VMware Virtual SAN and VMware Horizon 6: Cost Leadership Through Storage Savings

WHITE PAPER

vmware®

Table of Contents

Innovations from VMware have helped lower the cost of virtual desktop infrastructure (VDI) every year since 2012. VMware business case data[1] shows that the average annual cost of VDI per user for a typical 2,500-seat deployment in 2012—hardware, software, and labor—was approximately $595. With the introduction of VMware Horizon 6® Enterprise Edition, the average annual cost for the same deployment was reduced to $470. Most of these savings are a result of reductions in the cost and capacity of the hardware required. And VMware continues to drive innovation to bring down these costs even further.

Lowering Storage Costs with VMware Virtual SAN

Horizon 6, released in 2014, represented a significant milestone in the VDI cost journey. For the first time, the per-user capital cost of VDI fell below the cost of physical desktops. Storage contributed the primary cost savings, with VMware Virtual SAN™ delivering the long-term promise of low-cost clustered storage. Virtual SAN extends the power of virtualization to the storage layer by using industry-standard servers and attached commodity storage to enable scaling up and out with fine-grained adjustability. This combination of readily available x86 computing power, solid-state drives (SSD), and hard-disk drives (HDD) maintains linearly scalable I/O performance while reducing the unit cost per GB in comparison to previous network storage approaches. The savings are dramatic: For the same deployment model—2,500 users, each with a 30 GB image and 5 GB of data—we estimate that storage costs with Virtual SAN fell from around $140 per user to slightly over $50 per user.

Greater proximity of storage to servers and lower latency with direct-to-hypervisor I/O path (through Virtual SAN) improved performance, and consequently, increased the number of users each server could support. Proportionate reductions in related operational overhead also occurred. Overall, our benchmark 2,500-seat deployment showed an average per-user reduction of more than 25 percent in capital costs and 10 percent in operational costs.[2] VDI now holds cost leadership over physical desktops.

Why Is Virtual SAN Cost Effective?

Virtual SAN is a recent development in the evolution of storage. Conventional network storage (free-standing hardware on which storage can be centrally managed) has been reliable for critical data but at a high CapEx. Scaling traditional storage arrays involves costly stair-step upgrades in storage controllers and disk capacity.

Virtual SAN pools server-attached HDDs and SSDs to create a distributed shared data store that abstracts the storage hardware and groups discrete Virtual SAN nodes into a cluster that can scale up or out. The storage can be of any type and is attached to the server hosting the virtual machines, with the consequent advantages in latency and speed provided by the close proximity of the storage to the CPU. If needed, existing storage arrays can also be attached, enabling reuse.

Many conventional desktop environments use network-attached storage (NAS) rather than SAN, because historically it has been a more cost-effective, although less performant, solution. Recent VDI deployments have often used storage appliances, containing an array of HDDs and SSDs. Such appliances offer SAN levels of performance, but at a higher cost than NAS. Virtual SAN is an important component of the VMware VDI solution because it delivers the benefits of SAN storage at a total cost of ownership (TCO) that is significantly lower than NAS.

Virtual SAN cuts costs by using existing hardware (storage directly attached within the servers hosting the virtual machines) rather than purpose-built storage arrays. As a result, cost is predictable and linear, and support costs are lower. Virtual SAN is included in Horizon 6 Advanced and Enterprise editions, making it effectively free to customers.

1. Business cases built for over 700 VMware customers with their own cost data.
2. These are average figures and intentionally conservative. Many customers achieve even higher levels of savings.

With suitable SSDs added, Virtual SAN delivers performance almost equal to an all-flash storage system at a fraction of the cost. For environments that require the highest possible IOPS capacity, Virtual SAN also offers all-flash configurations that use SSD for both caching and buffering as well as data persistency tiers.

Flexible Storage with Virtual SAN

Virtual SAN offers a simple way to configure your storage to the specific requirements of your VDI users, without purchasing expensive storage arrays. Typically, the enterprise-class servers used to host virtual desktops have direct-attached storage built in that combine SSDs and near-line serial-attached SCSI (SAS) drives. Virtual SAN clusters these into pools of shared storage, which can then be allocated to the virtual desktops in any configuration based on the assigned storage policy and user requirements.

Virtual SAN offers the following advantages over conventional network storage.

- Lower cost from efficient, custom storage tiers.
- Optimized storage to support VDI workloads, which can be integrated and designed around the VMware product suite. For example, see the VMware Virtual SAN 6.0 Design and Sizing Guide.
- Operational simplicity with no new storage management tools to learn. IT admins manage Virtual SAN with the familiar VMware vCenter™ tools that they already use.
- Storage can be dynamically added in small increments, enabling the solution to scale up and out for capacity and performance as needed, with smaller and simpler purchases. Scale up by adding more embedded disks or flash. Scale out by adding more Virtual SAN hosts to the cluster.
- User experience matches, and in our testing often exceeds, conventional hybrid storage arrays, at a much lower TCO.

Storage Tiering

To obtain the full benefit of cost leadership while retaining optimum storage performance, it is important to use tiered storage in your virtual SAN. Storage tiering provides different grades of storage in an environment so that data with a particular access requirement sits on a device at an appropriate price point. Data is migrated from one type of storage to another based on patterns of access, performance requirements, or capacity constraints.

As an example of access patterns, consider a batch of high-resolution images taken 10 years ago and looked at only twice in the past 5 years. It doesn't make sense for these images to consume gigabytes of expensive, high-performance Fibre Channel storage that can be better used for frequently accessed or important data that needs to be read or written as quickly as possible. Instead, slower and less expensive HDD storage is more suitable for old and rarely accessed data. Conversely, an Oracle database accessed every second of the day and with extremely high-performance requirements should utilize a fast hybrid or all-flash storage solution that can accelerate the read and write requests created.

At the other end of the access speed scale is flash storage, or SSDs. These devices work in a fashion similar to memory (rather than rotating disks) and provide very fast data access.

Virtual SAN caches data based on frequency of access, ensuring quicker access to the data most often used. It uses a caching algorithm to store less frequently accessed data on less expensive HDDs. This process happens in the background and requires no intervention from support staff.

Storage Policies

The question of what proportion of data to allocate to less expensive storage is a difficult one. The traditional 80-20 rule—20 percent of data being accessed 80 percent of the time—might apply. However, this proportion is probably too conservative, because desktop users tend to keep old data forever and don't delete it until forced. A policy of ordering users to delete old data is expensive to implement. It consumes users' and support staff's time, and has the unfortunate side effect of decreasing goodwill.

How to split data between HDDs and SSDs depends on each organization's needs. It is best to experiment with a subset of users in a pilot program. If the proportion of SSDs is too low, the user experience suffers, because too high of a proportion of the data access is on the HDDs. If the proportion of SSDs is too high, the user experience is good—and probably better than needed—but the cost is unnecessarily high.

Tiered storage strategies are particularly suited for end-user computing, because its large quantities of rarely accessed data allow extensive use of slow, inexpensive storage. Conversely, application data, which is often stored in databases, uses files that are constantly being updated, making it more difficult to take advantage of tiered storage.

Virtual SAN Costs and Example Configurations

The following examples show how different organizations can implement VDI with Virtual SAN and the estimated costs compared to an equivalent network storage implementation.

We have assumed a storage appliance-type deployment for our comparisons, because these deployments are commonly used and deliver the performance required for VDI. If Fibre Channel SAN storage were used instead, the cost would be higher due to higher cost switches and the need to add host bus adaptors to each server. We have also taken a cautious approach to assessing server requirements, assuming no differences in the number of host or management servers required. As a result, servers are neutral to the comparisons. The calculations focus only on storage cost.[3] All estimates for Virtual SAN comprise failover disks, so that function will be maintained in the case of any single drive failure: Our Virtual SAN estimates are effectively over-provisioned.

In each case, the per-user storage costs are significantly lower using Virtual SAN with savings varying from 58 to 77 percent. The figures also indicate greater savings for more demanding workloads in both absolute and percentage terms. These savings reflect the consistently high performance that Virtual SAN delivers, which is difficult to match by conventional methods except at great expense. Our assumptions for the pricing are included after each example.

Example 1 – Domestic Call Center for a Retail Bank

In this example, the users are task workers carrying out a usually predictable set of functions that place only moderate demands on the infrastructure. The routine nature of the work means only a limited set of applications is required on the desktop. Virtual SAN can be implemented to satisfy these requirements without compromising the user experience.

This scenario consists of 5,000 users, each with a light image load. Nonpersistent images are used, with a 30 GB base image. There are two base images. A total of 31 TB of storage capacity is needed, to support replication of base images across servers. Since the images are nonpersistent, there is no user data storage requirement.

We have assumed 200 desktops are supported per host server. Each server can be fitted with four drives. For the network storage deployment, iSCSI connects to the shared storage array. For the Virtual SAN deployment, each host server houses four drives: two rotating platter and two SSD.

The savings compared to a network storage solution is $22 per user, or 58 percent.

SHARED STORAGE ARRAY COSTS		QUANTITY	COST
	Storage capacity (GB)	31,000	$155,000
	32 ports – 10Gb Ethernet switch	2	$28,978
Total			**$183,978**
Cost per user			**$37**

VIRTUAL SAN COSTS		QUANTITY	COST
	32 ports – 10Gb Ethernet switch	2	$28,978
	600 GB SAS drive	50	$19,950
	200 GB solid-state SAS drive	50	$28,450
Total			**$77,378**
Cost per user			**$15**

Table 1: Cost Comparison – Retail Bank Domestic Call Center

3. In many deployment scenarios, the greater proximity of servers to storage allows an increase in user density (number of users supported per server) due to the higher speeds of access to data and applications. In other deployment scenarios, the number of management servers might increase slightly.

Example 2 – Multisite Engineering Consultancy

In this example, the work involves a wide range of functions, often compute and storage intensive. Support staff perform many routine tasks. The knowledge workers, which comprise the majority of the workforce, use a demanding set of tools, such as business applications and data analysis systems. A wide variety of applications is used with a medium level of desktop customization, so a larger, more flexible virtual SAN setup is needed.

This scenario has 3,000 users, each with a medium image load. Persistent images are used, with a 30 GB base image and 20 GB per user for data. A total of 150 TB of storage capacity is needed, to store a full 50 GB image for each user. We have assumed 100 desktops are supported per host server and that each server can be fitted with six drives. For the network storage deployment, iSCSI connects to the shared storage array. For the Virtual SAN deployment, each host server houses six drives: four rotating platter and two SSD.

The savings compared to a network storage solution is $200 per user, or 77 percent.

SHARED STORAGE ARRAY COSTS		QUANTITY	COST
	Storage capacity	150,000	$750,000
	32 ports – 10Gb Ethernet switch	2	$28,978
Total			**$778,978**
Cost per user			**$260**

VIRTUAL SAN COSTS		QUANTITY	COST
	32 ports – 10Gb Ethernet switch	2	$28,978
	1.2 TB SAS drive	120	$69,960
	400 GB solid-state SAS drive	60	$82,140
Total			**$181,078**
Cost per user			**$60**

Table 2: Cost Comparison – Multisite Engineering Consultancy

Example 3 – Large Scientific Institution

Possible examples of this type of organization include national or multi-government centers responsible for research or leading-edge, science-based projects. The users are computer literate and often able to code. The desktop is of high specification and highly customized. The storage solution must be suitable for these variations and high demands.

This scenario has 1,000 users, each with a heavy image load. Persistent images are used, with a 30 GB base image and 50 GB per user for data. A total of 80 TB of storage capacity is needed, to store a full 80 GB image for each user. We have assumed 50 desktops are supported per host server and that each server can be fitted with five drives. For the network storage deployment, iSCSI connects to the shared storage array. For the Virtual SAN deployment, each host server houses five drives: three rotating platter and two SSD.

The savings compared to a network storage solution is $310 per user, or 72 percent.

SHARED STORAGE ARRAY COSTS		QUANTITY	COST
	Storage capacity	80,000	$400,000
	32 ports – 10Gb Ethernet switch	2	$28,978
Total			**$428,978**
Cost per user			**$429**

VIRTUAL SAN COSTS		QUANTITY	COST
	32 ports – 10Gb Ethernet Switch	2	$28,978
	1.2 TB SAS drive	60	$34,980
	400 GB solid-state SAS drive	40	$54,760
Total			**$118,718**
Cost per user			**$119**

Table 3: Cost Comparison – Large Scientific Institution

Cost Assumptions

The three scenarios are based on the following cost assumptions.

COMPONENT	MAKE AND MODEL	COST	UNIT
Shared storage array	Generic network storage[4]	$5	per GB
32 port Ethernet switch	Brocade ICX 6650 32 ports 10Gb switch with SFP+ transceivers	$14,489	each
1.2 TB SAS drive	10 K RPM 6 Gbps 2.5 inch hot-plug	$583	each
600 GB SAS drive	10 K RPM 6 Gbps 2.5 inch hot-plug	$399	each
400 GB solid-state drive SAS	6 Gbps 2.5 inch hot-plug	$1,369	each
200 GB solid-state drive SAS	6 Gbps 2.5 inch hot-plug	$569	each

Table 4: Cost Assumptions

All costs are list prices. Discounts would be applied,[5] we assume evenly for both the network storage and Virtual SAN options. List prices for the switch and drives were taken from public Web sites mid-June 2015. The cost per GB for shared storage is a conservative figure based on a combination of our business case data and analyst estimates.

In the third scenario, we could have selected a lower cost switch (fewer ports) for Virtual SAN and could have reduced the count of required drives by assuming a higher user-per-server density. Although our business case data suggests this would be valid, we wanted to ensure our cost comparison was not at risk of being seen as skewed in our favor. All scenarios assume a second, failover switch is deployed.

4. $5 per GB is a conservative estimate for the unit cost of network storage solutions based on our business case database. We also looked at equivalent hybrid storage array products and found similar costs. List price for a hybrid array with 30 TB of SAS HDD plus 10 TB of SSD is typically $200,000.

5. The actual discount levels achieved would depend on multiple factors, including (but not limited to) volume of overall purchase, status of existing relationships with suppliers, duration of contractual commitments, and external market conditions. Because "mileage may vary" widely, we have worked from list prices as a reliable and objective basis for cost estimates and comparisons.

Typical Implementation Methods for Maximum Cost Efficiency

Initially, you will probably have a large desktop environment, and it can be difficult to know how best to start your implementation of VDI. Two typical strategies are green field and brown field, and both can be used effectively. The green field method is starting fresh—deploying VDI completely from scratch. The brown field approach is an ongoing effort to update the desktop environment, perhaps replacing old equipment as it becomes end of life.

Green Field Method

An example of a green field project is several small groups of business users consolidating into a single data center and desktop environment. In a case such as this, a major implementation needs be undertaken whatever the desktop design.

This approach has the following advantages over the brown field method.

- Budget is already available, so justifying the capital expenditure isn't a problem.
- Sticking with the existing environment (if one even exists) is not likely to be an option, so an optimum solution can be designed without incorporating old equipment. VDI and conventional desktops do not have to coexist.
- It is easier to demonstrate to senior management the cost advantages of VDI, because different options can be proposed that show VDI as the clear winner both in terms of initial cost and ongoing maintenance.

This approach has one significant drawback: If it is the organization's first VDI implementation, limited data is available on which to base the design. For example, the optimal split between the drive types—HDDs and SSDs—can only be estimated.

If the proportion of HDDs is too high, the solution is less expensive, but I/O performance is negatively impacted. If the proportion of SSDs is too high, the user experience is excellent, but the system is overdesigned with unnecessary spending.

One way to mitigate these risks is to initially provide only enough storage for limited data growth, for example, equivalent to the first six months, holding back a portion of the storage budget until some tuning has taken place. An initial high estimate can be used for the proportion of fast storage with the intention of adding slower HDDs later to lower cost, assuming the performance—and thus the user experience—is acceptable. If it turns out that the proportion of fast storage is correct, storage types can be added in proportions similar to the initial setup. In the unlikely event of it being too low, more fast storage can be purchased.

Brown Field Method

With a brown field implementation, the project could involve replacing all existing desktops or be an ongoing effort to update the desktop environment using yearly budget allotments. Even if the goal is to replace all desktops, a phased implementation is advisable and is usually dictated by practical considerations.

A brown field project is more difficult than a green field one, in that both the physical and virtual infrastructures have to coexist, probably for an extended period. The design can be compromised by having to reuse existing components. These complications could prevent the full cost benefit of VDI from being realized, at least until the migration is complete.

On the other hand, this method has the following advantages.

- There is opportunity to start small, first with an experimental setup, and then with a pilot group of users.
- Some baseline data on I/O performance requirements is most likely available.
- Knowledge is gained as the project proceeds, so the final solution is likely to be more optimal, and hence more cost effective.
- Project deadlines might be less critical, because the lack of "big bang" makes slippage less visible or disruptive. Additionally, time taken at the start to ensure that the implementation is done correctly can make it easier to make up lost time later.

If the organization has a yearly budget line item for replacing the desktop infrastructure, this method can be used with no special CapEx allocation, because the VDI CapEx costs are lower than for conventional desktops. Money is saved immediately, although the expenditure type might be different. For example, data center space is required, and none had been used before.

Real-World Success: One Customer's Story

"After we finished virtualizing our server infrastructure, my team started to consider the desktop, which seemed like the next logical step.

"I initially thought it was going to be difficult to sell internally, because of the relatively large up-front investment in servers. When I looked more closely, by using Virtual SAN, the payback period was short enough so that we could see a positive return within the first year.

"The existing NAS was fully depreciated and coming to the end of its economic life, but we were able to use it as part of our storage grid in the new setup. As we replace it, we'll put in more attached storage, which is much cheaper than either new NAS or physical SAN. Performance is more than adequate, provided we take sensible decisions about proximity of the data to the virtual PCs.

"People have come to me and asked, 'How can this new setup be less expensive? You now have the cost of servers and storage in the data center, without eliminating your desktop PCs.' I tell them the total cost of the storage, servers, and software, plus the cost of the standardized, low-spec physical devices, is actually less than the cost of the physical desktops we used to buy—and that's before you even start to consider the lower maintenance cost and operational advantages of VDI.

"Some of the users were a little suspicious of the new setup to begin with, but were soon convinced of the advantages. The response from the virtual SAN is much better than they've been used to in the past. A failure of their desktop device no longer stopped them working, as they could access their virtual PC from anywhere—no more waiting around for their machine to be repaired or for a new one to be supplied. That's good for the IT support staff, too. It's easier and quicker to get people working again, so no more frustration from users on the shop floor or pressure from their managers if a target time is missed.

"There's a big advantage for business continuity planning, too. Disaster recovery exercises are no longer disrupted when we find there's some vital piece of data or software on someone's PC in a building to which we have no access. The Virtual SAN is set up with copies of the data at each site, and failover is all handled at software level.

"All in all, virtualizing the desktop makes total sense. I wouldn't want to go back to the way things were before, which now seems chaotic by comparison."

Conclusion

Cost-effective storage solutions often carry an element of compromise with respect to the solution adopted for end-user data. The databases associated with critical application data must have faster, more reliable storage to avoid expensive failures, so end users are typically allocated to older or cheaper storage.

This is not the case if you use VMware Virtual SAN.

Because of the flexibility in design when using Virtual SAN, an improved user experience can be delivered at a cost far lower than with traditional storage methods. The lower unit cost of attached storage provides an economical solution, while the caching mechanism ensures that the user experience is equal to that delivered by conventional SAN. VMware Virtual SAN is included in both the Advanced and Enterprise editions of Horizon 6, so comes at zero additional licensing cost. The greater simplicity of initial deployment, adjustment, and management also deliver ongoing operational cost savings compared with other storage solutions.

Virtual SAN should be the first choice for any organization serious about cost leadership in end-user storage solutions.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com