

Self-Driving Business on the Horizon? A Look at Enterprise AI

Summary

In evaluating recent artificial intelligence/machine learning (AI/ML) trends and innovation put forth by enterprise vendors, we see incremental benefit for CRM, NOW and SPLK. Though HDP is well-positioned, we believe the window of opportunity is narrow. Across the stack, we think the revenue opportunity lies with application vendors that address specific business problems or analytical tool providers that help enable a solution. The core AI engine layer of the stack will likely be dominated by large vendors such as Amazon.com, Apple, Facebook, Google and Microsoft; smaller players are likely to either get acquired or run into significant headwinds, in our view.

Key Points

Persistent trend toward enterprise automation. The current trend of accelerating AI/ML innovation continues the decades-long push to further automate enterprise workloads. We look at the evolution of key technologies over time as well as current leading innovations in the space from the large cloud players and emerging vendors, alike. Based on our conversations and industry analysis, we assess potential supporters and detractors of greater adoption in the near- to medium-term.

Standalone AI seems like a hard model to commercialize. As with those in the open-source space today, we believe business models reliant solely upon traction of core AI are likely to face similar scaling challenges. While we believe monetization strategies are very much in flux today with different models being tested, the near-term opportunity lies in tools that enable such workloads and end-user applications that solve specific business problems.

Mixed takeaways for core AI and analytics. In our view, 1) **AAPL** continues to embed AI functionality into its products, though incremental revenue opportunity remains limited; 2) big data capabilities required to support ML implementations is supportive of **HDP's** growth; however window of opportunity seems narrow for them; 3) **SPLK** seems to have multiple ways to win in this space and will be among the top vendors to gain from the trend; 4) **DATA** seems to have a very limited opportunity in the space.

CRM and NOW appear very well-positioned. We note that 1) **CRM** continues to be one of the biggest proponents of machine learning, with early organic investments plus robust M&A; Einstein is seeing strong interest from customers but production deployments could take some time; 2) **NOW** is likely to sustain momentum as it pushes for deeper enterprise automation across functional areas with investments in product.

Company	Symbol	Price (7/12)	Rating		PT
			Prior	Curr	
Apple Inc.	AAPL	\$145.74	-	Neutral	\$150.00
Hortonworks, Inc.	HDP	\$13.18	-	Neutral	\$14.00
Paycom Software, Inc.	PAYC	\$69.80	-	Neutral	\$70.00
Paylocity Holding Corporation	PCTY	\$46.81	-	Neutral	\$45.00
salesforce.com, inc.	CRM	\$90.40	-	Buy	\$100.00
ServiceNow, Inc.	NOW	\$112.26	-	Buy	\$125.00
Splunk Inc.	SPLK	\$60.74	-	Neutral	\$60.00
Tableau Software, Inc.	DATA	\$64.47	-	Neutral	\$68.00
Ultimate Software Group, Inc.	ULTI	\$215.74	-	Neutral	\$220.00

Source: Bloomberg and Mizuho Securities USA

Abhey Lamba
Managing Director, Americas Research
+1 415 268 5517
Abhey.Lamba@us.mizuho-sc.com

Parthiv Varadarajan
+1 415 268 5503
Parthiv.Varadarajan@us.mizuho-sc.com

PLEASE REFER TO PAGE 41 OF THIS REPORT FOR IMPORTANT DISCLOSURE AND ANALYST CERTIFICATION INFORMATION. Mizuho Securities USA LLC does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the Firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

Executive Summary

We assess the recent rapid developments in the fields of artificial intelligence (AI) and machine learning (ML) and evaluate their far reaching impact across industries. Our report takes investors through the evolution of AI/ML technologies over time, enterprise adoption patterns and potential commercialization strategies being implemented around by leaders in the space. We conclude our piece with takeaways for our coverage universe, delineating companies across the AI stack and evaluating near to long-term opportunities for vendors and their offerings.

Our key takeaways are highlighted below:

- **AI appears to be hitting a critical mass today, at least in terms of mindshare.** Trends in the hardware stack (rising GPU and NAND penetration, IoT connectivity) in addition to innovation on the software side (public cloud adoption, open-source acceptance, maturity of the data analytics stack) have elevated the prioritization of artificial intelligence and machine learning among technology vendors.
- **Consumer adoption leads the way but significant potential for enterprises exists.** Current penetration of consumer AI technologies is deepest for natural language processing (NLP), image detection and general AI (e.g. autonomous driving) use cases. On the enterprise side, we find that penetration is in its very early stages, with some early inroads from Robotic Process Automation (RPA) players followed by more sophisticated solutions now being launched by enterprise software vendors for, both, back-office (e.g. claims processing) and customer facing (chat bots) use cases. While there are some near-term challenges with respect to adoption, we think the enterprise space offers significant opportunities for software vendors.
- **Core AI capabilities could get commoditized over time.** Given the scale requirements and data moats that hyperscale providers such as Amazon, Microsoft, Google and Facebook enjoy, coupled with the trend of these vendors pursuing open-API and open-source strategies on the AI front, it will likely be very challenging for emerging vendors to compete in the core AI engine space. That is, an emerging vendor offering standalone core AI solutions (i.e. algorithms, models) is likely to face similar challenges that current open-source software vendors face.
- **Value likely to move upward in the stack toward analytical tools and applications.** Over time, we think value will shift to vendors offering analytical tools that can enable enterprises to build applications using frameworks supported by the public cloud vendors. In addition, application developers could create greater differentiation for their offerings and potentially improve product mix by tiering their solutions based on varying levels of AI capabilities.

- **Salesforce.com, ServiceNow and Splunk appear best positioned to benefit from enterprise AI in our coverage.** We look at takeaways for our coverage of hardware and software companies by respective positioning within the AI stack:
 - **Apple (Neutral, \$150 PT):** The company's AI/ML play is within the core AI layer, which will likely be commoditized. That is, while consumers will come to expect functionality enhancements via AI over time, they will not necessarily be willing to pay a large premium for this functionality. We think Apple has to offer this functionality to protect its broader franchise, though, we see limited routes for Apple to meaningfully capitalize on the trend.
 - **Hortonworks. (Neutral, \$14 PT):** Over time, we expect Hadoop to be one of the key enablers of machine learning technologies; that said, the architecture has a narrow window of opportunity to capitalize on current innovation trends, in our view. In the interim, our conversations continue to suggest that customers are looking to extract value from their early technology investments in the space, which could weigh on near-term spending patterns, or even the technology's credibility.
 - **Salesforce.com (Buy, \$100 PT):** The company remains one of the biggest proponents of enterprise AI in the application layer, having acquired significant IP over the recent past on its way to launching Einstein for use cases across sales, marketing and service. While production deployments remain in their early stages, we find that customer interest around the product's capabilities is robust.
 - **ServiceNow (Buy, \$125 PT):** ServiceNow's AI play lends itself well to the company's broader strategy of automating enterprise workloads. Recent M&A supports the company's ambitions in the space as management looks to penetrate functional areas outside of core ITSM (e.g. security, HR and customer service).
 - **Splunk (Neutral, \$60 PT):** We expect ongoing momentum of the company's solutions (supported by an expanding integration library of ML models for better machine analytics capabilities). We think Splunk can benefit as it competes, both, as an analytical tools provider in addition to an application vendor. For the stock, we believe the pace of the transition toward a ratable business model could continue to pressure financials and execution in EMEA needs to improve before we turn more constructive.

- **Tableau (Neutral, \$68 PT):** The company has built some AI/ML capabilities but we find them largely marginal in nature. As such, we do not expect Tableau to be a meaningful beneficiary of the trend.
- **Ultimate Software (Neutral, \$220 PT):** Ultimate Software has acquired assets such as Kanjoya (surveys for sentiment analysis) that are being integrated into the core product suite. While there could be an upsell opportunity as the product/use case matures, we expect limited contribution in the near-term.

Growing Interest in the Machine Learning Area

Automation has been a persistent enabler of societal progress over time. Since the industrial revolution, humans have sought to augment their own productive capabilities by deploying machines to handle manual, repetitive and error-prone tasks with an eye toward driving efficiency gains. This trend has intensified and has become more far reaching - the impact of machines, first felt in more traditional industries such as manufacturing and transportation during the early 1900s, then spread across industries, in some shape or form, in the latter half of the century. As technology has evolved and become more easily consumable in the past decade, the trend toward automation has gained steam, with significant implications for every single industry, in our view.

This dynamic is manifesting with respect to the artificial intelligence (AI) and machine learning (ML) areas today. These concepts are not new, having existed in different forms for several decades now; that said, AI/ML have noticeably expanded their potential applications off late owing to a number of factors. With the digitization of everything, for instance, data has quickly become one of the most valuable assets for companies enabling the trend toward cognitive and automated workflows across a variety of functions and industries. The concept of AI (dating back to the 1960s) has gradually transitioned from being a research-oriented intellectual exercise to a mainstream application in the consumer and enterprise worlds, with rising acceptance over the last few years. In fact, Diane Bryant, head of Intel's Datacenter Group, mentioned at the company's recent investor day that just 7% of servers in 2016 were dedicated to machine learning and deep learning workloads, and that AI compute cycles are expected to grow 12x by 2020. Data volume growth in addition to rising end-user appetite to draw insights from the data has enhanced interest in the domain. Further, notable advancements in various areas of the technology stack have made it more practical to process the data in large volumes for analytical workloads. For example, mainframe-type computing power was required for AI/ML workloads in the past, but technological progress in the hardware and software sectors has made it possible for a wide range of applications to operate on a commodity stack. We discuss some of these advancements below:

- **Hardware:** According to Mizuho's semiconductor research team led by Vijay Rakesh, the recent enablement of mainstream AI, among other things, was the result of: 1) the rising penetration of GPUs into the datacenter (which perform 50-100x better than CPUs for AI training given their significant parallel processing capabilities for multi-threaded applications) on the back of better per-unit-of-performance economics; 2) increasing mix of NAND (versus disk), that offers faster read-write performance for real-time data access, as prices move lower; 3) ongoing connectivity improvements as everyday devices incorporate 4G/Wi-Fi capabilities to share data with a central repository for real-time analysis and workflows.
- **Software:** Based on our field work and industry analysis, we believe the software space enables AI adoption via: 1) secular trend toward rising public

cloud adoption, allowing on-demand infrastructure consumption for hardware intensive AI workloads; 2) greater acceptance of open-source and open-API software architectures, backed by some of the largest proponents and early adopters of artificial intelligence in the industry; 3) the ongoing maturity of the data analytics stack, aided by open-source projects (e.g. Apache Spark), in dealing with high-volume unstructured and dark data; 4) rising acceptance of containers, allowing for quick and efficient spin-up/down of AI workloads.

Open-source Adoption Likely to Further Catalyze ML Traction

From a software perspective, we believe the rising acceptance of open-source solutions, especially by the enterprise, is very supportive of AI/ML adoption in the market. In particular, given the rising strategic value of massive repositories of customer data, open-source innovation in the analytics stack (from the likes of the Apache foundation) has been very influential with respect to innovation trends across the machine learning stack. Further, in addition to Open API strategies, leaders in the space have open-sourced their frameworks in order to drive adoption of their platforms. As a result, developers and software vendors do not need to expend effort building the core AI elements but, can focus on building applications that utilize the underlying AI infrastructure put forth by larger vendors that have the benefit of scale. More anecdotally, we point to data from the Apache foundation, ranking projects by the number of committers. As of June 2017, of the top 20 projects, 7 of them were in the big data or database buckets. We expect adoption of open-source software to continue as customers look to avoid lock-in with any one vendor over the longer term. Additionally, open source offers significant flexibility to developers when they are working on leading edge solutions.

Exhibit 1: Apache Project Ranking by Number of Committers

Project	Category	# Committers
Hadoop	Database	168
OpenOffice	Content	139
CloudStack	Cloud	112
Ambari	Big Data	90
Cordova	Library	89
Subversion	Build Management	81
Cocoon	Database	80
Geode	Database	79
MyFaces	JavaEE	77
Geronimo	Network-Server	72
Flex	Web Framework	69
Hive	Database	64
HBase	Database	63
Felix	Network-Server	61
ActiveMQ	Network-Client	59
Camel	Network-Client	58
CouchDB	Big Data	58
Struts	Web Framework	58
Maven	Build Management	57
Aries	Library	55

Source: Apache Foundation, Mizuho Securities USA

With a significant rise in funding from venture capitalists over the past few years, AI start-ups have built product platforms, with key dependencies on open-source projects from Apache such as Hadoop, Spark, Kafka and Hive. The Hadoop architecture, for instance, lends itself to use cases that require massive parallel processing in a distributed manner; although the execution of such architectures might be challenging, this capability presents value to an enterprise looking to, for instance, glean insights from its large store of transactional or customer usage/behavioral data (e.g. a company looking to optimize its marketing campaigns based on historical purchasing data). In addition, we increasingly find that the incumbent players in the technology industry are moving away from proprietary platforms toward open-sourcing their intellectual property, thereby, striving to be the platform upon which an AI/ML ecosystem develops. In 2015, IBM open-sourced its SystemML framework in order to allow programmers coding in R and Python to implement and scale their algorithms on popular big data architectures such as Apache Spark and Hadoop. Over the years, IBM has also made dozens of APIs available for developers to build cloud applications that take advantage of Watson at

the back-end (for instance, BuildOnMe's HR chat bot relies on Watson's reasoning/NLP engine to automate help desk requests). Similarly, in 2015, Google released its TensorFlow machine learning library to the open-source community, allowing developers to build and train AI systems using neural networks. Microsoft released its Distributed Machine Learning Toolkit (DMTK) training models to the public around the same time period, further enticing developers with the ability to deploy these models on the Azure cloud. Facebook's Caffe and Torch frameworks are popular, as well, given the social network's large user base. Amazon offers similar services that include its Natural Language Understanding (NLU), Automatic Speech Recognition (ASR) and text-to-speech (TTS) in addition to image recognition capabilities, with the option of deploying these workloads in the AWS cloud.

Expect Value to Move toward Applications and Away From Core Infrastructure

The key takeaway here, which we explore at greater length over the course of this report, is that as these frameworks make their way into the public domain, they largely become commoditized; instead the real value in these frameworks comes down, ideally, to functional or vertical-specific use cases that can be developed by other vendors or consultants for consumer and enterprise applications. The raw algorithm, itself, is unlikely to command a whole lot of value. In other words, customers are unlikely to be willing to pay for the use of these models. In terms of adoption, we are beginning to see some preliminary use cases take shape in the consumer and enterprise spaces over the past few years (e.g. Amazon Alexa and Apple Siri, which we will discuss later in the report in addition to AWS and Azure ML). Anecdotally, machines are able to beat humans in games such as Go and chess; that said, we believe that the hype around AI/ML far exceeds practical applications out there in the world, as it stands today. Rule-heavy applications such as games are likely significantly easier for machines to understand relative to solving practical business problems (for which there are fewer rules). However, if we look over the next 10 years as the cycle of adoption, we believe we are in the very early stages with current enterprise adoption largely being driven by the bleeding edge crowd for relatively simpler applications. We expect adoption to become more mainstream (and production worthy) as the ecosystem matures and some of the kinks around first generation applications are worked out.

Artificial Intelligence (AI) Versus Machine Learning (ML) and Deep Learning (DL)

With most enterprise software vendors vying for mindshare of AI/ML workloads, the marketing message can often cloud the product's actual capabilities. Before we attempt to understand the potential growth opportunities that artificial intelligence presents to software vendors, we begin by understanding: 1) the technologies underlying AI and the points of differentiation between them; 2) the building blocks required to be present before meaningful adoption can occur for each of these technologies. Our semiconductor research team has previously written extensively on the topic in its report titled [*Mi-Tech Vol. 46: AI & Deep Learning; Primer to a Revolution*](#) dated January 18, 2017. In addition, we draw our conclusions from industry analysis and our conversations with participants and analysts.

Basic Artificial Intelligence has Significant Limitations

Artificial Intelligence (AI) is the catch-all, umbrella term for the innovation occurring in the space today; since its earliest days going back to the 1960s, the term AI has been used to describe human-like cognitive capabilities (perception, reasoning, learning etc.) for machines. However, in its less sophisticated form initially, 'classic' AI came to represent automated behavior and responses that adhered to conditional 'if-then' rules that mimicked human behavior. While productivity was significantly enhanced with its advent, this type of machine behavior had a serious limitation – the AI was restricted to the rules that it was originally programmed with and was, thereby, limited in its ability to reason and learn on its own as data and circumstances changed. Robotic Process Automation (RPA) solutions (e.g. Blue Prism, UiPath) today are largely rules based and pre-programmed to perform certain tasks.

Advancements in Neural Networks Have Helped Overcome Basic AI Limitations

Simple Neural Networks, also known as Artificial Neural Networks (ANNs) evolved out of this limitation as a subset of AI. These networks were structured to operate like neurons in the human brain (i.e. transmitting information to one another and other parts of the system). Machine Learning (ML) is one such technology that was born out of ANNs over time, as the IT stack has developed to enable better data access and readability, machines can ingest information and make inferences about the future, with high levels of statistical accuracy, using correlations previously discovered in analyzing the data. In supervised learning, the owner 'trains' the model using vast amounts of 'labeled' data. Post training, the model is able to discern patterns in the data with a high degree of accuracy. The model is then exposed to 'unlabeled' data to execute certain types of workloads (e.g. image classification). Conversely, in unsupervised learning, the model is not trained for discrete workloads; instead, it is initially exposed to unlabeled (rather than labeled) data, with the expectation that the machine builds new learning models on the fly to glean patterns in the data autonomously.

There is a wide range of use cases for machine learning models. A few of the more common functional algorithms and their use cases are described below:

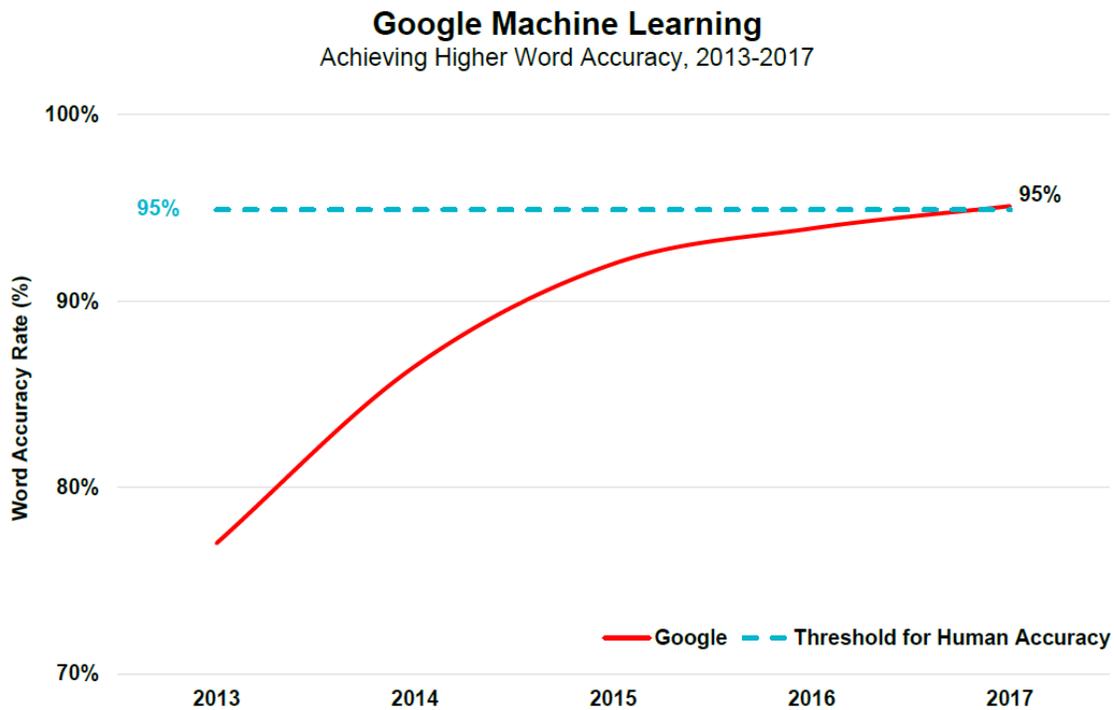
- **Classification:** The model is trained with discrete categorizations of groups that are specified by the trainer; ultimately, the machine is able to categorize ‘new’, unclassified objects into the optimal category with an acceptable level of accuracy.
- **Regression:** Based on prior correlations between variables that have been discovered in the data, the model attempts to predict outcomes using a set of discrete data inputs.
- **Clustering:** The model studies discrete objects with various characteristics in a data set; it then groups objects with similar characteristics into buckets such that objects in a bucket are more similar than dissimilar. The difference between clustering and classification is that the learning process, in the latter, begins with specified categories, to which, new objects are assigned, whereas, in clustering, the model creates and assigns categories on its own.
- **Association Rule Learning:** The model analyzes the underlying data set to understand the incidence with which two (or more) objects appear together. A common use case for this is in online shopping; the website recommends a purchase of another item based on what you currently have in your shopping cart.
- **Anomaly Detection:** The model determines patterns in the underlying data to flag discrete data points that are unusual relative to historical norms. A common use case for this algorithm is in fraud detection (insurance, banking etc.) or threat prevention (cyber security).

Recent Technological Innovation Has Paved the Way for Deep Learning Adoption

Deep learning (DL) is a variant of machine learning; relative to ML, DL technology requires a vastly greater number of ‘layers’ of information (i.e. deep neural networks) in training the machine, thereby, allowing more complex problem-solving with potentially better accuracy. In the past, the DL analytical process was very computationally intensive and largely out-of-reach for most practitioners given the significant infrastructure requirements. In recent years, however, with the availability of public clouds, rising use of GPUs with massive parallel processing capabilities and the ongoing maturity of the big data ecosystem, DL has become more viable as an artificial intelligence construct. Notably, one of the most prominent applications of DL today is in the field of Natural Language Processing (NLP), both for speech recognition and text-to-speech use cases. Projects such as Watson had the capability, early on, to take inputs of text/audio, establish a connection between the incoming data and an existing knowledge base and, finally, come up with an appropriate response. The penetration of deep learning for AI workloads on the back of enhanced infrastructure availability has driven improved accuracy of results over time.

Anecdotally, in the chart below, we illustrate the accuracy gains by Google’s machine learning algorithms for NLP use cases. The NLP model’s accuracy rate, based on the company’s definition, has increased from sub-80% levels in 2013 to around 95% today (Google’s threshold for the metric was 95%, with such levels indicating human-like speech or text).

Exhibit 2: Improving Accuracy of Google’s Word Accuracy



Source: KPCB Internet Trends (2017)

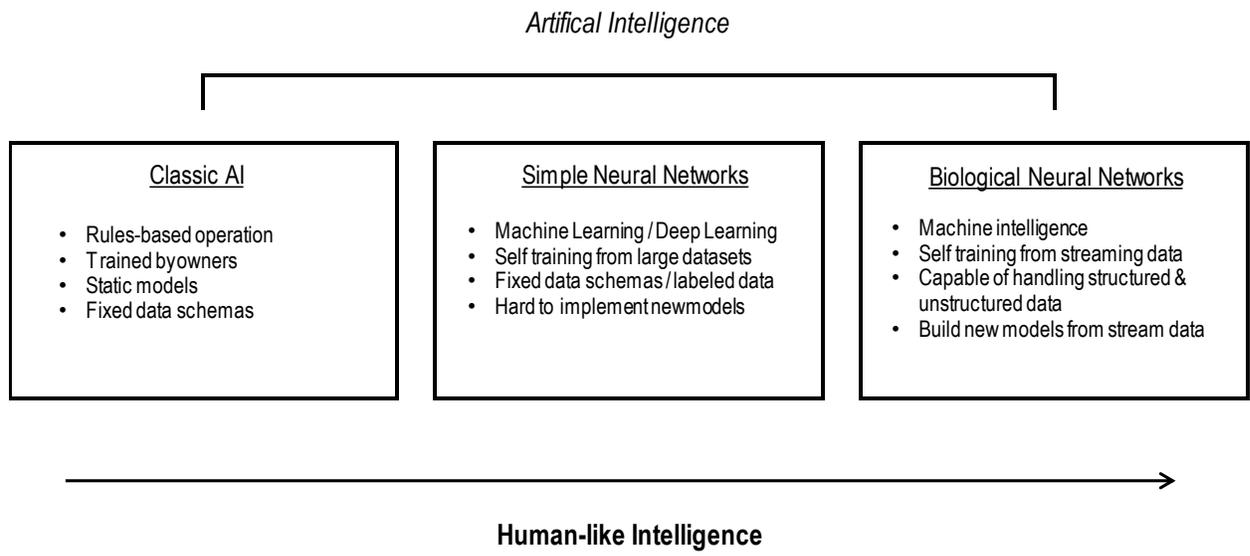
Biological Neural Networks Push Toward Human-Like Cognition

This brings us to the leading edge of AI technology today – Biological Neural Networks (BNNs), more commonly known as machine intelligence. The technology was born out of the limitations of previous AI constructs and brings machines closer to perceiving, reasoning, learning and reacting as the human brain does in an unstructured environment. Similar to ANNs, biological neural networks can train themselves on large datasets of pre-ordained models. However, BNNs can automatically build new models gleaned from constant data streams, using even unstructured data that does not fit a schema with which the program was previously trained. As the ‘new’ data streams in and collects in the repository, the program can then run historical analyses across the dataset to look for patterns, establish/classify various correlations and, ultimately, probabilistically predict certain outcomes. This ability to ingest ‘new’ information and build new models on-the-fly, with which the program trains itself, is a significant development toward human-like cognition in the field of AI. Whereas today’s technologies largely fall into the ‘supervised’ AI bucket,

these next-gen technologies potentially pave the way for machines to create new knowledge without the aid of humans (i.e. unsupervised learning). Companies such as Numenta and Vicarious are innovating on this front by developing core BNN technologies for a variety of use cases. For example, Grok is an IT operations analytics product (using core technology licensed from Numenta) that enables real-time monitoring IT infrastructure and applications, similar to Splunk. In another use case, Cortical.io offers a text analytics solution; the underlying NLP engine sits on top of Numenta’s HTM engine; example use cases include plagiarism detection in the legal and journalism spaces as well as text analytics for contract review purposes (once again in the legal realm). Lastly, Loop AI is making a play in unsupervised learning with its Loop Learning Appliance meant, combining hardware and software, for on-premise deployments in specific verticals (Mizuho’s semiconductor research team hosted a call with management with takeaways here: [Mi-Tech Vol. 47: NVDA/Loop AI-DL Transforms Enterprise Intelligence](#)).

The following exhibit summarizes various technologies available for artificial intelligence.

Exhibit 3: Classification of Various AI Technologies



Source: Mizuho Securities USA, KD Nuggets

Hurdles in Adopting AI for Enterprise Workloads

As previously mentioned, enterprise AI, in our view, is in the very early stages of a likely decade (or even multi-decade) long cycle. Our conversations indicate that enterprises are just about turning the corner with respect to deploying capital and resources toward AI-related projects for production workloads. To gauge the pace of adoption through the cycle, we start with Salesforce.com's evaluation of the current environment and assess the potential obstacles to widespread acceptance. Alongside IBM, Salesforce.com has been among the most vocal champions of the technology, with significant investments in the space in order to differentiate its offerings and to identify new customer use cases. The aggressive strategy on this front lends itself well to the company's significant presence as a repository of transactional data for its customers. As such, in our view, the company is one of the best positioned application software companies to meaningfully benefit from this trend.

Availability of Real-Time and Clean Data Important for Successful Implementations

According to Salesforce.com's chief scientist, Richard Socher, enterprise AI needs to progress through three phases before it achieves mass adoption. First, the underlying data needs to be usable and accessible. An ongoing issue that the BI ecosystem has been trying to solve is the limited usability of business data being collected from customers, suppliers, supply chain partners (among others) residing in legacy corporate data stores. Further, in addition to this transactional information, enterprises are being exposed to more unstructured data in the form of social media streams, images, graphs and sensor data. Today, sub-optimal data preparation and integration processes are a significant obstacle to greater near-term enterprise AI adoption. Many believe that until the underlying data architectures are built in a way that supports efficient consumption of new information by AI models, great strides in software-enabled enterprise automation will be unlikely. Granted, some industry participants do not believe it is a requirement that the data be completely 'clean' before it can be consumed by a model; in fact, raw data that contains a certain amount of 'noise' is instrumental for the models to discern the signal from the noise. Instead, we believe the underlying BI architecture needs to be capable enough to quickly consume structured or unstructured data from several different sources in order for the models to discern usable signals and insights.

To that end, companies such as Hortonworks and Cloudera provide support services and enterprise functionality (i.e. security, compliance, governance etc.) for several open-source projects in the Hadoop ecosystem aimed at solving this real-time data consumption problem. That said, we believe mainstream adoption for production workloads is likely a few years away. After initial deployments of Hadoop over the past couple of years, enterprise customers, now able to ingest and store various forms of data, are beginning to look at extracting value out of their early investments, which is proving harder to do than previously expected. As the ecosystem matures and new consumption models emerge, we believe the time-to-value for Hadoop implementations will come down; however, we also think we are a few years away from such an environment. To that end, we expect ongoing, gradual enhancements to

data prep and integration workloads over time (prominent vendors include Alteryx, Talend and SnapLogic) in support of more efficient enterprise data processing.

Algorithms Need to be More Robust

Second following the data issue is the required robustness of the algorithm in responding to changes in the environment around it. For an AI model to run unsupervised in the wild would necessitate controls around its potential behavior as the stream of data that it consumes changes. With respect to this, we believe there is a general sense of apprehension among most enterprises, especially for customer facing applications, given the potential for public relations mishaps and damages to the brand. For instance, Microsoft's conversational bot, Tay, which is meant to read social media sentiment and respond to tweets, had to be pulled offline in light of unexpected behavior; based on intentionally offensive tweets directed at the bot from the community, Tay, similarly, began replying with its own offensive remarks.

Human Factors and Trust Issues Will Need to Evolve

Lastly, according to Socher, effectively incorporating AI into operational workflows poses the most challenging hurdle to broad enterprise adoption. We concur – firstly, as with any emerging technology, the early hype around its promise can often cloud the true value that early adopters hope to extract out of the product. With respect to automating a workflow that would otherwise have been completed by a human employee, we think there are a couple of areas of skepticism that need to be overcome: 1) the employee's worry that she would be made redundant if her core workflows were automated by AI; 2) the trust required in the AI to complete the workflow effectively and accurately, even once the first hurdle is overcome. This second point is especially noteworthy; in our view, there is a certain inertia associated with the way core, operational (especially mission-critical) workflows are performed in large enterprises. Replacing current manual processes with more automated ones enabled by AI, while undoubtedly the way of the future, is likely to face an extremely high threshold for results before employees/management become comfortable deploying it with limited-to-no supervision (for e.g. self-driving cars). We believe current interest from customers is most likely in the early proof-of-concept phase with deployments for core operational workflows likely further out in time.

Importantly, the challenges around operationalizing workflows using AI remain top of mind for large enterprise customers. Based on our industry conversations, one potential way vendors can better help their customers more quickly ascertain the value proposition of their product, especially for AI where the use cases across functions and industries are vast, is to consider a functionalized/verticalized approach to product and go-to-market. We believe a model where, for instance, the vendor licenses the core AI model to the customer without participating in operationalizing the product, is unlikely to be a viable long-term business model, given the complexity of the implementation and supporting infrastructure. Instead, a potentially quicker route to near-term adoption lies with vendors that develop functional or vertical-specific applications by combining AI technology with domain

and industry expertise. Once this manner of AI has penetrated the enterprise and proven its capabilities for core workflows, it can then spread across functions toward wider adoption, in our view. As it stands today, we believe we are a few years away from that inflection in the adoption.

Potential Business Models for Commercializing Artificial Intelligence/Machine Learning

Next, we look at potential monetization strategies in commercializing ML product offerings. As has been the case with prior emerging technologies, we believe customers of ML products are, generally, wary of the near-term value they can extract out of current deployments. To be clear, we believe customers advocate for the immense value in applying artificial intelligence to enterprise workloads; however, the time-to-value is challenging to estimate at this point of the adoption cycle, thereby, potentially resulting in some hesitation in large-scale spending. In addition, as is typical for emerging technologies in their early days, talent capable of deploying such solutions tends to be limited, which further increases the cost of implementation and the necessary value hurdle required by customers. That said, as we look out over the next few years, we believe vendors, from start-ups to more established enterprise software players, have a few different options with respect to IP monetization. In our view, potential business models include:

- **Open API:** Initially, the early leaders in enterprise AI followed an open API model in order to drive adoption of their respective platforms. For simple use cases such as speech recognition and image detection, developers could tap into the underlying engines via APIs to enhance functionality of their own applications. This freed developers from having to develop their own back-end engines and allowed them to focus on innovation at the application layer. In addition, the pool of organizations that can successfully develop such capabilities is limited, given the tremendous scale requirements in product innovation. For instance, Amazon offers SDKs for its core AI services (natural language processing, automatic speech recognition, image recognition, text-to-speech, machine learning etc.) to developers in order to facilitate the creation of applications for specific use cases by the community. By not charging for the underlying technologies, Amazon hopes to draw developers to the AWS platform, likely, with the eventual goal of reaching a critical mass such that certain AI workloads become standardized on the platform. Similarly, Caffe is an open-source engine which is widely used for image recognition workloads, with a capacity to process up to 60 million images per day using NVIDIA GPUs, with potential to expand use cases over time. In our view, Microsoft, IBM, Google and Facebook appear to be approaching AI as a platform play so that a robust ecosystem of applications may develop on top of their core technologies. To that effect, Microsoft, at its 2017 Build Conference, indicated that its Cortana platform has 140 million end users, and that 130K developers had signed on to use the company's Bot framework.
- **Open-source and Machine Language-as-a-Service offerings:** As mentioned earlier in the report, we believe the trend toward open-source software is strong today, and will likely remain strong, over the foreseeable future. In our view, this trend is likely to be a key enabler of greater

workload automation as artificial intelligence penetrates deeper in the enterprises. To that end, we also note that the incumbent enterprise IT vendors have predominantly open-sourced their early generations of AI, ML and DL technologies with hopes of wider adoption, greater pace of innovation and, importantly, tighter customer lock-in. This includes frameworks such as DSSTNE from Amazon, Microsoft's Cognitive Toolkit, Google's TensorFlow and Caffe2 from Facebook.

That said, the current trend toward open-sourcing by the large vendors could render core AI technology a commodity over time. Looking at things from this perspective, in the absence of a specific horizontal or vertical application built on top of the core technology, the stand-alone AI model from any one of these vendors requires significant investment on the part of the customer in order to generate value out of the technology. With the industry moving in that direction, we believe start-ups and emerging players are unlikely to build sustainable business models purely around licensing of core AI intellectual property. A business model solely predicated on monetizing core AI intellectual property (via premium offerings, support or services) will likely face the same challenges that any other company attempting to commercialize open-source software faces. Numenta, for example, open sources its core Hierarchical Temporal Memory (HTM) machine intelligence technology via a general public license, with an agreement that any development work on top of the architecture must be released back to the community. Today, given the small size of the company, there is no premium support or services offering from Numenta; however, the company does have other models of monetization for commercial applications of its technology. Customers that develop proprietary applications on top of Numenta's HTM engine must purchase a commercial license that follows consumption or usage-based pricing. In all, similar to business models based on other open-source products, we believe one based on open-source AI could fall prey to growth and profitability pressures over time.

Interestingly, however, the incumbent large vendors are monetizing their early deep learning investments via their public cloud platforms at the IaaS layer. While the source code is available to the community, effective implementation of the model remains challenging for most developers (given that computing power on par with Google, for example, is hard to come by). As such, the public cloud vendors offer machine learning as-a-service, whereby, customers can deploy a TensorFlow model on Google's public cloud infrastructure (that is optimized for the workload) and pay for it on an hourly basis. The major public cloud vendors, all, offer a variant of this service (AWS ML, Azure ML, Google Cloud ML etc.). In the case of Google, the offering includes around 25 models for various use cases that package the data, train the model and perform an assessment of the results (to allow further iteration). The library of models is expected to grow over time to cover a broader array of use cases. Similarly, Microsoft charges

customers per hour of compute for its fully managed machine learning cloud service; however, the company further monetizes its assets by charging a fee ‘per prediction generated’ for the application that sits on top of the Azure ML platform. Pricing starts with a free tier (with limited experiment duration, storage space and the number of accessible nodes etc.) with a higher-end SKU that lifts these restrictions.

- **Commercial applications:** As ML becomes more mainstream and vendors look to monetize their offerings, some appear to be developing horizontal products for specific enterprise functions, while others seem to be going down the verticalization route. Our conversations with venture capital investors suggest that, early on, funds flowed through to both business models; however, more recently, venture investors appear to be targeting more vertical-specific applications given the potentially quicker returns (due to more controlled product development and sales and marketing costs for verticalized offerings). Further, verticalized offerings allow customers an accelerated payback period due to lack of significant customization needs.
 - **Horizontal applications:** Based on commentary from our software coverage in addition to a broader assessment of the industry, workloads across most (if not all) functional areas within the enterprise present an opportunity for AI in some shape or form. Among others, the more prominent enterprise use cases today involve sales force and marketing automation, customer support, business intelligence and security. For example, Infer offers predictive analytics for marketing use cases; the solution goes beyond traditional lead scoring to help marketers build customer profiles (using signal data from large datasets) in order to better segment and target prospects. On the security side, Deep Instinct is an example of a company that applies deep learning techniques in order to uncover threats prior to infection; most solutions offer detection once infected or prevention of the spreading of the infection after detection. Lastly, using our prior example of Numenta, we note that in addition to its open-source strategy, the company offers a commercial license which entitles Numenta a 2% royalty on all HTM-dependent revenue streams that its customers generate. In summary, we believe horizontal offerings can be monetized via traditional subscription plus support arrangements or more creative royalty agreements (though, we think the latter will remain niche given the potential cost burden on customers as the deployment scales). Within our coverage, we note that Salesforce.com and ServiceNow are already pursuing a horizontal AI product strategy for now, seemingly looking to penetrate their installed bases in their respective strongholds (e.g. ServiceNow with intelligent routing of incoming service requests). However, we would not be surprised if their AI offerings evolve toward becoming

more vertical specific in order to gain more rapid penetration while providing a potential pricing uplift.

- **Vertical applications:** Vertical software, as a business model, has advantages/detractors. On the positive side, given the industry focus, the product can be targeted very specifically to certain use cases, thereby, limiting development costs. For go-to-market, once the company develops an understanding and signs a few key reference customers, sales and marketing costs can scale very effectively as adoption picks up. That said, a common criticism of verticalized offerings is that the addressable market can be narrow. Although we think the points above hold true for enterprise AI offerings, the sheer diversity of potential use cases across industries for core technologies may necessitate verticalized offerings given the significant industry specific knowledge required to effectively operationalize workloads with artificial intelligence. In the absence of vertical offerings, vendors would need to heavily depend on their ecosystem of partners to develop feasible solutions in order to drive adoption (similar to an IoT play).

Looking at the current environment, again, while all industries face penetration by AI, headline verticals such as autos (driverless cars), financial services, retail/e-commerce, insurance and healthcare present some of the most active opportunities for early-stage investors today. Sentient Technologies, one of the highest funded AI start-ups (according to CB Insights) having raised over \$140 million from venture investors, is one such company pursuing a verticalized monetization model. The company's Sentient Aware product targets e-commerce use cases, enabling more personalized browsing experiences for website visitors. By studying a customer's historical browsing patterns and drawing similarities of the customer's behavior to 'cohort' spending patterns, the solution then customizes the layout of the landing page to incentivize a purchase. Similarly, the Ascend product, also for e-commerce use cases, offers A/B optimization techniques so that customers may offer different shopping experiences to website visitors depending upon various characteristics (e.g. browser or some other demographic indicator). What differentiates the product is the ability for marketers to try several campaigns on incoming traffic all at once and, in the process, dynamically devote more resources to the best performing campaigns.

In another example, Atomise, focuses on AI for healthcare use cases. The company's deep learning technologies aim to significantly reduce drug discovery cost and time by running the massively iterative process with limited supervision from humans; in the process, the algorithms can potentially uncover how existing

chemical compounds can be repurposed to treat new ailments. For the legal industry, Everlaw's NLP offering is able to search millions of documents, legal transcripts and other unstructured data to make attorney workflows more efficient. Similar examples include Civitas Learning and Climate Corporate for the education and agriculture industries, respectively. In all, given the difficulty in evaluating the return on AI investments (or a time-to-value), we believe potential customers could be more inclined toward such verticalized offerings of AI for specific use cases until the technology proves its worth. As such, we expect vertical-specific applications to be an important driver of enterprise AI adoption over the next few years.

Regulatory Impact on Enterprise AI Adoption

In our view, regulation around emerging technology trends has, historically, lagged innovation. With artificial intelligence and, more specifically, the rising penetration of robotic automation across industries, we believe regulatory bodies have begun formulating legislation and conferring with industry groups and participants, alike. Today, the areas that are most active in the legislative sense include autonomous vehicles and drones, which is not surprising, given the near-term implications for public safety as these technologies become more mainstream. In both cases, regulatory uncertainty, or even unfavourable rulings, has a very direct implication for near-term adoption of autonomous vehicles and drones. Looking at enterprise AI from a software perspective, we think the biggest area of concern from regulatory as well as customer policy point of view will be around data ownership and associated ability of large technology providers to process user information. For instance, Salesforce.com stores and processes customer information for all of its clients; however, there could be regulatory issues regarding the extent to which the company can potentially use this information in machine learning or analytical contexts. Further, some customers may be unwilling to share their client information with competitors. As such, we believe many such issues would need to be addressed in order to drive greater adoption of machine learning in the enterprise space. Laws could also develop around the necessity to audit enterprise algorithms and automated processes.

Current State of Machine Learning Adoption: Consumers versus Enterprises

Similar to various emerging technologies that preceded it, AI and ML have seen early penetration in the consumer realm, arguably, to a higher magnitude relative to enterprise adoption. However, also in-line with prior innovations, mainstream consumers are generally unwilling to pay a significant premium for AI-enabled services. Consumer facing companies, instead, often include functionality in product offerings as a way to differentiate themselves from the competition, usually without charging for the up-sell (for example, Siri in the iPhone). If effectively deployed, consumers are often unaware of the hefty AI processing working in the background. In the following section, we look at some of the key AI technologies that have penetrated consumer lifestyles:

- **Natural Language Processing (NLP):** NLP is one of the most pervasive use cases of artificial intelligence in the consumer space today. The technology, broadly, covers two areas: automated speech recognition, or ASR (which allows machines to understand human speech) and text-to-speech, or TTS (which involves machine generated voice from text). In particular, voice recognition is not a new concept, with cell phones including some sort of functionality since the early 2000s; that said, functionality has vastly improved with rising smartphone penetration and greater mobile data access. Apple Siri, Amazon Alexa, Microsoft Cortana and Google Assistant are all examples of consumer-facing conversational AI platforms today. The last two have the distinct advantage of being natively coupled with the companies' search engine technologies, thereby, allowing quick voice-enabled retrieval of information from the Internet. In fact, according to Gartner, 20% of Google searches on Android in the US are done by voice; similarly, Microsoft says that 25% of Cortana searches on the Windows taskbar are done by voice. Microsoft recently expanded Cortana's capabilities (similar to Alexa), allowing it to tap into third-party applications (weather, news, music, Skype etc.) rather than just the Bing engine. Apple, too, recently acquired Lattice Data (for \$200 million according to media reports) to enable better discovery and usage of unstructured data, likely, with an eye to enhance Siri's capabilities. NLP is gradually penetrating consumer lifestyles to a greater level – at the core of Google Home is Assistant, which enables a smart home environment for applications such as temperature control, lighting, regulation of power consumption etc. Over time, as these NLP engines mature (better conversational capabilities) and become more adept at learning from prior behavior, we expect penetration levels to rise across consumer lifestyle product offerings.

Another interesting use case of NLP on the consumer side (with vast potential to spread to the enterprise) is in proof-reading and anti-plagiarism use cases. As an example, Grammarly develops a plug-in that enables proof-reading of e-mails, text messages, social media posts etc. using more than 250 syntax/punctuation rules in the English language. On the plagiarism

side, the company’s offering compares text in the user’s document against text across over 8 billion web pages to identify content that appears to be plagiarized. The company recently raised \$110 million to build out its freemium offering across multiple use cases.

- **Image Recognition:** Similarly, image detection software has been around for a long time; however, it has penetrated consumer workloads in a significant way in the last decade or so enabled by advances in the AI ecosystem. Deep learning techniques (i.e. multi-layer analysis using deep neural networks) are particularly prevalent for image analysis workloads. Facebook, for example, runs facial recognition algorithms when users upload new pictures to their profile; by doing so, it is able to suggest the names of friends to tag in the picture based on its analysis of the user’s existing tagged albums. Such technology is also being used for security use cases; Android includes functionality for unlocking a smartphone using facial recognition using the device’s front-facing camera. A rapidly developing area of AI/ML involves facial/object recognition in video with use cases in security (i.e. scanning CCTV footage) or marketing (i.e. brands can advertise by adding tags to YouTube videos based on images detected in videos that are most closely associated with the brand). We expect the ecosystem to continue to mature over time, with incumbents likely to absorb attractive assets to build out their core offerings. According to Gartner, by 2025, deep neural networks will automate tagging of 80% of all image/video tagging workloads.
- **General AI/ML:** More generally, AI/ML underlies a significant portion of consumers’ daily interactions with technology-enabled products without us realizing it. Netflix’s suggestion engine relies on the user’s prior viewing history and crowd-sourced recommendations of similar shows based on analytics across its entire subscriber base. Amazon’s ‘other members purchased’ functionality works in the same way. Lastly, autonomous driving, pioneered by Tesla, is one of the consumer-facing use cases with the highest visibility today. Our colleague, Vijay Rakesh, has covered the underlying ADAS technologies in-depth in his recent report ([*Mi-Tech Vol. 45: ADAS 2017 Another Step Closer; LiDAR Comes of Age*](#)).

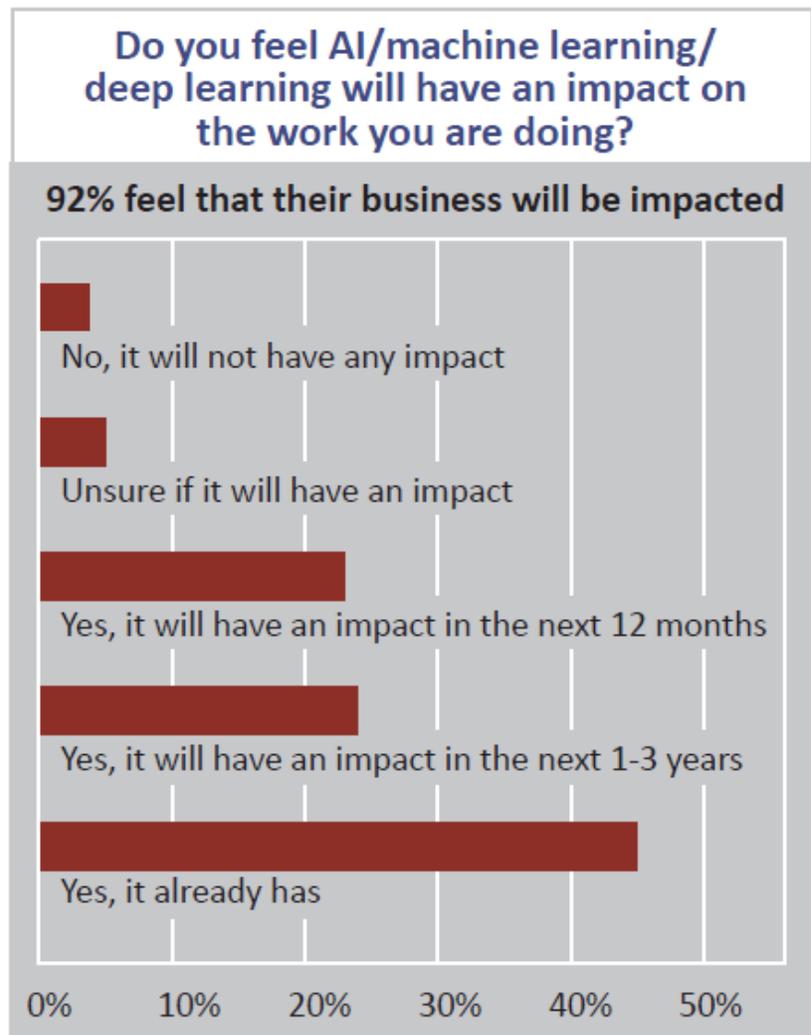
Enterprise Adoption Seems to be in Very Early Stages but Potential is Likely Significant

On the enterprise side, we believe adoption has been spreading across functions and use cases over the years. According to insideBIGDATA, over half of large enterprises globally are experimenting with AI in some shape or form. Some have recognized the secular trend and are investing aggressively. For example, GE has launched Predix, a PaaS/IoT play to collate data and extract insights from its industrial machinery installed base with a view to enhancing efficiency and minimizing downtime. Relatedly, Deloitte reports that a third of global enterprises are using bots of some kind within the finance and accounting functions; a fourth are

using AI/ML in their procurement or HR processes. An early iteration of a bot-enabled enterprise comes in the form of Robotic Process Automation (RPA). Vendors such as Blue Prism, UiPath and Automation Anywhere offer RPA solutions to automate highly repetitive back-office processes (for example, transferring information from a highly standardized purchase order into an ERP system). Though, we are very early in the process, PwC estimates that 45% of enterprise activity can be automated, producing an estimated \$2 trillion in workforce savings. RPA, in our view, is forcing the BPO industry to rapidly evolve by allowing enterprises to bring BPO-handled workloads back into the enterprise, while improving scale, cycle time and auditability along with a quick ROI. That said, RPA, as a technology is fairly ‘weak’ (i.e. task-specific); although it brings efficiency to highly repetitive, rules-based workflows, according to industry analysts, the technology is unlikely to mature into one that can learn on-the-fly by ingesting and analyzing unstructured data. Instead, the technology will need to be supplemented with more ‘intelligent’ AI so as to bring greater levels of efficiency to manual enterprise processes. In our view, collective innovation in various parts of the stack is likely to drive higher enterprise adoption over time. Other back-office areas such as claims processing are seeing inroads by AI as well; Shift Technology (based in France) is a SaaS offering that enables insurance providers to detect fraudulent claims by analyzing patterns across the customer’s repository of documents. Security, compliance and fraud detection workflows, collectively, present a large opportunity for enterprise AI.

In addition to back-office processes, we expect higher penetration levels for front-office processes such as marketing, sales and support as the underlying technologies and use cases mature over the next few years. A customer-facing area that has significant potential for enhancement via automation is customer service – chat bots (powered by NLP as part of offerings such as IBM’s Watson and Microsoft’s Cortana) are increasingly being deployed as the first layer of customer interaction. Though this workflow has existed for a long time, we believe the technology is seeing greater acceptance among enterprises as a means of two-way communication with customers (rather than the primarily one-way tool of the past that was used to collect information from the customer before being connected to a customer support representative). Zendesk’s offering, for example, includes an Answer Bot that responds to basic customer inquiries using NLP and an internally generated knowledge base of information (i.e. hours of operation, customer support numbers etc.). More broadly, we see evidence of ongoing enterprise adoption across functions; a recent survey by insideBIGDATA of 130 participants across functions (researchers, data scientists, IT administrators, business analysts, management etc.) indicated that 92% of respondents believed that their roles have already been or are likely to be impacted (over the next 3 years) by artificial intelligence.

Exhibit 4: Impact of AI/ML on Enterprise Workflows



Source: Guide to AI & Deep Learning by insideBIGDATA (2017)

Lastly, in the chart below, Topbots has listed the emerging (and to a lesser extent, legacy) players that are addressing functional enterprise workloads (customer service, engineering, security, marketing etc.) with their artificial intelligence offerings

Exhibit 5: Examples of Enterprise-focused AI Companies



Source: Topbots

On the flipside, while we maintain little doubt that this megatrend toward a more AI-enabled enterprise plays out over time, a few things could have an impact on the pace of adoption:

- Regulation:** As mentioned previously, the regulatory environment continues to respond to this nascent, yet rapidly evolving, technology trend. In essence, over time, legislation could develop around data usage; more specifically, laws curbing the potentially abusive usage of data by enterprises, particularly, for consumer facing purposes. In addition, stepping back, one of the biggest hurdles to toward greater AI penetration in the enterprise, in our view, is the sustained and latent lack of trust among humans in deploying such solutions (due to, both, the threat of job redundancy and AI capabilities in fulfilling the tasks assigned to it). Though regulation, in both instances, is unlikely to be able to assuage these concerns, adoption could be

slowed until industry standards around implementation develop in a manner that satisfies labor and other advocacy groups.

- **Skills shortage:** A commonly held view among technologists and industry observers, alike, is that the pace of technological innovation over the recent past meaningfully outstrips the ability to consume this innovation. Take, for example, the Hadoop ecosystem – with often several rapidly developing projects competing for mindshare as a means to tackle a single use case, it can often be challenging for consumers of this innovation to determine the ultimate winner over time. Given the nascence of many of these technologies, the bench of talent and skills tends to be shallow, with the result that development efforts are often hampered by a lack of knowledgeable heads despite investment dollars being allocated to assessing these technologies.
- **Cost and ROI:** There tends to be a general inertia with respect to reassessing, re-designing and re-operationalizing enterprise workflows to incorporate new technologies. Among the variety of reasons that existing processes are hard to rip-and-replace is cost – and similar to technologies that have preceded AI, very early investments in solutions may fail to garner the expected results. In fact, according to Gartner, by 2019, just 5% of organizations that have invested in AI will derive value from such technologies. In our view, it is incumbent upon vendors, especially with a technology that has such far-reaching use cases and capabilities, to offer a strong path toward value realization in order to drive rapid adoption.

Coverage Takeaways: Impact of AI to the Universe of Software Names

In this section, we offer investors a framework to think about opportunities in enterprise AI as they assess the technology’s impact over the immediate and medium terms. In segmenting this far-reaching opportunity, we maintain the traditional delineation of infrastructure from application technologies, as we illustrate in the diagram below.

Exhibit 6: Illustration of AI Stack



* Facebook implemented Caffe for its development efforts

Source: Mizuho Securities USA

At the lowest level of intelligent applications is the infrastructure stack that includes basic computing, storage and analytical framework to develop machine learning applications. Data and analytics solutions represent the next layer that relies on underlying compute engine for core processing. These solutions are also used to develop applications or supplement features within end user applications. The uppermost layer includes applications that utilize the underlying infrastructure to address some specific business workflows. These layers are described in detail below.

- Core AI / Open APIs:** As mentioned earlier in the report, these technologies include the training models and inference engines that support the development of more intelligent applications. Core AI offerings include those from Apple, Microsoft, Amazon, Google, Facebook etc. There are some smaller players also addressing specific user needs but it remains to be seen how many of them gain sufficient scale to build thriving businesses. We think this layer will eventually

be dominated by large cloud providers. These smaller players are likely to either get acquired or might get squeezed out over time.

- **Data & Predictive Analytics:** We believe the business intelligence stack (i.e. the infrastructure components that enable rapid, high-volume analysis of data) is a key enabler and driver of enterprise AI adoption. In our view, this area of our coverage universe represents the near-term opportunity for investors looking to play the trend toward a more intelligent and automated enterprise. Within this bucket, we will explore the opportunity set for data processing systems (Hortonworks) in addition to analytics/consumption (Splunk, Tableau). We believe vendors in this space could be attractive acquisition candidates since they enhance the value of, both the underlying platform providers as well as the application layers above them.
- **Applications:** Further up the stack, we believe enterprise application vendors are likely to rely on artificial intelligence to build horizontal as well as vertical capabilities within their solutions. As the technology matures from nascence, vendors are actively building out their product portfolios to incorporate capability, either via organic efforts or M&A. In our view, the trend is strong and secular enough to necessitate product functionality at some level given customer interest and general industry enthusiasm. That said, there remains significant uncertainty whether there is a pricing/revenue opportunity around the inclusion of AI functionality in enterprise applications. In our view, the appetite from customers to experiment with AI is likely to be offset by the limited near-term value realized during experimentation. As such, for the next few years, we believe customers will demand incremental functionality without necessarily being willing to pay for it. However, over time, the upsell opportunity exists as vendors look to differentiate their offerings. As such, vendors will likely continue to build out their product offerings in order to competitively differentiate themselves in the market. Within our coverage universe, we look at horizontal applications of AI across the following areas. Once applications reach a certain maturity level, vendors will likely start building vertical specific functionality in order to accelerate adoption.

 - **Sales, marketing and customer support:** This is one of the earliest customer-facing applications of AI that is already being deployed in the enterprise. Penetration areas include sales force automation to enable better deal scoring, marketing to aid more personalized engagement with prospects, and customer support to create efficiencies around low-level activities.
 - **Enterprise Service Management:** The core value proposition of ESM relies on higher automation of previously manual workflows across IT, security, customer support and HR. To the extent that vendors can enable more intelligent automation of workflows across these areas, we believe capabilities will grow over time.
 - **Human Capital Management:** As HR departments become more data savvy, the workforce analytics space has garnered significant attention to enable better hiring, engagement and learning/development.

- **Security:** This is one of the earliest domains that AI/ML has penetrated. Vendors offer various solutions that execute anomaly detection workloads in order to find intrusions of malware. More cutting-edge offerings use machine learning techniques to study patterns in the data to prevent intrusions before they occur.

Company Specific Takeaways

Key takeaways for the companies in our coverage universe are highlighted below.

Apple – Neutral-rated with \$150 PT: Tough to Directly Monetize the Opportunity Despite Making Significant Investments

Since 2011, when Siri was launched, the company has made ongoing investments in artificial intelligence and machine learning across the platform. In addition to penetration across platforms (i.e. from just iOS to MacOS and WatchOS), Siri's breadth of capabilities has expanded too, as the underlying models and technologies have become more robust. After initially restricting the engine to Apple's native applications, the company opened up Siri's functionality via APIs to allow integration with third-party apps as part of the iOS 10 release. This followed similar moves by Amazon, Google and Microsoft to spur development of AI-enabled apps on the platform. Besides NLP, Apple's AI play includes image recognition capabilities for intelligent photo tagging in the Memories application (where albums are auto-created based on the software's recognition of faces that appear more than once). More recently, at WWDC 2017, the company announced further enhancements to its AI toolkit. With iOS 11, SiriKit will enable on-device execution of machine learning models (rather than in the cloud previously) with an eye toward more optimized performance, better efficiency and improved security. This follows similar announcements from Google (TensorFlowLite announced at I/O a couple of months ago) and Facebook (Caffe2Go) to enable mobile-optimized machine learning models. Apple also announced functionality to allow translation of ML models from other platforms (for example Caffe) to iOS in a more seamless manner.

Overall, the company continues to position Siri as a text/voice capable virtual assistant with abilities to suggest workflows to the user based on browsing history, app usage and upcoming calendar appointments. Beyond the smartphone, we also point out the company's intentions to become a bigger part of the consumer lifestyle. HomePod, in our view, is an attempt to further penetrate the home with support for third-party 'smart home' appliances (i.e. bulbs, thermostats, doors etc.). Our key takeaway here is that, similar to Amazon and Google, we expect rising penetration of more advanced machine learning models into smartphones and consumer devices; however, we believe consumer willingness to pay a premium for these capabilities alone is largely non-existent. In addition, Apple appears to be late to the home assistant market. While it could face challenges in catching up with Amazon in the space, it likely needs to continue investing in the technology as an ecosystem play around its strong, core smartphone business. An important point to make – while Amazon can benefit, to a significant extent, from higher sales of its home assistant (search traffic with potential conversion to its retail business), we think Apple's upside is more limited. Although the company's services business continues to grow, we think the home assistant is unlikely to lead to meaningful upside to the segment. Instead, Apple needs a significant acceleration in smartphone growth or App Store revenues in order to monetize investments in areas such as the home assistant market. Overall, we do not expect Apple to be able to drive incremental monetization of ML

models through the developer community as the strategy will likely remain focused on developing and providing IP for free in order to enable greater adoption of the platform.

Hortonworks – Neutral-rated with \$14 PT: Well-Positioned for Now but Window of Opportunity Seems Narrow

Hortonworks, along with the larger Hadoop space, is well-positioned to benefit from greater penetration of AI and ML in the enterprise. The company falls into the ‘enabler’ bucket with solutions that not only allow faster and more efficient processing of new data types flowing into enterprise systems, but also provide better visibility of previously ‘dark’ data already ingested. The requirement of more usable data for purposes of training machine learning models across enterprise functions creates a natural tailwind for Hadoop-based solutions over the near-to-medium term.

Within the Apache ecosystem, developers have few different options when choosing between execution engines upon which to run their ML algorithms. Each framework comes with a similar library of algorithms for different use cases and workloads (for example classification, regressions, clustering, recommendation etc.). An early contender was Mahout, which saw initial success in the community given its basis in MapReduce that has gained in popularity as a processing engine. In fact, some early customers include AOL and Foursquare that use Mahout to power their recommendation engines. However, MapReduce is limited by its constant disk access for read-write cycles, with the result that Mahout was disadvantaged when running machine learning algorithms (which are heavily iterative and dependent on cycle times). Today, Spark has gained popularity in the community for its in-memory capabilities (particularly for inference rather than training workloads) and practitioners are shifting their attention to development on Spark from MapReduce while Mahout has expanded to include support for Spark. Spark MLlib offers a deep and fast growing library of algorithms for a variety of machine learning workloads, resulting in a significant uptick in investment dollars (which, in turn, drive ongoing development efforts). Another alternative to both Mahout and Spark MLlib comes from a company named H2O.ai – the company’s Sparkling Water offering combines an open-source library of ML/DL algorithms with Spark’s processing engine (see our semiconductor team’s recent report: [Mi-Tech Vol. 51: H2O and NVDA GPUs – A Look at AI in Healthcare](#)). H2O supports Windows, Linux as well as several distributions of HDFS (Cloudera, Hortonworks, MapR etc.).

As machine learning penetration of enterprise workloads rises over time, we expect sustained adoption of Hadoop solutions (given that the associated technologies lend themselves to big data workloads) and, thereby, an ongoing revenue opportunity for vendors that serve this market. In fact, leading up to its April IPO, Cloudera heavily marketed itself not as a company offering a Hadoop distribution, but as one that enables machine learning. In part, we believe this is due to the challenges involved with commercializing and building a sustainable business model around open-source Hadoop distribution. We expect this to persist – although Hortonworks continues to work toward positive free cash flow levels on a sustainable basis, we think there is

limited potential for strong near-term profitability in the space as investment needs remain high and skill shortages make it challenging to scale deployments. On the competitive front, given the rising prominence of Spark as the center of innovation in the ecosystem (and, in particular, for ML use cases), the framework appears to be taking workload share from other Apache projects. As a result, pure-play vendors focused on all things Spark, such as DataBricks, could see rapid uptake of their solutions at the expense of more general purpose Hadoop vendors such as Hortonworks and Cloudera. That said, the rapid innovation in the Hadoop ecosystem potentially exposes Spark to disruption by another superior project a few years down the line. Overall, we think the space will likely remain highly dynamic and while there are opportunities for Hadoop vendors, the window to capitalize on them seems narrow. As such, vendors such as Hortonworks need to move quickly or risk losing out to other smaller companies that are narrowly focused on the opportunity.

Salesforce.com – Buy-rated with \$100 PT: Leader in Embracing Machine Learning for Core Enterprise Use Cases

Salesforce.com has been one of the most active enterprise software vendors with respect to incorporating machine learning technologies into its product portfolio. In addition to organically building out its team of research scientists, the company's product strategy has been meaningfully supported by recent M&A: 1) PredictionIO (for close to \$60 million) allowing users to build customized predictive frameworks and recommendation engines; 2) RelateIQ (for around \$390 million) which enables automated scanning of emails and other customer interactions to make engagement recommendations; 3) MetaMind (amount unknown), or Einstein Vision, as it is known today, uses neural networks for image detection and classification use cases; 4) BeyondCore (for around \$110 million), known as Einstein Data Discovery today, enables pattern recognition of trends across data sources.

Earlier this year, at its FY18 kickoff event, the company provided a progress report with respect to the integration of these acquisitions into product portfolio; for example, Salesforce.com's Predictive Vision Services (PVS), which is based on MetaMind, allows users to upload large sets of image data in order to train the model toward an ultimate goal of being able to classify unstructured data on-the-fly. A use case was demonstrated at the event by Coca Cola – a store clerk could take a picture of a cooler in order to determine how much to re-stock, upcoming promotions to take advantage of, up-sell opportunities to hit volume discounts etc. (i.e. a marketing cloud opportunity. Similarly, a field service technician could take a picture of a damaged part in order to quickly identify it in order to procure and purchase a replacement part. Salesforce.com has exposed the APIs for PVS in order to allow customers and developers to build custom apps for specific use cases. Another offering, Predictive Sentiment Service (PSS), can be used for sentiment analysis workloads; for example, the PSS API can consume streaming social media data from Twitter to evaluate real-time changes to sentiment around a company's brand following an event.

Similarly, customer service workflows can be further automated in response to changes in sentiment. To support these real-time capabilities, Salesforce.com has also integrated a distribution of Apache Kafka into Heroku so that developers on its platform do not need to deal with setting up a stand-alone implementation of the open-source framework. Together, in our view, this constitutes a coherent medium-term product vision; however, we believe that the underlying technologies need to mature before it becomes a reality in production customer instances. More broadly, echoing the company's Chief Scientist, Richard Socher, we believe that the Einstein framework is unlikely to be a separate SKU, though, management is likely to use machine learning capabilities as a way to create wider differentiation in SKUs, thereby, driving more favorable product mix. Overall, we believe customers expect some level of AI functionality to be available on-demand across the product suite, especially given the premium pricing that the company already commands. Relatedly, we note that, according to a recent survey conducted by IDC, about 34% of Salesforce.com customers are early AI adopters; we expect this to rise over time as the company enhances its platform with embedded AI functionality.

On the competitive front, Microsoft has been actively investing in its CRM offering from a product and go-to-market standpoint, though, an uptick in adoption has yet to materialize likely due to lackluster messaging. Alongside other VCs, the company recently invested in a start-up called Tact that is developing artificial intelligence based solution specifically for CRM use cases. The company is looking to achieve this via a virtual assistant that integrates voice and messaging from smartphone with conventional enterprise apps such as calendar, e-mail, Slack etc. in order to recommend workflows to salespeople with the goal of making them more productive with their prospects. Relatedly, Optimize, a sales forecasting tool, applies machine learning techniques across a variety of data sources (ERP, CRM, economic data, sentiment data etc.) to help companies improve their long-range sales forecasting. Again, this is a very specific application of ML to sales operation use cases. In all, we expect machine learning to continue penetrating sales, marketing and customer support enterprise workloads over the next few years, though, the incremental revenue opportunity will likely be limited in the very near-term, while greater product differentiation could offer value over time.

ServiceNow – Buy-rated with \$125 PT: Significant Opportunity to Benefit from the Technology

ServiceNow, with its roots in enterprise automation, has delivered ongoing enhancements to its product offering with AI/ML capabilities, both organically and via M&A. Notably, the company's 'Intelligent Automation' engine allows customers to run training models on their own instances with the goal of layering predictive workloads on top of their implementations. With respect to product enhancement, management announced its acquisition of DxContinuum earlier this year; the deal gives ServiceNow technology to improve classification and, subsequently, the routing of incoming service requests to the appropriate group within the enterprise (with applications across IT operations, security, HR and customer service use cases). We note that the deal is consistent with the company's long-term product

strategy of enabling a more automated enterprise by defining structured workflows for routine tasks and anomalous instances alike. Later in May, the company announced its acquisition of Qlue, another technology deal that enhances the offering's chat bot capabilities (particularly for low-level service desk requests). Along with its simultaneous investment in BuildOnMe, the transaction is consistent with management's goal of automating enterprise service management workloads such as password reset requests or equipment orders from IT service desks. Interestingly, BuildOnMe's current commercial product, HR Bob, is powered by Watson's Conversation Service on the back-end (connected to ServiceNow's platform via APIs). This further reinforces our view that the underlying algorithms could get commoditized, and that an open API strategy is likely the primary means to driving adoption of the platform.

Management has been clear that pricing decisions around incremental AI/ML functionality have yet to be made, though capabilities are finding their way into the most recent release in May 2017. Over time, we expect inclusion and innovation around such capabilities to allow the company to maintain its pricing power with customers; similar to the case with other vendors, we do not expect the stand-alone AI offerings from ServiceNow. However, given enhanced capabilities, we expect the company to remain the product leader in the ESM marketplace, with the strong likelihood of sustained momentum over the next few years. Over time, similar to Salesforce.com, we think ServiceNow might be able to create better SKU differentiation with potential ASP benefit. In fact, we believe product innovation to be a key supporter of management's strategy and path toward its 2020 goals. As such, we remain positive on the name on potential upside to estimates and reasonable valuation.

Splunk – Neutral-rated with \$60 PT: Multiple Ways to Participate in the Market

With the release of Splunk Enterprise 6.5 in 2016, the company made a bigger push with respect to incorporating machine learning capabilities into the product portfolio. On-the-fly query capabilities on ingested data have been enhanced, enabled by 25+ pre-built ML algorithms (covering classification, clustering and regression) built into the platform; i.e. Splunk Enterprise plus premium apps such as ITSI, UBA and Enterprise Security. The company also offers a Machine Learning Toolkit, which includes 300+ models sourced from the open-source community that can be trained using ingested data for more specialized use cases. For example, the pre-built model package can enable detection and alerting of anomalies in data pulled into the platform, allowing security operations personnel to spot intrusions in real-time. In our view, Splunk can potentially straddle the dual roles of being a leading analytics provider in addition to being an application vendor. The company's solutions targeted at enterprise IT users will likely use AI and ML as key technologies to generate insights into core workloads. Additionally, Splunk can enable other applications in the IoT and technology infrastructure space; machine learning insights (e.g. pattern recognition) can be layered on top of Splunk's core competency of sifting through large volumes of log data. We point out that start-ups such as

Grok, which is enabled by Numenta’s HTM engine, are attempting to address the IT operations market from machine learning perspective; however, we believe Splunk’s product is likely to remain best-in-class over the near-to-medium term, which should drive ongoing new customer and installed base expansion activity. For the stock, we continue to believe near-term optics are likely to be muddled by the ongoing shift toward a more ratable revenue model. As such, we believe upside to estimates is likely more limited given the uncertainty around the pace of the ratable transition. On the execution front, recent hiccups in EMEA need to be addressed before investors can feel comfortable with its ability to deliver on near-term estimates.

Tableau – Neutral-rated with \$68 PT: Limited Opportunity to Participate

Tableau’s product has included NLP capabilities for a while now, allowing users to visualize their data sets by asking questions related to the information. With the 10.3 release, the company is adding further automation functionality by embedding machine learning more pervasively into the platform’s technology. At its recent analyst day event, Chief Product Officer, Francois Ajenstat, indicated that enhancements would include: 1) model-based suggestions of additional data sources/joins that could be useful to the user’s workflows; 2) recommendations for visualization templates based on patterns in the underlying data and; 3) customized alerts for the user once customizable thresholds in the data have been breached. While interesting, we view most of these product updates as largely incremental, yet, necessary to the offering’s evolution in order to remain best-of-breed. In our view, it is unlikely that the company can raise pricing with the inclusion of such features. Over time, we expect competition to continue to narrow to a mostly two-horse race between Tableau and Microsoft, with the latter likely narrowing the functionality gap over time while placing ongoing pressure on pricing. On the business model front, we remain of the view that Tableau’s transition to a more ratable revenue model is likely to pressure near-term growth, while ongoing investments to drive expansion are likely to limit profitability enhancement over the next 12-18 months. However, the company’s model shift seems to be well-understood at this point and likely baked into the stock at current levels. At current valuation, we believe the stock is appropriately priced given risks to the story.

Ultimate Software (Neutral-rated; \$220 PT), Paycom (Neutral-rated; \$70 PT) & Paylocity (Neutral-rated; \$45 PT)

For HR workloads, machine learning techniques are primarily finding their way into non-payroll applications and the workforce analytics space. The SAP SuccessFactors offering, for example, includes a HR virtual assistant for self-service help desk use cases; Workday’s WorkTalk is positioned similarly in the market in order to improve response times and replace manual inquiry response for low-level requests directed to HR. Workday’s Platfora acquisition in 2016 was an analytics play for the company to serve the workforce intelligence market (using data mining techniques across operational and machine data stores). For example, machine learning models can be trained to find correlations between payroll, performance management and learning

data so as to flag employees that are at a higher risk of leaving their employers. For learning applications, programs can be personalized to incorporate employee strengths and weaknesses, manager feedback and career goals using machine learning models. According to Bersin, the corporate learning and development market, sized at \$140 billion in 2017 (which includes human-led coaching and technology/non-instructor led solutions as well) represents a large area of opportunity for vendors looking to bring more productive programs to their customers. Chorus.ai is an interesting application of ML in the area of coaching; specifically sales coaching. The NLP engine embedded in the offering translates speech to text from rep calls with customer prospects and recommends courses of action to improve deal close rates. In another example, workflows around employee feedback and reviews are moving away from the conventional annual cadence to more real-time engagement. Pulse surveys, for instance, where Kanjoya competes (acquired by Ultimate Software) allow HR departments to collect close to real-time data from their employees; HR leaders can then run the data through machine learning models looking for potential correlations with data from performance management or payroll data stores (e.g. data indicating that an employee appears frustrated based on his/her survey responses could be cross-referenced with recent performance data to flag risk of quitting).

To that end, within our coverage, Ultimate Software launched UltiPro Perception, building on its acquisition of Kanjoya in 2016, to offer sentiment analysis capabilities to its installed base. For Ultimate, we believe ongoing penetration of the installed base with non-payroll modules is a key driver of near-term growth for the company. With respect to Paycom and Paylocity, though we expect AI/ML to incrementally enhance the product over time, given the smaller and less sophisticated average customer being targeted, we expect near-term growth to continue to be driven predominantly by cloud payroll penetration. In all, we expect deeper penetration of machine learning into HCM workloads over time as HR departments become more technologically savvy and workforce demands evolve to tend to millennial behavioral patterns. We do not expect material revenue contribution from AI/ML products with HR applications over the next few years.

Price Target Calculation and Key Risks

Apple Inc.

Price Target: Our price target is derived using an equally weighted three-pronged valuation approach, which includes a DCF, Enterprise Value (EV) to FCF, and EV to operating earnings analysis. Our assumptions are based on comparables in the information technology hardware universe. Given these inputs, we reach our 12-month price target of \$150.

Key Risks: On the downside, the biggest risk to Apple's stock could result from the company's inability to keep innovating. The current management team is very capable of delivering on Steve Jobs's vision, however, the stock could be materially impacted if the company were to lose some key executives. Additionally, more than expected slowdown in iPhone sales could push the stock materially lower. On the upside, a very strong iPhone 8 cycle combined with greater attach of services can drive materially higher upside to estimates, which can also help further expand the multiple.

Hortonworks, Inc.

Price Target: Our PT is based on 3.5x our 2018 total sales estimates.

Key Risks: We think lack of profitability is a big concern for investors and could weigh on the stock if the market becomes difficult. Although HDP completed a secondary offering, a tougher spending environment could weigh on cash profitability. Additionally, the technology is still in the early days and needs to mature for broader adoption. Lastly, execution challenges and demand conditions could impact the stock materially.

Paycom Software, Inc.

Price Target: Our \$70 price target is derived using an equally weighted three-pronged valuation approach, which includes a DCF, Enterprise Value (EV) to Sales, and EV to EBITDA analysis. Our assumptions use CY18 estimates and are based on comparables in the enterprise software universe. Based on an EV to Sales multiple of 8.0x, we come up with a share price of \$72. For our EV to EBITDA analysis, we have calculated a share price of \$70 based on a 28.0x multiple. Finally, our DCF assumes a discount rate of 8% and a terminal FCF multiple of 21x; based on these inputs, DCF analysis implies a price estimate of \$70.

Key Risks: Risks to our price target include an intensified competitive environment in the mid-market leading to a deceleration in growth rates. Given the company's growth model which depends on expansion of offices, sustained openings could limit near-term margin expansion. The company also faces the prospect of ACA being repealed; if so, the gap from lost revenue could be challenging to fill, leading to growth slowdown.

Paylocity Holding Corporation

Price Target: Our \$45 price target is derived using an equally weighted three-pronged valuation approach, which includes a DCF, Enterprise Value (EV) to Sales, and EV to EBITDA analysis. Our assumptions use FY18 estimates and are based on comparables in the enterprise software universe. Based on an EV to Sales multiple of 6.5x, we come up with a share price of \$46. For our EV to EBITDA analysis, we have calculated a share price of \$45 based on a 35.0x multiple. Finally, our DCF assumes a discount rate of 7% and a terminal FCF multiple of 30x; based on these inputs, DCF analysis implies a price estimate of \$45.

Key Risks: Risks to our price target include an a tougher competitive environment impacting near-term recurring revenue growth rates. Paylocity also has meaningful

revenue exposure to ACA which, in the event of repeal, could impact near-term growth rates. If the company's broker channel business sees sustained declines, sales productivity and, thereby, growth rates could be meaningfully impacted. Lastly, profitability expansion over the next 12-18 months could be impacted by the recent sales organization changes.

salesforce.com, inc.

Price Target: Our \$100 PT is derived using an equally weighted three-pronged valuation approach, which includes a DCF, EV/FCF, and Enterprise Value (EV) to Bookings analysis. Our assumptions use CY18E estimates and are based on comparables in the enterprise software universe. Based on an EV to bookings multiple of 5.0x, we come up with a share price of \$96. For our EV to free cash flow analysis, we have calculated a share price of \$98 based on a 25.0x multiple. Finally, our DCF assumes a discount rate of 10% and a terminal FCF multiple of 15x; based on these inputs, DCF analysis implies a price estimate of \$108.

Key Risks: Risks to our price target include any missteps in execution, which could lead to margin compression and depressed cash flows. Also, given the company's track record of strong growth, expectations have had a tendency to move ahead of reality. Investors have come to expect strong performance and any disappointment could put pressure on the stock. As the largest pure-play SaaS company, CRM gets lumped in with high multiple cloud names, despite valuation that's more reasonable. As these names experience volatility, so too does CRM. The company could also be at risk of disruption by any of the emerging cloud solutions in the marketplace. We see this as unlikely at this point as the company is very focused on extending its market leadership.

ServiceNow, Inc.

Price Target: Our price target is based on three-pronged approach that uses equal weighted DCF, EV to billings and price to CFFO analysis. On EV to CY18E billings basis, we derive a PT of \$124 based on a multiple of 7.5x. For EV to FCF, we get price of \$126 based on a multiple of 35x. Lastly, for DCF, we used discount rate of 9% and terminal multiple of 14x, implying perpetuity growth rate of 2%. Based on these assumptions, we get a price of \$126. As such, our equal weighted average price target amounts to \$125.

Key Risks: We believe the biggest risk to our positive thesis comes from execution missteps especially as the company is growing rapidly. Secondly, competition will become stiffer over the next couple of years.

Splunk Inc.

Price Target: Our price target is based on a three-pronged approach that uses equal weighted DCF, EV to Billings and EV to FCF analysis. On EV to FY18E billings basis, we drive a PT of \$64 based on a multiple of 5.5x. For EV to FCF, we get a price of \$58 based on a multiple of 30x. Lastly, for DCF, we used a discount rate of 10% and a terminal multiple of 11x, implying a perpetuity growth rate of 0.5%. Based on these assumptions, we get a price of \$60. As such, our equal weighted average price target amounts to \$60.

Key Risks: To the upside, we believe the biggest risk is potential re-acceleration of license revenue growth as the company works through the transition to the cloud, in addition to better-than-expected profitability gains. Further, in the event of an acquisition, the company could fetch a material premium over current levels given the strategic nature of the asset. To the downside, we believe the business model transition

could heavily weigh on license revenue growth and meaningfully slow profitability gains.

Tableau Software, Inc.

Price Target: Our price target is derived using Enterprise Value (EV) to Revenue analysis and EV to Recurring Revenue analysis. For EV to revenue, we use a target multiple of 5.0x on our CY18 revenue estimate to get a price target of \$69. For EV to recurring revenue, we assume a target multiple of 8.0x our CY18 recurring revenue estimate and arrive at a price target of \$67. Given these inputs, we reach our 12-month price target of \$68.

Key Risks: The ongoing risk entails the business model transition to a more ratable revenue model. Any execution missteps can cause a significant drop in the stock price. Upside to estimates could offer tailwind to the stock price.

Ultimate Software Group, Inc.

Price Target: Our \$220 price target is derived using an equally weighted three-pronged valuation approach, which includes a DCF, Enterprise Value (EV) to Sales, and EV to EBITDA analysis. Our assumptions use CY18 estimates and are based on comparables in the enterprise software universe. Based on an EV to Sales multiple of 6.0x, we come up with a share price of \$233. For our EV to EBITDA analysis, we have calculated a share price of \$218 based on a 22.0x multiple. Finally, our DCF assumes a discount rate of 8% and a terminal FCF multiple of 20x (implying perpetuity growth rate of 3%); based on these inputs, DCF analysis implies a price estimate of \$203.

Key Risks: Risks to our price target include an a tougher competitive environment as Ultimate makes a concerted effort to expand in the lower-end of the market. The product, which has typically been adopted by upper-mid-market and lower-enterprise customers could be perceived to be an overkill among lower-end customers, which would impact near-term growth. The move down-market could also weigh on near-term growth and profitability. Finally, the bear thesis on Ultimate could continue to weigh on sentiment in the absence of meaningful outperformance relative to consensus estimates.

IMPORTANT DISCLOSURES

The disclosures for the subject companies of this report as well as the disclosures for Mizuho Securities USA LLC entire coverage universe can be found at <https://msusa.bluematrix.com/sellside/Disclosures.action> or obtained by contacting EQSupervisoryAnalystUS@us.mizuho-sc.com or via postal mail at Equity Research Editorial Department, Mizuho Securities USA LLC, 320 Park Avenue, 12th Floor, New York NY, 10022.

Investment Risks and Valuation Methods can be located in the following section of this research report - Price Target Calculation and Key Risks.

Ownership Disclosures and Material Conflicts of Interest or Position as Officer or Director

As of the date of this report, the research analyst listed on the cover page of this report or household member beneficially owns or has a financial interest in None of the Apple Inc.Hortonworks, Inc.salesforce.com, inc.ServiceNow, Inc.Splunk Inc.Tableau Software, Inc.Paycom Software, Inc.Paylocity Holding CorporationUltimate Software Group, Inc..

Receipt of Compensation

Mizuho Securities USA LLC and or its affiliates makes a market in the following securities: Apple Inc., Hortonworks, Inc., salesforce.com, inc., ServiceNow, Inc., Splunk Inc., Tableau Software, Inc., Paycom Software, Inc., Paylocity Holding Corporation and Ultimate Software Group, Inc.

Mizuho Securities USA LLC and or its affiliates have provided non-securities-related services for Apple Inc. who is or was a client in the past 12 months. The compensation of the research analyst writing this report, in whole or part, is based on MSUSA's annual revenue and earnings and is not directly related to any specific investment banking compensation. MSUSA's internal policies and procedures prohibit research analysts from receiving compensation from companies covered in the research reports.

Regulation Analyst Certification (AC)

I, Abhey Lamba, hereby certify that the views expressed in this research report accurately reflect my personal views about any and all the subject companies. No part of my compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this research report.

Rating Definitions

Mizuho Securities USA investment ratings are based on the following definitions: Anticipated share price change is based on a 12-month time frame. Return expectation is based on the stock's total return. For disclosure purposes (for purposes of the FINRA requirements), we include the category headings of Buy, Hold, and Sell alongside our ratings of Buy, Neutral, and Underperform. (effective July 1, 2017).

Buy: Stock's total return is expected to appreciate above the average total return of the analyst's industry coverage universe, on a risk-adjusted basis, over the next 12 months.

Neutral: Stock's total return is expected to be within the average total return of the analyst's industry coverage universe, on a risk-adjusted basis, over the next 12 months.

Underperform: Stock's total return is expected to depreciate below the average total return of the analyst's industry coverage universe, on a risk-adjusted basis, over the next 12 months.

RS: Rating Suspended - rating and price objective temporarily suspended.

NR: No Rating - not covered, and therefore not assigned a rating.

Mizuho Securities USA LLC investment ratings are based on the following definitions. Anticipated share price change is based on a 6- to 12-month time frame. Return expectation excludes dividends (prior to July 10, 2017).

Buy:	Stocks for which the anticipated share price appreciation exceeds 10%.
Neutral:	Stocks for which the anticipated share price appreciation is within 10% of the share price.
Underperform:	Stocks for which the anticipated share price falls by 10% or more.
RS:	Rating Suspended - rating and price objective temporarily suspended.
NR:	No Rating - not covered, and therefore not assigned a rating.

Rating Distribution

(As of 7/12)	% of coverage	IB service past 12 mo
Buy (Buy)	43.24%	52.68%
Hold (Neutral)	50.97%	33.33%
Sell (Underperform)	5.79%	46.67%

For disclosure purposes only (NYSE and FINRA ratings distribution requirements), our Buy, Neutral and Underperform ratings are displayed as Buy, Hold and Sell, respectively.

For additional information: Please log on to <http://www.mizuhosecurities.com/us> or write to Mizuho Securities USA LLC 320 Park Ave, 12th FL, New York, NY 10020.

Disclaimers

This report has been prepared by Mizuho Securities USA LLC (“MSUSA”), a subsidiary of Mizuho Americas LLC, solely for the purpose of supplying information to the clients of MSUSA and/or its affiliates to whom it is distributed. This report is not, and should not be construed as, a solicitation or offer to buy or sell any securities or related financial products.

This report has been prepared by MSUSA solely from publicly available information. The information contained herein is believed to be reliable but has not been independently verified. MSUSA makes no guarantee, representation or warranty, and MSUSA, MHSC and/or their affiliates, directors, employees or agents accept no responsibility or liability whatsoever as to the accuracy, completeness or appropriateness of such information or for any loss or damage arising from the use or further communication of this report or any part of it. Information contained herein may not be current due to, among other things, changes in the financial markets or economic environment. Opinions reflected in this report are subject to change without notice.

This report does not constitute, and should not be used as a substitute for, tax, legal or investment advice. The report has been prepared without regard to the individual financial circumstances, needs or objectives of persons who receive it. The securities and investments related to the securities discussed in this report may not be suitable for all investors, and the report is intended for distribution to Institutional Investors. Readers should independently evaluate particular investments and strategies, and seek the advice of a financial adviser before making any investment or entering into any transaction in relation to the securities mentioned in this report.

MSUSA has no legal responsibility to any investor who directly or indirectly receives this material. Investment decisions are to be made by and remain as the sole responsibility of the investor. Investment involves risks. The price of securities may go down as well as up, and under certain circumstances investors may sustain total loss of investment. Past performance should not be taken as an indication or guarantee of future performance. Unless otherwise attributed, forecasts of future performance represent analysts’ estimates based on factors they consider relevant. Actual performance may vary. Consequently, no express or implied warranty can be made regarding future performance.

Any references in this report to Mizuho Financial Group, Inc. (“MHFG”), MHSC and/or its affiliates are based only on publicly available information. The authors of this report are prohibited from using or even obtaining any insider information. As a direct subsidiary of Mizuho Americas LLC and indirect subsidiary of MHFG, MSUSA does not, as a matter of corporate policy, cover MHFG or MHSC for investment recommendation purposes.

MSUSA or other companies affiliated with MHFG, Mizuho Americas LLC or MHSC, together with their respective directors and officers, may have or take positions in the securities mentioned in this report, or derivatives of such securities or other securities issued by companies mentioned in this report, for their own account or the accounts of others, or enter into transactions contrary to any recommendations contained herein, and also may perform or seek to perform broking and other investment or securities related services for the companies mentioned in this report as well as other parties generally.

Restrictions on Distribution

This report is not directed to, or intended for distribution to or use by, any person who is a citizen or resident of, or entity located in, any locality, territory, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to or restricted by law or regulation. Persons or entities into whose possession this report comes should inform themselves about and observe such restrictions.

United States: Mizuho Securities USA LLC, a subsidiary of Mizuho Americas LLC, 320 Park Avenue, 12th Floor, New York, NY 10022, USA, contact number +1-212-209-9300, distributes or approves the distribution of this report in the United States and takes responsibility for it. Any transaction by a US investor resulting from the information contained in this report may be effected only through MSUSA. Interested US investors should contact their MSUSA sales representative.

United Kingdom/European Economic Area: This report is distributed or has been approved for issue and distribution in the UK by Mizuho International plc (“MHI”), Mizuho House, 30 Old Bailey, London EC4M 7AU, a member of the MHSC Group. MHI is authorized and regulated by the Financial Services Authority and is a member of the London Stock Exchange. For the avoidance of doubt this report is not intended for retail clients. This report may be distributed in other member states of the European Union.

Japan: This report is distributed in Japan by Mizuho Securities Co., Ltd. (“MHSC”), Otemachi First Square Otemachi 1-chome, Chiyoda-ku, Tokyo 100-0004, Japan. Registered Financial Instruments Firm, No. 94 (Kinsho), issued by the Director, Kanto Local Finance Bureau. MHSC is a member of the Japan Securities Dealers Association, the Japan Securities Investment Advisers Association and the Financial Futures Association of Japan, and the Type II Financial Instruments Firms Association.

Singapore: This report is distributed or has been approved for distribution in Singapore by Mizuho Securities (Singapore) Pte. Ltd. (“MHSS”), a member of the MHSC Group, which is regulated by the Monetary Authority of Singapore. Any research report produced by a foreign Mizuho entity, analyst or affiliate is distributed in Singapore only to “Institutional Investors,” “Expert Investors” or “Accredited Investors” as defined in the Securities and Futures Act, Chap. 289 of Singapore. Any matters arising from, or in connection with this material, should be brought to the attention of MHSS.

Hong Kong: This report is being distributed in Hong Kong by Mizuho Securities Asia Limited (“MHSA”), a member of the MHSC Group, which is licensed and regulated by the Hong Kong Securities and Futures Commission.

Australia: This report is being distributed in Australia by MHSA, which is exempted from the requirement to hold an Australian financial services license under the Corporation Act 2001 (“CA”) in respect of the financial services provided to the recipients. MHSA is regulated by the Securities and Futures Commission under the laws of Hong Kong, which differ from Australian laws. Distribution of this report is intended only for recipients who are “wholesale clients” within the meaning of the CA.

If you do not wish to receive our reports in the future, please contact your sales person and request to be removed from receiving this distribution.

© Mizuho Securities USA LLC. All Rights Reserved 2017. This document may not be altered, reproduced or redistributed, or passed on to any other party, in whole or in part, without the prior written consent of Mizuho Securities USA LLC.