REAL-WORLD HYPERSCALED ENTERPRISE STORAGE

Global bank builds its own hyperscaled storage. By the SNIA Cloud Storage Initiative and the Data Protection and Capacity Optimization Committee.

TRADITIONALLY, enterprises have purchased highly available storage systems with built-in redundancy and dedicated storage controllers with proprietary firmware to meet enterprise storage demands, but these solutions are complex and expensive and growth of these types of systems has slowed. Some of the data is obviously moving to the cloud and that seems to be growing significantly. But few (Netflix mainly) have made the transition to total cloud-based IT. They principally cannot get all the services they need from today's public cloud.

SNIA research indicates that many of these enterprises, in addition to utilizing cloud storage, are also building their own storage systems using software defined storage (SDS) and best-in-class commodity components, assembled in racks. These enterprises are adopting methods from hyperscalers such as Amazon, Facebook, Google and Microsoft Azure to build their own storage systems while achieving the required service levels for internal projects. These private cloud storage systems are used to host the data for business critical applications and comply with regulations from multiple government jurisdictions.

Fortune 100 global bank becomes a hyperscaler

Large enterprises still building their own datacenters may be increasing their use of the public cloud, but they also have requirements which are difficult to meet with the services currently available. Banking institutions in particular are subject to many audits from multiple jurisdictions of regulators. They essentially create their own private storage cloud to ensure regulatory

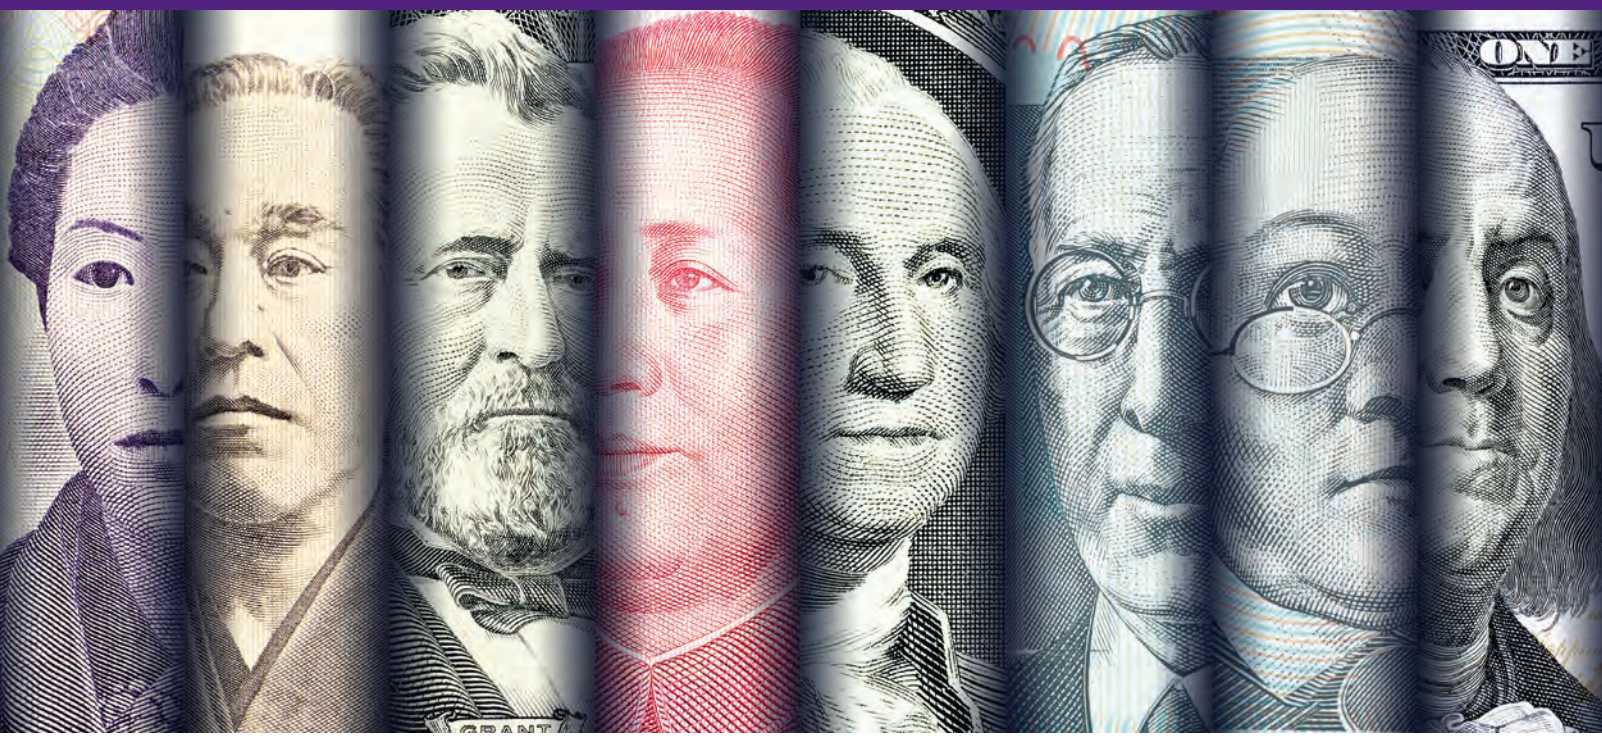

compliance with the many and varied requirements each government mandates. Hyperscale storage is built from best-in-class commodity storage components such as solid state disks and hard drives. These drives are typically delivered populated into racks and purchased as a "pod" with a dozen or more racks making up one storage "system." Around half of the racks are populated with server "heads" to perform the software-defined storage (SDS) in a scale-out manner. The disk trays are JBODs – just a bunch of disks, and the SSD trays are JBOFs – just a bunch of flash. Another technique is essentially densely populated compute nodes where all the storage is directly attached.

The SDS and other custom built software handle the resiliency by replicating or using erasure code across the datacenter and geographically between datacenters. Performance scales almost linearly at least within each pod. They also have a "fail in place" policy, letting failed components exist powered down in the datacenter until a certain percentage of trays or racks is unusable, and only then replacing these larger units.

In discussions with an IT vice president of a large multinational bank, SNIA has learned something of the scale and operation of their internal infrastructure. It matches that of the large Internet hyperscalers both in size and by using the same procurement and operation techniques.

The bank has over 20 datacenters around the world. They cannot use the public cloud as they currently need to comply with over 200 government regulations from different countries. Their solution was to create an internal private cloud for the entire bank's IT project usage. Their storage budget dwarfs the revenue of most medium sized storage vendors.

The bank's deployed storage has 10s of thousands of nodes with around 200 petabytes of active data and half an exabyte of inactive data. Their overall data footprint is growing at 45% annually. They process trillions of transactions daily, and downtime is very expensive.

The bank is big enough that vendors will custom build for them, but they also have a policy of no single source for any of their hardware. They buy storage in 6 PB pods that are half CPU and half storage drives, pre-assembled and sold by the original design manufacturer. They installed their first such pod in 2015. Their next pod will be all flash. Most importantly their cost savings is projected to be 50% over traditional storage.

The Role of software-defined storage

This institution uses SDS with the best-in-class commodity hardware to create a private cloud for internal IT projects and customer facing services. They are currently deploying about 11% of their storage as SDS today, but have a goal to grow this to 50% by 2020. They license their SDS from a major vendor (a site license) and are training up their staff in the new approach. The bank virtualizes the hardware to abstract away any differences between the multiple vendors. From the storage service they serve up an S3 compatible interface for new projects and mainly provide block services for existing applications. They plan to look at Ceph for SDS and non-volatile memory express for solid state disk interfaces in the future.

The hyperscale enterprise storage trend

The trend to build hyperscale storage infrastructures involves utilizing best-in-class commodity hardware and providing all the differentiation in the software to manage all aspects of the storage. Thus, all the "intelligence" is in the software, such that the infrastructure behavior is programmable via that software. Look for platforms running intelligent software on general-purpose hardware in the future. This is a trend enterprises should be closely following. Want to learn more about hyperscaled enterprise storage and this Fortune 100 global bank's implementation? Download the SNIA white paper, "Hyperscaled Enterprise Storage" at: **http://www.snia.org/hyperscaled**

About SNIA

The Storage Networking Industry Association (SNIA) is a non-profit organization made up of member companies spanning information technology.

A globally recognized and trusted authority, SNIA's mission is to lead the storage industry in developing and promoting vendor-neutral architectures, standards and educational services that facilitate the efficient management, movement and security of information.

For more information, visit: **www.snia.org**