EQUITY RESEARCH | March 6, 2019 | 3:02AM EST

Goldman Sachs

Cloud Platforms - Volume 6

Containers and Serverless Computing

Enabling the microservices journey

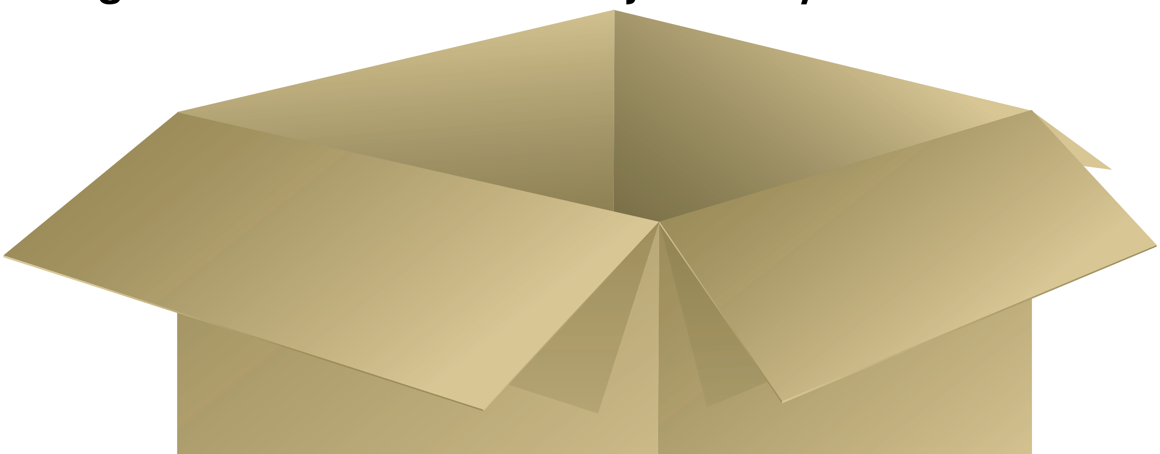

As workloads continue to migrate to the cloud where scalability is paramount, developers are shifting to a microservices software architecture that breaks down complex applications into individual components. Enabling the shift are containers, software packages that can isolate and house individual pieces of applications for more efficient scaling, repair and movement between environments. We estimate a $7bn market ($17bn in the bull scenario) for containers, capturing existing infrastructure spend and new web and IoT applications. Complementary to this is serverless computing, providing function-level isolation and rapid scaling ideal for batch processes and edge applications. Public cloud vendors look best positioned for the shift, namely Amazon, Microsoft, and Google. On-premise names like VMware and Pivotal Software will also benefit, but their onramp to container adoption looks longer and we expect public cloud to become the preferred approach.

Heather Bellini, CFA
+1 212 357-7710
heather.bellini@gs.com
Goldman Sachs & Co. LLC

Ted Lin
+1 212 357-6107
ted.lin@gs.com
Goldman Sachs & Co. LLC

Caroline Liu
+1 212 357-9111
caroline.liu@gs.com
Goldman Sachs & Co. LLC

Mark Grant
+1 212 357-4475
mark.grant@gs.com
Goldman Sachs & Co. LLC

Table of Contents

More in the Cloud Platforms Series

Workloads continue to migrate to the cloud, but for most large enterprises, this transition remains in its early innings. In our **Cloud Platforms series**, we map out the opportunities and implications of the move, and we profile markets enabled, and disrupted by the public cloud.

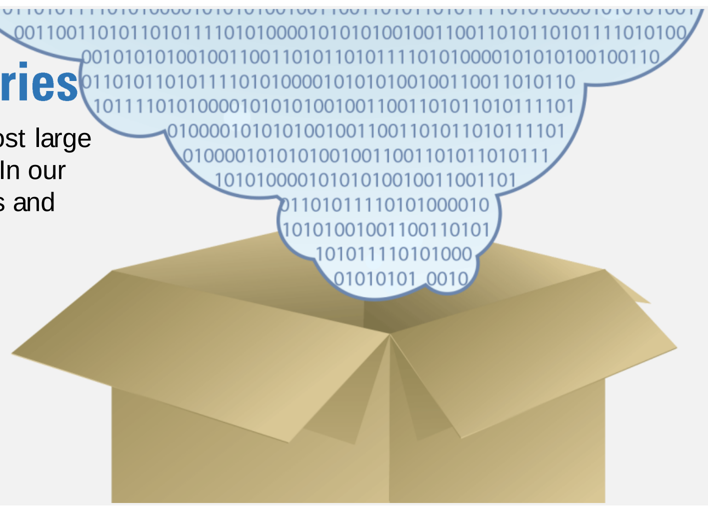

Vol. 1: Riding the Wave (Jan. 2015)
Vol. 2: NoSQL (Jun. 2015)
Vol. 3: Private Cloud (Feb. 2016)
Vol. 4: Public Cloud (Nov. 2016)
Vol. 5: Edge Computing (Oct. 2018)

PM Summary

Transition to a microservices architecture

Software developers are increasingly migrating from traditional, monolithic application architectures to microservices, motivated by the shift to the cloud that calls for more flexible, individually scalable applications. Microservices architecture is a software design style that breaks down a complex application into individual, isolated components (microservices) that have their own functions and can be reused by other applications. Traditionally, developers have been building applications with a monolithic architecture, which have interconnected components that must be run in chronological order and modified or scaled together. An IDC survey found that ~36% of applications are currently using a microservices architecture, which is expected to increase to ~45% in the next three years. Containers and serverless computing provide the isolation and scalability necessitated by the architecture, thus enabling the shift to the microservices design paradigm.

Why are containers important?

As workloads continue to migrate to the public cloud, we believe containers benefit from the increasing demand for application flexibility and portability, enabling enterprises to operate and scale in a multi-cloud environment. Containers are software packages that can house an application along with the tools and settings necessary to run the application. This grants developers the ability to build applications that can be easily shifted between server environments without disrupting its functionality. Furthermore, as applications increase in complexity, more developers are utilizing microservices architectures and the isolation necessary for microservices is enabled by containers, allowing errors to be isolated and addressed without having to shut down the entire application.

Refer to Exhibits 4-6 and 18-19 for examples of workloads ideal for containers or better suited for virtual machines or serverless computing.

We expect adoption of containers to accelerate and as such, demand for virtual machines to decrease over time, but not entirely be replaced. Unlike containers, virtual machines package an entire operating system with an application and its settings. As a result, containerized applications are more efficient and require less server space given that the application can be run without having to boot up an entire virtualized operating system. While we expect a portion of existing workloads to shift into containers, virtual machines are considered to be more secure and may be better suited for applications that have longer lifetimes and a complex, monolithic architecture. Net-net, virtual machine usage is expected to decline, but will not be completely replaced by containers in the near to medium term.

How big is this market?

In this report, we estimate the new and legacy software infrastructure spending that can be captured by containers, fueled by the deployment of new applications and migration of existing workloads from virtual machines and bare metal servers. Utilizing IDC's forecast of 1.63 billion paid container instances in 2021, we estimate that $5.5bn of new infrastructure spending and $1.1bn of existing infrastructure spending is captured by containers. This compares to the most aggressive model, with $13.9bn new infrastructure spending and $2.7bn from existing. We note that our estimates likely skew conservative as we assume a constant percentage of containerized applications will be migrated from existing infrastructure and we do not account for the on-premise container deployments which, while comprising of a smaller proportion of container instances, are less efficient and thus can drive up the container usage cost.

Serverless computing does not eliminate the need to containerize applications

Contrary to what its name suggests, serverless computing is still reliant on servers. Popularized by AWS Lambda, serverless computing depends on the cloud provider to maintain and manage the resources for servers. Serverless functions are event-triggered applications that are rapidly scaled in isolation in containers managed by the cloud provider and designed to be executed for a few seconds or milliseconds, allowing developers to run applications and services without having to manage how or where the applications are deployed – after the code snippet executes, the infrastructure is de-allocated. As such, there has been an on-going debate on the potential for serverless computing to disrupt the container market. We believe that organizations will still focus on containerizing their own applications given its greater flexibility and ability to support a variety of applications with longer run times. Early adopters have largely utilized serverless computing for batch processing and microservices applications. In addition to being inflationary, by enabling new applications and workloads, serverless computing does have the potential to be deflationary to existing cloud-based revenue given its cost efficiencies and elimination of the tendency for customers to over-provision resources.

Leaders & laggards of the container market

We expect the three largest public cloud vendors (Amazon's AWS, Microsoft's Azure, & Google's GCP) will be the primary beneficiaries of container adoption. On the other hand, over time it is possible that VMware could face a headwind as more workloads move to public cloud and leverage cloud native offerings such as Amazon Elastic Container Service, Azure Container Service, and Google Kubernetes Engine that are non-virtualized. We note that headwinds to VMware are somewhat mitigated by its diversified revenue base, revenue-sharing agreement with Pivotal, acquisition of Heptio (to assist in the implementation of Kubernetes), and partnership with AWS.

Based on our assessment of container services and survey results, we note that AWS is the most popular container services, offering the most diverse set of mature features, and we expect AWS to maintain its leadership position. However, Azure is rapidly gaining popularity among users and GCP, having created Kubernetes, can leverage the growing popularity of Kubernetes for orchestrating containers. Pivotal Container Service (PKS) could benefit in the longer term from container trends as it enables developers to build and efficiently run applications on both on-premise and cloud platforms. Pivotal's sales partnership with VMware leverages VMware's access to the enterprise market, although we are uncertain when Pivotal will realize this benefit as their products are still in early stages and users tend to run containers in a public cloud environment. That having been said, as with VMware we see this offering as having broader adoption potential on-premise versus on cloud native platforms.

Exhibit 1: We believe that certain vendors will benefit more than others as container adoption continues to grow

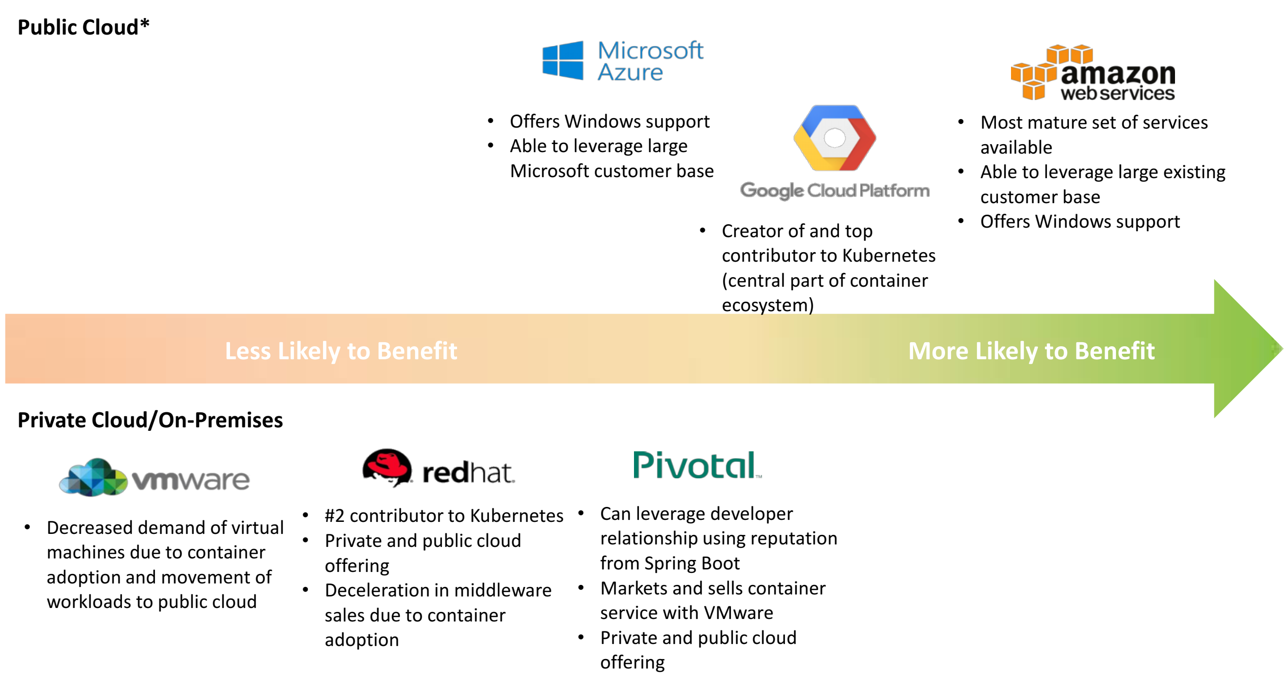

* July 2018: Google Cloud announced on-premises offering of Google Kubernetes Engine (not generally available yet)

Source: Company data, Goldman Sachs Global Investment Research

Pivotal Software: Downgrade to Neutral

While Pivotal Software could benefit from increasing adoption of containers and serverless computing, we believe that those trends are still in their early stages and increasingly more likely to benefit public cloud vendors. Pivotal does have an agreement with VMware to jointly market their container offering, Pivotal Container Service (PKS), with management recently stating that VMware's sales force is just starting to spin up on the offering, and benefit from VMware's acquisition of Heptio. Management has also highlighted their success with PKS within their installed base of 368 customers with greater than $50k in annualized recurring revenue, as of F3Q19. Likewise, Pivotal Function Service (PFS) has not been made generally available yet, making it difficult to assess if it is becoming a viable alternative to AWS Lambda. In general, given how new their product offerings are, it has been difficult to assess the traction they have garnered in the pipeline. On the other hand, users are increasingly mobilizing around public cloud vendors for container management and serverless functions deployment. Moreover, as the company scales, the lumpiness of their billings, variability in contract signings and collections, as well as their small base of subscription customers add to the volatility of their earnings results. We note that Pivotal continues to retain a multi-cloud advantage.

As such, we are downgrading PVTL from Buy to Neutral. At ~$22, PVTL is trading at an EV of ~7x CY19 and ~6x CY20 our sales estimates of $827mn and $1020mn, compared to consensus of $813mn and $1,000mn respectively, and at an EV of ~6.5x CY19 and ~5x CY20 our billings estimates of $921mn and $1,136mn, compared to consensus of $913mn and $1,123mn respectively.

VMware: Downgrade to Sell

While VMware could benefit from containers being run in virtual machines for security benefits, we expect that as container adoption matures, demand for VMs will face gradually increasing headwinds. While headwinds to VMware are somewhat mitigated by its diversified revenue base as well as its new product offerings and partnerships with AWS, we believe that product adoption and upside from VMware's partnerships are still in the early innings, with much of the potential upside to VMW already accounted for in its current valuation. Users are also increasingly mobilizing around public cloud vendors for container management.

Since 2017, VMW has benefitted from a resurgence in on-premise spending as customers recognize that they will remain in a hybrid cloud environment for longer than previously expected, driving an acceleration in billings from ELAs. For example, in CY17, signings from ELAs increased 27% versus the pool that was up for renewal in CY16 (+13%), in CY18 the increase in signings from ELAs was 26%, and in CY19, we estimate 17% growth as the company begins to face tougher comps. Given VMware's typical contract length of 3 years, we see VMware as facing a more normalized comp for CY20 vs. the comps for CY16-CY19. We note, however, that VMware's fundamentals continue to diverge from its ELA cycles over time, driven by 1) distance from the distinctive 2009/2010 ELA cycle and 2) a diversifying product portfolio that shifts focus away from renewals. As such, we would expect the ELA cycle to become less impactful for VMware over time. While we see the company as benefitting from its partnership with AWS and a potential deeper relationship with Microsoft, as reported by the Information, we see these as having a limited impact on financials this coming FY. As such, based on our view of more normalized comps and limited upside to topline over the course of the year, we downgrade the shares to Sell from Neutral as we see better risk reward in other names. To be clear, our rating is not a view on the company's execution, which has been excellent, but rather one where we see multiple expansion from current levels as unlikely.

At ~$173, VMW currently trades at 27x CY19 and 24x CY20 our non-GAAP EPS estimates of $6.49 and $7.31, compared to consensus of $6.58 and $7.29 respectively. This compares to MSFT which at ~$112, trades at 24x and 20X our GAAP EPS forecasts of $4.63 and $5.70.

Special thanks to Jessie Wang for her extensive contributions to this report.

What is a container?

A container is a software package that holds an application and everything it needs to run (i.e. its libraries, runtime, system tools, and settings). Containers became increasingly popular as the desire for application portability increased because they can be moved from one server to another without disrupting the application's ability to run. Containers are designed to function in isolation and to have the minimum resources needed to run an application. In essence, they operate similarly to a shipping container - multiple objects are packed into one shipping container for ease of transportation from one facility to another, a more efficient method than moving each individual object from one location to the next.

Despite being popularized in 2013 by Docker, the concept of containers existed long before then. Containers were introduced in the late 1970's, and one of the earlier technologies, the Linux-V Server, was introduced in the early 2000s. Early container technology did not disrupt the virtual machine market because its main advantage, application portability, had not been fully developed yet. As a result, these containers did not differ very much from virtual machines in terms of what they could offer. Thus, virtual machines were the go-to technology for hosting multiple applications on one server (and continue to be widely implemented today). Prior to virtual machines (server virtualization), companies had to purchase and maintain a server for each application they needed to run. Server virtualization enabled companies to fully utilize server capacity by allowing multiple operating systems – and therefore multiple applications – to run on one physical server. Companies were able to reduce capital and operating costs, as well as increase IT efficiency, productivity, and responsiveness.

Early approaches to containerization were heavily dependent on the operating system version; a containerized application running on Ubuntu (one version of Linux), for example, could not be moved to and run on Debian (a different version of Linux). Docker, however, provided a solution to this problem when it released its container technology. Docker containers could be moved between any operating system, as long as the kernel (the core of the operating system) was the same. This property gave container technology an edge, and contributed to the growth of the container market.

Containers vs. virtual machines

Containers package an application with its relevant tools and settings, whereas virtual machines package an entire operating system with the application and its settings. Virtual machines also rely on a hypervisor, which separates the host operating system from each VM and allocates computing resources to each machine. Because of this hypervisor-dependency, moving a virtual machine from one virtualization software to another (e.g. from VMware's vSphere to Microsoft's Hyper-V) requires fundamental changes. Containers, however, are designed so that they can be easily moved between any cloud platform and operating system with a container engine. However, containers built on one OS (e.g. Linux) cannot be run on another OS (e.g. Windows).

Containers are more lightweight and efficient than virtual machines, which take up more space given that they contain an entire copy of an operating system. As a result, given the same level of resources, a server can run more containers than virtual machines. Furthermore, containers can be spun up faster than virtual machines because on a VM, the virtualized operating system needs to be booted up first.

Exhibit 2: Virtual machines occupy more server space because they contain an entire operating system
Virtualization vs. containers on bare metal servers

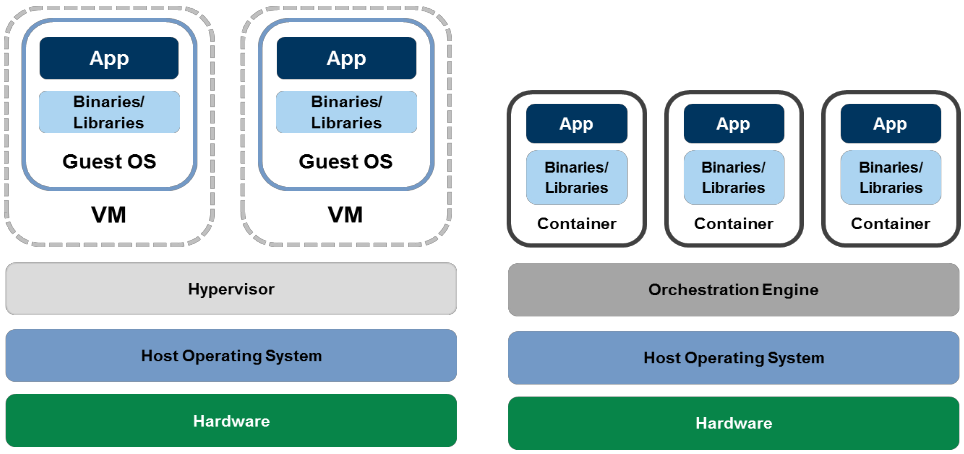

Source: Goldman Sachs Global Investment Research

Containers will not completely replace virtual machines

While demand for virtual machines will likely decrease as container adoption grows given the advantages containers offer over VMs, we expect VMs will still be used in conjunction with containers in the near to medium term. We outline advantages that virtual machines offer over containers below:

- **Security**: Although both containers and virtual machines are isolated, virtual machines are considered to be more secure than containers. Containers share an operating system, meaning a bug that affects the operating system will affect every container utilizing the OS. Conversely, each virtual machine runs its own OS, thus providing an extra level of security. As a result, a large number of containers are run inside virtual machines, taking advantage of both the portability and efficiency features of containers and the security benefits of virtual machines.
    - *Caveat:* In response to security concerns, container services have been evolving to reduce vulnerabilities and provide VM-level security. In 2016, for example, Canonical launched LXD, the "container hypervisor," which offers a user experience similar to a VM, but utilizing Linux containers instead.

Exhibit 3: A large number of containers are run inside of virtual machines because VMs provide an extra level of security
Containers + virtualization

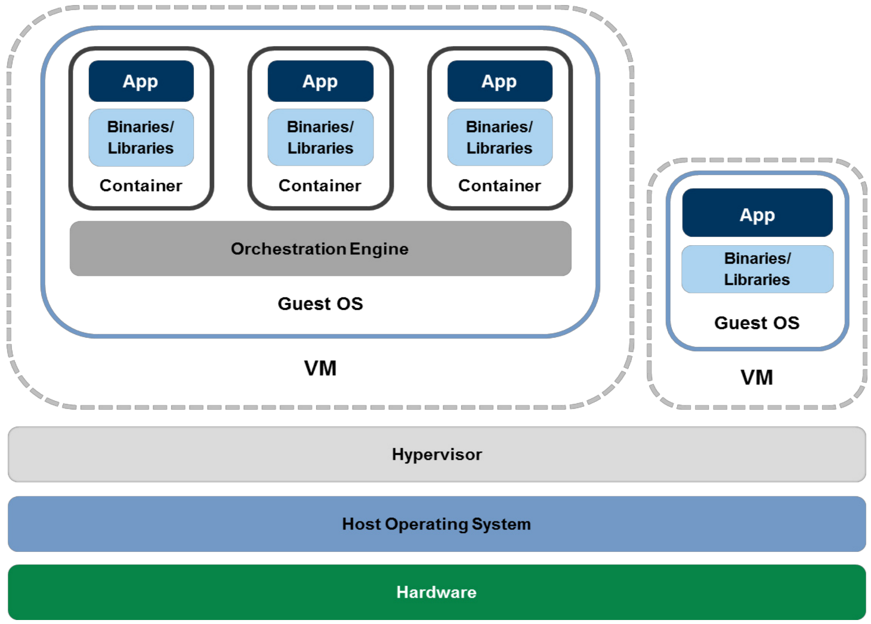

Source: Goldman Sachs Global Investment Research

- **Complexity of container technology**: Developing and managing containers can be complicated and requires new skills and talent. Additionally, many legacy applications have to be re-architected to suit a container environment, which also requires new skills and time.
    - *Caveat*: As containers are becoming more widely adopted, management tools are being developed and simplified. Azure Container Instances, for example, aims to offer a fast and simple way of running containers in Azure, while guaranteeing the security level of virtual machines.
- **Some stateful applications may remain virtualized**: Many existing applications are still structured using a monolithic architecture and are often stateful, meaning they store data on each user session and use the data every time a user makes a request. Thus, stateful applications require data that exists beyond the lifetime of a container. While it is possible to containerize a stateful application, the application must be refactored to make it suitable for containerization, breaking down the application into various functionalities that can be run in containers. Some stateful applications may be difficult to refactor, and doing so may not be worth the time and resource investment. As a result, these types of applications are more likely to remain in virtual machines.
- **Dependency**: Containers still depend on virtual machines to an extent, since containers must be compatible with the operating system they share. For example, Docker containers are Linux-based and therefore can only operate on Linux systems. Docker can support Windows and Mac OS X, but those containers must be spun up

in virtual machines on the OS in order to do so. Kubernetes supports both Linux and Windows.

Exhibit 4: Differences between a stateful and a stateless application

Stateful Server saves data generated from each user session for use in future requests		Stateless Server does not save data generated from a user session. Session data is stored in a database that is called for each request	
Benefits	**Drawbacks**	**Benefits**	**Drawbacks**
Lower latency as data on session is stored on the server	Not easily scalable as user sessions must be handled by the same server or an algorithm must be implemented to ensure consistency	Can be horizontally and independently scalable without concern over dependencies to accommodate load changes	Increased latency from having to call a database for each request
Sessions can be revoked any time Session data can be modified	Increased memory usage on the server side	Fast redeployment in the event of a failure Easily connected to other apps through APIs Enables consistency across various applications	Sessions cannot be revoked by the server Session data cannot be modified until it has expired
Examples			
Desktop applications Operating systems Distributed, multi-service environments for large-scale data processing		Web applications IoT devices	

Source: Goldman Sachs Global Investment Research

Exhibit 5: In deciding between virtual machines and containers, developers need to consider what qualities their applications have or need to have

Summary of applications best suited for containers vs. virtual machines

Factor to Consider	Use Containers	Use VMs	Explanation
Portability	Require portability		Containers can be moved between platforms more easily than virtual machines. VMs tend to be hypervisor-dependent, whereas containers can be moved between any Linux OS with a container engine.
State	Stateless	Stateful	Containers have short lifetimes, which is best suited for applications that do not need data stored for extended periods of time (i.e. stateless applications).
Architecture	Microservices	Monolithic	Containerizing a monolithic application would require refactoring the application. Microservices architecture depends on containers.
Lifetime	Flexible - updated or changed constantly and/or quickly	Constant with a long run-time	VMs have to load an entire OS, whereas containers rely on the already loaded host system, spinning up more quickly. Containers have short lifetimes; they are intended to run only for several minutes. VMs can run for days or months.

Source: Goldman Sachs Global Investment Research

Microservices as a driver of container adoption

Shift to microservices architecture from monolithic architectures

Adoption of microservices stemmed from the need to be able to fix and expand applications efficiently and independently. Microservices architecture is a software design style that breaks down a complex application into individual, isolated components. Each component or microservice has a discrete function, such as identification or authentication, and essentially becomes its own application. All of the individual microservices are connected via APIs to form the larger application. Furthermore, each microservice, given its discrete functionality, can be reused by other services and applications.

Netflix, for example, is a well-known example of a company that transitioned from a traditional software architecture to a cloud-based microservices architecture. Netflix needed to be able to support its fast growth and ensure constant availability. On a traditional software architecture, in which every application component is interdependent, a single coding error can shut down the entire operation. Adopting a microservices design paradigm for their applications offered a solution by isolating the components, meaning the failure of one component will not necessarily impact the performance of another. Furthermore, each component of an application can be quickly scaled depending on the number of requests sent to the application or component.

Exhibit 6: Microservices architecture breaks down an application into isolated, individual components that are connected via APIs
Example of an application under a microservices architecture

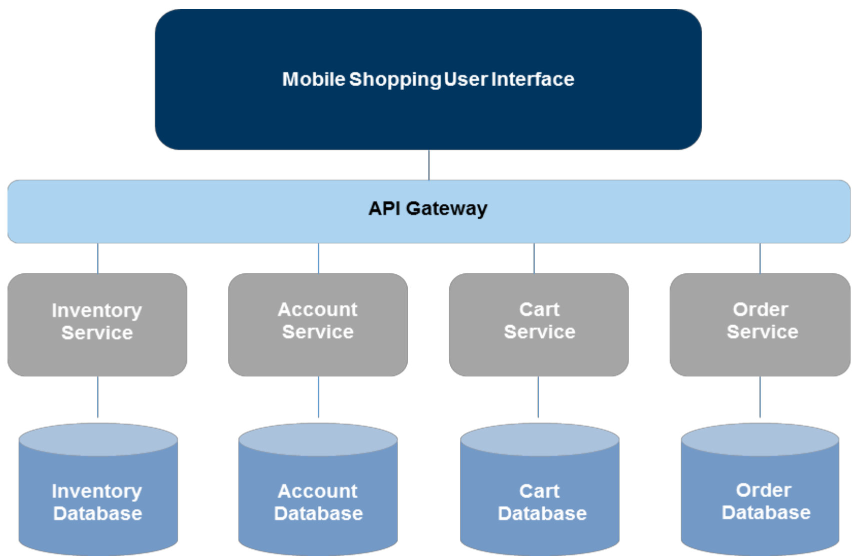

Source: Goldman Sachs Global Investment Research

Migration to the cloud driving increased usage of microservices

As workloads continue to migrate to the cloud, we expect the microservices architecture to become more mainstream, transforming the way applications are designed. According to a 2018 IDC survey of senior IT employees working with containers, respondents on average indicated that ~36% of their applications were currently using a microservices architecture and expect to increase their usage to ~45% of applications in three years. We note that when moving applications to the cloud, they often need to be refactored into cloud-friendly applications. Cloud-friendly applications take advantage of the platform's ability to optimize resource consumption, and to an extent, can recover if the server goes down. A step up from cloud-friendly is cloud-native, in which applications take full advantage of the cloud's scaling capabilities, and can recover quickly in the case of a server failure. Adopting a microservices paradigm is necessary to achieve cloud-native status. Some platforms like Pivotal Cloud Foundry require applications to be at least cloud-friendly, offering more structure and abstraction, allowing the user to focus on writing the applications vs. in the management and deployment.

Exhibit 7: Applications fall on a spectrum of cloud compatibility, and microservices are highly compatible with cloud computing
Overview of cloud compatibility levels

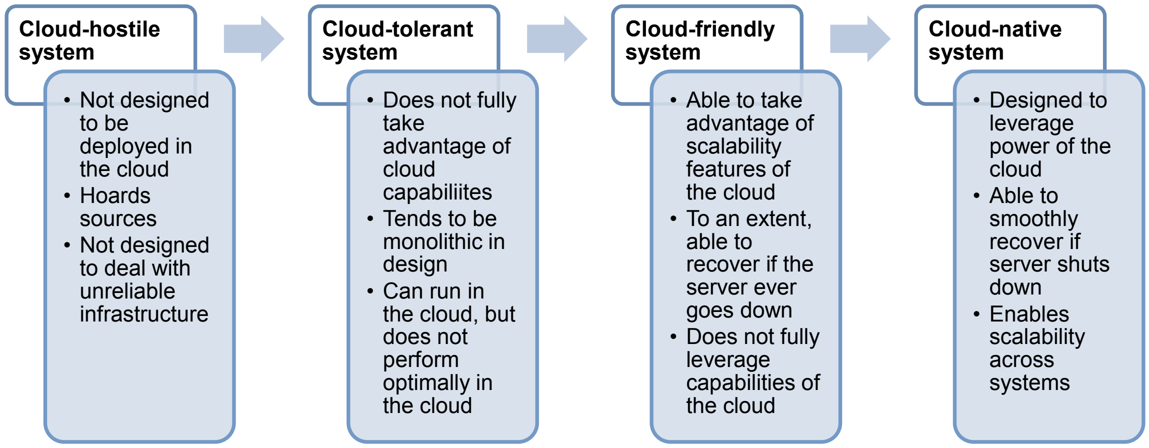

Source: Gartner, Goldman Sachs Global Investment Research

IDC expects that adoption of the microservices paradigm will continue to increase throughout 2020 due to factors such as the advancement of development tools, the gradual advancement of developers' skill levels in regard to microservices, the continued integration of DevOps, and the increasing adoption of practices that enable microservices (such as container and serverless computing usage). According to Gartner, industries currently using microservices include finance, retail, and digital businesses while insurance and manufacturing lag behind in adoption.

Advantages over monolithic architecture

An application with a microservices architecture is made up of smaller, independent applications that run in parallel with one another while traditional, monolithic applications contain interconnected components such that tasks must be completed in chronological order from beginning to end.

Exhibit 8: Monolithic applications have interconnected components

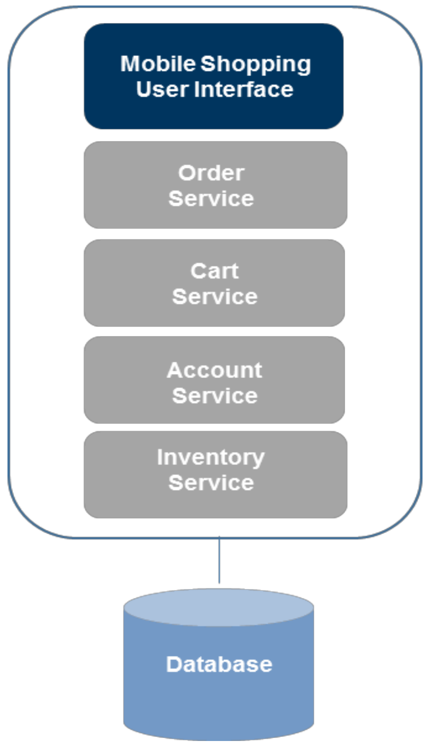

Source: Goldman Sachs Global Investment Research

Microservices architectures offer several advantages over monolithic architectures:

- **Independent components**: A monolithic system acts as a single application, meaning a change in one component of the application can break another. Thus, debugging one part of the application requires rebuilding, testing, and redeploying the entire application. A microservices system has parts that function independently, meaning each microservice can be developed, tested, and deployed without impacting other parts of the application.
- **Scalability**: The scalability of an application refers to its ability to process an increasing flow of requests per minute. As shown in Exhibit 9, in order for a monolithic application to handle more traffic, the entire application would have to be replicated and load balanced, even if only one component of the application needed to be scaled. On the other hand, microservices do not have this same issue because developers can scale the individual components of the application independently and as needed.

Exhibit 9: Microservices are more efficient at scaling than monolithic applications
Scaling with a monolithic vs. microservices architecture

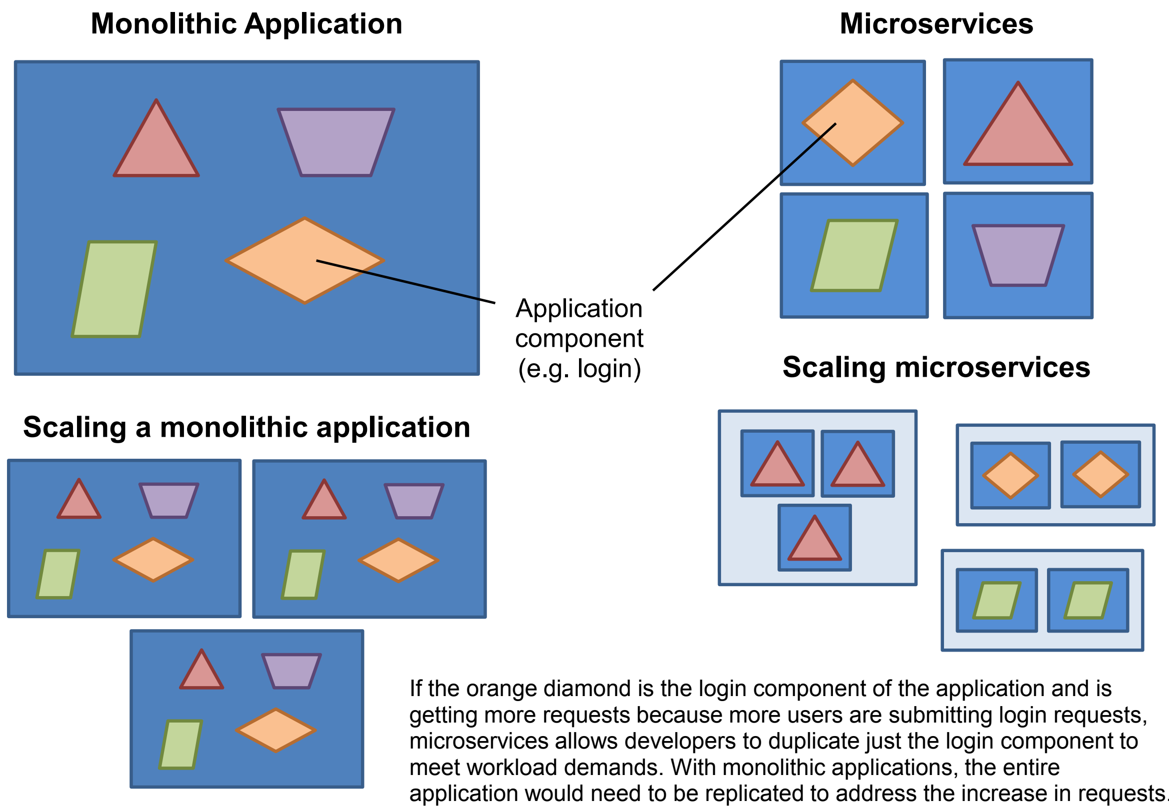

Source: Sourcesense, Goldman Sachs Global Investment Research

Although microservices offer flexibility and scalability advantages over monolithic applications, we believe that existing monolithic applications will continue to be used for the foreseeable future. Some barriers to adopting a microservices architecture include:

- **Complexity**: Microservices involve many moving parts, which can be complex and difficult to implement. Developers will have to be able to conceptualize and visualize an entire system of application components. Monolithic applications, on the other hand, are easier to develop.
- **Adoption of DevOps**: Developers are less familiar with the microservices architecture compared to the monolithic design structure, and need to adopt a DevOps design model. DevOps is a practice in which the development and operations teams work together (as opposed to traditionally being separated). Team members oversee the entire application lifecycle, which includes development, testing, and deployment. The skills of the team members are therefore not unique to a specific role or function. This is critical to the development of microservices, and the growing practice of DevOps is another driving factor in microservices adoption. IDC estimates that in 2017, ~28% of an organization's applications were built using DevOps practices, and expects that to increase to ~45% of applications by 2021.

- **Maturity of the technology**: Fewer mature platforms are available for running and managing microservices. Microservices currently rely on an amalgamation of technologies that developers need to run an application.

We expect that these factors will not deter microservices adoption and solutions are being developed to overcome these barriers to widespread adoption. Istio, for example, is an open source platform that is meant to provide a way to connect, manage, and secure microservices. Although it is relatively new (having been launched in 2017 by Google, IBM, and Lyft), Istio recently reached an important milestone: the release of version 1.0 in late July 2018. As platform offerings mature and as developers become more familiar with microservices, we believe that new applications will primarily be architected with microservices.

Microservices vs. Service-oriented architecture (SOA)

Microservices architecture is based on service-oriented architecture (SOA), and share the same goals and design fundamentals. They both separate an application's functions into smaller parts, in which each part is an individual service or application. The primary difference is that microservices are completely independent from one another whereas the services in SOA-designed applications are still somewhat dependent on each other. SOA components depend on an Enterprise Service Bus (ESB), a system that facilitates communication between each component.

The independence of microservices makes microservices-based applications more fault-tolerant than SOA-based ones. You can upgrade one microservice without upgrading the others, or you can shut down one microservice without affecting the others. Another differentiating factor is that the microservices architecture functions on a more granular level than SOA.

Exhibit 10: Microservices architecture vs. Service-oriented architecture

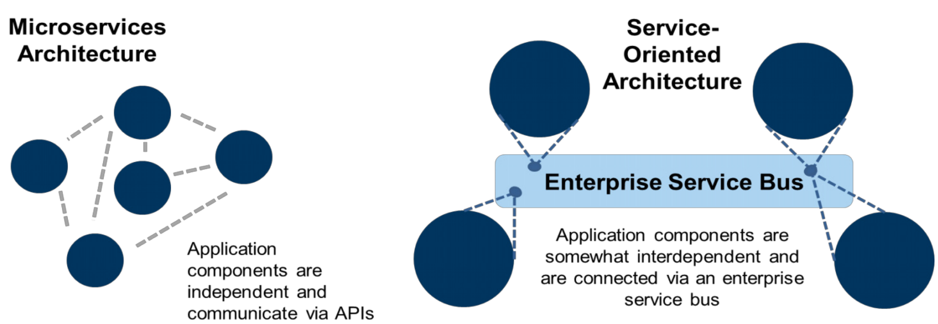

Source: Goldman Sachs Global Investment Research

Containers provide the isolation necessary for microservices

Microservices are heavily reliant on containers because they are comprised of everything an application needs to run, providing the isolation required by microservices architectures. Although microservices can run on a physical server or in a virtual machine, doing so can be risky and inefficient.

A major drawback of running microservices directly on a server (whether it be a physical or virtual one) is that there is little to no isolation of each service. If one service or component fails, it could affect the performance of another service. A potential solution would be to run one microservice per server in order to provide the isolation needed. However, dedicating an entire server to one component is an inefficient use of computing resources. Conversely, multiple containers can be run on one server, meaning multiple microservices can be run on one server, while providing the isolation needed by each microservice. As a result, most microservices-based applications are deployed within containers.

An additional benefit of using containers for microservices is the faster initialization and execution time. Unlike VMs, which have to reload an entire operating system, containers boot up in seconds as opposed to minutes. Thus, microservices in containers can be deployed faster and easily redeployed in the case of a network or server failure.

Sizing up the growth opportunity of containers

Acceleration in container deployments

We expect container adoption, while still in its early days, to accelerate over the next few years. In our December 2018 survey that tracked IT spending expectations of CIOs, we found that approximately 7% of workloads are currently run in containers, which is expected to increase to approximately 30% in three years. Reflecting the rapid adoption of containers, the proportion of respondents who have significant container deployments (i.e. at least 30% of their workloads are in containers) is expected to increase from roughly 7% to 35% in three years. Similarly, a 2018 IDC survey found that 42% of respondents' applications are currently running in containers, while an estimated 51% are expected to run in containers by 2021.

Exhibit 11: Container adoption is expected to increase to ~30% of workloads in the next three years
% of workloads today compared to expected in 3 years and longer term

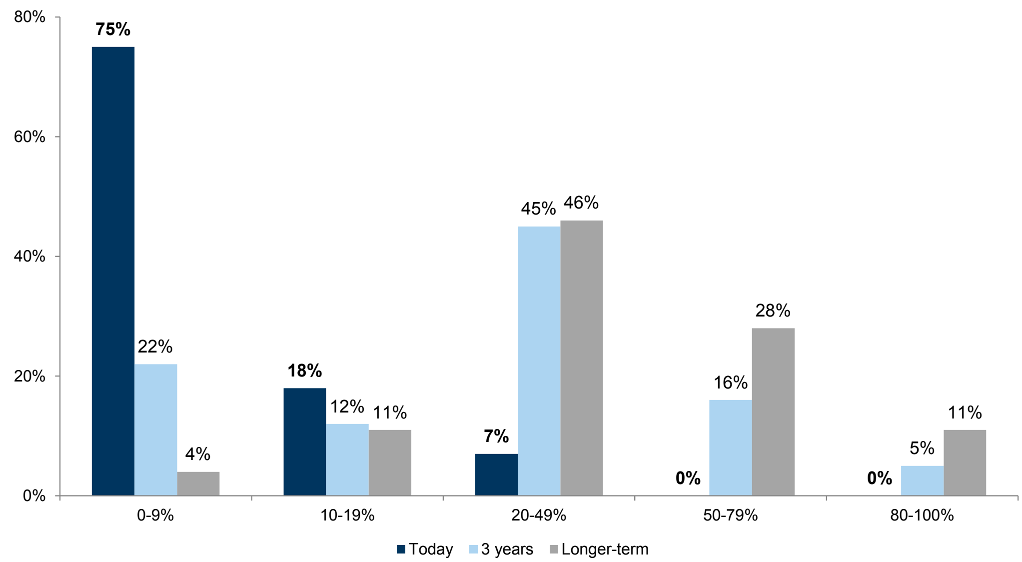

Source: Goldman Sachs Global Investment Research

The prevalence of virtualization highlights the potential for container-based virtualization to grow. In a 2016 report, Gartner estimated that many organizations virtualized (via hypervisor-based virtualization) more than 75% of their workloads. By 2020, Gartner expects that 95% of North American and Western European enterprises will have virtualized 80% of their on-premises x86 server workloads, and 50% of emerging market enterprises will have virtualized 60% of their workloads. As containers become more prevalent, it is possible that a fraction of the workloads running in virtual machines will be moved to containers running on bare metal. According to a 2018 IDC survey, approximately 53% of respondents' containerized applications were existing applications moved from a virtual machine or from bare metal. In addition to these existing workloads, containers are being used for new applications. New applications made up the other ~47% of containerized applications.

Companies are also beginning to anticipate greater spending on container software. In our December 2018 survey that tracked IT spending expectations across CIOs, we found that approximately 4% of respondents expect to accelerate their spending on containers/Docker.

Exhibit 12: The percentage of CIOs expecting to accelerate spending on containers is small but growing
% of GS IT spending survey respondents expecting accelerated spending on containers/Docker

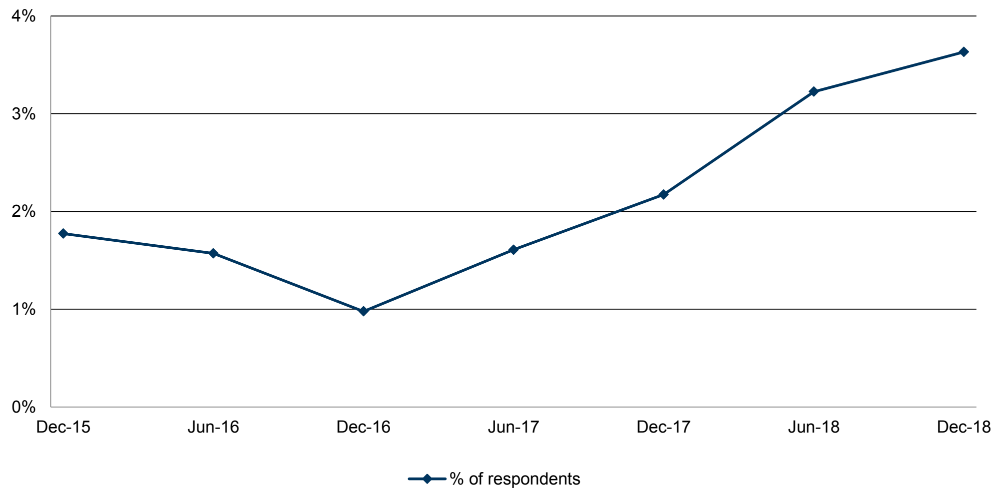

Source: Goldman Sachs Global Investment Research

Overall, container deployments are estimated to grow at an accelerated rate over the next few years. IDC estimates that worldwide container instance deployments will grow at a 65% '16-'21 CAGR. Excluding Web/SaaS provider deployments, this rate increases to almost 99%. Relative to all container instances, the number of paid instances will grow from about 12% in 2016 to about 54% in 2021. Public cloud deployments are also growing as organizations continue to build cloud-native applications. Furthermore, while more stateful applications are being deployed over time, the vast majority of deployments will involve stateless applications. We believe these container instance trends will have important implications for how vendors are positioned in the container market.

Exhibit 13: Worldwide container instances installed base forecasts

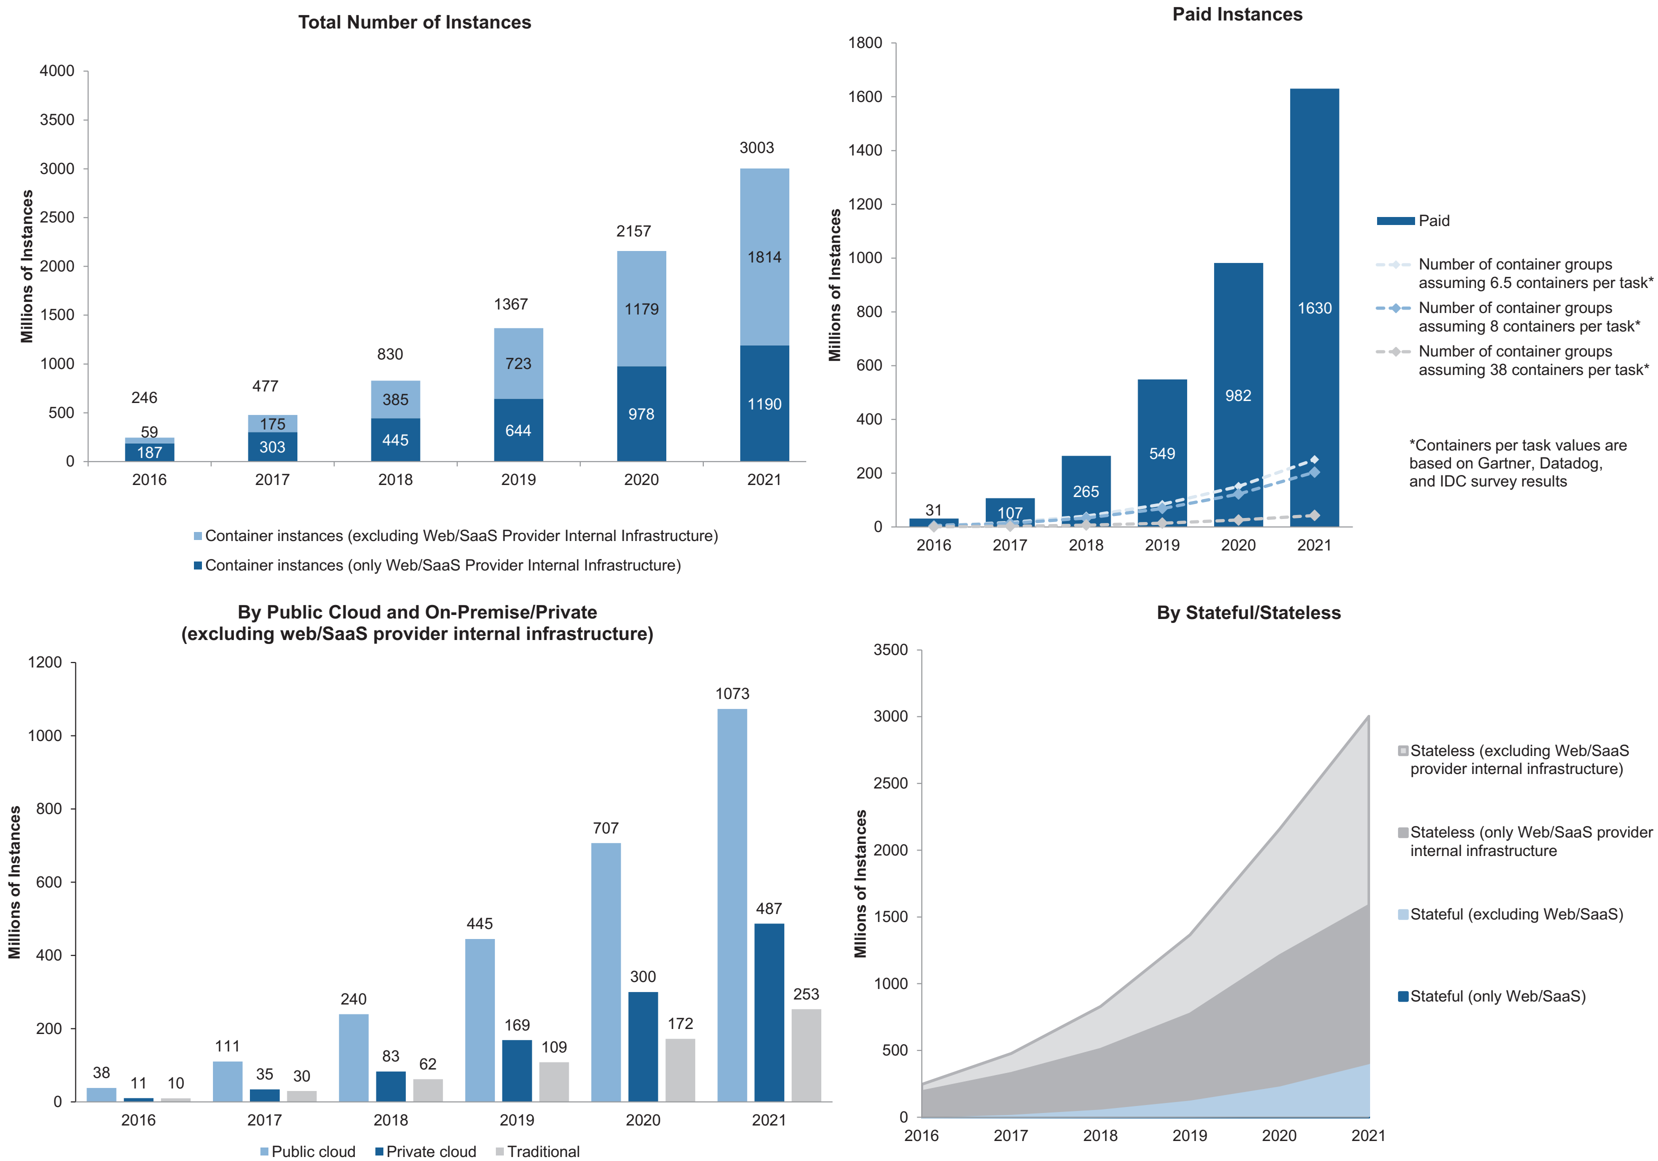

Source: IDC, Goldman Sachs Global Investment Research

Containers are shaping into ~$7bn opportunity by 2021

We deconstruct our review on the market size for containers into two segments 1) net new spending on containers for deploying new applications and use cases enabled by containers, and 2) reallocated spending on containers from moving existing applications to containers that were previously virtualized or on bare metal servers. Exhibit 14 details the base case for the size of the container market. Exhibit 15 details our methodology and a sensitivity analysis, with bull and bear cases. In the bear case (top left of the sensitivity table), the container market size is ~$1.7bn, with ~$1.2bn from new software infrastructure spending and ~$0.5bn from existing infrastructure spending captured by containers. In the bull case (bottom right of the sensitivity table), the container market could grow to ~$17bn, with new software infrastructure spending making up ~$14bn and legacy spending of ~$3bn.

Exhibit 14: We estimate that the container market will reach ~$7bn by 2021
Container market size by year, segmented by new and legacy software infrastructure spending captured by containers

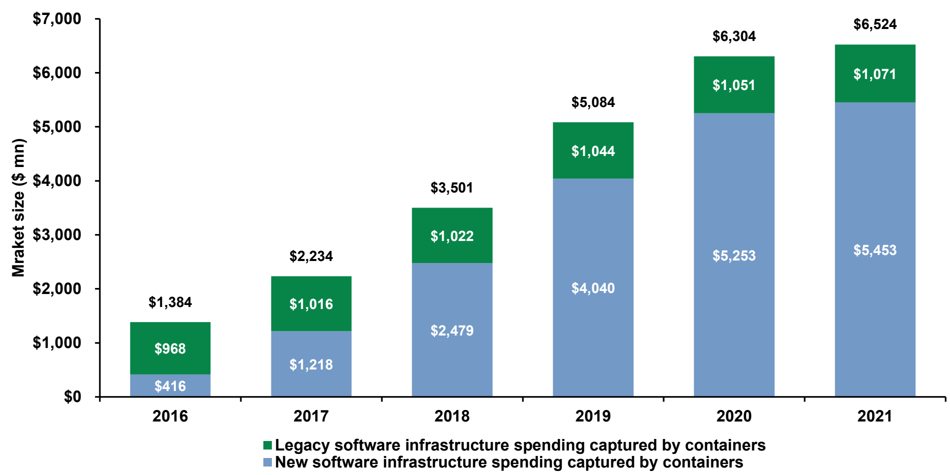

Source: Gartner, IDC, Goldman Sachs Global Investment Research

Using container instance data from IDC, pricing data from AWS, and container usage statistics from Datadog (a service that offers monitoring and analytics for cloud-scale applications), we were able to derive an estimate for new containerized application spending through 2021. Our process and considerations for estimating the total cost to run containers for each year are detailed below and in Exhibit 15.

Exhibit 15: Container market assessment methodology

	2016	2017	2018	2019	2020	2021	Source
New infrastructure spend captured by containers							
Paid container instances installed base	**31**	**107**	**265**	**549**	**982**	**1,630**	IDC
Average vCPU cores - base case	4.0	4.0	4.0	4.0	4.0	4.0	GSe
Average memory (GB) - base case	6.0	6.0	6.0	6.0	6.0	6.0	GSe
Average container lifetime (days)	2.0	1.7	1.4	1.1	0.8	0.5	Datadog - 2 days/container average; 0.5 days/container orchestrated
Cost per vCPU-sec	$0.00001406	$0.00001406	$0.00001406	$0.00001406	$0.00001406	$0.00001406	AWS Fargate
Memory cost per GB-sec	$0.00000353	$0.00000353	$0.00000353	$0.00000353	$0.00000353	$0.00000353	AWS Fargate
CPU cost	$302	$885	$1,801	$2,935	$3,816	$3,961	
Memory cost	$114	$333	$678	$1,105	$1,437	$1,492	
New software infrastructure spending captured by containers	**$416**	**$1,218**	**$2,479**	**$4,040**	**$5,253**	**$5,453**	
Legacy infrastructure spend captured by containers							
x86 Server Virtualization Infrastructure	$4,219	$4,486	$4,209	$4,289	$4,232	$4,262	Gartner
Application Platform Software	$4,218	$4,365	$4,702	$4,805	$4,925	$5,070	Gartner
Total legacy infrastructure addressable by containers	**$8,438**	**$8,851**	**$8,911**	**$9,094**	**$9,156**	**$9,332**	
% of total applications running on containers today	42.5%	42.5%	42.5%	42.5%	42.5%	42.5%	IDC
% of containerized applictaions that were migrated to a container from a VM or bare metal	54.0%	54.0%	54.0%	54.0%	54.0%	54.0%	IDC
% of total containerized applications migrated from VM or bare metal	**23.0%**	**23.0%**	**23.0%**	**23.0%**	**23.0%**	**23.0%**	Calc
Container efficiency over VMs	2.0x	2.0x	2.0x	2.0x	2.0x	2.0x	GSe
Legacy software infrastructure spending captured by containers	**$968**	**$1,016**	**$1,022**	**$1,044**	**$1,051**	**$1,071**	
Total container market size	**$1,384**	**$2,234**	**$3,501**	**$5,084**	**$6,304**	**$6,524**	

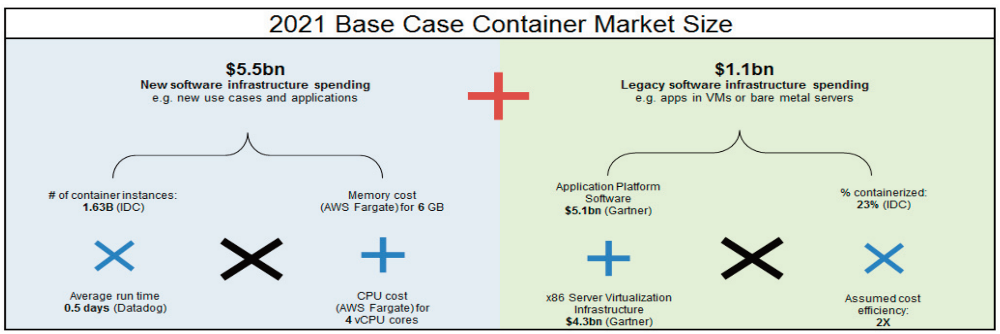

New software infrastructure spending on containers ($mn)

Average memory per container instance (GB) \ Average vCPU cores per container instance	1	2	3	4	5	6
1 GB	$1,239	$2,229	$3,220	$4,210	$5,200	$6,190
2 GB	$1,488	$2,478	$3,468	$4,458	$5,449	$6,439
4 GB	$1,985	$2,975	$3,965	$4,956	$5,946	$6,936
8 GB	$2,979	$3,970	$4,960	$5,950	$6,940	$7,931
16 GB	$4,968	$5,959	$6,949	$7,939	$8,930	$9,920
32 GB	$8,946	$9,937	$10,927	$11,917	$12,908	$13,898

Legacy infrastructure spend captured by containers ($mn)

% of total containerized applications migrated from VM or bare metal \ Container efficiency over VMs	3.5x	3.0x	2.5x	2.0x	1.5x	1.0x
19%	$507	$591	$709	$886	$1,182	$1,773
21%	$560	$653	$784	$980	$1,306	$1,960
23%	$613	$715	$859	$1,073	$1,431	$2,146
25%	$667	$778	$933	$1,166	$1,555	$2,333
27%	$720	$840	$1,008	$1,260	$1,680	$2,520
29%	$773	$902	$1,082	$1,353	$1,804	$2,706

Total container market size ($mn)

$1,745	$2,820	$3,929	$5,096	$6,382	$7,963
$2,047	$3,131	$4,252	$5,438	$6,755	$8,399
$2,598	$3,691	$4,824	$6,029	$7,377	$9,083
$3,646	$4,747	$5,893	$7,117	$8,496	$10,264
$5,688	$6,798	$7,957	$9,199	$10,609	$12,439
$9,720	$10,839	$12,009	$13,270	$14,712	$16,604

Source: IDC, Datadog, Gartner, Company data, Goldman Sachs Global Investment Research

New software infrastructure spending allocated to containers

New software infrastructure spending allocated to containers

We estimate ~$5.5bn of the new infrastructure spending in 2021 can be captured by containers. Our methodology is outlined below and in Exhibit 15.

- **Total vCPU charges** are based on the number of vCPUs used by each task and the duration of each task. It can be calculated by multiplying the total number of tasks by the price per vCPU-second, the number of vCPUs used by each task, and the total duration of all tasks.
  - **Paid container instances:** We use IDC data on the number of total paid container instances each year. An instance refers to the moment a container runs. The majority of container usage is for new applications and workloads, and as a result, we assume IDC container instance data to be representative of the opportunity for spending on containerizing new applications and workloads. For 2021, IDC forecasts 1.63 billion paid container instances. Importantly, we exclude unpaid instances, which are container deployments primarily by hyperscale web and SaaS providers. While these may account for a large proportion of container deployments today, these container instances likely have very little revenue impact for container vendors.
  - **Price per vCPU-second:** Prices are specific to AWS Fargate. The listed price per vCPU-second is $0.00001406.
  - **Number of vCPUs used:** Container instance pricing depends on the amount of memory needed. For the base case, we assume that all instances use 4 vCPUs and that deployments are for production purposes. In our sensitivity analysis, we flex the number of vCPUs between 1 and 6.
- **Total memory charges** are based on the amount of memory (GB) used and the duration of each container. It can be calculated by multiplying the total number of container instances by the price per GB-second, the amount of memory consumed by each container, and the duration of each container. The price per GB-second is $0.00000353. For the base case, we are assuming 8 GB of memory are used by each container, which can vary depending on if the deployment is low (1GB) to a high memory (32 GB). In our sensitivity analysis, we flex the amount of memory between 1GB and 32GB.
- **Task durations:** A Datadog survey found that the average container lifetime is 2 days and 0.5 days if the container is orchestrated. Given the increasing popularity of container orchestration systems like Kubernetes, we assume that as more containers are being orchestrated, the average container lifetime is decelerating from 2 days in 2016 to 0.5 days in 2021 as Kubernetes adoption becomes increasingly prominent and container usage becomes increasingly efficient.
- **Total container spending** is calculated by adding total vCPU charges and total memory charges.

For pricing metrics, we used AWS Fargate pricing; by using public cloud price per instance metrics, we believe we have a more conservative market size estimate.

Although a small proportion of container instances were deployed on-premises, we expect that on-premises deployments are less resource-efficient, and would therefore drive up the cost of container usage as users would have to provision more VMs in order to scale the application (as VMs are resource constrained). Our decision to use AWS-specific prices is based on a 2018 IDC survey which found that AWS captured the largest percentage of the total cloud container footprint among respondents. Furthermore, Fargate has a lower price per GB-sec than Azure Container Instance. We also assume that Fargate prices are comparable to or cheaper than AWS ECS, Azure Container Service, and Google Container Engine. Customers using Fargate do not need to provision virtual machine capacity (and potentially pay for excess capacity) and instead, only pay for the amount of consumed resources. We do not include additional charges associated with using other AWS services or data transfer charges.

We also do not adjust for potential price fluctuations. Looking at price fluctuations shown in Exhibit 16 of Amazon EC2 Instances, Microsoft Azure Instances, and Google Compute Engine usage, we note that virtual machine instance pricing has generally decreased from year to year. These prices are also used for Amazon’s Elastic Kubernetes Service, Azure Kubernetes Service, and Google Kubernetes Engine.

Exhibit 16: Public cloud compute prices have decreased over time
Snapshot of 2016 vs. 2019 compute prices for varying memory and vCPU types; prices as of 1/22/2019

Amazon EC2 Instances (Linux, General Purpose)				
	vCPU	Memory (GB)	Price	Discount vs. 11/6/16
t2.small	1	2	$0.023 / Hour	(12%)*
m4.xlarge	4	13	$0.20 / Hour	(25%)*
m4.2xlarge	8	26	$0.40 / Hour	(25%)*
m4.4xlarge	16	53.5	$0.80 / Hour	(28%)*

* Compared to similar CPU and Memory

Microsoft Azure Instances (Linux, Standard Tier)				
	vCPU	Memory (GB)	Price	Discount vs. 11/6/16
A1	1	1.75	$0.060 / Hour	-24%
A2	2	3.5	$0.120 / Hour	-24%
A3	4	7	$0.240 / Hour	-20%
A4	8	14	$0.480 / Hour	-20%

Google Compute Engine (Standard Machine Types)				
	vCores	Memory (GB)	Price	Discount vs. 11/6/16
n1-standard-1	1	3.75	$0.0475 / Hour	-47%
n1-standard-2	2	7.5	$0.095 / Hour	-47%
n1-standard-4	4	16	$0.190 / Hour	-47%
n1-standard-8	8	30	$0.380 / Hour	-47%

Source: Company data, Goldman Sachs Global Investment Research

Existing software infrastructure spending captured by containers

We estimate that ~$1.1bn of software spending is potentially redistributed to container spending, based on Gartner data and forecasts for infrastructure software markets (Exhibit 15). We note that the Application Platform Software and x86 Server Virtualization Infrastructure markets, as defined by Gartner, provide opportunities for spending redistribution.

Application Platform Software is largely comprised of middleware expenditures, which refers to software that lies between the operating system and the applications running on top of the OS. Application servers essentially provide an environment for applications to run on, which containers can do as well with an additional advantage of application portability, making them a viable substitute for traditional middleware software. As developers shift workloads to containers, we can expect to see a shift in middleware spending to container technology. According to Gartner, the middleware market is currently a ~$36bn market, of which Application Platform Software makes up $4.8bn (~8%), estimated to grow at a 3% '17-'22 CAGR.

A similar argument holds for server virtualization spending - for applications in which security is not a primary concern, developers can opt to use containers over virtual machines. Even if developers continue to run containers in virtual machines, the container market can still capture a portion of the server virtualization market. Fewer virtualization resources are needed to run the same workloads and containerizing applications that were running in VMs means developers can run more applications per server. Gartner currently estimates that x86 Server Virtualization Infrastructure is a $4.3bn market and is expected to remain constant or decline over time. We expect a portion of server virtualization spending can be reallocated to containerizing these existing applications.

Gartner's infrastructure software markets forecasts also include two additional areas that we do not include in our figures, but are relevant to container spending: Application Platform as a Service (aPaaS) and Delivery Automation.

- **Application Platform as a Service (aPaaS)** are cloud application infrastructure services which provide an environment to develop and execute applications. Given this category's general functional overlap with cloud container services, it is likely that this includes spending on containerizing new applications. This lends itself to the potential of double-counting a sizable portion of the container market. Gartner estimates that aPaaS is currently a $4.7bn market that is expected to increase to a $5.8bn by 2021.
- **Delivery Automation** refers to technology designed to automate (to some extent) manual processes, and included in this market is Container Management. The share of revenue that Container Management contributes is unclear as Gartner does not break down the Delivery Automation category further. As this category includes a long list of other subcategories, we estimate that the Container Management category only makes up a small percentage of the Delivery Automation market. Delivery Automation is currently a $8.5bn market that Gartner projects will increase to $9.8bn by 2021.

It is still possible that we have double-counted some portion of container spending. Application Platform Software, for example, also includes spending on container platforms such as Red Hat OpenShift Container Platform and Pivotal Container Service. This could include spending on containerizing new applications, but is likely minimal as application servers still constitute a majority of the Application Platform Software market.

We can estimate the total available market for containerizing existing applications by adding the market forecasts for x86 Server Virtualization Infrastructure and Application Platform Software. According to IDC survey results, approximately 43% of all respondents' applications run in a container today, and 54% of those containerized applications were moved from a virtual machine or from bare metal. Synthesizing these results suggest that 23% of all containerized applications were moved from a virtual machine or bare metal server. For the base case, we assume that this would imply that at most 23% of virtualization and middleware spending would be reallocated. Although this survey statistic is specific to 2018, we continue to use 23% through 2021 due to limited available data, but show scenarios in which the % of applications migrated could range from 19% to 29% in Exhibit 15.

Since applications can be more efficiently run on containers than in virtual machines or bare metal, we expect that it will cost less to run the same amount of workloads. For the base case, we assume 2x efficiency and consider cases in which containers are the same level of efficiency as VMs or up to 3.5x more efficient.

Serverless computing vs. containers

What is serverless computing?

Contrary to what its name suggests, serverless computing still relies on servers. Serverless computing is a form of cloud computing in which the cloud provider manages the server and resources needed to maintain the server. It was popularized by AWS Lambda, Amazon's serverless offering. Although serverless is largely associated with public cloud platforms, there are a small, but growing number of on-premises platforms. Apache OpenWhisk, for example, is an open source, serverless platform that can run in a private cloud using on-premises infrastructure.

Every action performed on a serverless platform is independent, which enables independent management and scaling. Thus, serverless environments are highly compatible with microservices architectures as they provide the isolation and scaling properties that enable microservices. As microservices become the norm, more IT teams may turn to serverless computing as a solution.

Serverless computing is implemented through Function-as-a-Service (FaaS), which falls in between Software-as-a-Service (SaaS) and Platform-as-a-Service (PaaS) models.

Exhibit 17: FaaS (also known as fPaaS) can be considered a subset of PaaS
FaaS vs. other as-a-Service models

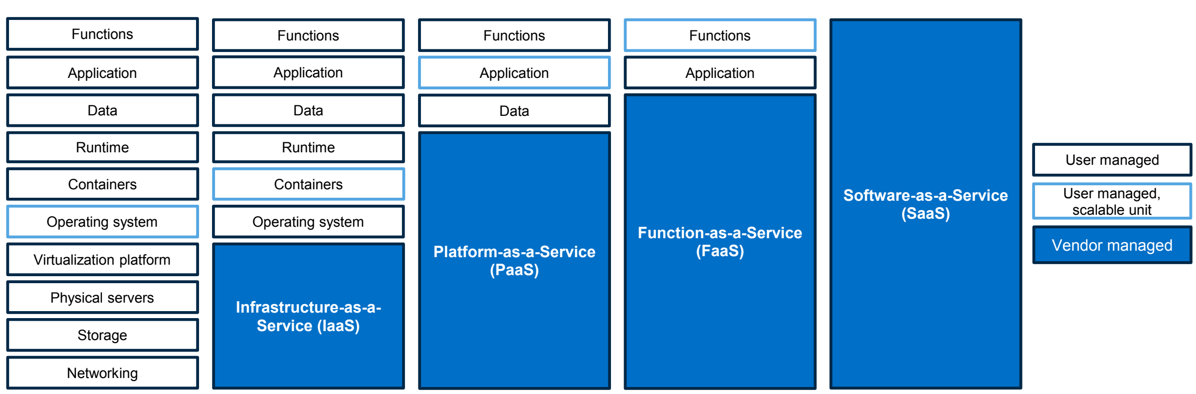

Source: Company data, Goldman Sachs Global Investment Research

Developers using FaaS are able to run applications by submitting functions to the serverless vendor. Functions comprise of a serverless application, and are packages that consist of some application code, its parameters, and its dependencies. These functions are executed when an event happens, such as an API (Application Programming Interface) call. The developer dictates what conditions will activate a function, and the platform checks for when these conditions are met. The FaaS model allows users to run applications and services without having to worry about how or where their applications are run.

For example, on a transit application running on AWS, when a user clicks on a subway schedule, an API is activated which calls the API endpoint, Amazon's API Gateway. The API Gateway activates a Lambda function, which is loaded into a container. The container runs code that pulls the transit schedule from the database and sends it back to the user. Once the task is complete, the container shuts down and the application developer is only charged for the amount of time the code was executed.

Exhibit 18: Serverless computing applied to a mobile transit application

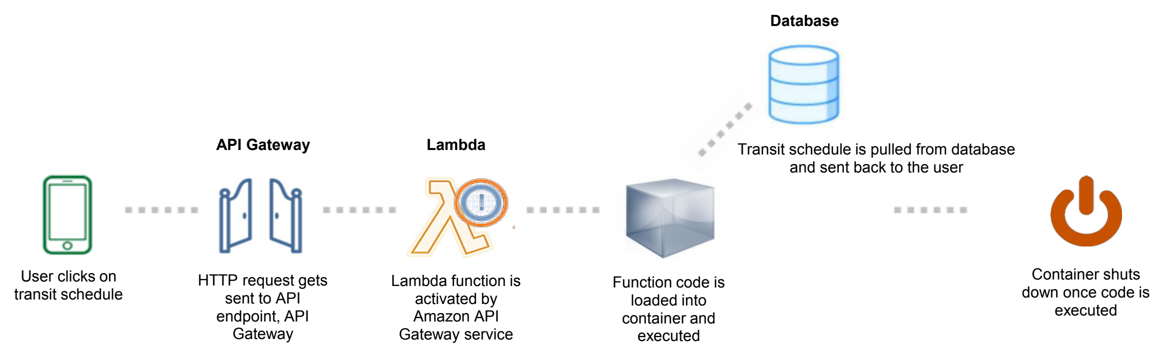

Source: Goldman Sachs Global Investment Research

Benefits of serverless computing include:

- **Ease of scaling**: Users of traditional PaaS need to estimate the amount of resources needed by an application ahead of time and provision accordingly. Comparatively, serverless computing scales applications automatically. The serverless provider takes care of replicating the application's functions and distributing computing resources as needed, reducing costs as users tend to over-provision resources.
- **Reduced operating costs**: Serverless computing is a pay-as-you-go service; costs are based on the resources consumed or the code run time as the infrastructure is de-allocated after the code runs. Companies no longer need to pay to have additional servers (physical or virtual) that are only waiting for larger workloads and therefore underutilized. Furthermore, companies are able to reduce infrastructure management costs as the underlying infrastructure is managed by the cloud provider.
- **Productivity**: Developers only need to focus on the code and the application itself; they do not need to worry about the server or operational issues

- **Decreased time to market**: Because developers do not have to spend time worrying about provisioning resources or the underlying infrastructure, producing an application takes less time.

Few companies, however, have actually adopted serverless computing; <5% of enterprises today are using serverless computing, which by 2020, is expected to increase to more than 20% (Gartner). Netflix, Adobe, and Airbnb are examples of early adopters that have partnered with AWS Lambda. Meanwhile, other companies are reluctant to turn to serverless computing because of its potential downsides:

- **Vendor lock-in:** A server and its underlying infrastructure are still needed to deploy a function, and these are vendor-specific. As a result, users depend on the vendor's other cloud services and have to write code that is catered to the serverless platform, making it difficult to switch services.
- **Lack of control:** Although one of the benefits of serverless computing is that users no longer need to manage servers, this can also be viewed as a downside. Users are no longer able to cater their computing environment to their application workloads. Compliance, monitoring, and debugging capabilities are also limited and still being developed.
- **Multitenancy:** Serverless platform resources are shared by multiple customers. If there are issues with one customer's code or if the application hogs resources, it can affect the performance of other customers' applications.
- **Limited run-time capabilities:** On serverless platforms, code is meant to be executed for seconds or milliseconds. On the other hand, containers can run for minutes and virtual machines can run for hours or days.

Exhibit 19: Virtual Machines vs. Containers vs. Serverless Computing

	Virtual Machines	Containers	Serverless
Scaling Unit	Virtual Machines (VMs)	Applications	Functions
Life Span	Days to months	Seconds to minutes	Milliseconds to seconds
Performance	Predictable	Less predictable	Least predictable
Operational Costs	High	Medium	Low
Operational Control	High	Medium	Low
Provider Lock-In	Medium to High	Low	High

Source: Gartner, Goldman Sachs Global Investment Research

Characteristics of serverless workloads

Due to the structure of serverless computing, serverless workloads should have the following properties:

- **Short runtime of code:** Functions are designed to execute code for seconds or milliseconds. Serverless computing is not useful for long-running applications.
- **Variable and infrequent workloads:** Applications are scaled automatically by the vendor; if demand for the application is too high, the vendor can quickly copy it and redistribute computing resources as needed. The serverless pay-as-you-go pricing structure can also help reduce costs in this scenario. Purchasing, configuring, and

managing an infrastructure is expensive and unnecessary for applications that do not need to be running all of the time.

- **Stateless and event-driven:** Functions are meant to be spun up and shut down quickly, and are usually triggered by events such as HTTP requests. Data from the execution of the function are neither saved nor made available for the next execution.

There are also several types of applications that may never be serverless, such as legacy applications that are monolithic because they contain interdependent parts and are often stateful. Refactoring a monolithic application into the appropriate architecture may not be worth the resource and time investment. Generally, applications that are stateful in that it generates a lot of data that is stored and recalled and is sensitive to performance and latency (such as transaction-processing applications or sites) are not suited for serverless computing. Furthermore, long-running jobs are not suited for serverless given that the functions expire after several minutes.

Assessing the serverless computing market opportunity

Generally available serverless platforms are relatively new to the computing scene, and as a result, very little data is available on serverless adoption usage and spending trends. We expect that the services provided by public cloud vendors such as AWS Lambda will continue to dominate the market as the base set of capabilities provided by them is generally good enough for customers. We note that there are on-premise open source products and an ecosystem of services being built on top of services like AWS Lambda (security, monitoring, debugging, etc), but are fairly nascent and most likely will comprise of a small portion of the market.

As serverless adoption continues to expand, we expect that AWS Lambda will likely maintain its leading position in the serverless market. AWS was the first of the three major cloud vendors to launch a serverless platform, and has since dominated the serverless space with AWS Lambda. AWS Lambda has continued to be the leader in the serverless market due to its large collection of products available for integration with Lambda functions. For example, Thomson Reuters started using AWS Lambda mainly because they were already leveraging AWS in other capacities and therefore, familiar with AWS capabilities.

Azure Functions is also quickly evolving and has continued to round out their assortment of cloud products. Their recent release in September 2018 (Functions 2.0) improved host run time and allows customers to run Azure Functions in more environments, such as locally on a Mac or Linux machine. Azure Functions can also be run on Kubernetes and on IoT Edge (public preview). Azure Functions on IoT Edge allows customers to deploy code that acts on Azure IoT Edge devices, such as deploying an Azure Function that filters sensor data from an IoT Edge device. To enhance visibility into the serverless application, a drawback as previously outlined, Azure Functions is now integrated with Azure Application Insights which provides more visibility and tracing of dependencies. Furthermore, new security features continue to be introduced to Azure Functions. Meanwhile, Google Cloud Functions seems to lag behind Azure Functions and AWS

Lambda, as Google has just made the service generally available in August. Much of the focus has been on Knative, which brings serverless to Kubernetes.

The future of Pivotal Function Service (PFS) and OpenShift Cloud Functions are still unclear as both were previewed less than a year ago and are still in early stages. Pivotal announced a preview of Pivotal Function Service (PFS) in Dec 2017, and Red Hat announced a preview of OpenShift Cloud Functions in May 2018. Both seek to provide customers the ability to run functions both on-premises and in the public cloud. Management is positive about PFS opportunity, which was expected to be GA at some point in 2018, but it still not yet available. Further information regarding OpenShift Cloud Functions have not been released yet.

Exhibit 20: AWS Lambda vs. Azure Functions vs. Google Functions features

	AWS Lambda	Azure Functions	Google Cloud Functions
Number of years in general availability	3 years	1.5 years	7 months
Supported languages	Largest number of supported languages C#, JavaScript, Java, Python, Go	Supports fewer languages than AWS Lambda C#, F#, JavaScript, Java, Python	Least number of supported languages JavaScript
Accepted code format	.zip and .jar uploads	Web editor	.zip uploads and Google repositories
Integrated services (supported event sources)	Large collection (~19) of AWS services, API gateway	Azure services (~7), GitHub, On-premises, Twilio	Internal event bus (Cloud Pub/Subtopics), Firebase, API gateway
Maximum execution time	300 seconds	300 seconds	540 seconds
Price	Requests: $0.20 per 1 million requests *First million requests per month are free Time: $0.00001667 for every GB-sec *First 400k GB-sec of compute time per month are free	Same as AWS Lambda	Higher price point than Lambda and Azure Functions Requests: $0.40 per 1 million requests *Initial two million requests per month are free Time: $0.0000231 for every GB-sec

Source: Goldman Sachs Global Investment Research

Knative – a potential open source alternative to AWS Lambda

Knative, announced at Google Next 2018, is an open-source project that is built on top of Kubernetes that enables customers to build and deploy container-based serverless applications in both on-prem and public cloud environments. Knative is being developed in close partnership with Pivotal Software, IBM, and SAP. The Knative components enable users to focus on writing the code while the components manage the orchestration, management, and scaling of the workload. More specifically, the serverless add-on component of Knative enables serverless functions to be easily run on Google Kubernetes Engine. We note that Knative could become a viable, open-source alternative to AWS Lambda given these features, the growing popularity of Kubernetes for orchestration, and its compatibility with on-premise environments (given that it is being built in conjunction with on-premise vendors).

Serverless computing reduces infrastructure management costs and enables new use cases

We note that serverless computing can be viewed as both deflationary due to lower application run costs, but also inflationary as it enables new use cases. One of the key benefits of serverless computing is how easily scalable its resources are without advance provisioning. As a result, serverless computing can be deflationary to traditional

IaaS as customers are only paying for the resources consumed while their applications are running, eliminating both costs associated with infrastructure management and the tendency for customers to over-provision resources. On the other hand, the key characteristics of serverless computing can enable new use cases and applications such as microservices applications, streaming and edge processing, batch processing, and automation of IT processes.

- **Microservices:** As previously mentioned, microservices architectures requires an environment where applications can be rapidly scaled independently. The ease of scaling offered by serverless computing is well-suited for a microservices architecture.
- **Stream & Edge Processing:** Stream processing requires real-time processing given the unpredictability of its high volumes of data traffic. Serverless enables edge processing, with use cases in voice-enabled devices like the Alexa or Google Home, which requires fast, real-time response rates. For more information on other edge applications that could benefit from serverless computing, please refer to our edge computing report (link).
- **Batch Processing:** Scheduled batch processing needs are well-suited for serverless because they are generally short-run and requires minimal intervention.
- **IT automation:** Automated IT tasks such as backup or patching can be initiated by serverless functions because they are event-driven.

We note that use cases such as using serverless for batch processing or as a part of applications with a microservices architecture are the most popular use cases, with edge/stream processing starting to gain traction, although still in its early days. According to Gartner, Autodesk and FINRA are two organizations that are early adopters of serverless computing. Autodesk utilizes AWS Lambda for batch processing and IT automation, automating the on and offboarding of accounts and the processing of log files. FINRA utilizes AWS Lambda to trigger data validation once a file has been uploaded and utilizes EC2 as a fallback in case the validation of the file exceeds the limits of a serverless function (five minutes).

Given how relatively nascent the technology is, we expect the serverless market opportunity to continue to expand as the offerings mature and more features such as debugging and monitoring are built out. We expect AWS Lambda to continue to remain a market leader given its maturity and increasing set of features. However, we note that Knative could become a viable, open-source alternative to AWS Lambda, although it is in its early stages of development.

Containers can coexist with serverless computing

When serverless platforms became available, there was debate over whether going serverless would eliminate the need for containers as cloud providers, instead of companies, would be deploying containers. This may be preferable for certain companies in which managing containers proves to be difficult; users may opt to use serverless computing to avoid having to deal with containerization, allowing them to focus solely on their application. In this extreme scenario, container usage would be limited to serverless computing providers. However, we believe this is unlikely to happen as in certain use cases, containers make up for what serverless computing lacks:

- **Management flexibility**: When going serverless, the user has to rely on how the vendor manages its servers and resources. With containers, the user can have control of the container system, including how to allocate and manage resources. Although this comes with an added level of complexity, it gives users more flexibility in managing their applications.
- **Application use-case variety**: Containers can support a wider variety of applications. Similar to serverless computing, containers are best suited for stateless applications since they do not have persistent storage and shuts down after the applications are run. However, unlike serverless platforms, containers can also handle stateful applications as platforms like Pivotal Container Service and Red Hat OpenShift support both stateless and stateful (we outline the difference between stateless and stateful in Exhibit 4) applications, providing persistent storage for containers such as databases for storing container state information. Developers are also able to utilize external storage tools which are mapped to the container and manage data outside of the container, allowing data to exist after the container terminates.
- **Code run-time**: Containers allow for longer code runtimes, while serverless platforms will limit the amount of time a function can be deployed. AWS Lambda, for example, imposes a 900 second limit after which, the function is aborted.
- **Vendor flexibility:** Users can take advantage of offerings from different vendors when running containers, avoiding the risk of vendor lock-in. For example, the Canonical Ubuntu operating system supports Docker containers, which can then be paired with a Kubernetes orchestration system. Conversely, serverless computing restricts the user to a single serverless platform.
- **Security:** Serverless functions are difficult to monitor, have a larger attack surface due to the increased number of events that can call the function, and rely on more third-party libraries, which enhance functionalities, but are also vulnerable to attacks and require monitoring.

Moving forward, we will likely see many IT teams using both computing approaches. The decision to use one over the other will vary depending on application type and team preferences. Serverless computing is more appropriate for more specialized use-cases and event-driven applications whereas containers are better suited for general use-cases and purposes.

Exhibit 21: Timeline of serverless computing product releases
Updated as of Feb 2019

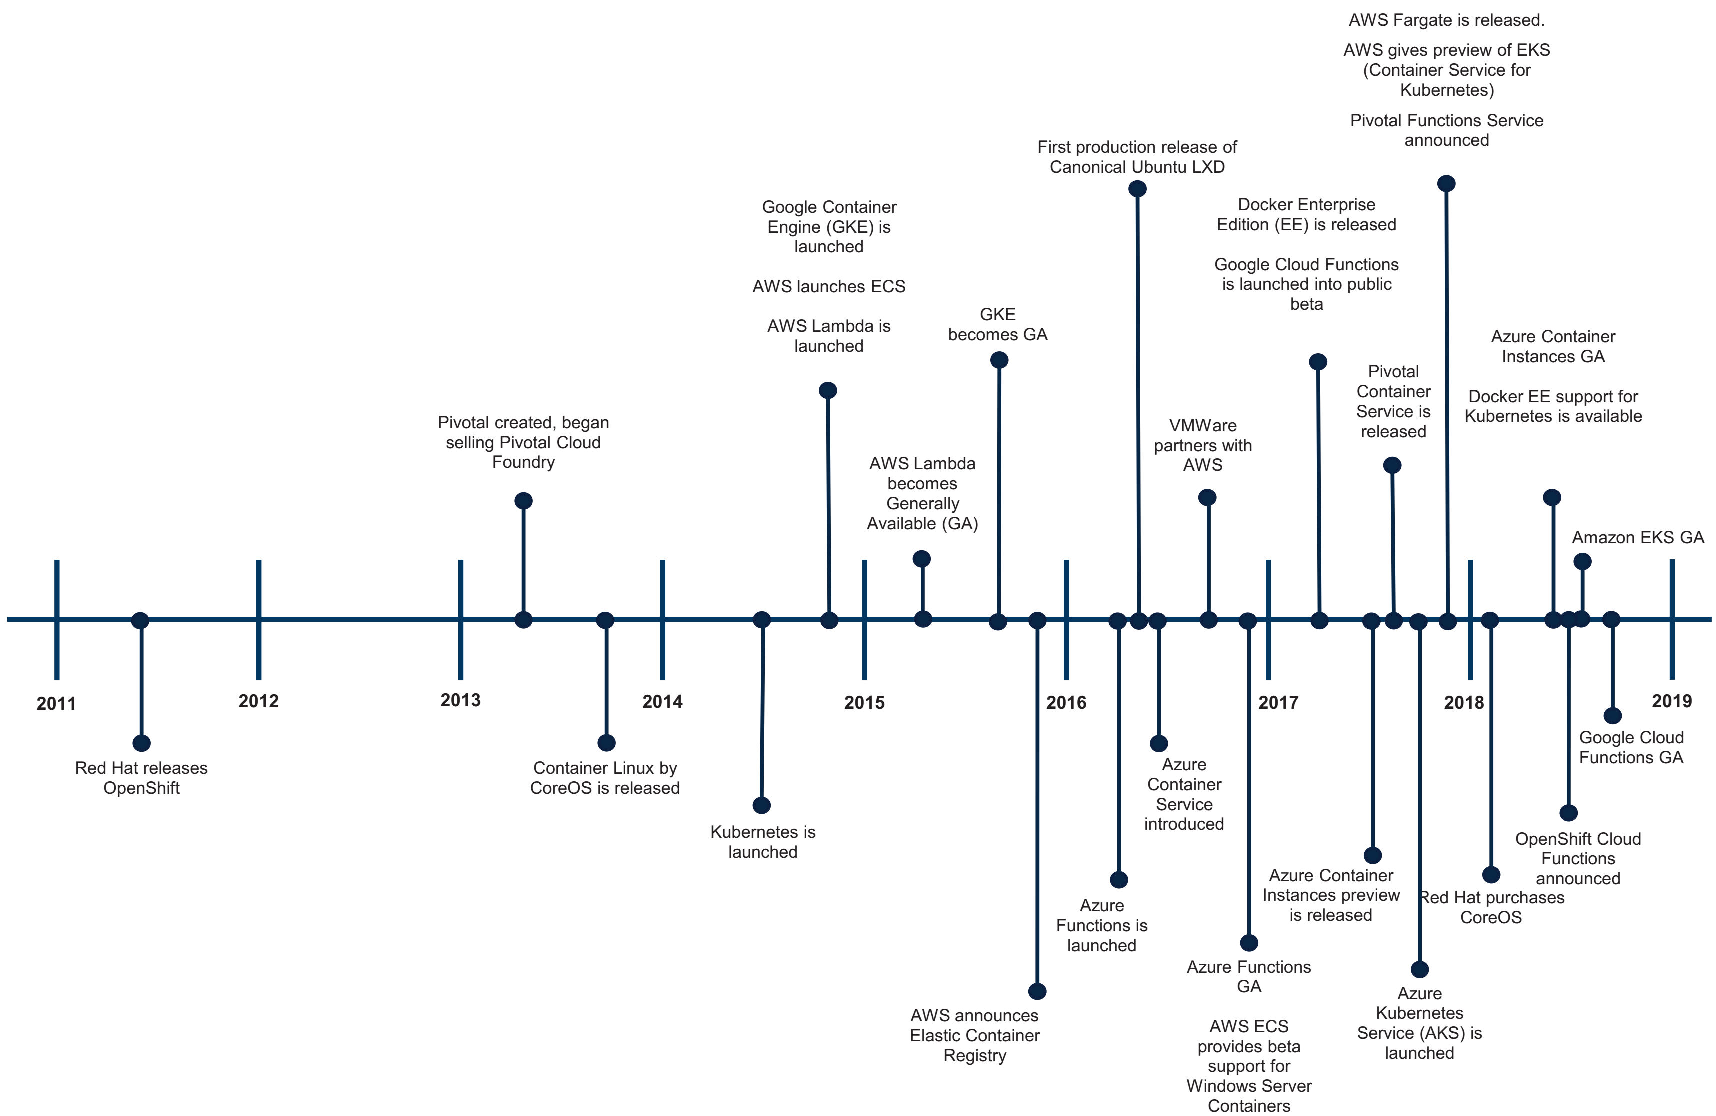

Source: Company data, Goldman Sachs Global Investment Research

Leaders & laggards of the container market

The core of the container framework lies with the container technology; Docker (the software) and CoreOS rkt are based on Linux technology, and all three are designed to run on Linux operating systems. Linux containers are more lightweight than virtual machines, but both are similar in that they can contain an operating system. Docker and CoreOS rkt contain applications.

Branching out, the container framework includes an ecosystem of services that companies can take advantage of when adopting containers: 1) guest operating system offerings, 2) orchestration and scheduling services, 3) PaaS platforms, and 4) deployment of containers via public cloud.

Exhibit 22: Container framework

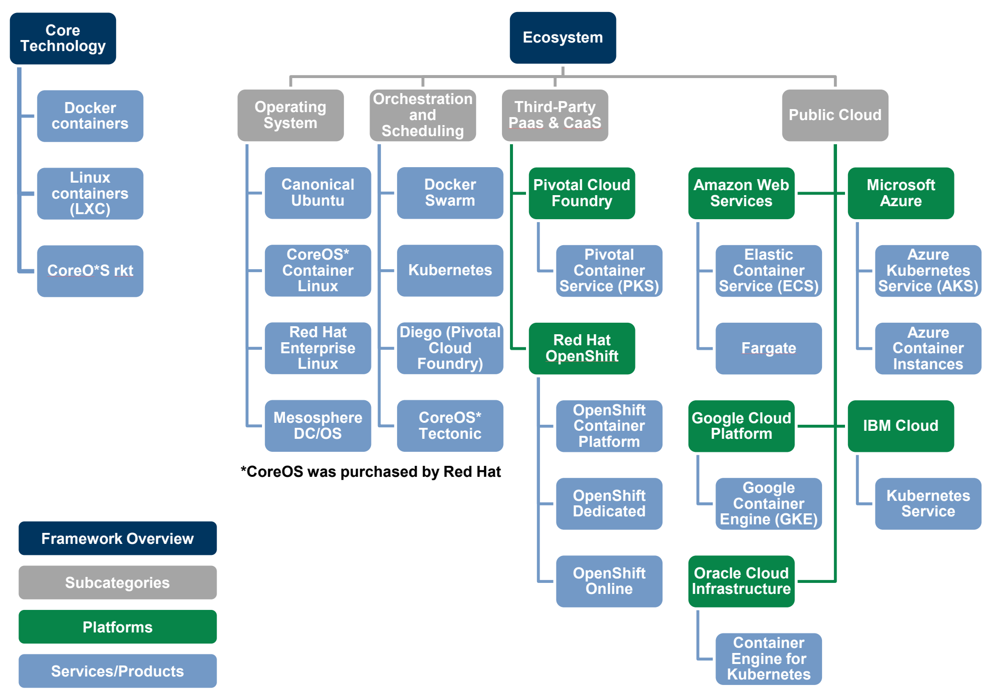

Source: Goldman Sachs Global Investment Research

- **Operating Systems:** OS offerings allow containers to run on a variety of platforms, regardless of whether it is a physical, virtual, or public/private cloud platform.
- **Orchestration and Scheduling:** Container orchestrators automate the deployment, scaling, and maintenance of containers. Orchestration is particularly useful as the number of containers and container hosts (i.e. the operating system a container runs on) grow. Orchestration systems allow developers to focus on running the application instead of the specific implementation details. They include a number of features, such as being able to provision container hosts, reschedule failed containers, and scale by adding or removing containers. Kubernetes (also known as K8s) is an open source orchestration system created by Google that has dominated

the container orchestration market; almost all container services offer some distribution of Kubernetes.

- **Third-Party PaaS & CaaS:** Third-party vendors such as Pivotal and Red Hat provide container services that can run on-premise and on the public cloud. In a PaaS approach, users only have to worry about running the application code; the vendor creates and packages the container for the user. This differs from CaaS (Containers-as-a-Service), in which the user is still responsible for creating the container image and related steps.
- **Public Cloud:** Public cloud vendors, such as AWS, Azure, and GCP, offer container solutions as a means of simplifying the container deployment and management process. Containers are launched on a cloud platform that provides orchestration support, in which vendors take on a large part of the complexity involved with containerization, allowing developers to focus more on their application. These cloud platforms also support and run third-party container services and products, such as those from Pivotal and Red Hat.

Pivotal's & Red Hat's multi-platform services

Exhibit 23: Summary of Pivotal Container Service vs. Red Hat OpenShift features

	Pivotal Container Service	Red Hat OpenShift
Platform	-vSphere -Google Container Engine (GKE)	-Red Hat Enterprise Linux -Red Hat CoreOS (previously CoreOS Container Linux)
Infrastructure	Multi-cloud	Multi-cloud -OpenShift Online (multi-tenant public cloud offering) -OpenShift Dedicated (single-tenant public cloud offering) -OpenShift Container Platform (private cloud offering)
Container Runtime	Cloud Foundry Container Runtime (previously Kubo)	Docker
Orchestration	Kubernetes	Kubernetes (starting in 2015), previously proprietary
Registry	Harbor (VMware)	-Proprietary -CoreOS Quay
Security features	NSX-T (VMware)	Proprietary
Storage	vSAN (VMWare)	Proprietary
Application States	-Stateful -Stateless	-Stateful -Stateless
Pricing	-Higher price point (approximately 2x OpenShift price) -Consumption-based pricing based on number of pods (collection of containers that share the same computing resources)	-Lower price point -Capacity-based pricing based on core count

Source: Gartner, Goldman Sachs Global Investment Research

Exhibit 24: Pivotal Container Services vs. Red Hat OpenShift platforms

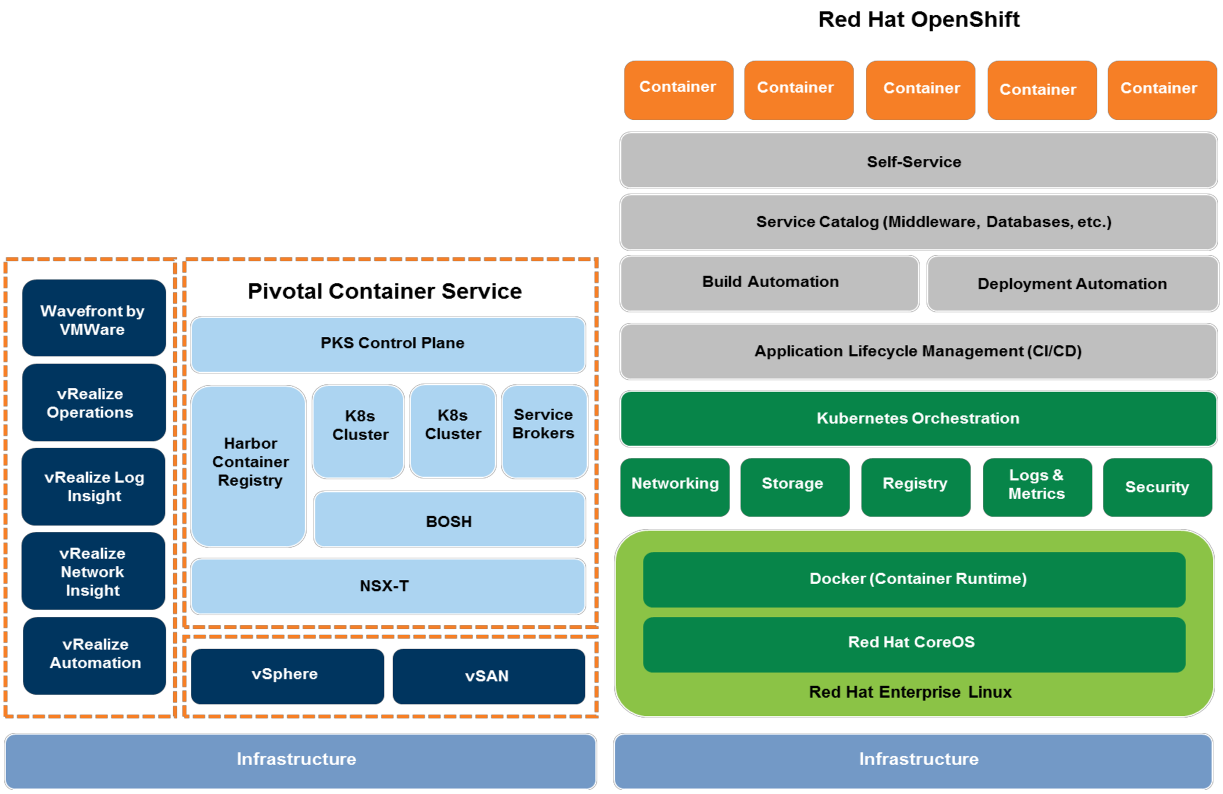

Source: Company data, Goldman Sachs Global Investment Research

Pivotal Container Service (PKS) partnership

Pivotal Container Service (PKS) is collaborative project between Pivotal, VMware, and Google Cloud that supports developers in building and efficiently running applications on multiple platforms. PKS is essentially an enterprise Kubernetes offering that is designed to use the same Kubernetes version as Google Container Engine (GKE), making it easy to migrate workloads between vSphere (VMware's server virtualization platform) and GKE. Due to its cross-compatibility with GKE, PKS will always have the latest, stable version of Kubernetes. Furthermore, PKS can be run on AWS and in its 1.3 release, now has support for Azure.

Pivotal jointly markets and sells its products and services with Dell/EMC and VMware, both of which have large preexisting footprints in most IT organizations. We note that transactions processed through the company's agency agreements with Dell/EMC and VMware accounted for 44% of the company's revenue in FY17 and 37% in FY18. Pivotal's agency fees to VMware range from 4% to 10% of the invoiced contract amount, incentivizing VMware to attach Pivotal offerings. Notably, PKS also includes VMware's NSX offering to manage software-defined virtual networks, and as a result, PKS is a 50/50 revenue share between Pivotal and VMware. During the F3Q19 earnings call, Pivotal noted that they are now seeing VMware's sales force starting to spin up and early traction within the customer pipeline for PKS. In November, VMware announced the acquisition of Heptio (closed in December), which was founded by two of the

Kubernetes creators and has a set of products to assist with the deployment and operation of Kubernetes. Management noted that VMware's acquisition of Heptio is also expected to help accelerate PKS.

Red Hat acquisition of CoreOS

Red Hat acquired CoreOS in early 2018 as a means of rounding out its container platform and increasing its influence in the container market. Red Hat seeks to integrate CoreOS Tectonic (a Kubernetes distribution), Container Linux (a lightweight operating system), and Quay (an enterprise container registry) into their existing ecosystem of container offerings by the end of 2018.

Red Hat is largely associated with Red Hat Enterprise Linux (RHEL), its general purpose, enterprise-ready, Linux-based operating system that was initially released in 2000. CoreOS released their open-source Container Linux in 2013, a Linux-based operating system that is a lightweight alternative optimized to run containers that also offers automatic, over-the-air system updates. Container Linux gained traction as containers grew in popularity. Two years later, Red Hat released its own lightweight OS for containers: a variation on RHEL called Atomic Host. However, our conversations with industry contacts suggest that Atomic Host did not garner much traction. Thus, the acquisition of CoreOS enables Red Hat to expand its existing container services and offer more flexibility to both existing and new users.

Red Hat's acquisition of CoreOS is also an attempt to better position itself as the market converges on Kubernetes for container orchestration. CoreOS Tectonic is an enterprise-ready Kubernetes distribution that features automated, over-the-air updates, which Red Hat intends to integrate into its existing Kubernetes system in OpenShift. The automated updates simplify the operational aspect of using Kubernetes, ensuring that users are running the most up-to-date version of Kubernetes.

Privates in the container space

There are a number of private companies in the container market that offer specialized solutions, such as Containers-as-a-Service (CaaS), operating systems, security features, and/or configuration management. We highlight several of them in Exhibit 25 below.

Exhibit 25: Categorical breakdown of private company offerings

Category	CaaS	Operating System		Security		Configuration Management	
Company & Product	**Docker Enterprise Edition**	**Canonical Ubuntu**	**Mesosphere DC/OS**	**Aqua Security - Container Security Platform**	**Twistlock**	**Chef Software**	**Puppet Labs**
Venture Capital Investors	AME Cloud Ventures, Benchmark, Greylock Partners		Andreessen Horowitz, Data Collective, Fuel Capital	Lightspeed Venture Partners, M12, TLV Partners	Dell Technologies Capital, ICONIQ Capital, Polaris Partners	Battery Ventures, Citi Ventures, Amplify Partners	EDBI, KPCB, Google Ventures
Latest Funding	Series E (Oct 2017) - $91.9M		Series D (May 2018) - $125M	Series B (Sep 2017) - $25M	Series C (Aug 2018) - $33M	Series E (Sep 2015) - $40M	Series F (June 2018) - $42M
Latest Valuation	$1.5bn (Oct 2017)		$402mn (Mar 2016)	$82mn (Sep 2017)	$76mn (Apr 2017)	$319mn (Sep 2015)	$550mn (June 2014)
Description	Enterprise-ready, on-premises CaaS platform	Linux-based operating system	Datacenter operating system that also acts as a container platform	Container security platforms that detect and protect containers from threats		Provides automated configuration management, ensuring that every system is configured correctly and consistently	
Supported by/Supports	-Windows -Amazon Web Services -Red Hat OpenShift -Azure -Oracle Linux -Ubuntu	-Docker EE -Pivotal Cloud Foundry	-Azure -Google Cloud Platform -Amazon Web Services	-AWS -Azure -GCP -Red Hat OpenShift -Pivotal Cloud Foundry (beta)	-AWS -Azure -Docker -GCP -Kubernetes -Mesosphere -Red Hat OpenShift	-GCP -AWS -Azure	-GCP -AWS -Azure

Source: Goldman Sachs Global Investment Research, CB Insights

CaaS: Since releasing their container software in 2013, **Docker** has been synonymous with the word containers. They offer Docker Enterprise Edition (EE), an enterprise-ready CaaS platform. Docker EE differentiates itself from other CaaS offerings through its ability to support Windows, Linux, and Linux-based mainframe applications on a variety of platforms. It can be installed on different operating systems, including Windows, Ubuntu, and RHEL, and can run on different cloud platforms as well. Docker EE integrates Docker Datacenter and Docker Swarm, its proprietary container orchestration service. Docker initially only supported Docker Swarm, but began supporting Kubernetes in 2018.

Operating system: Canonical offers an Ubuntu operating system that is used as a platform for numerous container products and services. Ubuntu can be integrated with Docker EE, and is also Pivotal's preferred Linux operating system. Canonical offers its own Kubernetes distribution, and also offers LXD, the "pure-container hypervisor." **Mesosphere** offers a different take on the traditional operating system with their Mesosphere DC/OS product. Datacenter operating systems are operating systems that run on several machines in a customer's datacenter or cloud and link them together so that they behave as one machine. DC/OS manages the collection of machines and provides services for deploying and running applications distributed throughout the collection. In addition to being an operating system, Mesosphere DC/OS also serve as a container orchestrator. Mesosphere DC/OS can be used with all three major cloud platforms, and is also partnered with vendors like Aqua and Twistlock.

Security: Container images are an inactive form of containers that can either be created or downloaded. Developers have to be cautious about potentially downloading a container image with malware or vulnerabilities. **Aqua** and **Twistlock** are container security platforms that scan container images for malware and potential threats, and also monitor and respond to threats against active containers. Both are intended to provide an end-to-end security platform. Aqua can be used on AWS, Azure, GCP, and Red Hat OpenShift, and is currently in beta for Pivotal Cloud Foundry. Twistlock provides solutions for AWS, Azure, Docker, GCP, Kubernetes, Mesosphere, and OpenShift.

Configuration management: Configuration management is a process that tracks and controls changes in system configurations (i.e. how different parts of a technology system are arranged and connected), and can be used for a variety of purposes, ranging from product management to debugging. With respect to debugging, configuration management can identify what changed from one test to the next, assisting users in determining the root cause of the bug. Configuration management enables DevOps, and in the container space, provides functionalities similar to that of Kubernetes. **Puppet** and **Chef** are two of the most well-known configuration management tools.

AWS emerges as the most popular platform for containers

Recognizing that implementing containers can be a daunting task, public cloud vendors offer Containers-as-a-Service (CaaS) offerings to simplify part of the adoption process, which are utilized by many companies and organizations. CaaS offerings shift the need to build and manage the underlying container infrastructure from the developer to the vendor. By mitigating many of the complexities associated with building the necessary technology to support containers, public cloud vendors are well-positioned in the midst of an increase in container adoption.

The three major public cloud vendors with container services are Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure. AWS was the first to release their initial container service, Elastic Container Service (ECS), followed by GCP with Google Kubernetes Engine (GKE), and then Azure with Azure Container Service (ACS). In addition to their own container services, they support third-party container products and services, such as those offered by Pivotal and Red Hat.

Exhibit 26: AWS vs. GCP vs. Azure container services

	AWS	GCP	Azure
Offerings	-Elastic Container Service (ECS) -Elastic Container Service for Kubernetes (EKS) -Fargate	Google Container/Kubernetes Engine (GKE)*	-Azure Kubernetes Service (AKS) -Azure Container Instances
Orchestration	-Proprietary -Blox (Custom) -Kubernetes (2018)	Kubernetes	-Mesos -Kubernetes -Docker Swarm
Registry	Amazon Elastic Compute Cloud (EC2) Container Registry (ECR)	Google Container Registry	Microsoft Azure Container Registry (ACR)
Storage	-Amazon Elastic Book Store (EBS) -Amazon Elastic File System (EFS)	Google Persistent Disks	-Azure Virtual Disk (AVD) -Azure Files -Azure Managed Disks
Monitoring and Logging	-Amazon CloudWatch -Amazon CloudTrail -Datadog, Splunk	-Stackdriver (Fluentd) -cAdvisor -Prometheus -Datadog, Splunk	-Microsoft Operations Management Suite (OMS) -Datadog, Splunk
Container Runtime	-Docker -Microsoft Windows	-Docker	-Docker -Microsoft Windows
Provisioning	-AWS CloudFormation -Terraform by HashiCorp	-Google Cloud Deployment Manager -Terraform	Azure Resource Manager
Supported Services/Products*	-VMWare vSphere -Red Hat OpenShift Container Platform -CoreOS,+ Tectonic -Mesosphere DC/OS -Docker Enterprise Edition -Pivotal Cloud Foundry	-Pivotal Cloud Foundry -Mesosphere DC/OS -Red Hat OpenShift	-Red Hat OpenShift Container Platform -Pivotal Cloud Foundry -Docker Enterprise Edition -Mesosphere DC/OS

*CoreOS is owned by Red Hat Google Cloud announced in July 2018 an on-premises offering of GKE. No further information is available yet. This list is not representative of all supported third-party services.

Source: Gartner, Goldman Sachs Global Investment Research

AWS currently has (and is expected to maintain) container market leadership

In 2014, AWS began offering Elastic Container Service (ECS), a proprietary container orchestration service that supports Docker containers and enables users to run and scale container applications that can leverage integrations with other AWS services. ECS supports both Linux and Windows containers. Datadog notes that as of November 2018, ~37% of workloads were orchestrated using ECS, closely trailed by Kubernetes with 35%. Given the growing popularity of Kubernetes, AWS launched Elastic Container Service for Kubernetes (EKS) in November 2017, a fully managed service for their customers that leveraged Kubernetes. In fact, Cloud Native Computing Foundation noted that in 2017, 63% of Kubernetes workloads run on AWS.

AWS also recently launched Fargate to address the maintenance issues and management of multi-layered access rules for Linux containers. By abstracting the underlying infrastructure, users can view all of their containers as a single machine and only have to specify the level of resources necessary for each container. Fargate manages the container hosts and determines how to launch the containers, essentially providing a serverless model for container orchestration. Currently, Fargate runs in 6% of AWS container environments (Datadog).

As a result of the numerous features and services that AWS offers for containers, AWS is the most popular platform for containers and we expect them to maintain their market leadership. According to our December 2018 CIO IT spending survey, as shown in Exhibit 27, over 40% of respondents indicated that they used AWS to run containers. AWS is expected to maintain this lead in the next three years, followed by Azure, which is expected to gain significant traction in the next few years. More specifically, 26% of respondents indicated that they utilized AWS ECS, followed by AWS EKS. Similarly, in a 2018 IDC survey of senior IT employees working with containers, roughly half of respondents indicated they used AWS to run containers and both Azure and GCP were leveraged by 45% of respondents. A little over 70% of IDC's respondents indicated that they used AWS ECS, 53% utilized AKS, and 42% leveraged GKE.

Azure expected to gain popularity as a container service

In 2016, Microsoft introduced Azure Container Service (ACS), which provided support for Mesosphere DC/OS, Docker Swarm, and Kubernetes orchestration systems, and supported both Windows and Linux containers. As the market converged on Kubernetes for orchestration, Azure relaunched with Azure Kubernetes Service (AKS), which shifted the focus from offering support for multiple orchestration systems to support for Kubernetes. Highlighting the popularity of Kubernetes, Microsoft announced in December that they will stop supporting ACS in January 2020, encouraging users to migrate to AKS. In 2017, the company announced Azure Container Instances (ACI) which, similar to AWS EKS, makes it easier for users to deploy containers as Azure would manage the deployments. Recently, the company released a public preview of AKS virtual nodes, which merges ACI's serverless capability with Kubernetes.

We note from our December 2018 CIO IT spending survey that while Azure is the second-most popular container service, it is expected to gain share in the next three years. 14% of respondents noted that they utilized Azure's container platform, although ~30% of respondents expect to utilize Azure in the next three years. Furthermore, Azure will also stand to benefit from widespread container adoption as interest in deploying containers on Windows servers is increasing. An IDC survey noted that 55% of respondents utilized Windows containers and predicts that Windows will have 32% of the share of worldwide container instances by 2021. Both Azure and AWS offer support for Windows systems, and are able to address the growing interest in Windows-based deployments.

GCP to benefit from the increasing popularity of Kubernetes

Launched in 2015, we note that Google's Kubernetes Engine (GKE) gives GCP an advantage, having created Kubernetes and continuing to be its top contributor. As a result, GKE is the most user-friendly and mature out of the three public cloud

Kubernetes offerings. As previously mentioned, Kubernetes for orchestration has been increasing in popularity, being the most popular on GCP with 85% of workloads (Datadog), followed by Azure (which has supported Kubernetes for two years) with 65% of workloads, and AWS with 35% (recently launched support for Kubernetes). While Datadog noted that AWS's ECS/Fargate orchestration is more widely used in AWS environments, Kubernetes is catching up quickly (35% of workloads in AWS vs. ECS/Fargate orchestrating 37%). We note that our December 2018 CIO IT spending survey showed that ~10% of respondents currently utilize or expect to leverage GCP's container offerings. However, Google also recently announced an on-premise offering of GKE, designed to function in data centers while also enabling users to make their applications "cloud-ready" and eventually migrate them to the cloud. Users can leverage a centralized management system that oversees both GKE and GKE On-Prem. GKE On-Prem is currently in early access mode, but depending on the traction the offering garners, it could provide GCP another competitive advantage in the container market.

Exhibit 27: AWS is the most widely utilized platform for containers...
% of respondents

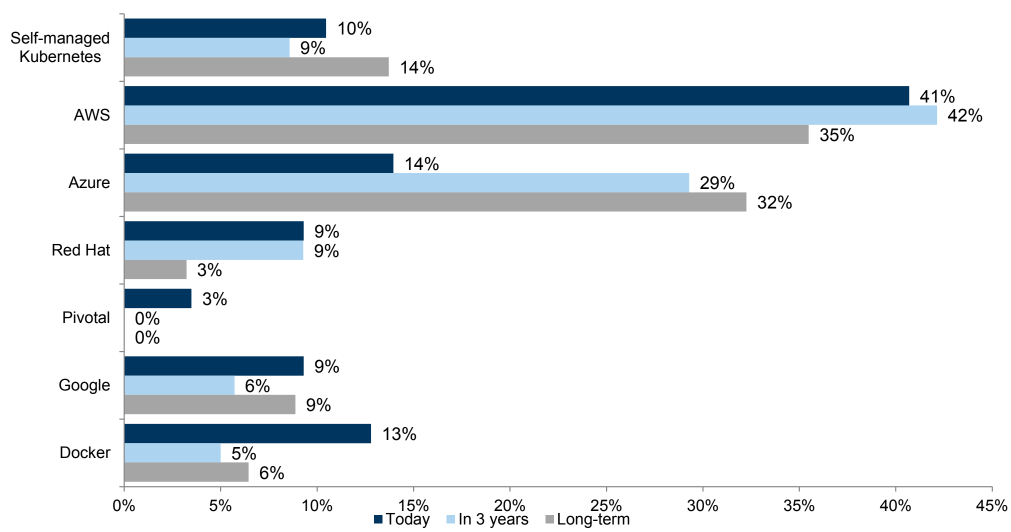

Source: Goldman Sachs Global Investment Research

Exhibit 28: ...with Amazon Elastic Container Service (ECS) being the most popular offering
% of respondents

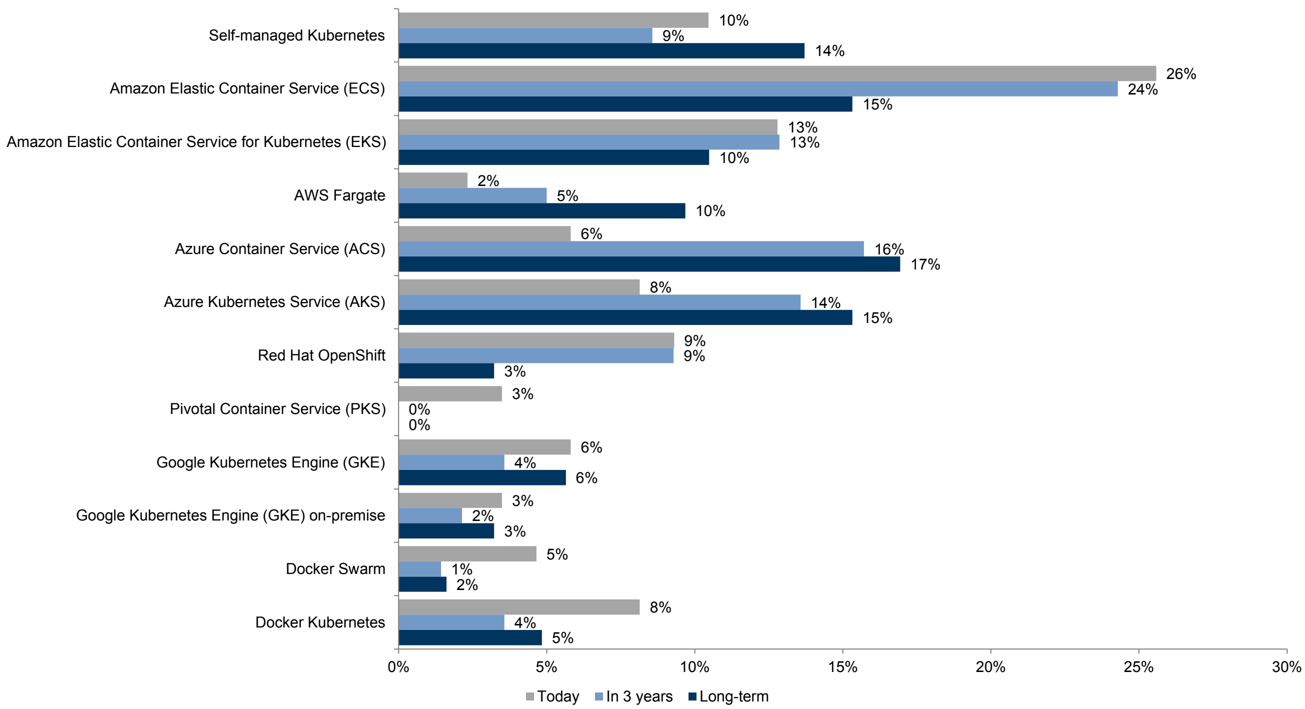

Source: Goldman Sachs Global Investment Research

Exhibit 29: Multiple trends benefit all three public cloud container services

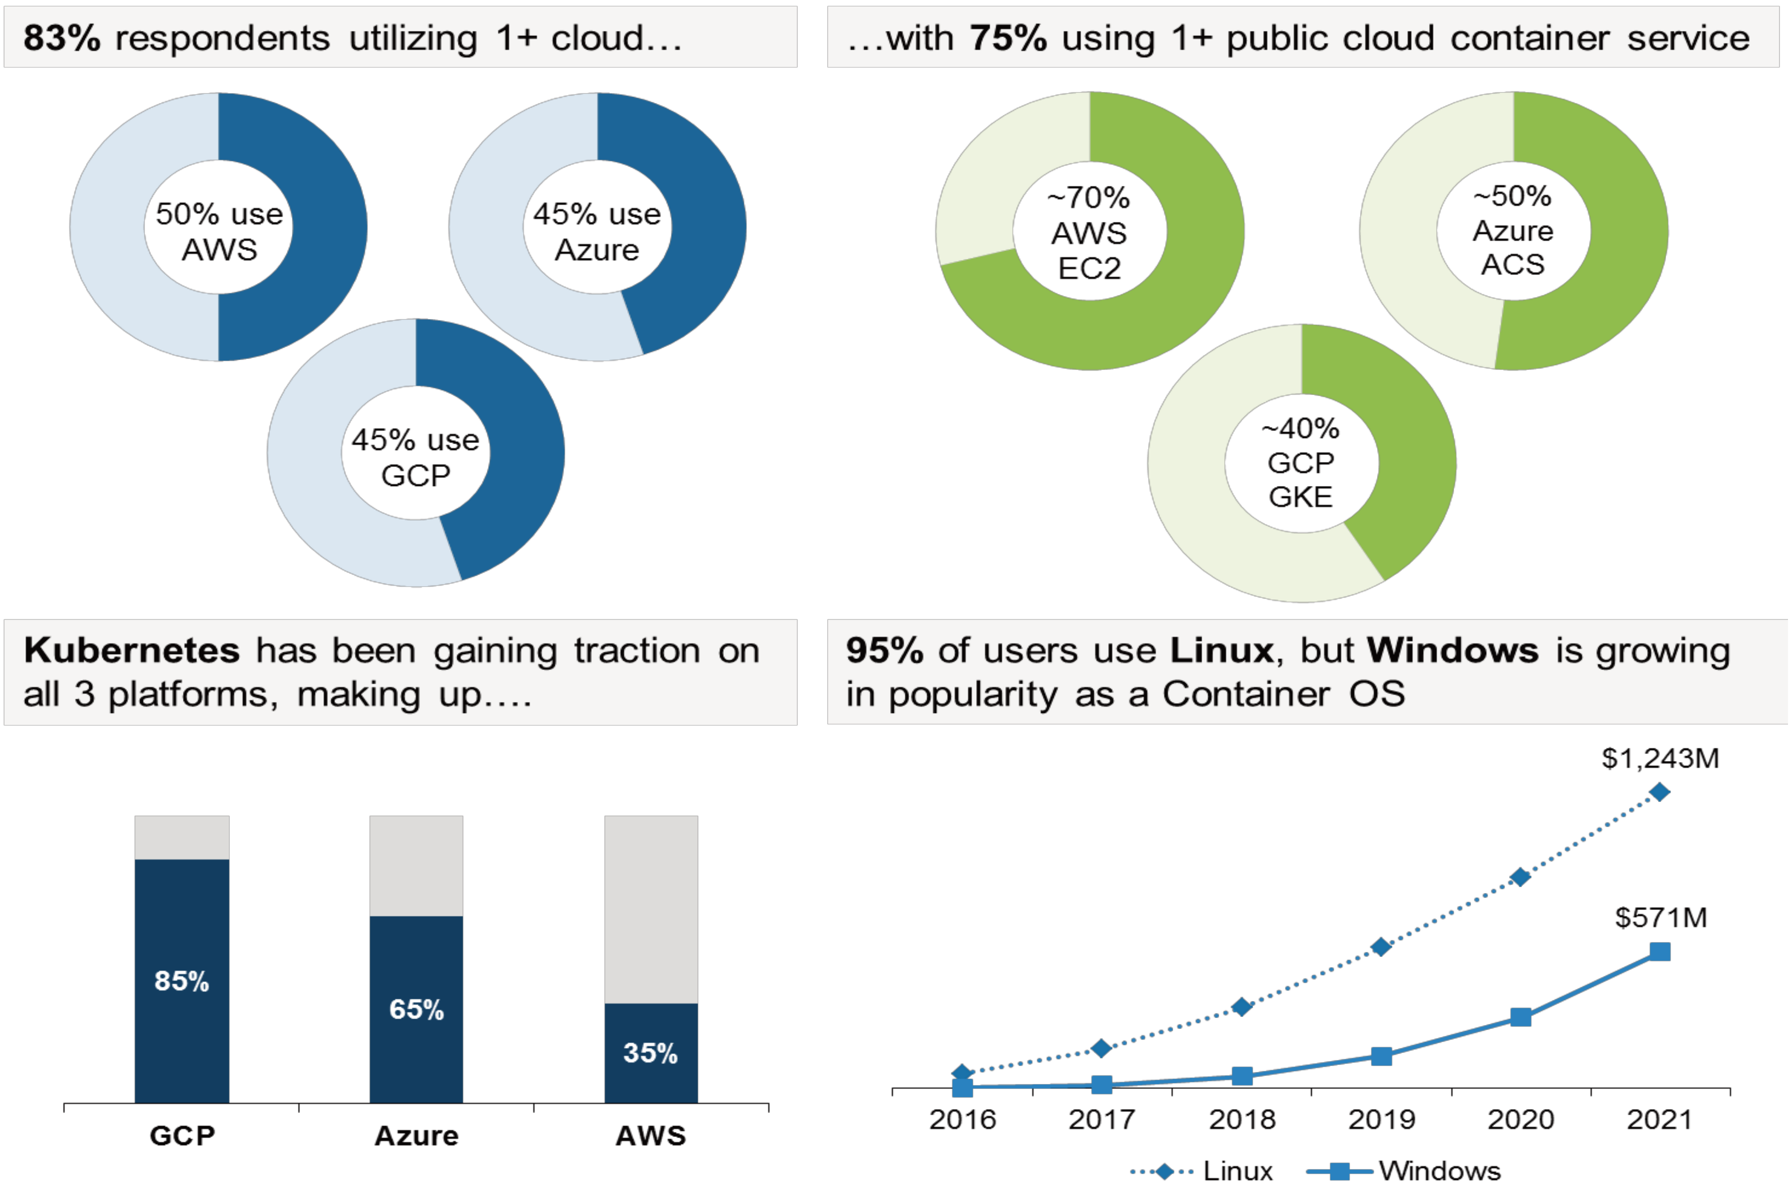

Source: IDC, Goldman Sachs Global Investment Research

Pivotal Software – Downgrade to Neutral

Containers and serverless computing are fairly nascent trends and we expect those workloads to be increasingly deployed in public cloud environments.

While we are seeing a general trend of workloads migrating to the public cloud, we expect the majority of workloads will still remain on-premise over the next 3-5 years. Based on our December 2018 IT spending survey results, most CIOs still expect to have a majority of their workloads run on-premise by 2021 although a smaller percentage than they do today.

Exhibit 30: Although CIOs are expecting to move a larger percentage of workloads to public cloud platforms, the majority of applications are still expected to run on-premises in the near future
% of applications moved to public cloud platforms today vs. expected in three years

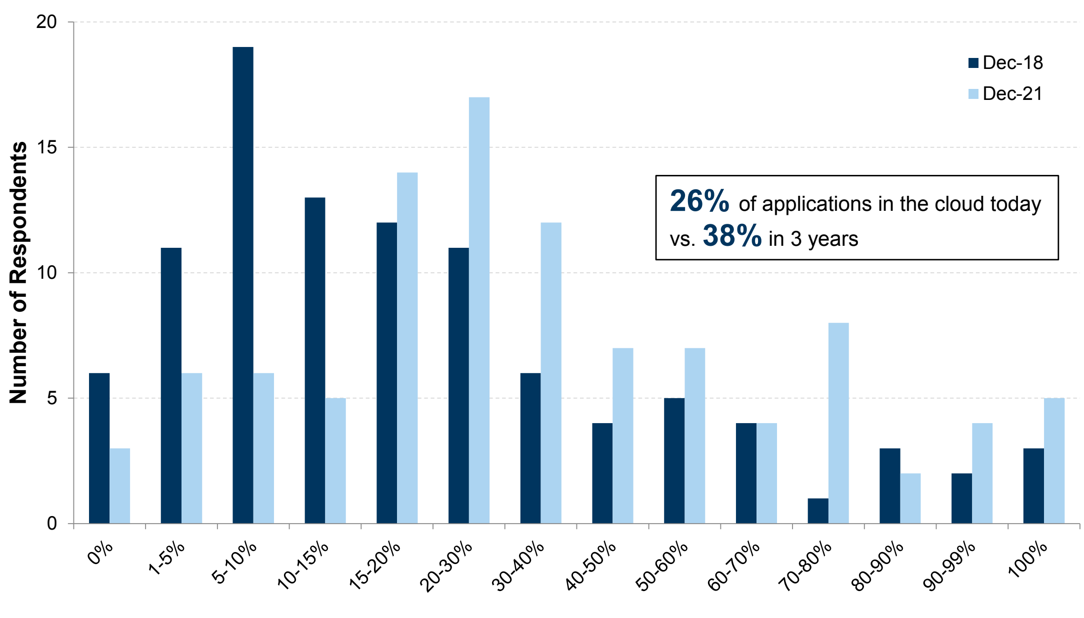

Source: Goldman Sachs Global Investment Research

While this could benefit multi-cloud platforms like Pivotal Container Service, which offers the ability to run containers in both on-prem and in a public cloud environment, the majority of containers and serverless workloads appear to being deployed more extensively in public cloud environments. Deploying containers or serverless functions in an on-premise environment has the potential to diminish the characteristics core to their purpose. Customers would have to leverage their own developer capabilities to orchestrate and manage their containers in their on-prem environments and not be able to leverage cloud native orchestration and management services. Likewise, the elasticity necessary for serverless functions would be limited in an on-premise, resource-constrained environment. Users would have to provision more virtual machines in order to scale those serverless applications, increasing infrastructure costs for the user which would have otherwise been reduced or eliminated in a public cloud environment. Furthermore, core to AWS Lambda's value is its data integration feature enabled by Kinesis. By connecting Kinesis Data Streams to AWS Lambda, users can collect and process data from a variety of streams of data sources. Conversely, on-premise environments do not have the same level of data integration as it would be difficult to connect multiple sources of data extending from different locations.

Exhibit 31: Roughly 60% of container workloads are expected to be in public cloud environments, growing at a 65% '18-'21 CAGR

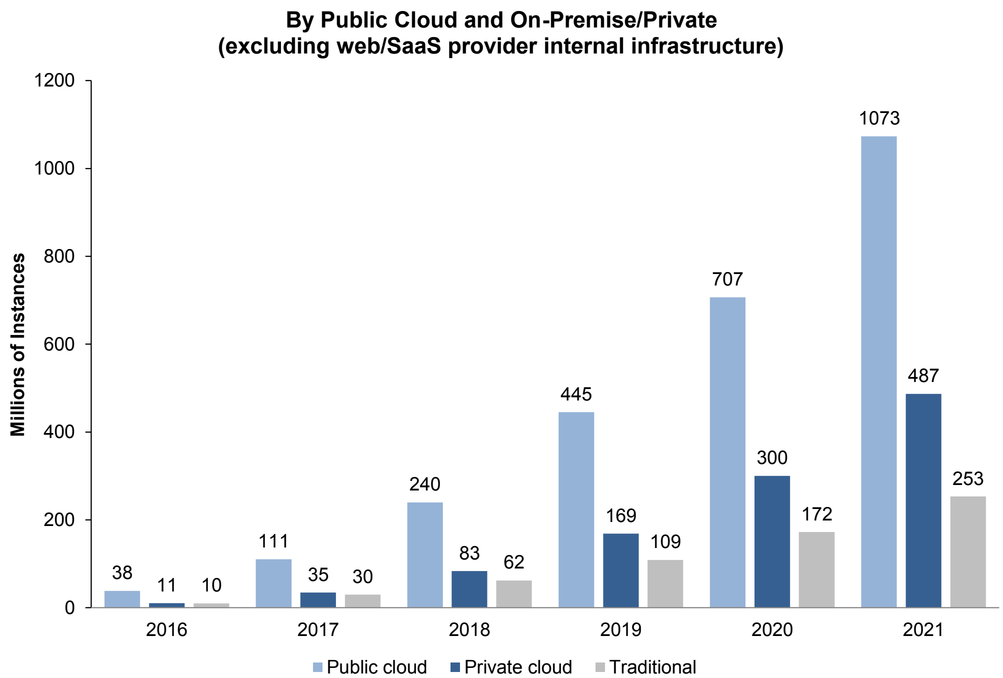

Source: IDC

Traction in Pivotal's products is still in its early days.

Pivotal Container Service (PKS), made GA in Feb 2018, leverages services and technologies from Pivotal Cloud Foundry and Kubernetes to run containers. Pivotal has revenue–sharing agreements with Dell/EMC and VMware to sell PKS, leveraging their existing enterprise customer base, which recently started seeing some early traction in the customer pipeline. Likewise, management has expressed optimism in Pivotal Function Service (PFS), Pivotal's FaaS platform, although it was expected to be made available at some point in 2018, but is still currently not GA. Given the clear market leadership and the growing number of features on AWS Lambda, it is unclear if PFS can serve as a viable alternative option. Our December CIO IT spending survey results suggest that interest in Pivotal services such as PKS is low and expected to remain that way over the next few years. Currently, 4% of respondents indicated they utilize Pivotal for container management, with a proportion of respondents expecting a decline in the next three years as more users opt to utilize public cloud container services.

Exhibit 32: More respondents indicated they utilize public cloud vendors for container management; a percentage that is expected to increase as services like Pivotal declines
% of respondents

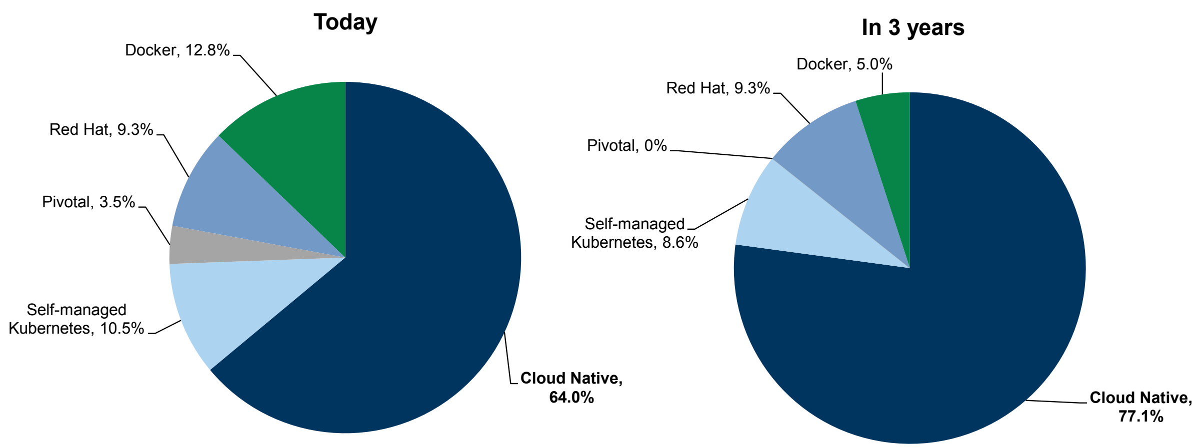

Source: Goldman Sachs Global Investment Research

Volatility of billings

As of F3Q19, Pivotal had 368 subscription customers with ARR over $50k. Combined with its relatively small number of net new customer additions in any given quarter (an average of 16 per quarter in FY19), these factors contribute to volatility around its subscription billings growth. We estimate subscription billings growth by taking subscription revenue + the change in short-term deferred revenue, adjusted for our estimate that ~10% of their short-term deferred revenue balance is related to professional services. In the exhibit below, the volatility in the growth from quarter to quarter can be quite significant. If we try and normalize for this and instead look at short-term subscription billings growth on a rolling 12 month basis we see that the growth smooths out, but the trend is still inconsistent.

Last quarter, management noted that they expect RPO to peak in Q4 due to seasonality. As they continue to scale, they expect RPO growth rates yoy to decline compared to the prior period.

Exhibit 33: Short-term Subscription Billings Growth (adjusted for professional services)
$ in mn

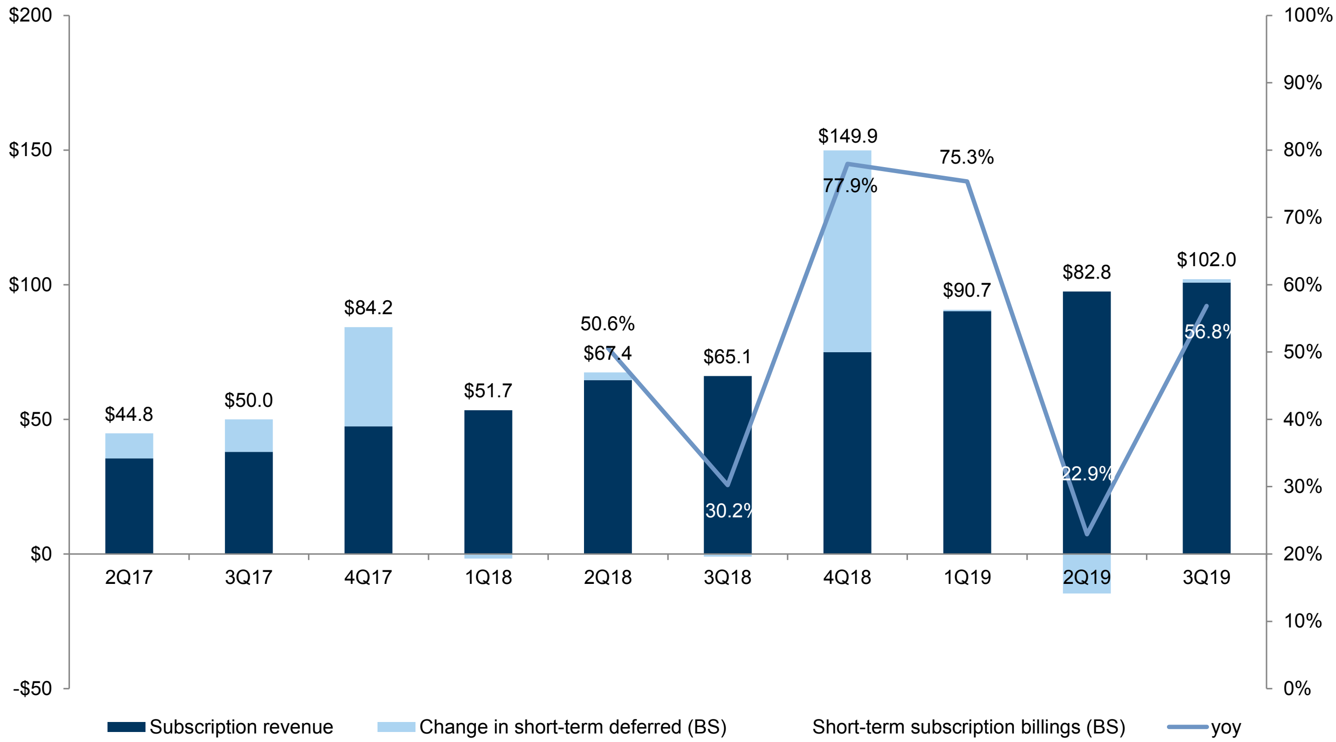

We adjust for professional services by assuming 90% of short term deferred revenue is subscription.

Source: Company data, Goldman Sachs Global Investment Research

While we expect new customers with ARR greater than $50k to start to improve as the sales force focuses on both new and renewals, we believe the volatility from quarter to quarter will remain for some time. Their customer count grew 16% yoy in FY18 and we expect it to grow 23% and 21% yoy in FY19 and FY20 off of a small base.

Exhibit 34: Short term deferred revenue trends (FY17 through F3Q19)
$ in mn

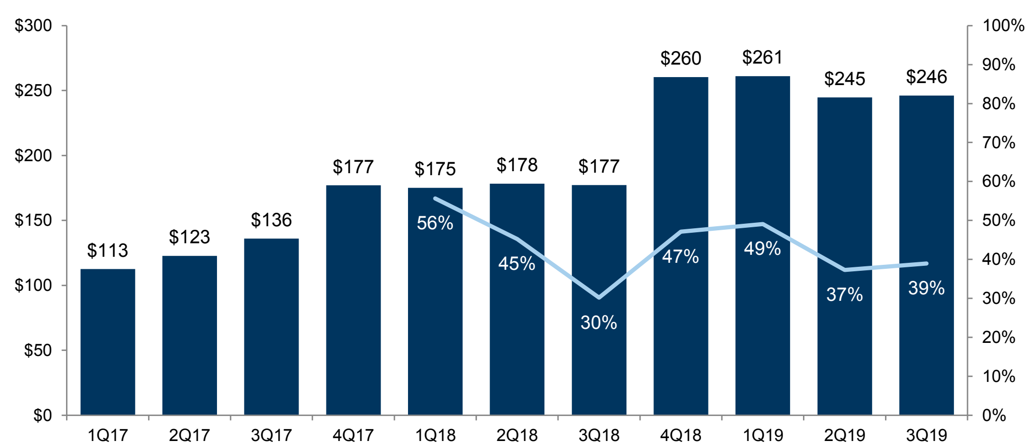

Source: Company data, Goldman Sachs Global Investment Research

Valuation & Key risks

We downgrade PVTL from Buy to Neutral and lower our 12-month price target to $24 from $25 as we lower our forecasts based on confidence in out year estimates. At ~$22, PVTL is trading at an EV of ~7x CY19 and ~6x CY20 our sales estimates of $827mn and $1020mn, compared to consensus of $813mn and $1,000mn respectively, and at an EV of ~6.5x CY19 and ~5x CY20 our billings estimates of $921mn and $1,136mn, compared to consensus of $913mn and $1,123mn respectively. Please refer to the Exhibit 35 below for our estimate revisions. Our price target is derived from an equal-weighting of our EV/Sales (8.5x CY19 sales, unchanged) and a DCF (~5% perpetuity growth rate, unchanged) analyses. Since we initiated on PVTL with a Buy rating on May 15, 2018, the stock is up 13.5% versus the S&P 500 up 3%.

Exhibit 35: PVTL estimate revisions

$ in mn except per share items

All figures in $ mns	1Q20 (E) Guidance	1Q20 (E) Prior Est.	1Q20 (E) New Est.	1Q20 (E) Δ	1Q20 (E) Consensus	FY19 (E) Prior Est.	FY19 (E) New Est.	FY19 (E) Δ	FY19 (E) Consensus
Subscription revenue	$110 - $111	$117.9	$103.6	($14.2)	$120.1	$398.6	$398.2	($0.3)	$398.4
YoY	*22% - 23%*	*31%*	*15%*		*33%*	*54%*	*54%*		*54%*
QoQ	*9% - 10%*	*7%*	*(6%)*		--				
Services revenue	$60 - $61	$74.8	$74.8	$0.0	$69.2	$259.9	$259.9	$0.0	$260.0
YoY	*(9%) - (8%)*	*14%*	*14%*		*5%*	*4%*	*4%*		*4%*
QoQ	*(12%) - (10%)*	*25%*	*25%*		--				
Total revenue	**$169.0 - $171.0**	**$192.7**	**$178.4**	**($14.2)**	**$191.7**	**$658.5**	**$658.1**	**($0.3)**	**$658.4**
YoY	*9% - 10%*	*24%*	*15%*		*23%*	*29%*	*29%*		*29%*
QoQ	*1% - 2%*	*13%*	*5%*		--				
Gross profit (non-GAAP)		$127.2	$114.1	($13.1)	$125.8	$431.2	$430.9	($0.3)	$430.7
YoY		*28%*	*15%*		*27%*	*47%*	*47%*		*47%*
Gross margin		66.0%	64.0%	(2.1%)	65.6%	65.5%	65.5%	(0.0%)	65.4%
Operating expenses (non-GAAP)		$145.8	$145.8	$0.0	$143.5	$507.1	$507.1	$0.0	$506.5
YoY		*21%*	*21%*		*19%*	*20%*	*20%*		*20%*
QoQ		*7%*	*7%*		--				
Operating income (non-GAAP)	**($26.0) - ($25.0)**	**($18.6)**	**($31.7)**	**($13.1)**	**($17.7)**	**($75.9)**	**($76.2)**	**($0.3)**	**($75.8)**
YoY	*24% - 19%*	*(12%)*	*50%*		*(16%)*	*(41%)*	*(41%)*		*(41%)*
QoQ	*75% - 69%*	*(27%)*	*23%*		--				
Operating margin	*(15.4%) - (14.6%)*	*(9.6%)*	*(17.7%)*	*-809 bps*	*(9.2%)*	*(11.5%)*	*(11.6%)*	*-5 bps*	*(11.5%)*
EPS (non-GAAP)	**($0.10) - ($0.09)**	**($0.07)**	**($0.12)**	**($0.05)**	**($0.06)**	**($0.31)**	**($0.31)**	**$0.00**	**($0.30)**
Net expansion rate		**142%**	**135%**	**(7%)**	**NA**	**148%**	**146%**	**(2%)**	**NA**
New customer additions		**26**	**26**	**-**	**NA**	**74**	**74**	**-**	**NA**
Billings (via CF)		$177	$163	($14)	$191	$746	$746	($0)	$745
YoY		*1%*	*(7%)*		*8%*	*28%*	*28%*		*28%*
QoQ		*(34%)*	*(39%)*		--				
Short-term billings		$176	$161	($14)	$185	$740	$740	($0)	$740
YoY		*12%*	*3%*		*18%*	*25%*	*25%*		*25%*
QoQ		*(34%)*	*(39%)*		--				
ST subscription billings		$101	$87	($14)	$113	$480	$480	($0)	$480
YoY		*11%*	*(5%)*		*25%*	*40%*	*40%*		*40%*
QoQ		*(51%)*	*(58%)*		--				
ST deferred revenue		**$325**	**$325**	**$0**	**$335.3**	**$342**	**$342**	**$0**	**$342.0**
YoY		*25%*	*25%*		*28%*	*31%*	*31%*		*31%*
QoQ		*(5%)*	*(5%)*		--				
LT deferred revenue		**$65**	**$65**	**$0**	**$60.8**	**$63**	**$63**	**$0**	**$64.0**
YoY		*(16%)*	*(16%)*		*(21%)*	*10%*	*10%*		*12%*
QoQ		*3%*	*3%*		--				
Deferred revenue (total)		**$390**	**$390**	**$0**	**$405**	**$405**	**$405**	**$0**	**$403**
YoY		*15%*	*15%*		*20%*	*28%*	*28%*		*27%*
QoQ		*(4%)*	*(4%)*		--				
OCF		**($8)**	**($3)**	**$4**	**$20**	**($40)**	**($40)**	**$0**	**($37)**
YoY		*(272%)*	*(178%)*		*343%*	*(66%)*	*(66%)*		*(68%)*
FCF		**($12)**	**($8)**	**$4**	**$16**	**($52)**	**($51)**	**$0**	**($47)**
YoY		*(578%)*	*(404%)*		*522%*	*(60%)*	*(60%)*		*(63%)*

All figures in $ mns	FY19 (E) Guidance	FY20 (E) Prior Est.	FY20 (E) New Est.	FY20 (E) Δ	FY20 (E) Consensus	FY21 (E) Prior Est.	FY21 (E) New Est.	FY21 (E) Δ	FY21 (E) Consensus
Subscription revenue	$398 - $399	$563.7	$552.2	($11.5)	$532.1	$779.0	$734.8	($44.2)	$694.3
YoY	*54% - 54%*	*41%*	*39%*		*34%*	*38%*	*33%*		*30%*
QoQ									
Services revenue	$259 - $260	$275.0	$275.0	$0.0	$274.4	$285.3	$285.3	$0.0	$290.9
YoY	*3% - 4%*	*6%*	*6%*		*6%*	*4%*	*4%*		*6%*
QoQ									
Total revenue	**$657 - $659**	**$838.7**	**$827.2**	**($11.5)**	**$812.8**	**$1,064.3**	**$1,020.1**	**($44.2)**	**$1,000.1**
YoY	*29% - 29%*	*27%*	*26%*		*23%*	*27%*	*23%*		*23%*
QoQ									
Gross profit (non-GAAP)		$587.4	$576.8	($10.5)	$548.1	$788.6	$747.7	($40.9)	$708.9
YoY		*36%*	*34%*		*27%*	*34%*	*30%*		*29%*
Gross margin		70.0%	69.7%	(0.3%)	67.4%	74.1%	73.3%	(0.8%)	70.9%
Operating expenses (non-GAAP)		$604.7	$604.7	$0.0	$599.4	$700.9	$700.9	$0.0	$687.6
YoY		*19%*	*19%*		*18%*	*16%*	*16%*		*15%*
QoQ									
Operating income (non-GAAP)	**($77) - ($76)**	**($17.3)**	**($27.9)**	**($10.5)**	**($51.3)**	**$87.6**	**$46.7**	**($40.9)**	**$21.4**
YoY	*(40%) - (41%)*	*(77%)*	*(63%)*		*(32%)*	*(606%)*	*(268%)*		*(142%)*
QoQ									
Operating margin	*(11.6%) - (11.5%)*	*(2.1%)*	*(3.4%)*	*-130 bps*	*(6.3%)*	*8.2%*	*4.6%*	*-365 bps*	*2.1%*
EPS (non-GAAP)	**($0.32) - ($0.31)**	**($0.07)**	**($0.10)**	**($0.03)**	**($0.17)**	**$0.23**	**$0.14**	**($0.10)**	**$0.09**
Net expansion rate		**138%**	**135%**	**(3%)**	**NA**	**136%**	**130%**	**(6%)**	**NA**
New customer additions		**83**	**83**	**-**	**NA**	**84**	**97**	**+12**	**NA**
Billings (via CF)		$933	$921	($12)	$913	$1,180	$1,136	($44)	$1,123
YoY		*25%*	*24%*		*22%*	*27%*	*23%*		*23%*
QoQ									
Short-term billings		$932	$920	($12)	$917	$1,177	$1,133	($44)	$1,154
YoY		*26%*	*24%*		*24%*	*26%*	*23%*		*26%*
QoQ									
ST subscription billings		$657	$645	($12)	$636	$892	$848	($44)	$848
YoY		*37%*	*34%*		*33%*	*36%*	*31%*		*33%*
QoQ									
ST deferred revenue		**$435**	**$435**	**$0**	**$446.3**	**$548**	**$548**	**$0**	**$600.1**
YoY		*27%*	*27%*		*30%*	*26%*	*26%*		*34%*
QoQ									
LT deferred revenue		**$64**	**$64**	**$0**	**$58.4**	**$66**	**$66**	**$0**	**$55.3**
YoY		*1%*	*1%*		*(9%)*	*4%*	*4%*		*(5%)*
QoQ									
Deferred revenue (total)		**$499**	**$499**	**$0**	**$502**	**$615**	**$615**	**$0**	**$630**
YoY		*23%*	*23%*		*24%*	*23%*	*23%*		*26%*
QoQ									
OCF		**$4**	**($8)**	**($13)**	**($15)**	**$49**	**$33**	**($16)**	**$56**
YoY		*(111%)*	*(79%)*		*(59%)*	*1003%*	*(495%)*		*(464%)*
FCF		**($15)**	**($28)**	**($13)**	**($28)**	**$25**	**$10**	**($15)**	**$47**
YoY		*(70%)*	*(45%)*		*(41%)*	*(261%)*	*(134%)*		*(266%)*

Source: Company data, Goldman Sachs Global Investment Research, FactSet

VMware - Downgrade to Sell

We downgrade VMware from Neutral to Sell and maintain our 12-month price target of $177. While we see the company as having done an excellent job diversifying its revenue base away from core virtualization, we see the benefit from the resurgence in on-premise spending, which has helped the growth in ELA renewal signings, as starting to face more normalized comps. Although we are optimistic on the revenue potential from partnerships with platforms such as AWS, we see the current multiple as factoring in material upside to revenue and non-GAAP EPS in CY19 and CY20. At ~$173, VMW currently trades at 27x CY19 and 24x CY20 our non-GAAP EPS estimates of $6.49 and $7.31, compared to consensus of $6.58 and $7.29 respectively. This compares to MSFT which at ~$112, trades at 24x and 20X our GAAP EPS forecasts of $4.63 and $5.70.

As shown in Exhibit 36, Gartner projects a deceleration in the worldwide x86 server virtualization infrastructure market, and given VMware's dominant market position in this market, we would expect a deceleration to have a larger impact on VMware. Likewise, our IT Spending Survey results suggest that the majority of companies are expecting to either maintain or decrease spending with vendors offering virtualization services in the future. Although virtual machines will still be used with containers in the near to medium term, we believe that it is possible that companies can provision fewer virtual machines to run the same amount of workloads as they increasingly leverage containers to run application instances. Additionally, longer-term, if developers begin transitioning their containerized workloads to run on bare metal versus in a virtualized environment, this could further impact server virtualization growth, in our view. As shown in our container market sizing analysis, we believe it is possible that $1.1bn of legacy infrastructure spending (spending on VMs and bare metal servers) could be shifted to containerizing applications.

Exhibit 36: VMware could face headwinds to growth as the server virtualization market continues to stagnate
Worldwide x86 server virtualization infrastructure market size 2016-2022

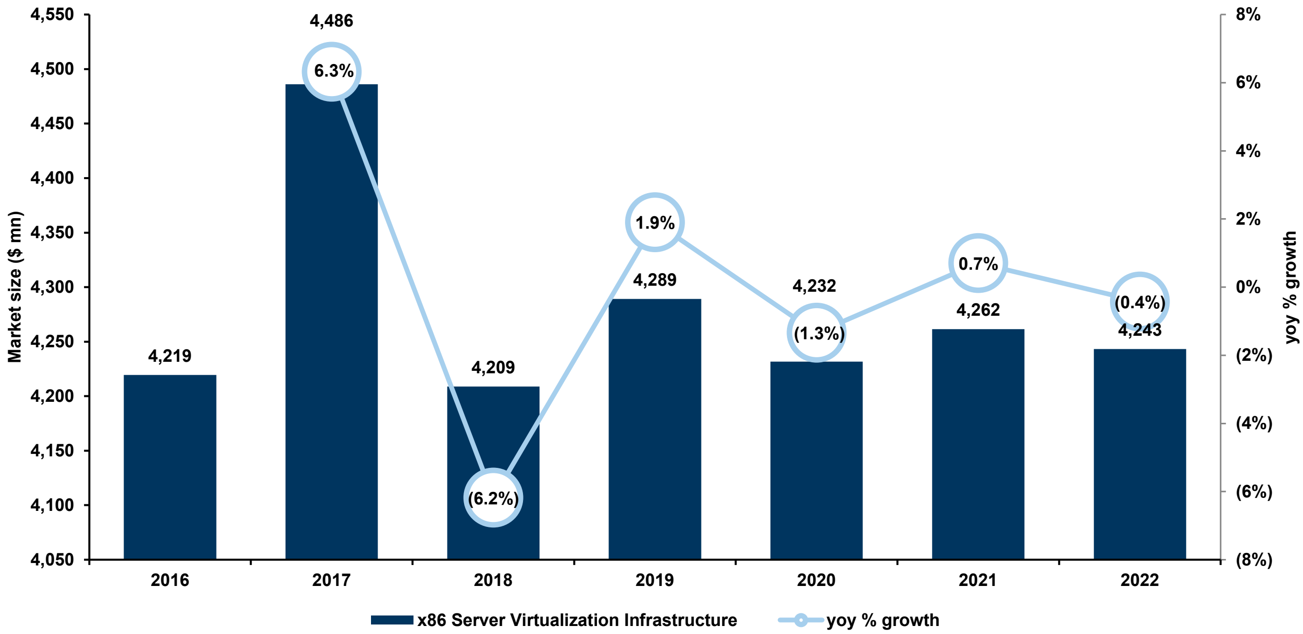

Source: Gartner, Goldman Sachs Global Investment Research

Tailwind from resurgence in on-premise spending starting to slow

Recall that VMware's largest customers typically sign enterprise license agreements (ELAs) which tend to be three year contracts. As such, we evaluate the pool of ELAs that were signed three years ago that will likely be up for renewal. We note, however, that VMware's fundamentals continue to diverge from its ELA cycles over time, driven by 1) distance from the distinctive 2009/2010 ELA cycle and 2) a diversifying product portfolio that shifts focus away from renewals. 2009 ELA renewals were underwhelming and 2010 renewals were strong, and with three-year contracts, historically, renewals on the 2009 cycle (2012, 2015) have been weak while renewals on the 2010 cycle (2013, 2016) have been strong. However, as time progresses, VMware generally sees a "reversion to the mean" effect, where sharp historical oscillations are dampened – although three-year contracts are the standard, in any given period, VMware sees some renewals pulled forward and some pushed out, and with the 2009/2010 ELA cycle now nearly a decade in the rearview mirror, this has smoothed out the ELA cycles. Furthermore, while VMware's ELAs have historically been compute-focused, the company's increasingly diverse product portfolio (i.e. compute license bookings have been below 50% of total license bookings for several years at this point), including NSX, vSAN, and hybrid cloud offerings have resulted in the company cross-selling to its customer base, further reducing the dependency and significance of historical ELA patterns.

Although the significance of VMware's ELA cycle is declining, we continue to monitor it, given that we estimate that ~50% of overall VMware bookings are still standalone vSphere, though this proportion has generally declined over time. Looking ahead, we estimate that the ELA renewal opportunity in CY19 is 13% greater than CY18's opportunity, as we estimate that CY16 ELAs (for CY19 renewal) were $2,828mn vs. CY15 ELAs (renewed in CY18), which were $2,504mn.

VMW has had very strong ELA renewal cycles thus far, with ELAs driving a larger portion of billings each year, increasing from ~41% in CY17 to ~45% of total billings in CY18. However, we expect that to become tougher as the comp becomes more normalized. For example, we saw an acceleration in billings from ELAs starting in CY17 (up 27% yoy vs. 13% in CY16) and have seen continued strong growth through CY18 (26% yoy) and expect growth to decelerate to 17% in CY19 and 13% in CY20.

In our view, the sharp increase in billings from ELAs in CY17 was driven by corporations recognizing that the migration of production workloads to public cloud was going to take longer than anticipated and as such on-premise spending received a considerable boost as CIOs found that they had underestimated the required on-premise capacity they would need and as a result, underspent with on-premise infrastructure software vendors. As shown in our biannual IT spending survey, expectations in terms of the pace to the shift of workloads to public cloud became more tempered in CY17 and reversed its trend of acceleration.

Exhibit 37: IT Spending Survey - Public cloud adoption today vs. expected in 3 years
% of workloads

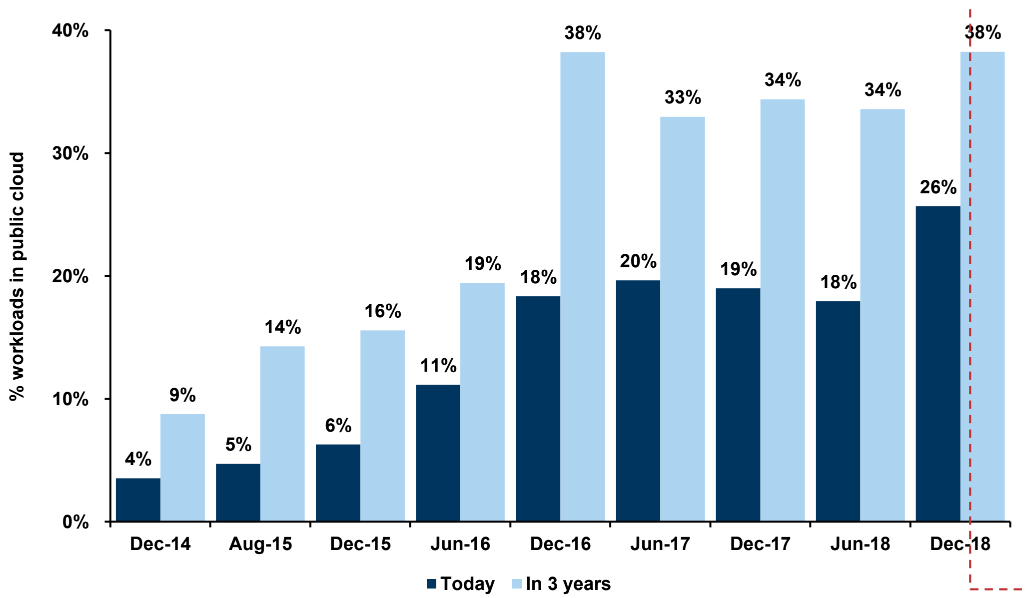

Source: Goldman Sachs Global Investment Research

As a strategic on-premise vendor, we believe that VMware benefitted from this catch-up spending in infrastructure, starting in CY17. While CY18 likely continued to benefit from this dynamic, given that VMware generally has three-year contracts, this benefit starts to become a more normalized comp in CY19, with the uptick in spending becoming more normalized in CY20 (as it will be the elevated CY17 pool up for renewal once again).

Exhibit 38: ELA billings and renewal opportunity trends
$ in mn

Fiscal year Calendar year	FY11 (A) CY11	FY12 (A) CY12	FY13 (A) CY13	FY14 (A) CY14	FY15 (A) CY15	FY16 (A) CY16	FY18 (A) CY17	FY19 (A) CY18	FY20 (E) CY19	FY21 (E) CY20
ELA renewal opportunity base year	*CY08*	*CY09*	*CY10*	*CY11*	*CY12*	*CY13*	*CY14*	*CY15*	*CY16*	*CY17*
ELA billings in year (GSe)	1,177	1,484	2,105	2,245	2,504	2,828	3,585	4,506	5,262	5,963
ELA billings renewal opportunity	441	433	724	1,177	1,484	2,105	2,245	2,504	2,828	3,585
ELA billings yoy increase		**26.1%**	**41.9%**	**6.6%**	**11.5%**	**13.0%**	**26.8%**	**25.7%**	**16.8%**	**13.3%**
ELAs % billings	25.5%	27.7%	35.3%	33.4%	36.3%	37.0%	41.0%	44.6%	46.8%	48.3%
ELA renewal opportunity yoy increase		(1.9%)	67.1%	62.6%	26.1%	41.9%	6.6%	11.5%	13.0%	26.8%

Source: Company data, Goldman Sachs Global Investment Research

Broadening its cloud footprint but revenue opportunities likely take time to build

As outlined in our recent ratings reinstatement note (link), VMware has successfully transitioned to a more diversified revenue base, with non-compute business segments such as infrastructure management, end user computing, and software-defined networking & security driving a larger portion of license billings. Outside of vSphere, management is likely the top contributor to revenue, followed by EUC. Our partner checks have also noted increased traction for VMware's NSX business, driven by microsegmentation use cases, with management noting on the F4Q19 earnings call that bookings for NSX for the quarter ended the year at $1.3bn. The company's hybrid cloud initiatives and HCI storage software (vSAN) business are also starting to emerge as larger opportunities as well. However, we estimate that virtualization is still the main driver of ELA renewals based on our estimate that their standalone compute (vSphere) business is 25% of total license billings and compute continues to make up ~50% of total billings.

VMware's partnerships with AWS and Pivotal (to a lesser extent) may help mitigate potential revenue impacts from the anticipated decrease in demand for virtual machines. We are also positive about the partnership with AWS and the potential for deeper engagement with Azure, as reported by the Information, although we expect revenue contributions to be relatively small in 2019, as management is focused in the near term on customer adoption, and we see limited multiple upside. Moreover, VMware acquired Heptio, which enhances support for the Kubernetes container management system and has the potential to help drive on-premise container revenues to their service. However, as companies move workloads to public cloud platforms, they may not find VMware's public cloud offerings as attractive. Users already integrated with public cloud platforms can leverage the proprietary container services without having to provision VMware-specific virtual machines. Regarding Pivotal, VMware's 50/50 revenue split from PKS revenues places VMware in a position to benefit from container adoption, although uptake of this offering could also be impacted by the desire to run containers on that of native cloud platform providers.

VMware Cloud (VMC) on AWS is now available across multiple regions in US, Europe and APAC (Sydney and Tokyo), and more recently in AWS Gov Cloud as of F4Q19, after being initially available in the US West region in F3Q18. Management commented in F3Q18 that they don't expect it to be material in 2017 or 2018 and will continue to focus on the net billing aspects of VMC and customer growth. Most recently, during the F4Q19 call, management noted that they continue to see momentum in VMC on AWS across all three geos and closed their largest deal so far at $20mn during the quarter. As such, while we see the company as benefitting from the partnership with AWS and the potential for a deeper relationship with Microsoft, as reported by the Information, we see these as having only a limited impact on financials this coming FY.

What would make us more positive on the stock?

We would be more positive on VMW if 1) the company introduces more integrations and services tied to VMC on AWS, driving more customer adoption and expansion of current deal sizes, 2) the sales force is more ramped on their newer products and services, such as PKS Essentials (Heptio) and CloudHealth, and able to drive more uptake in their hybrid cloud and SaaS products, 3) gain leverage faster than expected in operating margins 4) drive more and larger deals with their non-standalone vSphere products such as NSX and EUC.

Valuation

Given how well the stock has performed relative to the comp group over the past few years and at its current valuation level, we are downgrading VMW from Neutral to Sell as we prefer other companies in our coverage and we view its current valuation as largely reflecting its diversified revenue base and potential upside from their partnerships and newer offerings. Our 12-month price target of $177 is based on equal weight of our P/E (28x CY19 EPS), EV/FCF (20x CY19 FCF), and a DCF analysis (~1% perpetuity growth rate). Our price target has ~2.5% upside vs. ~5.0% upside in our coverage universe.

Investment Risks

Pivotal Software (PVTL)

Downside:

- **Unique ownership structure.** Dell Technologies is Pivotal's majority shareholder and controls a majority of the combined voting power of both of Pivotal's classes of common stock.
- **High mix of services.** Pivotal's revenue mix has historically been heavily weighted towards services. In FY16, Pivotal's business was two-thirds services, moving to an even split between its subscription software business and services business in FY18 and towards 40% in FY19. Although the mix is expected to move increasingly towards subscription software, the current 40% mix of services is still substantially higher than many of Pivotal's peers, which typically have just 10-15% of revenue as services.
- **Competitive landscape**. When aggregating Pivotal's TAM of application infrastructure and middleware, application development, and PaaS, the marketplace

includes several scaled competitors, including IBM, Oracle, Microsoft, and Salesforce.com, as well as large public cloud incumbents (i.e. AWS, Microsoft Azure, and Google Cloud Platform)

Upside:

- **Faster than expected customer additions.** Ending F3Q19, Pivotal had 368 customers paying over $50k in ARR. Pivotal could accelerate their customer additions and adoption of their products.
- **IT spending.** Our recent IT spending survey showed that overall CIO spending intentions ticked down meaningfully in December 2018 relative to June 2018, suggesting deceleration of IT spending expectations for CY19 vs. CY18. If IT spending intentions accelerate, this would lead to more spending on PaaS vendors like PVTL.
- **PKS.** The VMware sales force is just starting to spin up on the product and PVTL is starting to see early traction in the customer pipeline. As the sales force continues to ramp, this could drive more customer adoption of PKS and benefit PVTL's revenue.

VMware (VMW)

- **IT spending.** While our recent IT spending survey suggested a deceleration of IT spending expectations for CY19 vs. CY18, and continued decline in spending in virtualization software, an uptick would benefit VMW's business as most of their business is still driven by compute.
- **Growth in non-standalone vSphere segments.** VMW's non-standalone vSphere segments could accelerate as their newer products gain more traction and they gain more market share.
- **AWS partnership.** VMW's VMC on AWS product could see faster customer adoption and as a result, greater and/or earlier than expected monetization of their partnership with AWS.
- **Heptio & PKS adoption**. Customer adoption of PKS and PKS Essentials (Heptio) accelerates, mitigating the potential headwinds they face from increased container adoption and decreased VM spending.
- **Hybrid cloud & SaaS revenues**. As of F4Q19, hybrid cloud & SaaS revenues made up 10% of total revenues (50/50 split between maintenance and subscription revenue). An acceleration in those revenues would signal increased uptake in their cloud and SaaS-based product offerings.

Last close: Pivotal Software Inc. ($21.84) and VMware Inc. ($172.69)

Disclosure Appendix

Reg AC

I, Heather Bellini, CFA, hereby certify that all of the views expressed in this report accurately reflect my personal views about the subject company or companies and its or their securities. I also certify that no part of my compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed in this report.

Unless otherwise stated, the individuals listed on the cover page of this report are analysts in Goldman Sachs' Global Investment Research division.

GS Factor Profile

The Goldman Sachs Factor Profile provides investment context for a stock by comparing key attributes to the market (i.e. our coverage universe) and its sector peers. The four key attributes depicted are: Growth, Financial Returns, Multiple (e.g. valuation) and Integrated (a composite of Growth, Financial Returns and Multiple). Growth, Financial Returns and Multiple are calculated by using normalized ranks for specific metrics for each stock. The normalized ranks for the metrics are then averaged and converted into percentiles for the relevant attribute. The precise calculation of each metric may vary depending on the fiscal year, industry and region, but the standard approach is as follows:

Growth is based on a stock's forward-looking sales growth, EBITDA growth and EPS growth (for financial stocks, only EPS and sales growth), with a higher percentile indicating a higher growth company. **Financial Returns** is based on a stock's forward-looking ROE, ROCE and CROCI (for financial stocks, only ROE), with a higher percentile indicating a company with higher financial returns. **Multiple** is based on a stock's forward-looking P/E, P/B, price/dividend (P/D), EV/EBITDA, EV/FCF and EV/Debt Adjusted Cash Flow (DACF) (for financial stocks, only P/E, P/B and P/D), with a higher percentile indicating a stock trading at a higher multiple. The **Integrated** percentile is calculated as the average of the Growth percentile, Financial Returns percentile and (100% - Multiple percentile).

Financial Returns and Multiple use the Goldman Sachs analyst forecasts at the fiscal year-end at least three quarters in the future. Growth uses inputs for the fiscal year at least seven quarters in the future compared with the year at least three quarters in the future (on a per-share basis for all metrics).

For a more detailed description of how we calculate the GS Factor Profile, please contact your GS representative.

M&A Rank

Across our global coverage, we examine stocks using an M&A framework, considering both qualitative factors and quantitative factors (which may vary across sectors and regions) to incorporate the potential that certain companies could be acquired. We then assign a M&A rank as a means of scoring companies under our rated coverage from 1 to 3, with 1 representing high (30%-50%) probability of the company becoming an acquisition target, 2 representing medium (15%-30%) probability and 3 representing low (0%-15%) probability. For companies ranked 1 or 2, in line with our standard departmental guidelines we incorporate an M&A component into our target price. M&A rank of 3 is considered immaterial and therefore does not factor into our price target, and may or may not be discussed in research.

Quantum

Quantum is Goldman Sachs' proprietary database providing access to detailed financial statement histories, forecasts and ratios. It can be used for in-depth analysis of a single company, or to make comparisons between companies in different sectors and markets.

GS SUSTAIN

GS SUSTAIN is a global investment strategy focused on the generation of long-term alpha through identifying high quality industry leaders. The GS SUSTAIN 50 list includes leaders we believe to be well positioned to deliver long-term outperformance through superior returns on capital, sustainable competitive advantage and effective management of ESG risks vs. global industry peers. Candidates are selected largely on a combination of quantifiable analysis of these three aspects of corporate performance.

Disclosures

Coverage group(s) of stocks by primary analyst(s)

Heather Bellini, CFA: America-Software. Ted Lin: America-Software. Mark Grant: America-Software.

America-Software: Adobe Systems Inc., Akamai Technologies Inc., Alphabet Inc., Anaplan Inc., Atlassian Corp., Autodesk Inc., Box Inc., Citrix Systems Inc., Docusign Inc., Dropbox Inc., Elastic NV, Endurance International Group, Facebook Inc., GoDaddy.com Inc., Microsoft Corp., MongoDB Inc., Okta Inc., Oracle Corp., Pivotal Software Inc., Red Hat Inc., RingCentral, Salesforce.com Inc., SolarWinds Corp., Twilio, VMware Inc., Wix.com, Workday Inc..

Company-specific regulatory disclosures

Compendium report: please see disclosures at http://www.gs.com/research/hedge.html. Disclosures applicable to the companies included in this compendium can be found in the latest relevant published research

Distribution of ratings/investment banking relationships

Goldman Sachs Investment Research global Equity coverage universe

	Rating Distribution			Investment Banking Relationships		
	Buy	Hold	Sell	Buy	Hold	Sell
Global	35%	54%	11%	65%	58%	56%

As of January 1, 2019, Goldman Sachs Global Investment Research had investment ratings on 2,945 equity securities. Goldman Sachs assigns stocks as Buys and Sells on various regional Investment Lists; stocks not so assigned are deemed Neutral. Such assignments equate to Buy, Hold and Sell for the purposes of the above disclosure required by the FINRA Rules. See 'Ratings, Coverage groups and views and related definitions' below. The Investment Banking Relationships chart reflects the percentage of subject companies within each rating category for whom Goldman Sachs has provided investment banking services within the previous twelve months.

Price target and rating history chart(s)

Compendium report: please see disclosures at http://www.gs.com/research/hedge.html. Disclosures applicable to the companies included in this compendium can be found in the latest relevant published research

Regulatory disclosures

Disclosures required by United States laws and regulations

See company-specific regulatory disclosures above for any of the following disclosures required as to companies referred to in this report: manager or co-manager in a pending transaction; 1% or other ownership; compensation for certain services; types of client relationships; managed/co-managed public offerings in prior periods; directorships; for equity securities, market making and/or specialist role. Goldman Sachs trades or may trade as a principal in debt securities (or in related derivatives) of issuers discussed in this report.

The following are additional required disclosures: **Ownership and material conflicts of interest:** Goldman Sachs policy prohibits its analysts, professionals reporting to analysts and members of their households from owning securities of any company in the analyst's area of coverage. **Analyst compensation:** Analysts are paid in part based on the profitability of Goldman Sachs, which includes investment banking revenues. **Analyst as officer or director:** Goldman Sachs policy generally prohibits its analysts, persons reporting to analysts or members of their households from serving as an officer, director or advisor of any company in the analyst's area of coverage. **Non-U.S. Analysts:** Non-U.S. analysts may not be associated persons of Goldman Sachs & Co. LLC and therefore may not be subject to FINRA Rule 2241 or FINRA Rule 2242 restrictions on communications with subject company, public appearances and trading securities held by the analysts.

Distribution of ratings: See the distribution of ratings disclosure above. **Price chart:** See the price chart, with changes of ratings and price targets in prior periods, above, or, if electronic format or if with respect to multiple companies which are the subject of this report, on the Goldman Sachs website at http://www.gs.com/research/hedge.html.

Additional disclosures required under the laws and regulations of jurisdictions other than the United States

The following disclosures are those required by the jurisdiction indicated, except to the extent already made above pursuant to United States laws and regulations. **Australia:** Goldman Sachs Australia Pty Ltd and its affiliates are not authorized deposit-taking institutions (as that term is defined in the Banking Act 1959 (Cth)) in Australia and do not provide banking services, nor carry on a banking business, in Australia. This research, and any access to it, is intended only for "wholesale clients" within the meaning of the Australian Corporations Act, unless otherwise agreed by Goldman Sachs. In producing research reports, members of the Global Investment Research Division of Goldman Sachs Australia may attend site visits and other meetings hosted by the companies and other entities which are the subject of its research reports. In some instances the costs of such site visits or meetings may be met in part or in whole by the issuers concerned if Goldman Sachs Australia considers it is appropriate and reasonable in the specific circumstances relating to the site visit or meeting. To the extent that the contents of this document contains any financial product advice, it is general advice only and has been prepared by Goldman Sachs without taking into account a client's objectives, financial situation or needs. A client should, before acting on any such advice, consider the appropriateness of the advice having regard to the client's own objectives, financial situation and needs. A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests and a copy of Goldman Sachs' Australian Sell-Side Research Independence Policy Statement are available at: https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html. **Brazil:** Disclosure information in relation to CVM Instruction 598 is available at http://www.gs.com/worldwide/brazil/area/gir/index.html. Where applicable, the Brazil-registered analyst primarily responsible for the content of this research report, as defined in Article 20 of CVM Instruction 598, is the first author named at the beginning of this report, unless indicated otherwise at the end of the text. **Canada:** Goldman Sachs Canada Inc. is an affiliate of The Goldman Sachs Group Inc. and therefore is included in the company specific disclosures relating to Goldman Sachs (as defined above). Goldman Sachs Canada Inc. has approved of, and agreed to take responsibility for, this research report in Canada if and to the extent that Goldman Sachs Canada Inc. disseminates this research report to its clients. **Hong Kong:** Further information on the securities of covered companies referred to in this research may be obtained on request from Goldman Sachs (Asia) L.L.C. **India:** Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (India) Securities Private Limited, Research Analyst - SEBI Registration Number INH000001493, 951-A, Rational House, Appasaheb Marathe Marg, Prabhadevi, Mumbai 400 025, India, Corporate Identity Number U74140MH2006FTC160634, Phone +91 22 6616 9000, Fax +91 22 6616 9001. Goldman Sachs may beneficially own 1% or more of the securities (as such term is defined in clause 2 (h) the Indian Securities Contracts (Regulation) Act, 1956) of the subject company or companies referred to in this research report. **Japan:** See below. **Korea:** This research, and any access to it, is intended only for "professional investors" within the meaning of the Financial Services and Capital Markets Act, unless otherwise agreed by Goldman Sachs. Further information on the subject company or companies referred to in this research may be obtained from Goldman Sachs (Asia) L.L.C., Seoul Branch. **New Zealand:** Goldman Sachs New Zealand Limited and its affiliates are neither "registered banks" nor "deposit takers" (as defined in the Reserve Bank of New Zealand Act 1989) in New Zealand. This research, and any access to it, is intended for "wholesale clients" (as defined in the Financial Advisers Act 2008) unless otherwise agreed by Goldman Sachs. A copy of certain Goldman Sachs Australia and New Zealand disclosure of interests is available at: https://www.goldmansachs.com/disclosures/australia-new-zealand/index.html. **Russia:** Research reports distributed in the Russian Federation are not advertising as defined in the Russian legislation, but are information and analysis not having product promotion as their main purpose and do not provide appraisal within the meaning of the Russian legislation on appraisal activity. Research reports do not constitute a personalized investment recommendation as defined in Russian laws and regulations, are not addressed to a specific client, and are prepared without analyzing the financial circumstances, investment profiles or risk profiles of clients. Goldman Sachs assumes no responsibility for any investment decisions that may be taken by a client or any other person based on this research report. **Singapore:** Further information on the covered companies referred to in this research may be obtained from Goldman Sachs (Singapore) Pte. (Company Number: 198602165W). **Taiwan:** This material is for reference only and must not be reprinted without permission. Investors should carefully consider their own investment risk. Investment results are the responsibility of the individual investor. **United Kingdom:** Persons who would be categorized as retail clients in the United Kingdom, as such term is defined in the rules of the Financial Conduct Authority, should read this research in conjunction with prior Goldman Sachs research on the covered companies referred to herein and should refer to the risk warnings that have been sent to them by Goldman Sachs International. A copy of these risks warnings, and a glossary of certain financial terms used in this report, are available from Goldman Sachs International on request.

European Union: Disclosure information in relation to Article 6 (2) of the European Commission Delegated Regulation (EU) (2016/958) supplementing Regulation (EU) No 596/2014 of the European Parliament and of the Council with regard to regulatory technical standards for the technical arrangements for objective presentation of investment recommendations or other information recommending or suggesting an investment strategy and for disclosure of particular interests or indications of conflicts of interest is available at http://www.gs.com/disclosures/europeanpolicy.html which states the European Policy for Managing Conflicts of Interest in Connection with Investment Research.

Japan: Goldman Sachs Japan Co., Ltd. is a Financial Instrument Dealer registered with the Kanto Financial Bureau under registration number Kinsho 69, and a member of Japan Securities Dealers Association, Financial Futures Association of Japan and Type II Financial Instruments Firms Association. Sales and purchase of equities are subject to commission pre-determined with clients plus consumption tax. See company-specific disclosures as to any applicable disclosures required by Japanese stock exchanges, the Japanese Securities Dealers Association or the Japanese Securities Finance Company.

Ratings, coverage groups and views and related definitions

Buy (B), Neutral (N), Sell (S) -Analysts recommend stocks as Buys or Sells for inclusion on various regional Investment Lists. Being assigned a Buy or Sell on an Investment List is determined by a stock's total return potential relative to its coverage. Any stock not assigned as a Buy or a Sell on an Investment List with an active rating (i.e., a stock that is not Rating Suspended, Not Rated, Coverage Suspended or Not Covered), is deemed Neutral. Each regional Investment Review Committee manages various regional Investment Lists to a global guideline of 25%-35% of stocks as Buy and 10%-15% of stocks as Sell; however, the distribution of Buys and Sells in any particular analyst's coverage group may vary as determined by the regional Investment Review Committee. Additionally, each Investment Review Committee manages Regional Conviction lists, which represent investment recommendations focused on the size of the total return potential and/or the likelihood of the realization of the return across their respective areas of coverage. The addition or removal of stocks from such Conviction lists do not represent a change in the analysts' investment rating for such stocks.

Total return potential represents the upside or downside differential between the current share price and the price target, including all paid or anticipated dividends, expected during the time horizon associated with the price target. Price targets are required for all covered stocks. The total return potential, price target and associated time horizon are stated in each report adding or reiterating an Investment List membership.

Coverage groups and views: A list of all stocks in each coverage group is available by primary analyst, stock and coverage group at http://www.gs.com/research/hedge.html. The analyst assigns one of the following coverage views which represents the analyst's investment outlook on the coverage group relative to the group's historical fundamentals and/or valuation. **Attractive (A).** The investment outlook over the following 12 months is favorable relative to the coverage group's historical fundamentals and/or valuation. **Neutral (N).** The investment outlook over the following 12 months is neutral relative to the coverage group's historical fundamentals and/or valuation. **Cautious (C).** The investment outlook over the following 12 months is unfavorable relative to the coverage group's historical fundamentals and/or valuation.

Not Rated (NR). The investment rating and target price have been removed pursuant to Goldman Sachs policy when Goldman Sachs is acting in an advisory capacity in a merger or strategic transaction involving this company and in certain other circumstances. **Rating Suspended (RS).** Goldman Sachs Research has suspended the investment rating and price target for this stock, because there is not a sufficient fundamental basis for determining, or there are legal, regulatory or policy constraints around publishing, an investment rating or target. The previous investment rating and price target, if any, are no longer in effect for this stock and should not be relied upon. **Coverage Suspended (CS).** Goldman Sachs has suspended coverage of this company. **Not Covered (NC).** Goldman Sachs does not cover this company. **Not Available or Not Applicable (NA).** The information is not available for display or is not applicable. **Not Meaningful (NM).** The information is not meaningful and is therefore excluded.

Global product; distributing entities

The Global Investment Research Division of Goldman Sachs produces and distributes research products for clients of Goldman Sachs on a global basis. Analysts based in Goldman Sachs offices around the world produce equity research on industries and companies, and research on macroeconomics, currencies, commodities and portfolio strategy. This research is disseminated in Australia by Goldman Sachs Australia Pty Ltd (ABN 21 006 797 897); in Brazil by Goldman Sachs do Brasil Corretora de Títulos e Valores Mobiliários S.A.; Ombudsman Goldman Sachs Brazil: 0800 727 5764 and / or ouvidoriagoldmansachs@gs.com. Available Weekdays (except holidays), from 9am to 6pm. Ouvidoria Goldman Sachs Brasil: 0800 727 5764 e/ou ouvidoriagoldmansachs@gs.com. Horário de funcionamento: segunda-feira à sexta-feira (exceto feriados), das 9h às 18h; in Canada by either Goldman Sachs Canada Inc. or Goldman Sachs & Co. LLC; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in India by Goldman Sachs (India) Securities Private Ltd.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs New Zealand Limited; in Russia by OOO Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); and in the United States of America by Goldman Sachs & Co. LLC. Goldman Sachs International has approved this research in connection with its distribution in the United Kingdom and European Union.

European Union: Goldman Sachs International authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, has approved this research in connection with its distribution in the European Union and United Kingdom; Goldman Sachs AG and Goldman Sachs International Zweigniederlassung Frankfurt, regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht, may also distribute research in Germany.

General disclosures

This research is for our clients only. Other than disclosures relating to Goldman Sachs, this research is based on current public information that we consider reliable, but we do not represent it is accurate or complete, and it should not be relied on as such. The information, opinions, estimates and forecasts contained herein are as of the date hereof and are subject to change without prior notification. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Other than certain industry reports published on a periodic basis, the large majority of reports are published at irregular intervals as appropriate in the analyst's judgment.

Goldman Sachs conducts a global full-service, integrated investment banking, investment management, and brokerage business. We have investment banking and other business relationships with a substantial percentage of the companies covered by our Global Investment Research Division. Goldman Sachs & Co. LLC, the United States broker dealer, is a member of SIPC (http://www.sipc.org).

Our salespeople, traders, and other professionals may provide oral or written market commentary or trading strategies to our clients and principal trading desks that reflect opinions that are contrary to the opinions expressed in this research. Our asset management area, principal trading desks and investing businesses may make investment decisions that are inconsistent with the recommendations or views expressed in this research.

The analysts named in this report may have from time to time discussed with our clients, including Goldman Sachs salespersons and traders, or may discuss in this report, trading strategies that reference catalysts or events that may have a near-term impact on the market price of the equity securities discussed in this report, which impact may be directionally counter to the analyst's published price target expectations for such stocks. Any such trading strategies are distinct from and do not affect the analyst's fundamental equity rating for such stocks, which rating reflects a stock's return potential relative to its coverage group as described herein.

We and our affiliates, officers, directors, and employees, excluding equity and credit analysts, will from time to time have long or short positions in, act as principal in, and buy or sell, the securities or derivatives, if any, referred to in this research.

The views attributed to third party presenters at Goldman Sachs arranged conferences, including individuals from other parts of Goldman Sachs, do not necessarily reflect those of Global Investment Research and are not an official view of Goldman Sachs.

Any third party referenced herein, including any salespeople, traders and other professionals or members of their household, may have positions in the products mentioned that are inconsistent with the views expressed by analysts named in this report.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients. Clients should consider whether any advice or recommendation in this research is suitable for their particular circumstances and, if appropriate, seek professional advice, including tax advice. The price and value of investments referred to in this research and the income from them may fluctuate. Past performance is not a guide to future performance, future returns are not guaranteed, and a loss of original capital may occur. Fluctuations in exchange rates could have adverse effects on the value or price of, or income derived from, certain investments.

Certain transactions, including those involving futures, options, and other derivatives, give rise to substantial risk and are not suitable for all investors. Investors should review current options disclosure documents which are available from Goldman Sachs sales representatives or at http://www.theocc.com/about/publications/character-risks.jsp. Transaction costs may be significant in option strategies calling for multiple purchase and sales of options such as spreads. Supporting documentation will be supplied upon request.

Differing Levels of Service provided by Global Investment Research: The level and types of services provided to you by the Global Investment Research division of GS may vary as compared to that provided to internal and other external clients of GS, depending on various factors including your individual preferences as to the frequency and manner of receiving communication, your risk profile and investment focus and perspective (e.g., marketwide, sector specific, long term, short term), the size and scope of your overall client relationship with GS, and legal and regulatory constraints. As an example, certain clients may request to receive notifications when research on specific securities is published, and certain clients may request that specific data underlying analysts' fundamental analysis available on our internal client websites be delivered to them electronically through data feeds or otherwise. No change to an analyst's fundamental research views (e.g., ratings, price targets, or material changes to earnings estimates for equity securities), will be communicated to any client prior to inclusion of such information in a research report broadly disseminated through electronic publication to our internal client websites or through other means, as necessary, to all clients who are entitled to receive such reports.

All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client websites. Not all research content is redistributed to our clients or available to third-party aggregators, nor is Goldman Sachs responsible for the redistribution of our research by third party aggregators. For research, models or other data related to one or more securities, markets or asset classes (including related services) that may be available to you, please contact your GS representative or go to http://research.gs.com.

Disclosure information is also available at http://www.gs.com/research/hedge.html or from Research Compliance, 200 West Street, New York, NY 10282.

MINDCRAFT: OUR THEMATIC DEEP DIVES

Innovation & Disruption

Edge Computing | B2B: Next Payments Frontier | Precision Farming | Silicon Carbide | Why Technology is not a bubble | 5G

Extended Reality | The Genome Revolution | eSports | AI Hardware | Digital Divergence | Digital Health

GS Sustain

The PM's Guide To The ESG Revolution

37 GS Sustain Gems

Macro Insights and Policy

Top of Mind | Bear Necessities | Macro Factors | Making Cents | Trade Policy

PM Resources

A taste of the high (vol) life | Healthcare's Holy Grail | Corporate Simplification

Closing the Gender Gaps

The Low Carbon Economy

IMO 2020 | More Lean, More Green | Digitalisation and Utilities | Offshore Opportunity | Power Shift 2.0 | From Pump to Plug | NextGen Power: Solar

Music's Return

Commodity Corner

Top Oil & Gas Projects | Copper Top Projects | Re-Imagining Big Oils

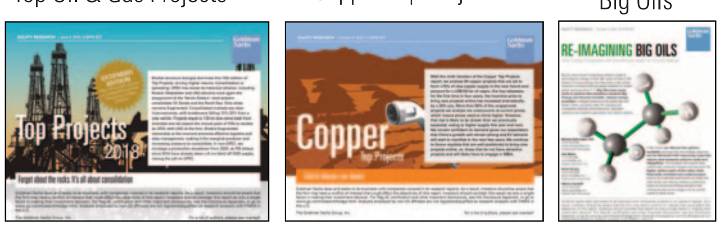

Venture Capital Horizons

Rising Asia

The Rise of China FinTech | China's Rise in AI | Mapping China's Credit - Vol.2

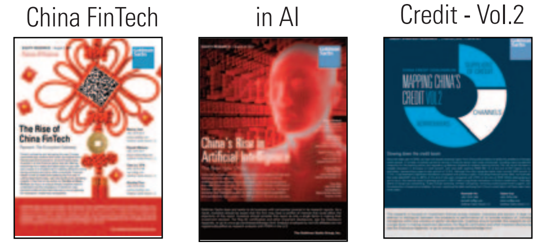

Quantamentals

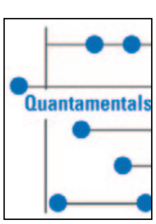

Consumer Currents

Supermarket Shift: Expect Measured Disruption | Millennials | World of Games: Consoles | The Great Wall of SKUs | The Store of the Future

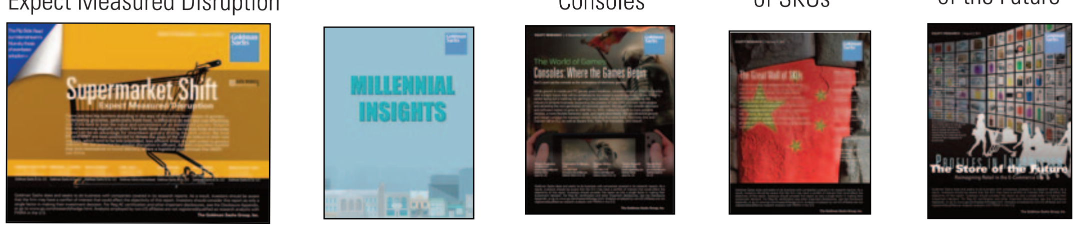

...and more