KeyBanc Capital Markets

April 29, 2018

Technology: Enterprise Infrastructure

Industry Update

Enterprise Infrastructure: Rubrik's Polaris Release Aims at Emerging Hybrid Cloud Software Marketplace

Rubrik, a privately held company with $292M in funding, has recently announced a strategic expansion of its product portfolio with the release of Polaris—a SaaS-based platform that will enable improved visibility and workflow orchestration of primary and secondary data assets. We see this as part of a broader market development as IT organizations look for improved control of application resources across public and private clouds with solutions from emerging vendors like Rubrik, Cohesity, Datrium, NTNX, PSTG, and ANET, as well as established vendors that are introducing API offerings into their core products (CVLT, NTAP, FFIV, and CSCO).

Key Investment Points

Rubrik impacting the secondary storage landscape. Rubrik has emerged as a leading provider of next-gen data management software, delivering the scale-out Cloud Data Management platform through appliances, as well as cloud and reference architectures. Rubrik has grown to nearly $300M in run-rate bookings (February 2018) by simplifying how IT organizations manage their data, including backup, recovery, archival, and compliance use cases, all from a central management console that enables search and analytics across all data sets. Channel feedback on Rubrik's Cloud Data Management highlights the simplicity of its UI functions and easy scale-out capabilities. As a result, Rubrik has demonstrated strong adoption in our quarterly channel survey (here) and has evolved into an important go-to-market partner with Pure Storage and Nutanix.

New Rubrik Polaris SaaS platform extending backup capabilities. Rubrik's new Polaris SaaS platform allows IT organizations to create a unified system of record across a company's data, including on-premise and in the cloud. Polaris provides visibility, workflow orchestration, and data intelligence across all of a customer's global data managed by the Rubrik Cloud Data Management platform. Polaris utilizes an open API architecture to expose the platform's underlying functionality and enable higher-level applications provided by both Rubrik and third-party developers. The first of these higher-level applications provided by Rubrik is Polaris GPS, a centralized dashboard for global policy management and monitoring.

Commvault simplifying architecture with HyperScale. Commvault's recent HyperScale product releases (HyperScale Software, HyperScale Appliance, and ScaleProtect with Cisco) are part of the Company's strategic push to simplify the core functions of Commvault's architecture in a scale-out form factor targeted at the midmarket. Our feedback from several channel partners suggests positive initial takeaways from customer trials, and we see a race among Commvault, Rubrik, and Cohesity (among others) to secure incremental channel mindshare from the Dell EMC shakeup. While Rubrik has grown quickly in the U.S., we do not see this impacting Commvault's near-term outlook or results (Tuesday before the open) as we would highlight a recently closed $1M+ Commvault deal out of northern EU as a signal of the Company's broad distribution and reach.

Stock Implications ▲ Positive, ▼ Negative, — Neutral

	Company	Stock	Rating	Target	Price	Mkt Cap.
—	Commvault Systems, Inc.	CVLT	OW	$75.00	$69.80	$3,298.7M

Enterprise Infrastructure

Alex Kurtz
503.821.3898
akurtz@key.com

Steve Enders, CFA
503.821.3978
senders@key.com

Hybrid Cloud and the Market Evolution of Data Management Software

The term hybrid and multi-cloud has come to encapsulate a wide range of strategic initiatives across vendor product roadmaps in addition to how end customers are now deploying applications and data across on-premise and off-premise environments. Providing this set of tools and functionality has become a core requirement, in our view, as end customers are increasingly looking for longer-term assurances that their data management platforms can leverage public and private cloud assets.

From our perspective, we view the core change as one that allows software providers like Rubrik, Commvault, Nutanix, VMware, Red Hat, et al., to extend their core capabilities/IP into public cloud platforms like AWS or Azure. Given the relatively nascent stage of many hybrid cloud product initiatives (and how they are priced and monetized), we believe the market opportunity for this functionality is still too early to quantify, although would note these providers are competing in multibillion-dollar markets.

Rubrik: Quick to Scale in the Secondary Data Management Market

Rubrik has quickly established itself as a leader in the next-gen data management and secondary storage markets, growing to nearly $300M in run-rate bookings in the less than four years since the Company introduced its first Brik for data management in late 2014 (Rubrik was founded in January 2014). From our channel discussions, as well as speaking with management over the last several years, core competitors are typically Dell EMC (primarily the Data Domain/Avamar product families), HPE, Veritas, Commvault, and IBM. Also included in this competitive set are emerging backup vendors including Cohesity, Veeam, Zerto, Actifio, and Datrium, but this set will vary depending on how well distributed these companies are in specific markets.

Rubrik last raised $180M in April 2017 at a $1.3B valuation, led by IVP, Greylock Partners, and Lightspeed Venture Partners, and brought total funding to $292M. The Company was founded by Bipul Sinha (CEO of Rubrik, former partner at Lightspeed), Arvind Jain (Google), Soham Mazumdar (Google), and Arvind Nithrakashyap aka Nitro (Oracle/Exadata). Also, sales is led by Mark Smith (Arista Networks), and the Company has recently hired Murray Demo as CFO (Atlassian).

Rubrik has continuously expanded its core product offering to include support for all major hypervisors (VMware, Hyper-V, Nutanix AHV), SQL databases (Oracle, Microsoft SQL Server), storage protocols, including file (iSCSI, NFS, SMB), and cloud providers (AWS, Azure, Google Cloud), allowing Rubrik to expand its customer base by more than 4x last year and grow to over 800 employees.

Rubrik's core Cloud Data Management Platform has found strong traction with enterprise and government customers by simplifying the backup and recovery market. The Cloud Data Management product has traditionally been delivered as a converged appliance with a simple UI and quick-start functionality (up and running in less than an hour). Rubrik's core product offering provides snapshots for backup, recovery, and archival use cases with a near-zero recovery time objective (RTO, or how quickly an app/data can be brought back after failing) that enables additional data use cases for analytics, compliance, dev/test, and search, no matter where the data resides (cloud vs. on-premise).

Alex Kurtz 503.821.3898 / akurtz@key.com

Exhibit 1: Traditional Incumbent Backup Environments

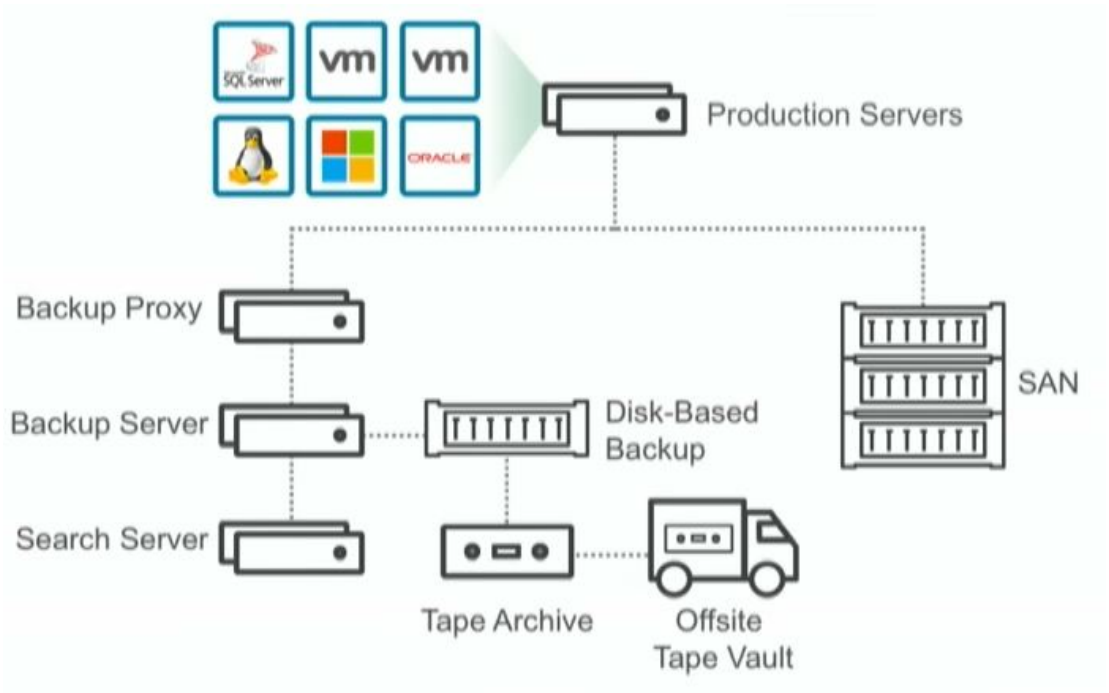

Source: Company reports

Rubrik has improved upon incumbent backup solutions by reducing the architectural complexity and simplifying both the setup and management process. Rubrik has developed a software-defined solution, collapsing the traditional secondary storage environment with multiple servers (backup proxy, backup server, search server, etc.) into a scale-out converged offering (both compute and storage in a single box) delivered in a 2U appliance. Also, Pure Storage and Nutanix have become important go-to-market partners as they look to dislodge primary and secondary storage assets in a combined floor-sweep type of sales program.

Exhibit 2: Rubrik Backup Environment

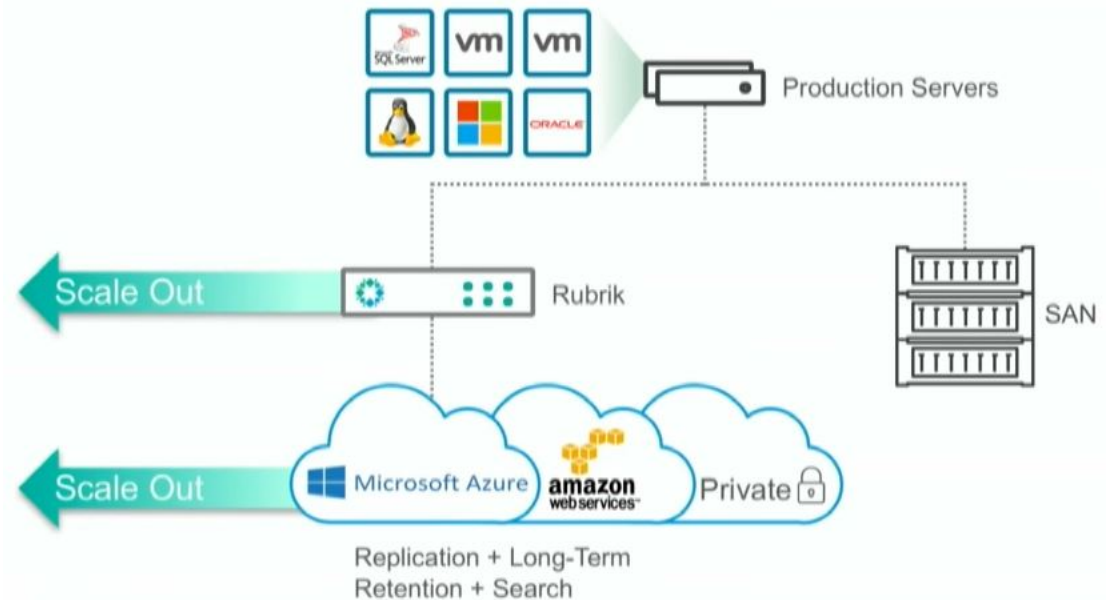

Source: Company reports

Rubrik's converged software scales linearly (scale-out, or just add a node and performance scales in line) with no additional backup software or servers required as has typically been the case in legacy systems. Customers can set SLA policy and interact with Rubrik's APIs to automate service delivery, limiting the scope that is usually needed to manage a backup system. Similarly, Rubrik built its platform with cloud APIs enabled from the beginning so cloud connections are not bolted-on like in some incumbent solutions.

The latest Rubrik release includes "CloudOn" services in all major cloud providers, allowing server-less conversion of VMs to cloud instances. This can lower the cost of cloud-based backup by only utilizing cloud services when needed. As well, this release included SQL Live Mount, allowing Rubrik customers to power on SQL Server with near-zero RTO at any point on Rubrik, providing instant access to data without moving the data.

Rubrik's Architecture a Differentiator

Rubrik's core Cloud Data Management Platform has three layers in its tech stack:

- The interface layer, including Infinity integrating the system via APIs and the Crystal UI and API interface.
- The logic layer, including Cerebro, the brains of the Rubrik platform, which has abstracted the data control plane from the underlying infrastructure.
- The core layer, highlighted by the Atlas cloud-scale file system that is masterless (enabling scale-out architecture with no single choke point) and self healing.

Exhibit 3: Rubrik Cloud Data Management Platform Architecture

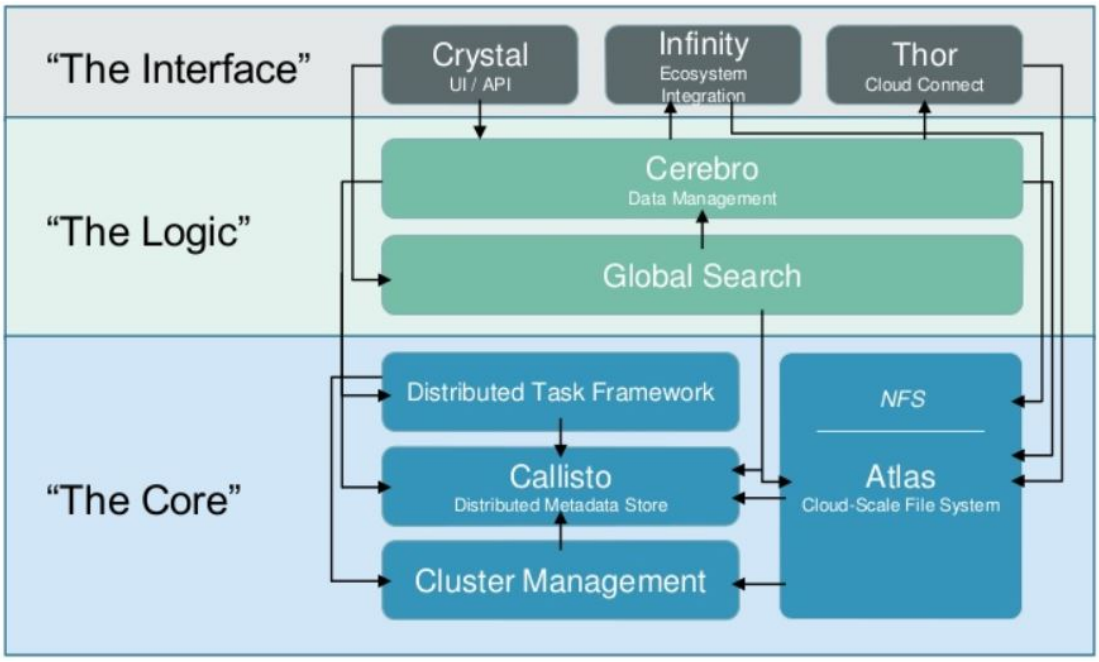

Source: Company reports

Interface Layer

Rubrik's interface layer is an API-based architecture that is designed to abstract away the underlying platform complexity. Rubrik's APIs can deliver granular control for customers to utilize the Rubrik platform in the way that best fits their needs. Rubrik's APIs can be tied to various external and internal systems, including: weather-triggered DR to the cloud via an API call from OpenWeatherMap; or integration with ServiceNow to automate protection for new workloads, self-service recovery, or enable developers to provision clones for dev/test.

Logic Layer

The core of Rubrik's logic layer is Cerebro, Rubrik's data control plane that is detached from the underlying infrastructure, enabling the platform to manage data across on-premise systems to the cloud with a single fabric. This decoupled system is software-defined and resilient and enables data mobility across the fabric.

Importantly, Rubrik's software utilizes a declarative model, meaning the user defines what the end-state should look like and where data should reside, and the underlying software executes on this vision. Most incumbent backup models use an imperative model, where the user describes every step needed to achieve a desired end-state for each workload. This model typically means a user has to schedule jobs manually and is unnecessarily complex and time consuming. Comparatively, Rubrik is simpler, letting the end user set just the SLA policy and the desired end-state for the data, and allowing the Cerebro engine to allocate resources and schedule tasks, eliminating unnecessary manual tasks.

A key aspect of Cerebro is the underlying Rubrik Blob Engine, a distributed version-control system that provides many of Rubrik's core data management services, including immutability, deduplication, retention, replication, and archival, and manages the data lifecycle. The Blob Engine takes the user-defined SLA policy as an input to schedule tasks behind the scenes and automates the data lifecycle management process, including compression, deduplication, consolidation, and garbage collection.

The Blob Engine was also designed to minimize fragmentation issues related to deduplication by constantly analyzing snapshots to defragment/optimize to ensure fast recovery times (minimize RTO) for tier-one apps. Rubrik delivers lower RTO through Live Mounts by issuing parallel requests to the cluster nodes and underlying storage concurrently and utilizes parallel synthesis of the data to improve the amount of time in which data is recovered by the system.

Exhibit 4: Rubrik Utilizing Parallel Reads

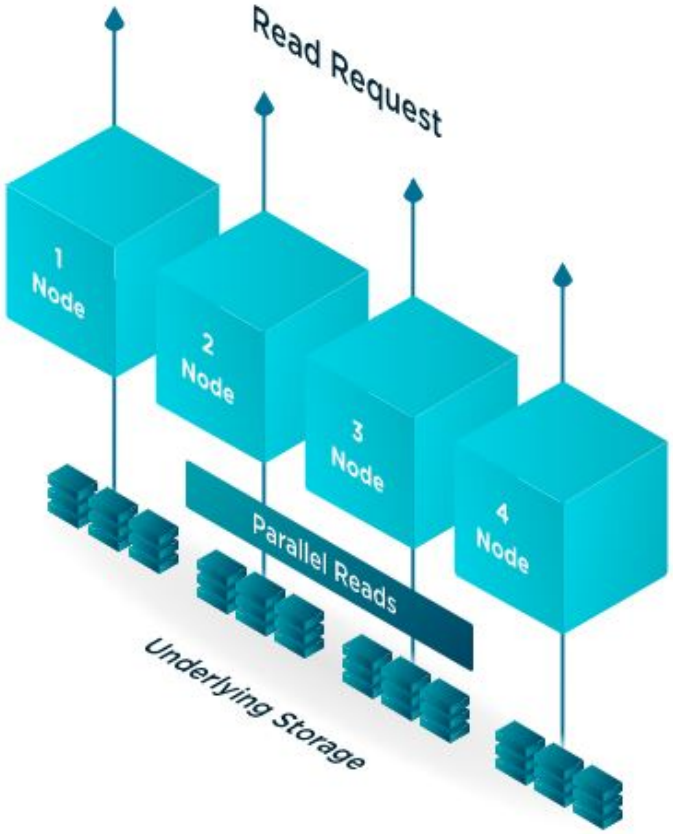

Source: Company reports

Core Layer

Directly connected to Cerebro in Rubrik's core is the platform's Distributed Task Framework. The Task Framework schedules and assigns the jobs/tasks across the system to enforce the user-defined SLA policies in a fault-tolerant and efficient manner. The framework has two components: a Task Scheduler, which evenly assigns tasks across the cluster; and Task Maintenance, which enforces activities to uphold the assigned SLA policies (which are then distributed by the scheduler). Workloads can be prioritized by a QoS component within the system as dictated by the declarative nature. The Distributed Task Framework abstracts away complexity for the end user by automatically distributing jobs across the cluster so the user does not have to deal with this scheduling complexity.

Rubrik's main file system, Atlas, was built from the ground up to be scale-out, masterless, and fault tolerant (no single point of failure and no choke point), so that if the system fails, it can replicate whatever data was lost. Rubrik was able to create an app-aware file system and optimize for its data management application. Atlas understands it is storing snapshot chains, which is what the Blob Engine uses to back its snapshots, and ensures data integrity so the system knows that the snapshots are immutable and that data has never been changed (immutability is key to preventing ransomware and to roll back to a prior version).

Atlas has focused primarily on restore efficiency and customized its apps and features to support blob engine and snapshot chains. As a result, Rubrik can do a nearly instantaneous recovery (one second or less to build a workload while other vendors can take hours). Atlas is able to do this because it understands snapshot chains and uses a metadata operation instead of a data read—no data has to move and reads are served on the fly by merging the snapshot chain.

Rubrik Polaris Enabling SaaS Delivered Data Management

Rubrik's recently announced Polaris SaaS platform allows IT organizations to create a unified system of record across all of a company's data, including on-premise and in the cloud. Polaris provides a centralized platform for data visibility, workflow orchestration, and data intelligence across all of a customer's fragmented data that is managed within the Rubrik Cloud Data Management platform. The Polaris platform manages data by tracking the data's metadata, allowing the underlying data to stay where it resides.

Polaris utilizes Rubrik's open API architecture to expose the platform's underlying functionality and enable higher-level applications provided by both Rubrik and third-party developers.

Exhibit 5: Rubrik Polaris Platform

Source: Company reports

The first of these higher level applications provided by Rubrik is **Polaris GPS**, a centralized dashboard for global policy management and monitoring across all Rubrik environments. Polaris GPS provides a simple and interactive global dashboard to track all apps and data with deep search and filtering to better identify and understand incidents.

Through Polaris GPS, enterprises can produce reports to demonstrate global SLA compliance for all of an organization's data. We expect Rubrik to continue expanding Polaris' functionality into adjacent data management markets.

Hybrid & Multi Cloud: A Lot of Different Things to Different People

Over the last two-to-three years, there has been an industry-wide push from enterprise-focused infrastructure providers to introduce solutions that can extend their core platforms (IP) into public clouds, as many of their end customers have started their own journeys to deploy applications and data tiers into AWS, Azure, and Google Cloud, to name a few.

The technical differences between hybrid and multi-cloud is starting to get a bit murky, but in essence, a customer may see an environment first through the prism of private (on-premise) and public clouds as a hybrid cloud approach, and secondarily, if leveraging multiple public cloud platforms (AWS for compute and Azure for SQL, as an example), this could be seen as a multi-cloud strategy.

Ultimately, we see this technology integration driving a host of new product introductions into the market through two high-level approaches:

(1) **API integrations** that enable infrastructure providers like NetApp, Commvault, or Rubrik to enable the movement/visibility of application data across clouds (both public and private). Customers benefit by having access to a more dynamic set of compute/storage pricing across various platforms and the ability to burst into and out of public cloud resources (although more limited to date).

(2) **Cloud-agnostic distribution** of applications/run-time platforms, where a platform like Red Hat OpenShift can enable a customer to deploy containers both on-premise, as well as in a public cloud. Another example would be VMware's initiative to enable vSphere/VMs to run on bare-metal servers at AWS.

The following exhibit aims to outline the various categories of hybrid and multi-cloud integration. We would caution that it is still early innings as far as end-customer adoption of approaches 1 and 2 (above), but is quickly becoming a standard requirement for end customers.

Exhibit 6: Hybrid and Multi-Cloud Vendor Landscape

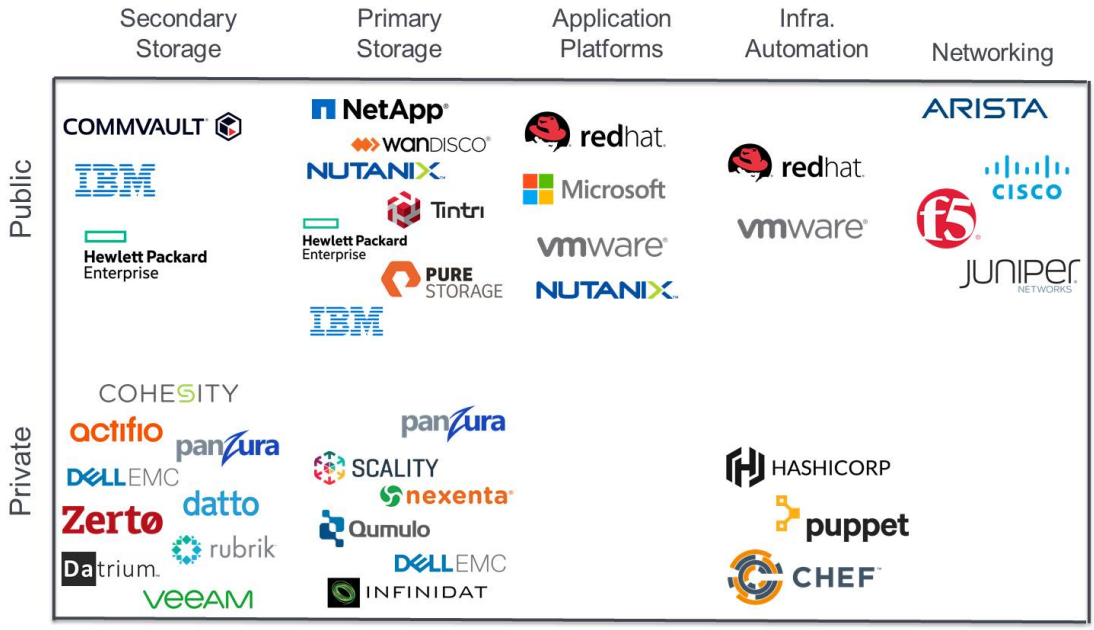

Source: Company reports

Also of note in this discussion is the various complexities involved with creating a hybrid cloud product. As recently reported, Nutanix has experienced delays with its Xi Cloud Services solution due to the engineering complexities involved with building a multi-cloud product. Nutanix's first cloud-based service, Xi Disaster Recovery, a DRaaS product, is attempting to solve the difficult problem of replicating an on-premise application in public cloud in real time. We believe these delays demonstrate the complexity of both creating a public cloud service and of solving this novel problem for core, tier-1 applications. While we are not modeling Xi currently, we expect this to become a service that provides additional differentiation for Nutanix in the hyperconverged market.

Commvault – Adapting Quickly with HyperScale Launch

Commvault's HyperScale product offering, announced in October 2017, improves upon Commvault's existing software-only offering by providing additional agility, resiliency, and availability to on-premise data and apps with simpler setup, management, and upgrades compared to Commvault's prior delivery model. HyperScale modernizes Commvault's software delivery model by creating a product with modular scale and non-disruptive capacity and performance expansion – this includes built-in resiliency with erasure-coding and no single point of failure, creating a true scale-out architecture. We believe, if the Company can execute, this product has the potential to evolve into a $100M revenue platform given Commvault's strong core IP, large installed base of active customers (26K), and years of channel investments.

Additionally, HyperScale utilizes a software-defined storage architecture to leverage server-based storage instead of traditional controller-based or object-oriented storage devices. The total HyperScale solution includes the HyperScale Software, Commvault Data Platform (including Protection and Archive), Red Hat (RHEL and GlusterFS with the appliance also including RHEV), and Windows 2016.

Exhibit 7: Commvault HyperScale Architecture

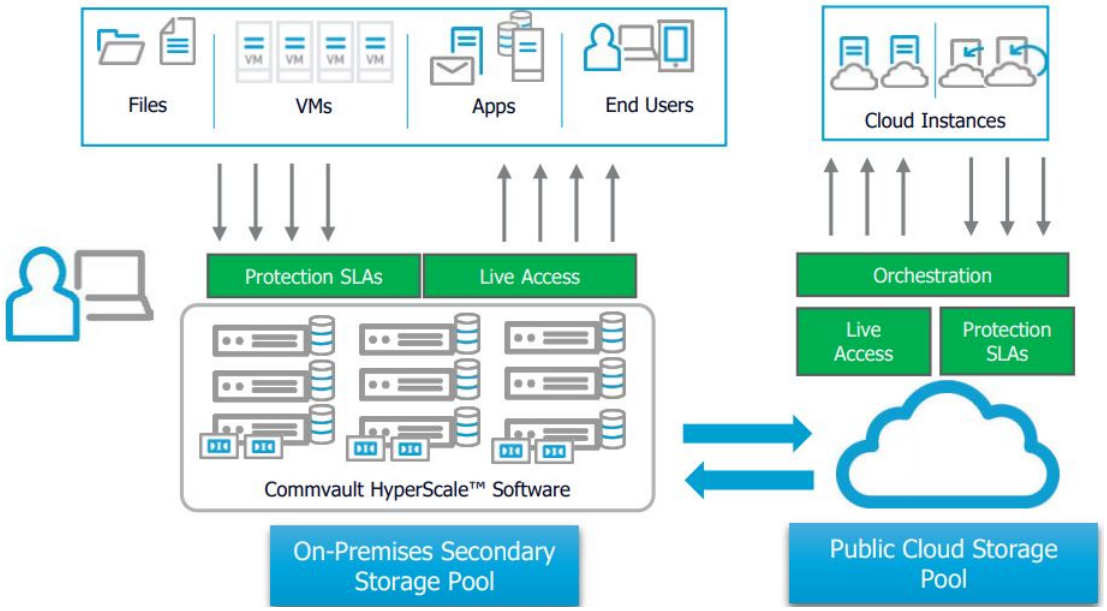

Source: Company reports

The biggest difference with HyperScale is that Commvault has updated its go-to-market strategy, now providing software that is scale-out as well as an integrated appliance with both Cisco (ScaleProtect) and Fujitsu (HyperScale Appliance – converged offering that is most comparable to what Rubrik provides). The solution also allows for a unique consumption-based pricing model and an ability to start with a small offering and scale up to multiple petabytes (scale to 10s of PBs)

HyperScale is more targeted at mid-market customers with the simpler setup approach and it remains a longer-term question of how the Company's larger enterprises will respond to the product launch. HyperScale software provides reference architectures with key technology partners including Fujitsu (including HyperScale Appliance), Cisco (including ScaleProtect with Cisco UCS), Lenovo, HPE, Super Micro, Huawei, and DellEMC. The product integrates with AWS, Google Cloud, Oracle Cloud, and Azure, and includes REST APIs to automate complex tasks and policies. Importantly, HyperScale enables increased simplicity, flexible pricing, automation, orchestration services, programmable APIs, and direct data access through standard interfaces.

Commvault's updated HyperScale offering closes some of the simplicity and ease of use "gaps" that we have heard in our channel discussions, and initial feedback has been positive as partners are training on the product and performing initial customer trials.

Valuation & Investment Risks

Commvault Systems, Inc. (CVLT)
Our price target of $75 is based on an EV/FCF multiple of 25.0x our FY19 estimate, which compares to the two-year range of 16-24x. This equates to an EV/sales multiple of 4.0x vs. other mature software companies at 4-6x and the takeout multiples for TIBCO and Veritas at 3.7x and 2.5x, respectively. This also compares to the historical EV/sales NTM range of 2-6x. CVLT currently trades at 22.8x FY19E EV/FCF.

Market and macroeconomic conditions could interfere with the realization of this price target, as could risks such as any large-deal delays that negatively affect a given quarter, management turnover, potential acquisitions, heightened competitive pressures, success of alternative technologies, and currency fluctuations.

Important Disclosures

Important disclosures for the companies mentioned in this report can be found at https://key2.bluematrix.com/sellside/Disclosures.action.

Please refer to the analysts' recently published reports for company-specific valuation and risks.

Reg A/C Certification

The research analyst(s) responsible for the preparation of this research report certifies that:(1) all the views expressed in this research report accurately reflect the research analyst's personal views about any and all of the subject securities or issuers; and (2) no part of the research analyst's compensation was, is, or will be directly or indirectly related to the specific recommendations or views expressed by the research analyst(s) in this research report.

Rating Disclosures

Distribution of Ratings/IB Services Firmwide and by Sector

KeyBanc Capital Markets			IB Serv/Past 12 Mos.		Technology			IB Serv/Past 12 Mos.	
Rating	Count	Percent	Count	Percent	Rating	Count	Percent	Count	Percent
Overweight [OW]	284	46.63	74	26.06	Overweight [OW]	108	50.70	22	20.37
Sector Weight [SW]	315	51.72	56	17.78	Sector Weight [SW]	105	49.30	5	4.76
Underweight [UW]	10	1.64	3	30.00	Underweight [UW]	0	0.00	0	0.00

Rating System

Overweight - We expect the stock to outperform the analyst's coverage sector over the coming 6-12 months.

Sector Weight - We expect the stock to perform in line with the analyst's coverage sector over the coming 6-12 months.

Underweight - We expect the stock to underperform the analyst's coverage sector over the coming 6-12 months.

Alex Kurtz 503.821.3898 / akurtz@key.com

Important disclosures for the companies mentioned in this report can be found at **https://key2.bluematrix.com/sellside/Disclosures.action**.

Please refer to the analysts' recently published reports for company-specific valuation and risks.

Other Disclosures

KeyBanc Capital Markets is a trade name under which corporate and investment banking products and services of KeyCorp and its subsidiaries, KeyBanc Capital Markets Inc., Member NYSE/FINRA/SIPC ("KBCMI"), and KeyBank National Association ("KeyBank N.A."), are marketed.

KeyBanc Capital Markets Inc. ("KBCMI") does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

This report has been prepared by KBCMI. The material contained herein is based on data from sources considered to be reliable; however, KBCMI does not guarantee or warrant the accuracy or completeness of the information. It is published for informational purposes only and should not be used as the primary basis of investment decisions. Neither the information nor any opinion expressed constitutes an offer, or the solicitation of an offer, to buy or sell any security. The opinions and estimates expressed reflect the current judgment of KBCMI and are subject to change without notice. This report may contain forward-looking statements, which involve risk and uncertainty. Actual results may differ significantly from the forward-looking statements. This report is not intended to provide personal investment advice and it does not take into account the specific investment objectives, financial situation and the specific needs of any person or entity.

No portion of an analyst's compensation is based on a specific banking transaction; however, part of his/her compensation may be based upon overall firm revenue and profitability, of which investment banking is a component. Individuals associated with KBCMI (other than the research analyst(s) listed on page 1 of this research report) may have a position (long or short) in the securities covered in this research report and may make purchases and/or sales of those securities in the open market or otherwise without notice. As required by FINRA Rule 2241(C)(4)(A), financial interest, if any, by any research analysts listed on page 1 of this report will be disclosed in Important Disclosures, Company-specific regulatory disclosures located above in the Disclosure Appendix. KBCMI itself may have a position (long or short) in the securities covered in this research report and may make purchases and/or sales of those securities in the open market or otherwise without notice. As required by FINRA Rule 2241(C)(4)(F), if KBCMI, or its affiliates, beneficially own 1% or more of any class of common equity securities in the subject company(ies) in this research report, it will be disclosed in Important Disclosures, Company-specific regulatory disclosures located above in the Disclosures Appendix. This communication is intended solely for use by KBCMI clients. The recipient agrees not to forward or copy the information to any other person without the express written consent of KBCMI.

Alex Kurtz 503.821.3898 / akurtz@key.com