



White Paper

Thin Provisioning Data Services with Silk

Thin provisioning of disk capacity has been a valuable enterprise data services feature that IT teams have been using for years in their on-prem data centers to simplify application and database deployments and to gain greater efficiencies in resource utilization. But today's teams are struggling when they try to move some of those workloads to the public cloud. In many cases, migrating a VM from on-premise to the cloud requires configuring the same capacity footprint in the cloud as was provisioned on-prem. This can cause significant capacity overprovisioning in the cloud since much of that provisioned capacity is almost always not actually being used. In the cloud, you will pay for every byte of provisioned capacity regardless of the true level of usage, or what is known as "allocated capacity". Provisioned capacity is the total amount of capacity that you purchase up front for your VMs. Allocated capacity is what the VM actually writes onto that capacity. But the cloud does not differentiate between provisioned and allocated capacity when charging for usage—you always pay the full provisioned amount, regardless of allocation levels. In the cloud, overprovisioning results in pure wasted resources that you always pay for but never use.

One could argue that every VM deployed in the cloud is somewhat overprovisioned as there is always some amount of capacity allocated to it that is not being fully used. This may not be much of a concern if you only have a few VMs but could get quite costly if you have hundreds or even thousands of VMs, especially if each has a good amount of overprovisioning. Cost is not the only challenge that exists when provisioning capacity in the cloud.

Resource availability can also be a challenge since not all regions and zones have the same resources available in them. Remember: underlying cloud infrastructure is architected to be a shared environment. While the individual virtual resources (CPU, Network and Disk) that you provision for your usage are not shared, the underlying hardware they are running on is heavily shared, usually with many, many other clients. What resources that are available one day may not be there the next and vice versa. It is possible that you might need to quickly scale up your capacity for a mission-critical database but not be able to due to resources not being available in that zone or region. Some cloud users have been forced to scale up their VMs to a larger instance type just to support adding additional capacity. This will negatively impact costs even further since you are then paying for even more resources you do not need and are not going to use, and it could even increase your database licensing costs since you will be using more vCPUs. These are considerations that users did not need to worry as much about when their applications were on-premise since they were using a private shared storage platform that provided all the great benefits of a shared resource architecture (something very powerful and cost effective) but that also was dedicated for use to just that single set of users or a single company (a private cloud, actually)—which is something the public cloud providers do not offer today as a standard offering.

The Silk Cloud Platform

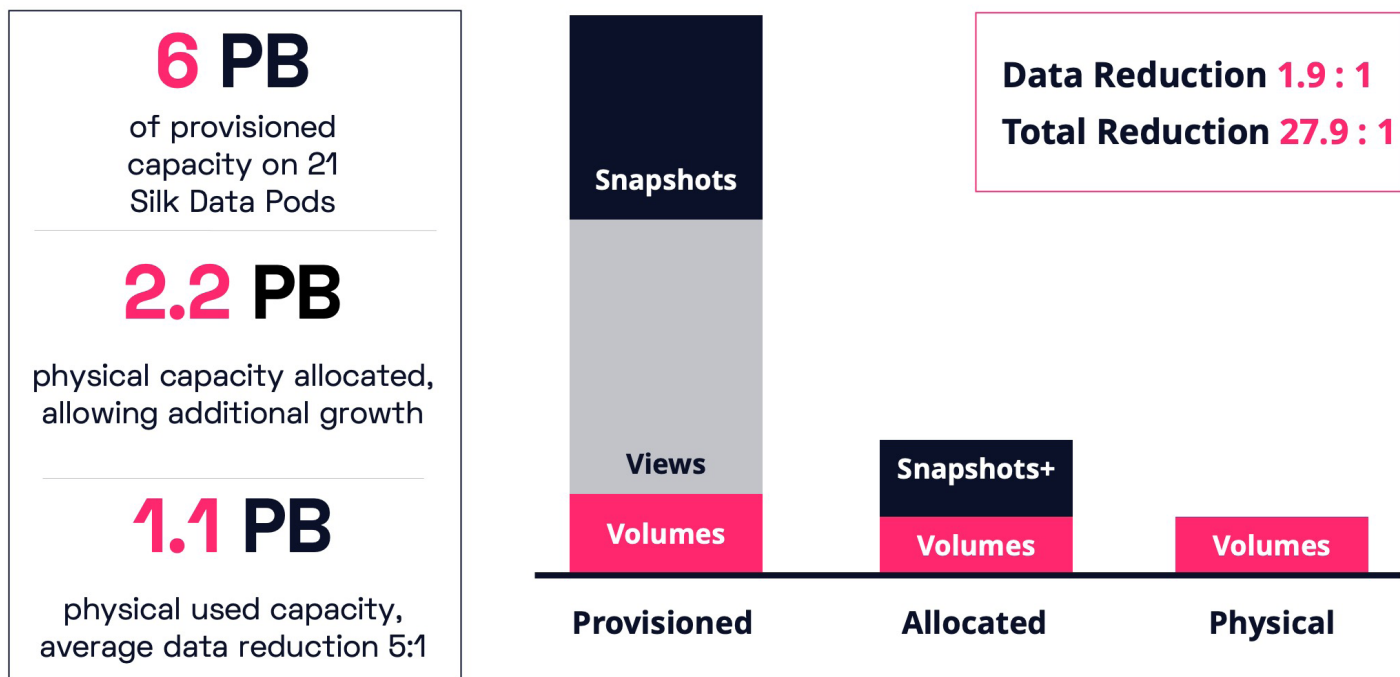
Silk provides a virtualized cloud platform whose resources and data services can be shared by multiple VMs. This platform is decoupled from the performance and capacity limitations the cloud imposes with different VM types and allows users to maximize the vCPU and network performance of each VM without having to overprovision anything. Users can provision multiple volumes of any size and map them to any VM type regardless of its vCPU count or memory footprint. When user provision volumes on Silk, those volumes are always thin-provisioned. The volume size can be any size up to the maximum allowed by the particular version of operating system being utilized (often 64TB or even larger). However, capacity is only allocated on a volume when it has data being written to it. This has the net effect of users being able to provision literally hundreds of Terabytes and even Petabytes of capacity without having to incur any additional expenses. Users can provision all the “thin” capacity they need for their application’s requirements without having to purchase any more cloud resources. The Silk Cloud Platform also makes it very easy to dynamically scale up or down the capacity of a volume with just a couple of UI clicks or via API-driven automation tools. Using thin provisioning can massively reduce physical capacity usage, typically by 50-60%, but 90% reduction is not uncommon. No more paying for overprovisioned cloud resources that are not being used—including vCPUs, memory, networking and storage!

The Silk platform provides inline data compression, deduplication, and pattern removal without impacting performance. These additional services can significantly reduce your cloud data footprint, with a 50%-75% reduction being most common. The combined effect of thin provisioning and inline data reduction can result in a 10x, 20x or more increase in effective usable capacity.

Another advantage of the Silk Cloud Platform’s enterprise data services is inline data reduction -- something else the cloud vendors don’t offer. This has been a common feature in on-premise systems for years. Losing these powerful data services features when applications and databases get migrated to the cloud results in many users seeing significant increases in their capacity provisioning footprints compared to the on-prem sizing, especially for data sets that were benefitting from high data reduction ratios. The result: grossly overprovisioned resources and higher-than-expected cloud costs. The Silk platform provides inline data compression, deduplication, and pattern removal without impacting performance. These additional services can significantly reduce your cloud data footprint, with a 50%-75% reduction being most common. The combined effect of thin provisioning and inline data reduction can result in a 10x, 20x or more increase in effective usable capacity.

In the example on the next page, Silk achieved a 2800% increase in usable capacity, delivering major cost and efficiency benefits, not to mention the significantly upgraded agility and scalability capabilities that go along with this compared to cloud native alone. It should also be briefly mentioned that all clones or snapshots taken on the Silk platform also do not use any real capacity and are “free” of utilization, until those snapshots are turned into writable copies and data is written to them. This means users can take dozens or hundreds of snapshots that take up zero capacity, instantly and for free, with no performance impact, and will only use a small amount of capacity as they are truly used and being written to, benefitting from the same inline data reduction as normal volumes do. The overall benefit can be extremely large, especially for database or app dev shops who need to iterate continually against multiple copies or who need to refresh copies of large datasets frequently and quickly in the course of normal operations. This is extremely beneficial for those users.

Data Reduction



An example of what is possible is shown here, where a furniture and home goods e-commerce company had a corporate objective and a strict deadline to get all of its data out of their on-premise data centers. The IT teams started migrating their dozens of Microsoft SQL Server hosts to the cloud but kept hitting throughput limitations no matter how many cloud resources they assigned to them. With the deadline bearing down, they didn't have time to refactor all of their SQL databases. They needed a way to get faster performance and streamline the deployment of their cloud resources with a common configuration for each SQL instance rather than every instance being customized based on peak workload requirements.

With Silk, they were easily able to exceed their throughput requirements and get better performance with even lower latencies than their previous on-premise solution because of Silk's ability to dynamically scale performance as needed. More importantly, Silk's enterprise data services -- such as thin provisioning, compression and zero footprint instantaneous snapshots -- allowed the company to standardize their mission critical SQL server deployments without having to worry about massively overprovisioning cloud resources. This chart shows how much space this company is saving compared to what would have been required with native cloud, making their cloud resources go much further with Silk.

About Silk

Silk gives demanding workloads 10x faster performance on the cloud compared to native cloud alone. The Silk Cloud Platform is a virtualization layer that sits between the underlying cloud infrastructure and customers' workloads. Without refactoring, workloads such as Oracle, Microsoft SQL Server, and industry-specific applications can move onto the GCP and Azure cloud and massively improve user experience. Industry leaders in e-commerce, software publishing, FinTech, and healthcare, trust Silk with their mission-critical workloads to get the ultra-fast speeds their customer's demand. Silk is headquartered in Needham, MA. To learn more, visit silk.us.