

Disaggregation: a long word for a simple and key concept


The 451 Take

Over the last two to three years, disaggregation has become a hot topic in multiple areas of IT, including data storage. However, the concept is far from new, and it has been a foundation of datacenter architectures for many years. The term means the separation of infrastructure resources into independent layers or domains, and a major example is the disentanglement of data storage from compute. In datacenters, that process began about 20 years ago when storage inside servers began to be replaced with external storage systems, connected to the servers via storage area networks (SANs). This disaggregation of storage and compute offered advantages in both cost and operational flexibility, and as a result, it quickly became the dominant approach to datacenter storage.

However, this approach could not be applied to all workloads because external storage could not meet the needs of some applications that required extremely fast access to data – even when the storage systems were powered entirely by flash. Fast-forward to the present and two important developments that have taken place. The first is an increase in the number of those performance-sensitive applications in areas including but not restricted to AI and machine learning and analytics. The second is the emergence of external storage systems that deliver the performance needed by those applications. Based in part on NVMe and NVMe over Fabric (NVMe-oF) protocols, these systems are bringing the benefits of disaggregation to workloads that were previously left out in the cold, able only to use server-based storage. And by doing that, they are also bringing the concept of disaggregation back into the limelight.

On-Premises Infrastructure Attributes to be Improved to Meet the Needs of Customers Source: 451 Research's Voice of the Enterprise: Storage, Workloads and Key Projects 2018

Q. Which of the following attributes of your on-premises infrastructure are you looking to improve to meet the current and future needs of your customers? Please select all that apply.



451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 120 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

Business Impact

DISAGGREGATION IS A FAR-RANGING VIRTUE. The benefits of moving data out of servers and into external storage systems include dramatically simplified and seamless operations and more efficient use of resources, resulting in reduced capital spending. However, these benefits do not just apply to the disaggregation of storage from compute. Server vendors are currently promising to disaggregate memory from server CPU because that will allow memory to be used more efficiently and flexibly.

STORAGE DISAGGREGATION REDUCES CAPEX. One of the drawbacks of using storage within servers is that it is difficult to enable one server to share its storage capacity with another. This results in trapped storage, unused capacity and wasted capital spending. Server-based storage can also cause unnecessary spending when there is a need to boost either server processing power or total storage capacity because it limits the ability to alter those two resources independently. Since NVMe is one of the most expensive types of flash storage, the issue is compounded further. Disaggregated storage avoids both of these problems.

STORAGE DISAGGREGATION REDUCES OPEX. The data in external storage systems can be shared or accessed by multiple applications running on different servers, and this ability can be of great value for multiple types of applications. Disaggregated storage also significantly simplifies the implementation and management of data protection and disaster recovery mechanisms designed to provide resilience and handle hardware failures. It does this by eliminating multiple silos of server-based storage, for which those mechanisms must be managed independently.

Looking Ahead

Until recently, the performance needs of many applications could not be met by external or disaggregated storage because of the networks that are needed to connect servers to those storage systems. The networks impose overhead that reduces the speed of access to data. When that causes unacceptably low performance, IT organizations are forced to store data within servers, where it can be directly accessed by the servers with no intervening networks.

The NVMe protocol is changing this situation by making stand-alone storage systems faster, which helps compensate for the network impact on overall performance. NVMe replaces the slower SAS and SATA protocols that were developed during the disk era, and until now were the dominant interfaces used within storage systems. However, NVMe does not directly address the problem of the network impact on performance. As a result, many applications still need to store data within servers, on NVMe flash drives. Enter NVMe-oF, which is a networked version of NVMe and a more recent development. When NVMe-oF is used in the networks connecting servers to storage systems, it eliminates virtually all of the network overhead, eliminating the performance penalty of using stand-alone storage.

Incumbent storage vendors are modifying their existing storage systems to work with NVMe and NVMe-oF. However, there are challenges in doing this, and several leading-edge vendors argue that retrofitted, legacy systems cannot fully unleash NVMe performance because they are based on architectures that date back to the disk era. Instead, these independent suppliers have designed systems from the ground up to work with NVMe and NVMe-oF, with the aim of eliminating architectural bottlenecks and, therefore, fully unlocking the performance benefits of the two protocols. As adoption of these purpose-designed NVMe storage systems grows, the virtues of storage disaggregation will spread to high-performance applications that until now have been forced to suffer all the disadvantages of local or server-based storage.



Disaggregation is ushering in a 2nd wave of NVMe adoption where radical performance impact with cost savings cannot be ignored. The parallels of NVMe-oF to SAN adoption in the late 1990s are uncanny. Pavilion Data is leading this 2nd wave with a storage platform designed specifically to scale linearly with an open choice of NVMe drives and fabrics. Organizations worldwide are catching this wave with Pavilion Data. Surf's up! To learn more, please visit www.paviliondata.com.