TVP Strategy

White Paper:

SDDC Software Defined Storage Coverage Report

for DataCore SAN Symphony, Hedvig, and ioFABRIC

Edward L. Haletky

Analyst - Secure Hybrid Cloud

TVP Strategy

Version 1.0 June 2017

Abstract

Our software defined storage(SDS) analysis focuses on the ever-growing list of requirements of the modern data center. We look for those requirements that address the future, not the past, and classify them into major categories, adding categories as needed.

Each time we talk with an enterprise, peer, evangelist, end user, or vendor, we expand and clarify our requirements list. We look at each SDS use case, tool, option, and mechanism during our conversations and either match it to an existing requirement or add a new one. Our categories are broad by nature, but our requirements are specific by design.

The goal of our research is to demonstrate the extent to which each vendor's offerings address our requirements, as well as how they stack up against each other. Where there is little overlap in requirements coverage between two or more offerings, you may be able to use them together to form a comprehensive SDS architecture. Where they do overlap, your choice of tool will depend entirely on the specific needs outlined

by your SDS architecture.

Our goal is to determine how each vendor's offerings address the SDS needs of the secure and protected hybrid cloud.

I. Coverage Graph

In the following radar graph(s), we represent our ideal requirements coverage in each category using the light gray shadow. The colored areas demonstrate how well the offerings of one or more vendors stack up against our ideal and against each other. No single offering meets all of the ideal requirements; a suite of tools is often necessary to meet one's needs. The multi-axis graph indicates where various tools overlap, where each may be lacking, and where they may complement each other.

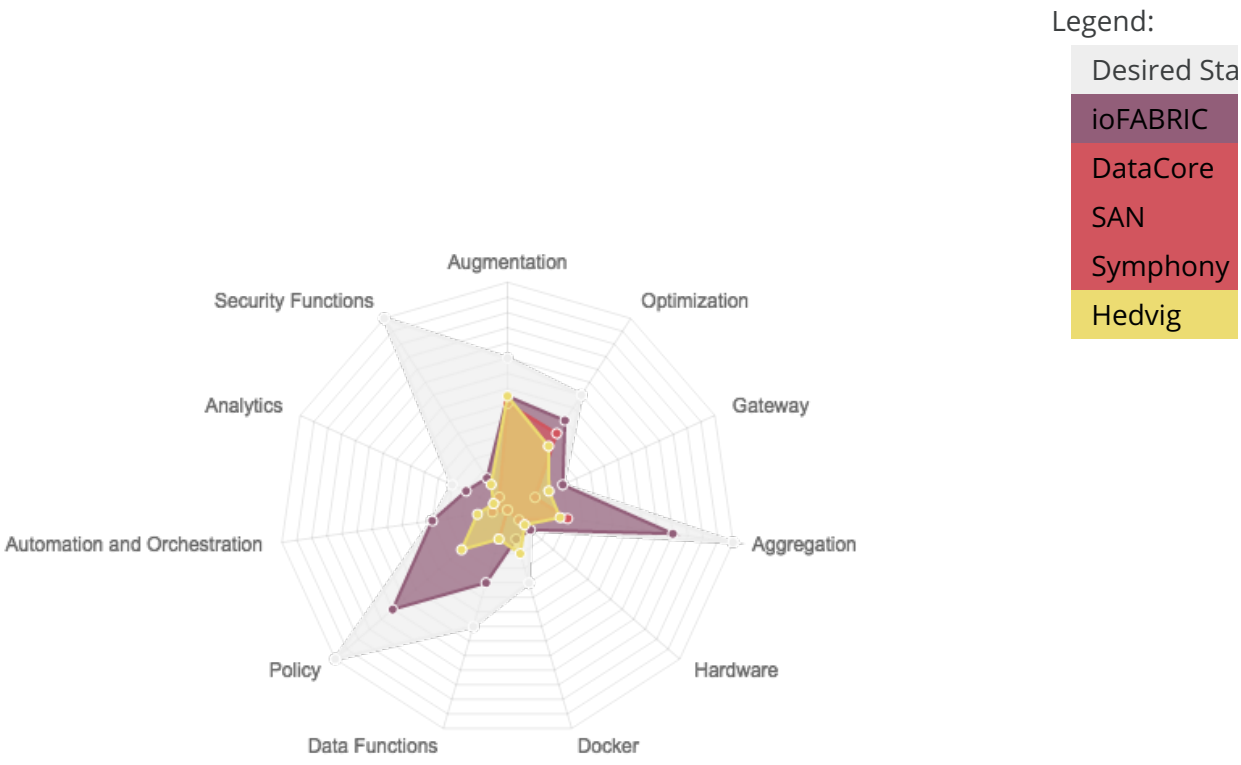

The following multi-axis graph indicates the coverage of our desired state by products from just ioFABRIC:

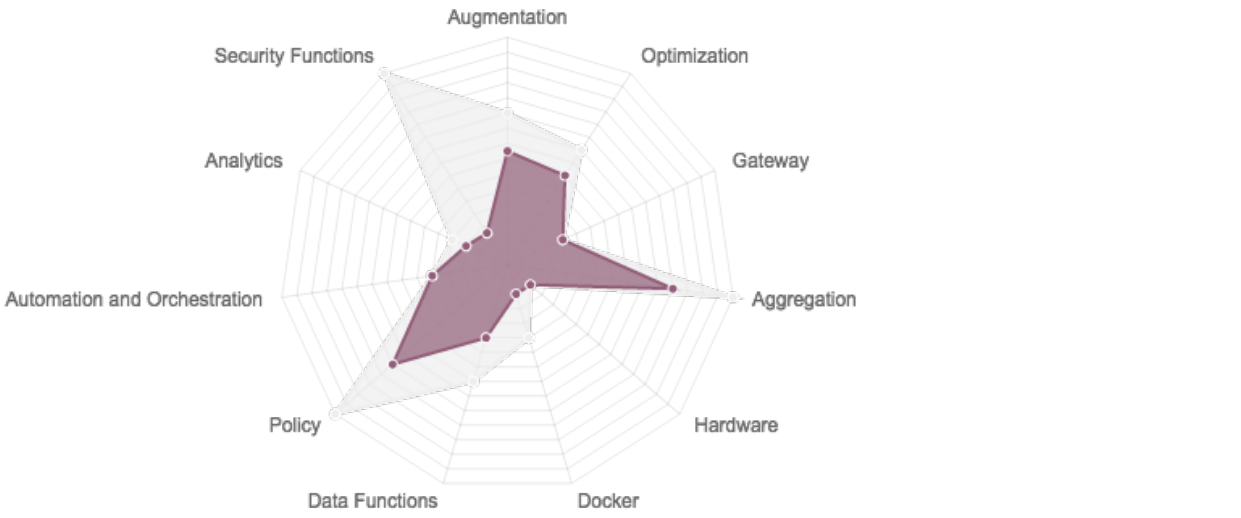

The following multi-axis graph indicates the coverage of our desired state by products from just DataCore SAN Symphony:

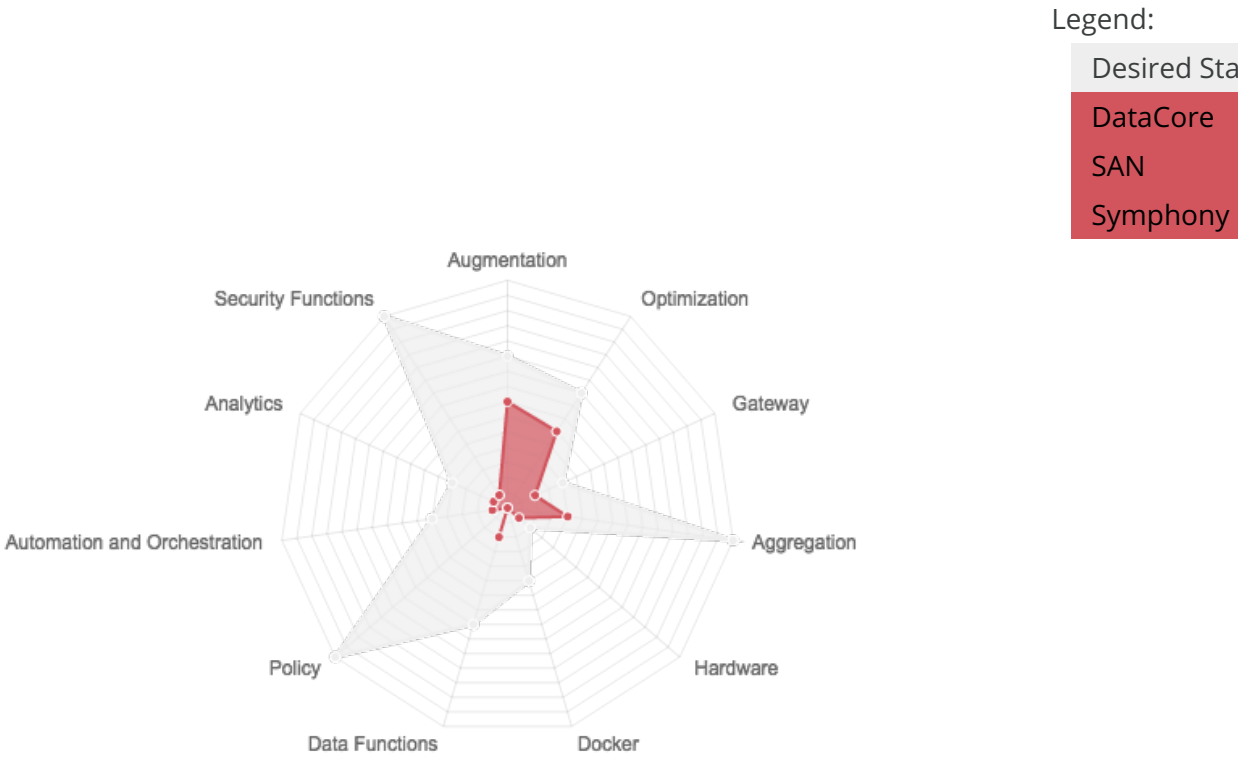

The following multi-axis graph indicates the coverage of our desired state by products from just Hedvig:

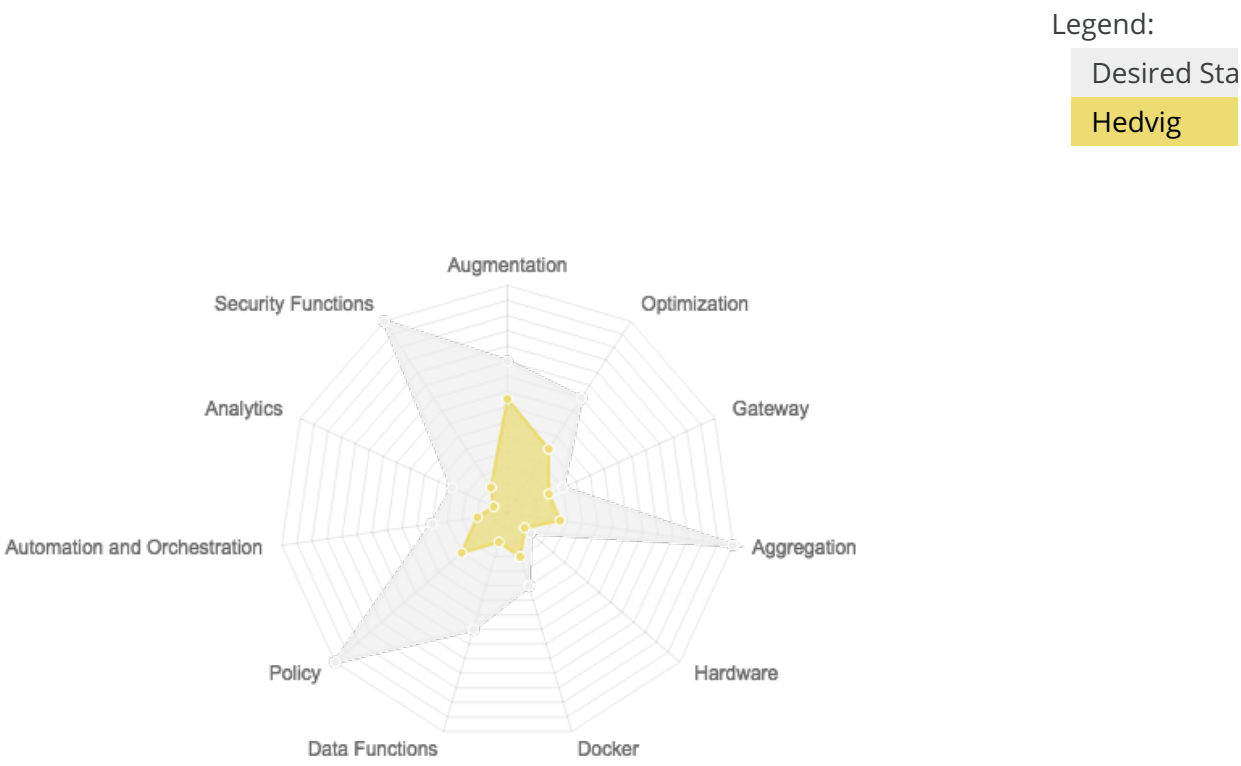

II. Terminology

Terms used in the multi-axis graphs are defined as follows:

Augmentation - Augmentation is the addition of data services and features to existing storage that may not already have those services. In essence, if deduplication, for example, is not available, then we can augment the storage by adding it in the data path. We could also add data path optimization and other services.

Optimization - SDS fabrics optimize existing storage. Optimization includes read and write caching, write coalescing, auto-tiering (optimizing for yesterday's workload today), and other elements that improve the overall quality of service (QoS) and performance of the storage subsystems.

Gateway - SDS fabrics provide gateways to other services such as cloud storage as well as translation between disparate protocols. For example, translating from iSCSI protocols to Amazon S3 and object storage, in the case of a cloud gateway.

Aggregation - Aggregation is the combining of other storage services into one control plane. It can also combine unlike storage into a larger storage pseudo device. Aggregation ingests various storage protocols (iSCSI, NFS, FC, SMB, Memory, Object, SCSI) and presents as another storage protocol (usually iSCSI, NFS, or SMB).

Hardware - Where we look at specific hardware requirements for the SDS fabric.

Docker - Where we look at Docker or container support within the SDS fabric

Data Functions - An SDS fabric adds higher functionality to the storage environment, such as automated data motion for migrating workloads, tiering, and achieving data locality. Data locality involves moving data closer to compute to reduce latency, and other related processes. Additional functionalities are data inspection, data tiering, and adaptive optimization or micro-tiering.

Policy - Software defined storage requires various policies to control the delivery of the storage to meet specific security, delivery, protection, performance, and other goals.

Automation and Orchestration - Software defined storage makes use of automation and orchestration to improve delivery of storage services. We look at what forms of automation and orchestration is available.

Analytics - Storage and workload analytics ensure that storage objectives and goals are being met with some level of status reporting, alerting, and adaptive actions, such as calling automation and orchestration tools to remediate the situation. In general, analytics should be looking not at the data, but at the underlying storage itself. Looking within the data for patterns, transforming the data, etc., would be considered data functions.

Security Functions - Security, which these days is much more than just encryption. We extend security to include layered role-based access controls, namespace control, encryption in motion, encryption at rest, and use of a key manager, as well as data tokenization and investigation.

III. Conclusion

The graphs show that the available tools do not completely address our ideal requirements. This is intentional on our part—our requirements and categories represent what we believe all SDS tools should address as they enter the realm of the next generation of IT. There is a growing need for SDS to provide even more analytics, and data functions. SDS becomes the key to keeping storage in check.

SDS in the hybrid cloud is mostly about presenting data functions across distance but it is also about knowing where our data resides at all times. We are just beginning to address this need.

IV. About TVP Strategy

Established in 2009, TVP Strategy is an analyst organization focused on analyzing, researching, and reporting on the virtualization and cloud trends that impact businesses today. Our analysts all “come from the trenches” and therefore have unique insight into how technology can deliver business value. We also have the battle scars that show we know how not to do things. This experience means we provide our clients with a variety of analytical content that can really move the needle for their business.

Edward L. Haletky, aka Texiwill, is an analyst, author, architect, technologist, and out-of-the-box thinker. As an analyst, Edward looks at all things IoT, big data, cloud, security, and DevOps. As an author, he has written about virtualization and security. As an architect, Edward creates peer-reviewed reference architectures for hybrid cloud, cloud-native applications, and many other aspects of the modern business. As a technologist, Edward creates code prototypes for parts of those architectures. Edward is solving today's problems in an implementable fashion.