IBM Spectrum Fusion Focuses on Software-Defined Storage and Reduced Opex Costs

By Jean S. Bozman

At the beginning of March, IBM announced its latest version of IBM Spectrum Fusion storage-management software, addressing a blended environment of containers, distributed data and Kubernetes orchestration – and leveraging the growing customer trend toward managing data across hybrid clouds and software-defined storage (SDS). The new Spectrum Fusion release consistently manages data – file data and object data – across on-premises data centers, the edge, and the hybrid cloud.

In a software-centric world of hybrid clouds, customers are at various stages of adoption of software-defined storage and containers. In recent years, customers have been looking to reduce the complexity of deploying stored data across hybrid clouds, with the option to extend them by reaching into multi-cloud deployments, as well. With software-defined storage, the enabling software no longer has to be re-invented for each project. Rather, data can be managed in a consistent way across the hybrid cloud, including data centers and a range of public and private cloud services from cloud service providers.

IBM Spectrum Fusion Specifics

What IBM Spectrum Fusion offers is an abstraction layer that is largely hardware-agnostic – provided that the appropriate APIs and data-management policies are in place. In short, this software-defined approach to storing data provides application data portability across the enterprise. It is designed to reduce administrative time associated with managing and protecting distributed data in a hybrid cloud.

The new IBM Fusion release, supporting Red Hat Open Shift Container Platform (OCP) and Kubernetes orchestration, increases support for containerized applications and access to multi-site data. The release became available in 1Q22. Customers can manage the full range of IBM Spectrum Fusion capabilities by using a single management portal.

Specifically, IBM announced last fall that IBM Spectrum Fusion contains multiple storage technologies for data management – and that it is expanding Spectrum Fusion now with the support for Red Hat Open Shift Data Foundation Essentials. Earlier, IBM had combined functionality from IBM Spectrum Scale's General Parallel File System (GPFS); IBM Spectrum Discover, which uses an organization's metadata to locate data; and IBM Spectrum Protect Plus for backup/restore capability in the hybrid cloud. Previously, those IBM Spectrum products had been offered as separate products in a six-part software suite.

Copyright © 2004-2022 ***Robert Frances Group****, all rights reserved*
46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; Contact: inquiry@rfgonline.com

Impact on Storage Administrators

This approach to software-defined storage supports customers' business model for cloud-based services, with a pay-as-you-go approach to consumption of storage infrastructure. In this case, Spectrum Fusion is taking on the traditional tasks of storage administrators who are focused on tweaking and tuning management software that controls the storage infrastructure. This approach is designed to improve efficiency of data management, while reducing data-management-related operational expenses (Opex).

Regarding Opex considerations, IBM's offer is aimed at application developers and IT administrators who work with software containers and Kubernetes, often in DevOps units throughout the enterprise. Using Spectrum Fusion, they can focus on DevOps, without refactoring applications for data-placement purposes. With the new Spectrum Fusion release, automation and software tools are enabled by Red Hat OpenShift Data Foundation Essentials, supporting tight data services integration via Red Hat OpenShift software.

Application Mobility in a Changing Software Landscape

As customer adoption of containers has increased in recent years (2015-2021), the realization has grown that moving the applications to the data – rather than moving the data across the network – is the most efficient way to reduce network-related delays in data transfer. If the data can be located, and managed in place, then DevOps programmers and data administrators will have less work to do when adapting enterprise applications for hybrid cloud deployments.

The COVID-19 pandemic, accelerated the rate of application migrations to cloud-based consumption for IT services, accessing servers, storage, and networking as cloud services, beginning in 2020. While many customers continue to use both virtual machines (VMs) and containers in their software environments, customer migrations are leaning heavily on containers to support mission-critical and business-critical workloads.

This approach to customer adoption of software and cloud migrations was seen most strongly when VMware released its Tanzu technology, which directly embedded Kubernetes support into the VMware flagship vSphere 7 virtualized software in fall, 2020. This allowed many enterprise customers to move more quickly to containers and hybrid clouds that use Kubernetes orchestration to manage containers. VMware, with a large installed base of VMware and vSphere customers worldwide, supports the IBM Cloud and multiple CSP cloud-services for customers' multi-cloud deployments.

After IBM acquired Red Hat in 2019, IBM placed more focus on containers and Kubernetes. Subsequently, IBM developed a number of software initiatives designed to grow the company's

Copyright © 2004-2022 ***Robert Frances Group****, all rights reserved*

opportunities in hybrid cloud, artificial intelligence (AI), application development and automation. Today, hybrid cloud, multi-cloud and AI are three leading IBM initiatives, as stated, clearly and repeatedly, by IBM CEO Arvind Krishna.

This was, in part, due to IBM's recognition that customer adoption patterns were changing from hardware-centric purchases to cloud-centric software-development and cloud-based subscription models for servers and storage. With the COVID-19 pandemic, IBM moved to support customer acceleration of digital transformation in the data center – investing more heavily in software, AI, and management of data and storage in 2020 and 2021.

Meanwhile, competition in this space has grown, with other container/orchestration offers from Dell Technologies, Hewlett-Packard Enterprise (HPE), and Oracle Corp., among others. In the cloud provider services (CSP) arena, IBM Cloud competes with Amazon Web Services (AWS); Microsoft Azure; and Google Cloud. However, in recognition of the shift to hybrid cloud and multi-cloud, other software vendors now offer third-party support for multi-cloud access – and CSPs offer connections to other CSP cloud services. This is, in short, a pragmatic co-opetition model for cloud services providers, based on customers' cloud requirements.

Software-Defined Infrastructure and Multi-Cloud Adoption

For customers, the evolution of hybrid clouds requires rapid access to multiple data sources, which are distributed across the network – linking on-premises data centers, public clouds, and private clouds. IBM is clear about saying that the Spectrum Fusion release supports mobility, caching and replication – along with strong support for Kubernetes orchestration.

This approach gains customers the flexibility to tap multiple clouds and a multi-cloud strategy – as needed – based on application type, service type, or capabilities within specific clouds (e.g., processor types supported; storage capacity; network speeds; security levels, data protection and backup/recovery services).

Reasons for multi-cloud adoption include mergers/acquisitions, resulting in many data silos, and previous decisions made by multiple lines of business (LOBs) within one corporation. Data access by containers supports customers' move to improve application mobility and to reduce costs and staff time associated with moving data across the network, and repeated data migrations.

Key Takeaways

Given cyberattacks and ransomware threats, it's clear to us that data management across the hybrid cloud is a Priority-One challenge for most businesses. Customers must find ways to

Copyright © 2004-2022 ***Robert Frances Group****, all rights reserved*

leverage more data, reaching across the entire organization, while getting accurate results for better data analytics and better business outcomes.

We're already seeing the impact of distributed data on traditional "core" enterprise issues, such as corporate security, application availability, and the accurate use of production data. But many obstacles stand in the way of more efficient data management – including network latency, large data-stores in isolated data "silos" and aging applications that must be re-written.

The next frontier for enterprises that are moving more applications to the hybrid cloud will be the emergence of simpler and easier management of large volumes of distributed data – from the Core to the Cloud to the Edge. Achieving customers' goals for efficiency, productivity, security, and Opex cost reductions will require software tools that are able to identify, manage and protect large volumes of distributed data more productively and efficiently.

Copyright © 2004-2022 ***Robert Frances Group****, all rights reserved*
46 Kent Hills Lane, Wilton, CT. 06897; (203) 429 8951;
http://www.rfgonline.com/; Contact: inquiry@rfgonline.com