SPEED, EFFICIENCY, & SCALE FOR DATA-INTENSIVE WORKLOADS

CISCO'S MULTI-PRONGED STRATEGY IS DESIGNED TO ADDRESS EVOLVING CUSTOMER NEEDS

EXECUTIVE SUMMARY

The way companies think about data is changing dramatically—from exponential increases in the amount of available data to the way in which it is structured and used. How effectively a company turns its data into actionable insights is now seen as a competitive differentiator. A new class of data-intensive applications has emerged to allow organizations to effectively process, analyze, and store all of their available data. Data-intensive applications require a paradigm shift in the way the industry thinks about storage, because traditional storage solutions do not provide the flexibility, scalability, or performance to align with these needs. Leading vendors are responding with new solutions to align with this new generation of workloads, a range of hybrid cloud deployment options, and enhancements in intelligent systems management capabilities.

Cisco has developed a multi-pronged strategy to target the data-intensive solution market. The company is extending its UCS product line and UCS management framework with scale out storage solutions for data-intensive computing that use Cisco UCS S-Series Storage Servers. For companies looking to deploy hybrid clouds for their data-intensive workloads, Cisco offers Cisco CloudCenter, a hybrid cloud management platform that provisions infrastructure resources and deploys applications to on premise datacenters, private clouds, and public cloud environments. In addition, Cisco is investing in software, ecosystem partnerships, and go-to-market capabilities to bring optimized data-intensive workload solution stacks to market. As the market demand for data-intensive applications evolves and expands, Moor Insights & Strategy (MI&S) believes Cisco will continue to enhance this strategy with more solution stacks, product enhancements, and additional systems management capabilities.

MARKET TRENDS

The confluence of mobile applications, cloud computing, and the Internet of Things (IoT) is causing massive increases in the volume of data that needs to be moved, processed, and stored by IT. While there is much uncertainty in growth expectations for data processing and analytics over the coming years, IT organizations must be prepared to scale quickly as demands evolve and grow over time.

Copyright ©2016 Moor Insights & Strategy

Not only is the volume of data exploding, but the data itself, sources for data, and usage models are transforming. Below are a few major changes that are happening today.

- **More Unstructured Data**: By many estimates, the majority of data being generated today is considered unstructured, meaning a schema does not exist or is hidden. Traditional file and block-based storage systems are based on a hierarchical structure that can limit their overall scalability and impact performance for unstructured data sets. Demand for object-based storage approaches have risen to address the scaling needs of unstructured data sets like static web content, multimedia, data backups, and archival images.
- **More Data Sources**: As the world becomes more data-driven, organizations must prepare for data to be in a constant state of motion. Data may originate at the "edge" of the network (from the sea of emerging IoT smart devices) or be provided by an independent business partner. Regardless of the source, data analysis will likely happen in several layers and at multiple locations.
- **More "Active Data"**: Intelligent analytics has become central to business success. A company's ability to turn available data into actionable insight is now a competitive differentiator. With this approach in mind, organizations are transitioning from data management strategies focused on infrequent access of warehoused information ("cold", archival data) to models that rely on continuous analysis and action (with "warm / hot" or active data). Applications that operate in this manner are the ones providing analytics, insight, and security in real time. As an example, video surveillance operated on a capture / store / retrieve model in the past, but it now is evolving to include real-time analysis for facial recognition or retail traffic analysis. With more applications creating and requiring "active data" than ever before, wholesale increases in the effectiveness (cost / performance) of supporting storage platforms are required.

The data explosion and desire for intelligent data analytics are driving a new class of applications that requires data to be closer to compute resources. Some refer to this class of workloads and the requirements it drives on the underlying infrastructure as **data-intensive computing**.

For data-intensive computing to be effective, MI&S believes IT can no longer rely on traditional storage products that require very specific expertise and experience. The future must embrace a decentralized shared-nothing architecture based on flexible building blocks that can be managed using a common set of standards.

Copyright ©2016 Moor Insights & Strategy

VENDOR DIRECTIONS TO ADDRESS MARKET NEEDS

As storage needs change over time, many IT organizations look to deploy solutions that provide flexibility to meet the range of needs for data-intensive workloads. However, choosing the right platform approach for each workload can be daunting. There are a wide range of needs and tradeoffs related to both performance (computation, capacity, latency, *etc*.) and cost (per I/O, GB, transaction, *etc*.). Figure 1 depicts the storage solution categories available to meet the needs of data-intensive workloads today.

FIGURE 1: STORAGE SOLUTION APPROACHES TODAY

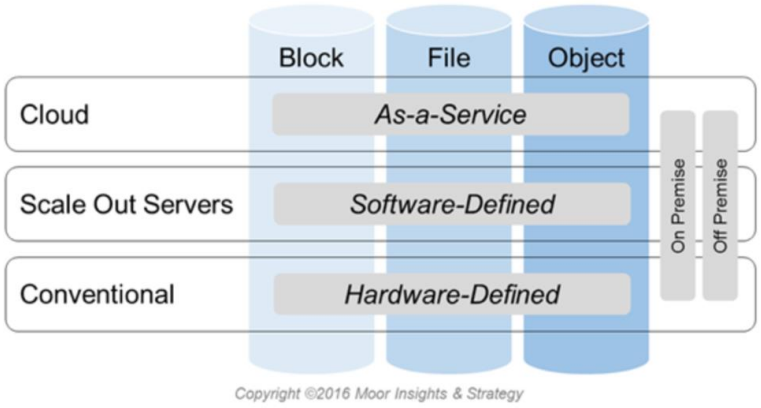

(Source: Moor Insights & Strategy)

As IT organizations evaluate these solutions, we see several notable trends.

- **"Right Sizing" the Solution for the Workload**: As IT evaluates the wide range of storage options available, companies are developing storage strategies that help align the right solution type with each workload and use case. IT organizations are also getting a better understanding of whether the data they are using is hot, warm, or cold to determine how quickly and often data needs to be accessed. For example, high-value, data-intensive workloads that require frequent data access may require solutions optimized for performance metrics such as compute capability, latency, or storage throughput. Other workloads that require lower performance and need a significant amount of capacity for infrequently accessed-data can be optimized for lower cost and improved efficiency. An optimal system design provides the right balance of CPU performance, number of CPU cores, amount of system memory, network latency and bandwidth, and storage capacity for a specific workload. Table 1 provides some general workload examples and areas that require particular attention for a balanced system design.

Copyright ©2016 Moor Insights & Strategy

TABLE 1: DATA-INTENSIVE WORKLOADS & CHARACTERISTICS

Workload Type	Notable Workload Characteristics
Software-Defined Storage	Solutions vary widely. Storage feature / function requires significant processing. High capacity HDDs provide capacity combined with SSD(s) for reduced latency. External network bandwidth must match system capability.
Data Protection (BURA)	Varied complexity from simple flat file to complex snapshot catalog. Flat file schemes depend on number / capacity of HDDs with sufficient network bandwidth. Limited processing requirements.
Big Data (Hadoop, Log Analytics)	Parallel processing with distributed data. Maximum CPU threads (sockets & cores). Data split among HDDs. Modulo split size for distributed file system.
Content Distribution	Receives content & redistributes as needed. Network ingress / egress requires balance. Maximize HDDs for capacity with SSDs for caching frequent titles. Data translation or transcoding (if needed) requires additional processing.
Video Analytics	Maximum CPU threads (sockets & cores). In-memory data requires maximum memory. CPU / HDDs make up distributed file system.

- **Hybrid Cloud Adoption**: As IT organizations look to "right size" their storage resources, many look to adopt a hybrid cloud strategy that gives them the flexibility to choose the optimal solution for each workload. The public cloud provides an opportunity to increase flexibility and quickly scale as demands require. However, IT should also seek solutions that provide the flexibility to move data to and from the public cloud as needed. In addition, organizations should understand their data sovereignty and compliance requirements and ensure their cloud resources meet these needs.
- **New Platform Choices**: IT has more choices than ever before for storage solutions. Within the traditional SAN space, hybrid or all flash arrays have gained traction as cost/GB of flash has come down closer to spinning media. In addition, there has been a steady move away from traditional SAN products for secondary or warm data to software-defined storage (SDS). SDS unbundles storage services from the actual hardware by using storage-optimized server hardware based on industry standard x86 architecture combined with software-defined storage stacks. For the appropriate workloads, SDS has the potential to improve scalability, lower costs, reduce latency, and streamline operations. Increasing numbers of scale out server products targeted at data-intensive workloads are designed specifically to be optimized for cost/GB. While the market for software-defined storage is small today, we expect significant growth in storage array capacity in enterprise datacenters deployed with SDS or hyperconverged system architectures based on x86 hardware over the coming years.
- **Unified Intelligent Systems Management**: As the world of IT becomes increasingly complex, there is a desire to break down system management silos between storage, servers, and networking platforms. In addition, as IT looks to

Copyright ©2016 Moor Insights & Strategy

implement more sophisticated data tiering capabilities, they desire a consistent management framework for all data whether it is hot, warm, or cold. System management solutions are becoming more intelligent and include capabilities such as embedded automation and orchestration, logical device consolidation, and zero-touch configuration for adds / moves / changes.

Vendors that provide storage solutions to the market can no longer operate in silos. As IT organizations seek to move to a hybrid cloud model, they are looking to the storage ecosystem to provide the path for deployment. Integrated storage strategies must encompass each of the areas shown in Figure 1 above.

CISCO'S STRATEGY FOR DATA-INTENSIVE COMPUTING

In 2009, Cisco introduced its UCS product line targeted at simplifying the technology silos necessary to support virtualization. As a result, UCS has enjoyed considerable market success by simplifying and automating tasks notorious for complicating and slowing solution deployment. Cisco is extending the UCS management framework to a set of scale out storage solutions for data-intensive computing that use Cisco UCS S-Series Storage Servers. By leveraging the simplicity of UCS Manager, Cisco's goal with these solutions is to allow its customers to "scale to petabytes within minutes" and to experience "zero-touch, policy driven management" for ongoing operations. In addition, Cisco recently entered the hyperconverged solution space with its HyperFlex product family. Hyperconverged solutions are designed to bring web scale economics to enterprise datacenters and private cloud. The hyperconverged line is complementary to scale out secondary storage, as both are designed to improve economies of scale.

Modern workloads require wide variation in the ratio of compute, storage, and I/O, which makes unique form factors with limited modularity an inefficient choice. Cisco is taking a modular approach for its data-intensive solutions to align with the range of data and storage needs for target workloads. In addition to the UCS product expansion, Cisco is helping enable hybrid clouds through Cisco CloudCenter (via its acquisition of CliQr in April 2016). CloudCenter is a hybrid cloud management platform that provisions infrastructure resources and deploys applications to on premise datacenters, private clouds, and public cloud environments. Also, Cisco is investing in storage ecosystem partnerships including a system validation program with partners. For example, Pure Storage and Cisco have developed a Cisco Validated Design for FlashStack (Cisco UCS servers with Pure Storage flash arrays) with VMware Horizon for Enterprise VDI deployments

Copyright ©2016 Moor Insights & Strategy

CISCO UCS S3260 STORAGE SERVER OVERVIEW

The Cisco UCS S3260 Storage Server is the flagship data-intensive computing and storage platform within the UCS family of products. Unlike many of the novel form factors emerging in the industry, the UCS S3260 is extremely compact at only 32 inches (813mm) deep, which means it can fit in a standard depth 900mm common rack with about 3.5” (87mm) of rear clearance for cabling and cooling. Figure 2 provides an overview of the system attributes.

FIGURE 2: CISCO UCS 3000 SOLUTION ATTRIBUTES

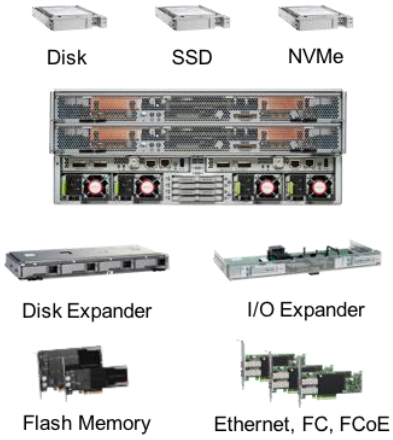

(Source: Cisco)

Key attributes that position this product for data-intensive workloads include:

- **Modularity**: The UCS S3260 is a modular design with disaggregated compute, network, and storage resources. This design allows IT to refresh each subsystem independently for multi-generational investment protection. The UCS S3260 is positioned for warm-tier storage for unstructured data and archiving. Target workloads include software-defined storage, object store, media streaming, content distribution, video surveillance, email, BURA, and big data analytics.
- **Flexibility**: The modular design of the UCS S3260 makes the product suitable for both traditional and emerging scale out data-intensive workloads, essential for eliminating unique storage islands. Cisco has outlined a number of optimized configurations with varying node counts, acceleration technologies, storage media options, and networking capabilities to meet the needs of different data-intensive workloads. For capacity intensive workloads, the UCS S3260 can deliver over half a petabyte (600TB) of total raw data storage in a single chassis and up to 4.8 petabytes per rack. For applications that require high performance,

Copyright ©2016 Moor Insights & Strategy

two complete drive rows (28 of the 56 available LFF slots) can be outfitted with solid-state storage (SSD) for caching or data tiering. The latest server nodes will include Intel Broadwell CPUs available in one node or two node configurations with support for NVMe. Cisco also offers expansion card options. One is an optional disk expander card that provides an additional 4 LFF slots to achieve its 600TB total capacity. Another is an optional I/O expander card that enables a variety of unified connectivity for Ethernet, Fibre Channel, or FCoE to bridge disparate storage systems. The I/O expander card can also accept third-party flash memory cards for application acceleration of direct-attached local caching.

- **Simplified Management**: IT organizations are looking for ways to eliminate silos by managing their infrastructure using a "single pane of glass". The UCS S3260 uses the Cisco UCS Manager for system management of blade, rack mount, or storage optimized servers and the networking that connects them all. Leveraging the same system management plane is beneficial for existing UCS bare metal, converged infrastructure, or hyperconverged infrastructure customers who are already using UCS Manager in their environments. S-Series could be used for seamless addition of tiered storage under the same systems management umbrella. Cisco UCS Manager includes a full API based management framework that provides access, identity, and configuration management capabilities.

Potential TCO Benefits of UCS S-Series Over Public Cloud

Cisco has developed a Total Cost of Ownership (TCO) model to demonstrate potential cost savings of the Cisco S-Series compared to the Amazon S3 public cloud offering. The example in Figure 3 illustrates the 3-year TCO of the UCS S3260 in an on premises datacenter versus Amazon S3 for a workload that requires 420TB of usable hot storage. Results will vary for each workload and customer; this example shows the Cisco S-Series solution providing a 56% TCO advantage with a breakeven period of 14 months.

Figure 3: TCO Comparison: Cisco UCS S-Series vs Amazon S3

(Source: Cisco)

Copyright ©2016 Moor Insights & Strategy

FUTURE MARKET DIRECTION & POTENTIAL OPPORTUNITIES

Over the next several years, MI&S expects to see more advanced sophistication in data-intensive computing workloads. One of the key areas of enhancement will be increased use of tiered analytics models that drive demand for analytics capabilities from edge devices all the way through to the cloud. To meet the various performance and cost needs of next generation workloads, we expect to see an increased use of flash in enterprise-class storage solutions along with the introduction of products that include storage-class memory. However, hybrid media strategies will continue to be commonplace for the next several years, as HDDs will be used to meet the cost and capacity requirements for most workloads. We also expect to see continued increases in the adoption of hybrid cloud strategies for storage and a desire from IT to have a unified way to manage both public and private cloud resources.

Cisco is a relative newcomer to the data-intensive solutions market, and others are likely further along in terms of feature sets and software optimizations. We expect to see Cisco round out its data-intensive product portfolio over the next several years with additional offerings in the Cisco UCS S-Series family. As the market evolves, Cisco should look to develop solutions for more advanced big data analytics applications like Apache Spark. Cisco should also provide new products that meet next generation requirements like support for “all-flash” configurations and storage class memory.

The modular design approach for the UCS S-Series means that Cisco can upgrade the compute, storage, and networking capabilities on separate paths as technology enhancements become available in each area. In the future, we expect to see additional processor upgrades, additional acceleration options, and new storage capabilities. Cisco UCS Manager is continuously evolving; over time, we expect to see additional enhancements at the management layer, such as more intelligent automation features.

CALL TO ACTION

IT organizations are in a state of transition as they look to modernize their infrastructure to keep up with the changing needs of data-centric workloads. Many are looking to deploy hybrid cloud strategies that combine traditional storage platforms, server-based storage, and cloud-based storage. Using a combination of solutions allows IT to have the flexibility to meet the needs of each workload.

Cisco is investing in platforms, software, ecosystem partnerships, and go-to-market capabilities to respond to the market demands for data-intensive computing. The Cisco

Copyright ©2016 Moor Insights & Strategy

UCS S-Series uses Cisco UCS Manager for systems management; this makes it a good match for existing UCS customers looking to deploy data-intensive computing solutions under the same systems management umbrella as their other UCS workloads. Cisco is also investing in software ecosystem partnerships in the storage space and helping customers deploy and manage hybrid clouds with Cisco CloudCenter. Cisco and non-Cisco customers should talk to Cisco about their plans for data-intensive computing and understand how their strategies align with specific application needs over the long term.

Copyright ©2016 Moor Insights & Strategy

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR

Gina Longoria, Senior Analyst at Moor Insights & Strategy
Jimmy Pike, Technologist at Moor Insights & Strategy

PUBLISHER

Patrick Moorhead, President & Principal Analyst at Moor Insights & Strategy

EDITOR / DESIGN

Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES

Please contact us here if you would like to discuss this report, and Moor Insights & Strategy will promptly respond.

CITATIONS

This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

Cisco is a research client of Moor Insights & Strategy and this paper was commissioned by Cisco. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2016 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.