to cloud: hybrid cloud, multi cloud, or single cloud. Hybrid and multi cloud are not mutually exclusive. The hybrid cloud approach simply means that a customer will run some workloads with a cloud provider and keep some on premise. At the moment, a majority of customers are opting for this model. However, we believe that the long-term equilibrium will be predominantly cloud with some exceptional workloads remaining on premise. For the sake of simplicity, the cloud offering is usually broken down into three components: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). In order to examine the shift from a legacy to a modern IT stack, it is important to understand these three different parts of the cloud offering.

Cloud Computing

Infrastructure as a Service (IaaS) is the core offering and is purely focused on replacing on-premise hardware with cloud provider hardware. Hardware is often misperceived as being just storage, while the more important service it provides is compute. Compute refers to the processing power required to run code on hardware. With the IaaS offering, an organization rents storage and compute capacity on the cloud provider's servers and pays based on usage. The IaaS customer continues to manage the workloads, but no longer has to manage physical servers or the facilities that store these servers. The result is that capex shifts to opex and the organization gains flexibility as it is no longer necessary to invest in and maintain capacity for peak demand, because capacity can be adjusted with demand. Simultaneously, most of the IT staff in charge of managing the hardware are made redundant. There are a variety of monetization models today, but the standard format for compute is that a customer pays for compute instance used per minute. At AWS, for pure server compute usage, a customer can pay as little as $0.0225 per hour and as much as $7 per hour, depending on the capabilities of the server. The same model applies for storage, but the costs are significantly lower than for compute.

Platform as a Service (PaaS) is the cloud offering of platform layer services, which aims to replace on-premise software licenses. The best example are databases. Instead of purchasing an Oracle database license, installing it on the cloud providers' servers, and managing that license, an organization can now purchase a database directly from their cloud provider and use it without any installation. The benefits to the organization are clear: save costs compared to the traditional license, reduce time to deployment, and reduce the complexity of their IT stack. However, one should note that platform services offered by the cloud providers are normally tied to their infrastructure offering. In addition, platform services are very sticky since applications are built on top of them and replacing a platform service usually requires a complete rebuild of the application. This means that every platform layer service purchased by an organization makes it more difficult to leave the cloud provider. On the other hand, the cloud provider benefits from selling higher margin software products and can increase the stickiness of their offering, which incentivizes them to offer increasingly compelling platform services. There is an alternative to traditional platform services and the cloud offering. These are cloud-agnostic and cloud-based third-party platform services including database providers like MongoDB or data warehouse providers like Snowflake. One can consider these to be co-opetition for the cloud providers, since they offer competitive products, but host their software on the cloud providers' infrastructure. For simple use cases we think the benefit of cloud platform services increasingly outweigh worries about vendor lock-in, based on conversations with implementation partners and consultants. For more complex tasks we think it can be helpful to use best-of-breed third-party platform solutions.

On top of the platform services, there are also software services. Software as a Service (SaaS) originated in the early 2000s with companies like Salesforce. The concept is that instead of purchasing software licenses, installing the software on their own servers, and then paying for updates and maintenance, an organization can purchase the software as a subscription and the software is hosted by the software provider. This means that the software is continuously updated by the provider and the customer requires very limited hardware to host the software. An example of a SaaS application would be the messaging service Slack. Slack began hosting their service on AWS shortly after they were founded in 2009. Hence, any organization that uses Slack, is running it on the AWS servers. This enables you to use Slack simultaneously from your browser, a desktop app, an iPad app, and a mobile app, since the information is being processed on AWS. SaaS is revolutionizing enterprise software by offering best-in-class solutions to specific problems, with short implementation times, continuous updates, and hardware agnostic deployment. SaaS applications also require much fewer IT staff to manage, install, and update applications. Currently, Amazon, Google, and Microsoft sell very few SaaS application outside of G Suite and Office 365. Microsoft turned their license Office business into a SaaS offering, with the launch of Office 365 in 2011. So, for now, they are focused on hosting most of the world's SaaS applications on their infrastructure, but going forward we think they will start offering more of their own SaaS applications.

The Modern IT Stack

Now that we understand what the cloud offers, it is possible to explain the modern IT stack. In the modern IT stack, the layers are still hardware, virtualization, operating system, platform layer, and application layer. However, now most of these layers are managed by a cloud provider. If we assume that an organization goes full cloud, rather than hybrid, then the organization will no longer need any server hardware, data centers, or maintenance staff. The compute and storage requirements of the organization are managed by the cloud and paid for on a usage basis. Virtualization can be purchased directly from the cloud providers, or one can continue to use providers like VMware. The same goes for operating systems, although

Microsoft Azure customers are more likely to go with the Windows server operating system. On the platform layer, most companies will probably start using some platform services from the cloud providers over time and decide to what extent the benefits of using these outweigh the potential vendor lock-in, while the remaining platform requirements will be satisfied by cloud-based third-party providers. Finally, on the application layer the main change from a purchasing perspective is that there will be fewer software licenses and more SaaS products, while the development team will have the ability to build applications with the flexibility of the cloud.

The ability for developers to build applications on the cloud also fundamentally changes the enterprise software architecture. In the legacy IT stack, software development requires working on large monolithic applications that solve a multitude of problems. An old ERP system is a good example: for developers to make small fixes, downtime for the whole system is required and every change puts the entire application at risk. In addition, the application becomes one large collection of software patches, which means that it becomes increasingly difficult to work with, both for the business users and the developers. Moving applications to the cloud also means modernizing these applications for the cloud environment. This is an important point, since some companies that just move their existing legacy software to cloud infrastructure do not reap most of the benefits. In some cases, modernizing means completely rewriting legacy applications and in others, it means replacing parts of applications with existing solutions. In either case, developers are rearchitecting their applications so that they are splitting the large monolithic applications into a collection of microservices. These microservices each perform one specific part of the larger application and are connected to each other with application programming interfaces (APIs). APIs are software intermediaries that allow two applications to communicate with each other. This is commonly referred to as a microservices architecture. All tech companies work with a microservices architecture, because it allows developers to work without any required downtime, it gives greater flexibility to make improvements, and it limits the risk of damaging the functionality of the entire application. The result is better applications and more effective deployment of software development resources.

More nimble deployment of software requires a different organizational structure and operational model. In the legacy IT stack, the majority of employees are IT operators. This includes all the personnel that ensures the hardware is working, the data center facilities are being managed, licenses are being installed and maintained, updates are being implemented, applications and changes to applications are being deployed, and more. The minority of employees are software developers that are solely responsible for writing code. Furthermore, IT is considered to be a cost center and therefore frequently reports to finance, rather than being considered a growth enabler with access to the CEO. This structure results in a clear distinction between software developers and IT personnel. A cloud environment makes a lot of these IT functions redundant. It is no longer necessary to manage the hardware, and deployment does not require a separate skillset from software development. The result is the DevOps organizational model. DevOps combines the functions of software development (Dev) and IT operations (Ops) into one way of working. This means that software developers now write the applications and deploy these themselves. Evidently, this is a reduction in complexity, which allows for faster deployments and better allocation of resources. A modern IT stack requires a DevOps operational model in order to take full advantage of the cloud shift.

The Operational Impact

Every year, Google Cloud together with DevOps Research & Assessment (DORA) release a State of DevOps report.[2] The report incorporates six years of research and data from 31,000 professionals in the space. Based on this industry feedback, the team identified four key metrics to measure software delivery performance. The four metrics are 1) code deployment frequency (which measures the rate at which changes to software are being deployed in the development pipeline), 2) lead time for changes to be implemented, 3) time to restore a service when it goes down, and 4) the percentage of changes that result in a failure.

They then segregate the companies into elite, high, medium, and low performers. To be an elite performer you have to 1) deploy code multiple times a day, 2) take less than one day for changes to be implemented, 3) take less than one hour to restore a service, and 4) have a failure rate of 0-15% for changes. The researchers compared elite performers with low performers and found that on average elite performers had 1) 208x more frequent code deployments, 2) 106x faster lead time from finishing code to deployment, 3) 2,604x times faster to recover from incidents, and 4) 7x times lower change failure rate.

The report also outlines five essential characteristics for cloud-based organizations: 1) they have on-demand self-service of the infrastructure for developers, 2) they

2 *https://services.google.com/fh/files/misc/state-of-devops-2019.pdf*

have broad network access to the infrastructure, 3) they can pool IT resources, 4) they can rapidly scale up or down usage, and 5) they can measure their service effectively. This definition of cloud-based organizations overlaps with our modern IT stack. Elite performers were 24x more likely to meet all five of these characteristics, which implies that achieving elite performance is nearly impossible without a cloud first approach. In addition, companies that met all five characteristics were 2.6x more likely to accurately estimate their cloud cost, 2x more likely to correctly identify their most expensive application, and 1.65x as likely to stay under their software budget. Finally, organizations that run on large chunks of legacy code were 1.6x less productive, and elite performers were 1.4x less likely to run on legacy code. This shows that one can quantitatively measure the improvement from a legacy IT stack to a modern one.

The Financial Impact

This level of organizational and operational change also has to be reflected in the IT budget. We spoke to multiple cloud deployment consultants, interviewed a number of former cloud provider employees, and surveyed over one hundred CTOs and CIOs to better understand the impact on IT budgets. First, despite the improved productivity and efficiency of the modern IT stack, everyone expects IT budgets as a share of their revenues to continue to grow significantly over the next five years. However, the composition of the IT budget will change significantly. In the legacy IT stack, the majority of spend went towards IT operations. This includes hardware, data center costs, and all the IT staff. The data we collected suggests that in the legacy IT stack 40-50% is spent on IT operations. Approximately, 10-15% is spent on the cloud, since even legacy IT stacks have some of their workloads in the cloud today (email servers, newer applications, etc.). For the remainder, 10-15% is spent on software development, 5% on cyber security, and 10-25% on software licenses and other expenses. The shift from self-managed to cloud-managed is also reflected in the modern IT budget. Our research indicates that in the modern stack, IT operations will only represent 15-20% of the IT budget, cloud will represent 30-35%, software development will represent 15-20%, cyber security 5-10% (although we have heard of cases where cyber security spend quadrupled), and software and other expenses will represent 15-20%. Furthermore, the majority of software purchased will be in the form of SaaS rather than license based, with 83% of our respondents stating they will increase the number of SaaS applications they are using in the next three years.

The legacy IT budget composition and the dramatic change required to attain a modern IT stack, also shows one of the biggest hurdles for digital transformation: the incumbent IT staff and misaligned decision makers. IT personnel are aware that the cloud will make most of their positions redundant and that a majority of the spend will be reallocated to the cloud providers. Hence, there are a lot of decision makers within these organizations whose incentives are directly opposed to a successful shift to the cloud. In addition, even the remaining employees will be given new KPIs, since many of the old KPIs will no longer be attainable during the process of transformation. This means that to successfully transform an organization the decision must be embraced by the highest form of leadership, since it requires the willingness to sacrifice the legacy system. The difficulty of this decision, the cost associated with the transformation, and the extensive duration of a full transformation (5-7 years), means most executives will try to avoid taking on this task. This results in partial transformations and small projects that can be presented as quick wins to the board and shareholders. However, we have witnessed multiple times that partial transformations are just kicking the can down the road. Often, a crisis is required for management teams to realize they need to rebuild completely. For some organizations, COVID could be this triggering event.

Nike Case Study

Public companies rarely provide information about their IT infrastructure. However, to showcase the impact of a successful digital transformation we can look at two exceptions: Nike and Equifax. Nike used to have a very traditional IT infrastructure, which was the starting point for their transformation.[3] Back in 2013 the company had most of their IT in one data center and two distinct IT and software development teams.[4] The infrastructure was organized in a way that all IT solutions, such as Nike.com and Nike apps, were running on the same servers and databases. The result was that any change had to be approved and then deployed with the next release. It was a very manual process, depended on a number of different vendors, and had to be approved by a waterfall process involving both the software and the IT teams. As of 2018, the company has four AWS regions, 150 software engineers, three development locations, and multiple data center locations. In the process, the company decided that they would not just lift and shift their existing applications from their own servers to the public cloud, but instead decided to rethink every single component of their IT organization. The results

3 *https://chainstoreage.com/technology/nike-gets-its-footing-cloud*

4 *https://www.youtube.com/watch?v=6A1tOFqvgek&ab_channel=AmazonWebServices*

show that this transformation worked. The organization went from one software deployment every two months to 2.6 deployments per day. Nike went from 90% manual software testing to 100% automated testing, which freed up a lot of developer time. They managed to reduce the time to make small changes on the website and apps from 3 hours to 5 seconds, which means they could react to sports and similar live events. In the past it took more than six months to add a new experience to their digital services, and today it takes one day. In the past they would have a 3-month lead time for new hardware and today they can scale and deploy without any lead time.[5] The IT infrastructure now supports 50+ commerce countries versus 6 in 2012, supports 25 languages versus 7, and enables the e-commerce site to access the inventory of 500+ retail stores.

The early move to the cloud and the willingness to adapt to the new environment also allowed Nike to benefit from some significant learnings. For instance, the company first used the Cassandra database when they moved to the cloud.[6] However, due to many technical limitations, it would not allow them to scale for peak demand.[7] Peak demand was becoming a big problem because the Nike SNKRS App would launch products with very limited availability, which meant that millions of people would access the app at the same time. Nike then decided to move to the AWS DynamoDB database (a platform offering), which allowed them to scale up prior to these launches, and thereby spend 98% less than with Cassandra, while offering the same service.[8] In addition, they managed to monitor the launches in real time, which allowed them to react to problems and error messages within seconds. The vast amount of data that is generated within this very short period is now analyzed with machine learning techniques to improve the stability, reliability, and optimization of future launches. The company is working on a number of other efforts that benefit from the cloud environment, such as the implementation of RFID whose data output is managed through the AWS IoT offering.

The benefits of this transformation to the consumer are clear. Nike can now deal with higher demand, deal with sudden spikes in orders, offer better product recommendations, offer more customization, provide better product fulfillment, and more. In addition, the company benefits from a leaner and more efficient IT organization, better product conversion, more feedback data from customers, social integration into products, and ultimately a more satisfied costumer. We think Nike's continued investment into their modern IT stack will be a key differentiator for their competitive positioning.

Equifax Case Study

Unlike the proactive approach at Nike, the digital transformation at Equifax was triggered by a crisis. For context, Equifax is an American consumer credit reporting agency that competes with the likes of Experian, TransUnion, and, to some extent, FICO. Equifax collects consumer and business data and then sells this data to corporations, government agencies, and consumers. In 2017 Equifax suffered a security breach as part of which the private records of 143m Americans and 15m UK citizens were exposed.[8][9] Given the confidential nature of these records, this was a detrimental blow to the company's credibility. Those affected sued the company and many corporations shifted business from Equifax to their competitors. The new management team realized that in order to save the business they would have to rearchitect their IT and security stack. At the beginning of 2019 they announced that as part of this reorganization they would move their entire IT infrastructure to the Google Cloud Platform (GCP).

The transition to the cloud is being led by Bryson Koehler, the former CTO of IBM Watson and IBM Cloud Platform.[10] Koehler joined IBM in 2015 when they acquired The Weather Channel. He was the CTO responsible for moving the business to the cloud and thereby transforming it into the leading global weather data provider. He joined Equifax in June 2018 and was responsible for determining the cloud-first strategy and choosing Google as the main provider. There are some interesting case studies on Koehler's work at The Weather Channel, including two from Harvard Business School.

The ongoing technology transformation for Equifax is based on three principles. First, the company will become cloud-native, which means that all of their infrastructure will be moved to the cloud. Second, all applications will be built in a way that services and components can

5 *https://www.youtube.com/watch?v=u_7xMcZl3D0&ab_channel=AmazonWebServices*

6 *Cassandra is an open-source NoSQL database managed by the Apache Foundation and the basis of Amazon's Dynamo database offering.*

7 *https://www.youtube.com/watch?v=f7FSpT7jrX4&ab_channel=AmazonWebServices*

8 *https://www.slideshare.net/AmazonWebServices/becoming-a-nimble-giant-how-amazon-dynamodb-serves-nike-at-scale-dat320-aws-reinvent-2018*

9 *https://www.ft.com/content/c70d723a-941f-11e7-a9e6-11d2f0ebb7f0*

10 *https://www.ciodive.com/news/ibms-bryson-koehler-becomes-equifax-cto/525741/*

be easily assembled and connected using standard APIs (microservices architecture). This enables developers to reuse parts of their code and make sure that all applications can run on the same data infrastructure. The third principle is rationalization. In 2018 the company had six to eight versions of the same applications, running across different systems and customers. In the cloud there will only be one version, which is continuously updated. This means that development resources can now be focused on one version, rather than split across eight.

A transition of this sort is only possible with the right team in place. CTO Koehler replaced 50% of his leadership team within the first six months of arriving, by recruiting talent from the best technology companies. This exemplifies the difficulty associated with a true digital transformation and the likely employee resistance that organizations face.

The three principles are executed across five different tracks. First, the company will build what they call a common "data fabric." The idea is to build one centralized data warehouse on the Google Cloud Platform that allows them to ingest, govern, enrich, and manage all of this data. This will replace a multitude of current purpose-built systems and siloed databases. It will allow customers access to data in real-time versus days or weeks based under the current system. They will also be able to take advantage of Google's industry leading querying, analytics, and machine learning tools. Most of the US data will be fully migrated by the end of 2020. Second, all of the customer applications will be rebuilt for or migrated to the cloud in a way that they can be delivered in a software as a service (SaaS) format. All new applications will be built to this standard. Equifax already initiated this process in 2016, when they launched their InterConnect SaaS product. Third, Equifax's 4,000 global customers will be migrated from legacy systems and integrations to the SaaS products. Fourth, customer support software will be moved to the public cloud. As part of this transition, the company will also deploy SaaS solutions for customer support, such as Salesforce. Fifth, some of the operational business systems will be moved to SaaS applications. For instance, Equifax will move to Gmail for email and collaboration, move their Oracle financial systems from on-premise to AWS, and their sales management applications will be moved to the Salesforce cloud.[11]

The financial impact Equifax expects from this cloud transition is twofold. The simpler of the two is the additional revenue growth they anticipate. The management team believes that through their ability to launch more data products at a faster rate (100 products implemented in 2020 vs. 10 in 2019) they will be able to sell more to existing customers and attract additional ones. In addition, the increased speed at which the data and analytics will be available to customers should increase their data consumption. Finally, their access to Google's machine learning and analytics capabilities should further accelerate their ability to sell more of the higher cost data insights rather than the raw data. In essence, they aim to offer a product that customers will purchase more of and at higher prices.

The second financial impact is the anticipated cost reduction from the transition. The two areas of cost savings will be technology costs and development expenses. Technology costs represent approximately 45% of the COGs and management expects to reduce these by 15%, which implies a 7% reduction in total COGs or $90m applied to the 2019 cost base. The company also projected a 25% reduction of product development expenses. Based on the 2019 numbers the combined cost saving would be $125m or 3-4% of revenues.[12] Equifax also expects a 35% reduction in capital spending. Applying this 35% to the 2019 numbers results in a $115m saving. Consequently, the total pretax cash savings on 2019 numbers would be $240m or 7% of revenues.[13]

Equifax is a good case study because they provide an unusual level of detail with regard to their cloud

11 *Equifax Q1 2019 Earnings Call*

12 *Equifax Q1 2020 Earnings Release Presentation*

13 *Equifax Q1 2020 Earnings Call*

Dynamo Cougar x IBX x Ibovespa Performance up to October 2020 (in R$)

Period	Dynamo Cougar	IBX	Ibovespa
60 months	166.4%	108.6%	104.8%
36 months	86.9%	29.1%	25.6%
24 months	72.7%	11.7%	8.1%
12 months	24.3%	-12.0%	-13.3%
Year to date	5.7%	-17.9%	-18.8%

NAV/Share on October 31 = **R$ 1,405.166503600**

DYNAMO COUGAR x IBOVESPA
(Percentual de Rentabilidade em US$)

Period	DYNAMO COUGAR* Year	DYNAMO COUGAR* Since Sep 1, 1993	IBOVESPA** Year	IBOVESPA** Since Sep 1, 1993
1993	38.8%	38.8%	7.7%	7.7%
1994	245.6%	379.5%	62.6%	75.1%
1995	-3.6%	362.2%	-14.0%	50.5%
1996	53.6%	609.8%	53.2%	130.6%
1997	-6.2%	565.5%	34.7%	210.6%
1998	-19.1%	438.1%	-38.5%	91.0%
1999	104.6%	1,001.2%	70.2%	224.9%
2000	3.0%	1,034.5%	-18.3%	165.4%
2001	-6.4%	962.4%	-25.0%	99.0%
2002	-7.9%	878.9%	-45.5%	8.5%
2003	93.9%	1,798.5%	141.3%	161.8%
2004	64.4%	3,020.2%	28.2%	235.7%
2005	41.2%	4,305.5%	44.8%	386.1%
2006	49.8%	6,498.3%	45.5%	607.5%
2007	59.7%	10,436.6%	73.4%	1,126.8%
2008	-47.1%	5,470.1%	-55.4%	446.5%
2009	143.7%	13,472.6%	145.2%	1,239.9%
2010	28.1%	17,282.0%	5.6%	1,331.8%
2011	-4.4%	16,514.5%	-27.3%	929.1%
2012	14.0%	18,844.6%	-1.4%	914.5%
2013	-7.3%	17,456.8%	-26.3%	647.9%
2014	-6.0%	16,401.5%	-14.4%	540.4%
2015	-23.3%	12,560.8%	-41.0%	277.6%
2016	42.4%	17,926.4%	66.5%	528.6%
2017	25.8%	22,574.0%	25.0%	685.6%
2018	-8.9%	20,567.8%	-1.8%	671.5%
2019	53.2%	31,570.4%	26.5%	875.9%

2020	DYNAMO COUGAR* Month	DYNAMO COUGAR* Year	IBOVESPA** Month	IBOVESPA** Year
JAN	-0.1%	-0.1%	-7.1%	-7.1%
FEB	-13.0%	-13.0%	-13.1%	-19.3%
MAR	-41.2%	-48.9%	-39.3%	-51.0%
APR	10.6%	-43.5%	5.6%	-48.3%
MAI	9.9%	-37.9%	8.6%	-43.9%
JUN	12.1%	-30.3%	7.8%	-39.5%
JUL	18.0%	-17.8%	13.9%	-31.1%
AUG	-3.5%	-20.7%	-8.2%	-36.7%
SEP	-5.4%	-25.1%	-7.0%	-41.1%
OCT	-1.3%	-26.1%	-3.6%	-43.2%

Average Net Asset Value for Dynamo Cougar
(Last 12 months): R$ 5,008.5 milhões

(*) The Dynamo Cougar Fund figures are audited by Price Waterhouse and Coopers and returns net of all costs and fees, except for Adjustment of Performance Fee, if due.

(**) Ibovespa closing.

transition. We concede that the company is more data reliant than the standard enterprise and therefore the benefits from moving to the cloud are easier to understand. Notably, we spoke with the former CTO of a >$100bn market cap enterprise about the Equifax transition and he was adamant that the focus of a digital transformation should be a better customer proposition and the subsequent revenue growth that follows, rather than cost savings. We agree with this view, but since Equifax was already a data and technology-based operation prior to the transformation, they are also likely to benefit on the cost side. Furthermore, revenue synergies are harder to quantify and sell to the market than cost synergies, given the complexity of these transformations, which might explain why they have been more vocal about these to date.

Conclusion

In our view, Nike and Equifax provide two tangible examples of the dramatic impact that a digital transformation can have on a business. The table stakes of a digital transformation are no longer an ability to compete incrementally against peers, but rather the long-term viability of the business. In order to continue to make concentrated investments with a multi-year time horizon, it is crucial to closely monitor the digital transformation efforts of the companies in the portfolio as well as the latest developments at the leading technology providers. This impacts our research and potential investment opportunities in two ways. First, it allows us to better determine whether the state of transformation of a traditional business poses a risk or an opportunity that is not understood by the market. Second, the foundational understanding of these transformation efforts has enabled us to focus some of our research efforts on finding investment opportunities in the new enterprise software market.

Rio de Janeiro, November 23, 2020.

Please visit our website if you would like to compare the performance of Dynamo funds to other indices:

www.dynamo.com.br

DYNAMO

DYNAMO ADMINISTRAÇÃO DE RECURSOS LTDA.
Av. Ataulfo de Paiva, 1235 / 6° andar. Leblon. 22440-034. Rio. RJ. Brazil. Phone: (55 21) 2512-9394. Fax: (55 21) 2512-5720