



How SC//Platform Lowers the Total Cost of Infrastructure

A Discussion of Soft,
Hidden, and Ongoing
Infrastructure Costs



Table of Contents

Introduction	3
Operational Cost Considerations	3-7
Deployment.....	4
Training.....	5
Software Licensing.....	5
Scaling.....	6
Downtime	6
Management	7
Data Center vs. Edge Computing	8
The Multiplier	8
Going the Distance.....	8
Right-Sizing.....	8
Summary	9



Introduction

When considering a new IT infrastructure solution, the acquisition cost of the hardware and software to standup the infrastructure is only the starting point for cost analysis. It is important to look into the operational costs of deployment, training, licensing, scale out, downtime, and management. All of these considerations contribute to the total cost of ownership (TCO) of the solution.

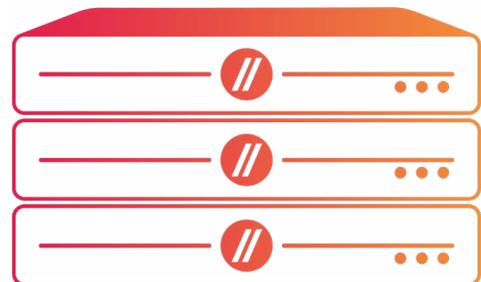
This document will dive into the TCO of Scale Computing Platform and discuss how it compares to traditional architecture models. The combined considerations of different areas of costs will give a broader and more complete picture of how SC//Platform affects IT costs. With a focus on serving small and midmarket IT, Scale Computing is committed to making virtualization infrastructure both accessible and affordable for any size organization.

Operational Cost Considerations

Acquisition costs of hardware and software are generally easy to determine based on vendor quotes, particularly if allocated as a CapEx (a Capital Expense). Acquisition costs can also be calculated as fixed OpEx (Operational Expense) which is also easy to calculate based on direct vendor quotes. Other types of operational expenses are not as easy to calculate and are often overlooked as part of TCO. Such costs that will be discussed here are:

- Deployment
- Training
- Software Licensing
- Scale Out
- Downtime
- Management

With each area of consideration, organizations will need to evaluate their own needs, own processes, and specific costs to determine the actual numbers. This will serve as a guide to calculating those costs.





Deployment

Traditional infrastructure deployment usually consists of deploying a SAN or NAS appliance, multiple virtual host servers, and then layering on a virtualization hypervisor. These three components are usually from different vendors meaning that each component is deployed separately and then configured and integrated into the complete infrastructure solution. With SC//Platform, the deployment is almost turnkey.

Because Scale Computing HyperCore™ combines server, storage, hypervisor, and management in a single appliance, all supported by a single vendor, a cluster of three or more nodes can be quickly deployed in under an hour. There is no integration of components other than assigning IP addresses to cluster nodes to create network connections for clustering. Storage is automatically striped and pooled across the entire cluster so there is no storage configuration. The hypervisor is included and preinstalled on each node, so there is no added configuration or integration of software needed.

From shipping container to rack to running live virtual machines (VMs), even a novice administrator can have SC//Platform deployed and ready for production in well under one hour. Compare that to traditional infrastructure that may take days or even weeks to deploy and test before putting into production. Getting up and running is only the first phase of deployment. The second phase is migrating your workloads to the infrastructure.

The migration process is going to be more or less the same on SC//Platform as on any virtualization platform. All hypervisors have tools to import VMs from other hypervisors and convert them to their own format. There are third-party tools that can assist in migrating workloads from any physical or virtual server with practically no downtime. It involves some planning and at least minimal planned downtime to account for any potential issues that may pop up.

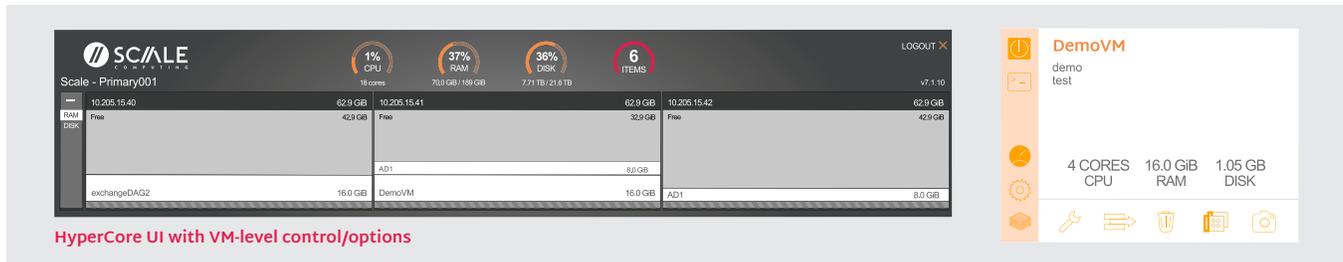
Because the process is almost the same on any hypervisor, migration alone is not a differentiator, however support and services for migration can be a big difference. Scale offers a number of migration services. These services range from migrating every workload onto SC//HyperCore and handing every aspect of migration to comprehensive migration walkthrough and training with users, enabling them to repeat the process on their own. Aside from service offerings, migration is supported by ScaleCare support services including how-to documentation along with an optional migration tool called Scale Computing Move. All ScaleCare customers also enjoy 24/7 phone and email support to assist through any critical issues.

So to sum up the deployment and migration costs, it is a simple matter of calculating the costs of administrator hours, third-party migration tools (if used), and downtime (if applicable). Compare the rapid delivery within a couple hours for SC//Platform to traditional virtualization infrastructure that can take days or weeks to fully deploy.

Training

Training and certification is often required to administer many technologies traditionally deployed in virtualization infrastructure including SAN or NAS solutions, hypervisors, and virtualization management tools. This kind of training can require dedicated offsite classroom learning or self-driven training through online resources, requiring days of time to complete. The cost of training can be high in terms of purchasing trainer classroom time, travel, and administrator time spent. Even with classroom training, administrators are often required to self-train hours of each week on updates and changes to ensure all components of the infrastructure remain compatible.

SC//Platform requires no specific training and has been designed so that it can be administered easily by IT generalists or even novices. Although there are some minimal training offerings around migration and deployment, most users are able to become familiar with the management interface rapidly during deployment without the need for any training resources. In fact, it's so simple that part of the included installation service is a HyperCore UI introduction and first VM creation walkthrough as part of the installation time. There are no certifications required. Training costs are usually calculated as zero for SC//Platform administration.



HyperCore UI with VM-level control/options

Software Licensing

One of the most oppressive costs of virtualization can be renewing hypervisor licensing costs. In particular, VMware can come with high annual license renewal rates both for the hypervisor and for vCenter management. Software licensing for virtualization can vary depending on the license type and implementation size. Free versions of hypervisors are sometimes available but often lack many of the features of the full hypervisor version that comes with the hefty price tag. Licensing can become an issue when scaling out infrastructure as well, when adding a new host server can require new licensing and sometimes move the whole solution into the next tier of licensing and significantly higher cost.

SC//Platform does not require any software licensing fees or renewals. All SC//Platform customers receive the full version of the built-in HyperCore hypervisor and all updates are available at no extra charge to customers under active support contracts including any new or existing features. Also, SC//Platform does not require any additional management tools at extra cost. This results in zero software licensing cost for SC//Platform. When compared to many VMware implementations, the cost of SC//Platform is often absorbed simply by the savings from eliminating the software licensing renewals from VMware.

Scaling

In traditional architecture built on servers, SAN/NAS, and hypervisors from different vendors, scaling out can be challenging because of the unknown. It is not always clear how or even if a new piece of hardware will interact with the existing implementation. Usually testing must be planned and executed before rolling out the new hardware into production. Often the existing infrastructure must be brought offline to add resource capacity. This is followed by crossing fingers that a latent incompatibility is not present which could cause new problems during implementation or at any later time, including inadvertently breaking out of a supported configuration.

It is also difficult to properly calculate or estimate the resources needed over a given period of time. As a result, organizations often over-provision, which means overspending. Between the cost of over-provisioning and spending days or weeks testing new hardware, and downtime for implementation, the costs add up.

With SC//HyperCore, scale out couldn't be simpler. A new node can be added to the cluster within minutes, automatically making new resource capacity instantly available and adding storage capacity to the existing storage pool without taking any workloads offline. No testing needs to be performed because the hardware has been pre-validated.

There is no need to over-provision with SC//HyperCore because scaling out can be done so quickly and easily, only the resources required at the time need be purchased. Future needs can be easily met with more cluster nodes. Furthermore, nodes can be mixed and matched adding resources as needed. Need more storage but not much more CPU? Need more RAM but not much more CPU and storage? New nodes can be configured with whatever you need and can be added to a cluster seamlessly.

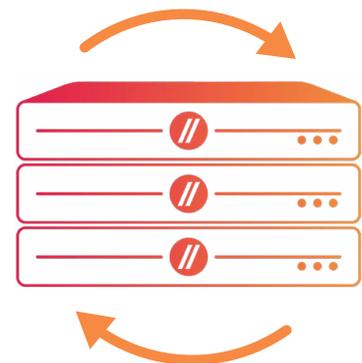
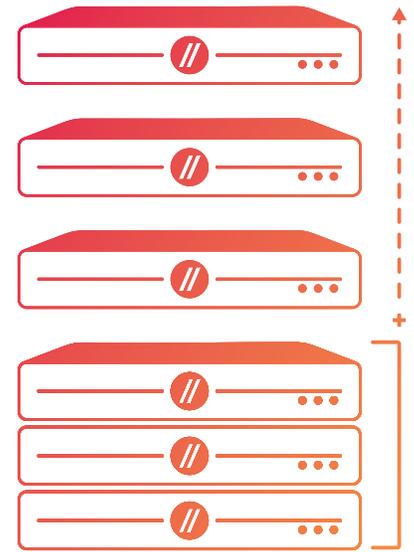
The ability to scale out at will, while mixing and matching nodes, nearly eliminates the costs associated with testing hardware, taking workloads offline, over-provisioning, and dealing with incompatibility issues. The cost of scaling out with SC//Platform is basically just the cost of a new node or two at whatever capacity or configuration your scale out requires.

Downtime

Downtime costs can cripple organizations, especially when man-made or natural disasters strike. Without sufficient high availability, disaster recovery, and backup strategies in place, such events can cause hours or days of downtime that may lead to critical revenue loss. But downtime can be both planned or unplanned and planned downtime can be costly as well.

The hourly downtime cost of each organization will vary depending on the type of business operations being performed. At the very least, downtime erodes credibility with customers both internal and external. For some businesses, an hour of downtime can be measured in the tens of thousands of dollars, especially if occurring during a seasonal rush. Planned downtime is generally much less costly than unplanned, but still carries costs of loss of productivity and administrators working extra hours on nights and weekends when the downtime is planned.

SC//Platform practically eliminates all planned and unplanned downtime for infrastructure. With built-in redundancy, high availability, backup, replication, snapshots, failover, and failback at no extra cost, SC//HyperCore provides a comprehensive disaster recovery strategy to ensure practically no unplanned downtime. Additionally, non-disruptive rolling updates and non-disruptive scale out eliminate nearly all planned downtime.



Within any SC//HyperCore cluster, VMs are automatically highly available and will failover to another node of the cluster if an entire node fails. Then, by enabling replication between a production cluster and a remote cluster or single node appliance, VM workloads can be backed up with multiple recovery points that can be used for failover within minutes of entire cluster failure. There is no need for a third-party DR solution with SC//HyperCore; the capabilities are built-in.

After a failover to a remote site, data can be quickly restored to the recovered production cluster by sending only the data that changed. The VMs can then be failed back within a matter of minutes. This reduces unplanned downtime from hours or even days to mere minutes in an incredible savings for business operation costs.

The redundant hardware and storage architecture also protects SC//Platform from various failures ranging from failed network ports or failed drives to failed power supplies. The resilient appliances continue to operate normally with self-healing abilities and administrators can replace failed components days after failure. Replacement components are shipped for next business day delivery whenever possible.

All of these combined abilities nearly eliminate downtime costs of any kind, dramatically reducing the TCO of the solution. Eliminating the cost of a third-party solution for DR capabilities is yet another way SC//Platform reduces the cost of ownership.

Management

Management costs generally refer to any administrator time spent dealing with the solution and therefore many of the areas already discussed here include elements of management time from deployment to dealing with downtime. There are more general areas of management that include monitoring and break/fix that can also be significant.

With traditional architecture that contains multiple technologies from multiple vendors, infrastructure often needs a lot of babysitting, in many cases requiring full time infrastructure administrators to monitor infrastructure performance and resource consumption, perform break/fix, and to plan and deploy new infrastructure for refresh or scale out. Management time is a significant cost especially when calculated over the lifecycle period of infrastructure, often between three to five years.

By eliminating complexity from infrastructure, SC//Platform significantly reduces management costs. One of the most significant savings comes from providing a single vendor solution. Rather than having to manage servers, storage, and hypervisor as separate components, with SC//Platform, the entire system is managed from a single management interface. The need to either manage through multiple consoles or purchase additional management tools is not necessary. The single management console alone reduces the management cost of monitoring the infrastructure by allowing status to be checked at a glance.

Another significant cost of infrastructure management is dealing with support issues ranging from incompatibilities between vendor solutions, software or firmware errors, hardware failures, and performance issues. In traditional architectures that combine multiple solutions, multiple vendor support organizations are needed to troubleshoot and resolve issues. Troubleshooting is complicated when dealing with multiple vendor technologies and often results in vendor run-around without any vendor willing to take responsibility, but instead blaming the other vendors whose technologies make up part of the solution.

With Scale Computing, there is only one vendor supporting the storage, hypervisor, and server compute resources. ScaleCare support services troubleshoots all issues and offers solutions to all issues in the virtualization infrastructure, significantly reducing the time spent by administrators dealing with support. While user experiences may vary based on the management costs tied directly to support issues, one can reasonably compare support costs between solutions based on the number of vendors involved. A single vendor solution can be expected to cost approximately one half to one third of a three-vendor solution in terms of support management costs.

One of the key value propositions of SC//Platform is decreasing management time. By nearly eliminating downtime, providing a single vendor infrastructure solution, and automating many typical IT tasks, SC//Platform dramatically reduces management time and costs. This allows administrators who would generally be employed full time to monitor and maintain infrastructure to instead focus on applications and activities that drive business forward instead of merely keeping the ship afloat.

Data Center vs. Edge Computing

Edge computing happens outside the data center in remote sites such as retail stores, branch offices, and distributed facilities and it has all of the cost considerations of data center computing. With edge computing, however, the costs can take on new dynamics, particularly being stretched across dozens, hundreds, or even thousands of sites.

The Multiplier

Extra costs that may seem inconsequential in a data center project such as a few minutes of management time or a few extra hours of deployment, can become monumental when multiplied across sites.

Consider

- An extra 10 minutes of management time per site x 240 sites adds up to a 40 hour work week
- An extra 4 hours of deployment time per site x 240 sites is an extra 6 months added to deployment
- An extra 24GB of RAM resource overhead per site x 240 sites is 5.7TB of resource overhead across sites

Going the Distance

Management and maintenance costs for remote sites can also be dynamically different from data center costs. Without onsite dedicated IT resources, performing routine maintenance and management tasks can be exponentially costlier when resources must travel to sites. Computing systems that require less onsite management and maintenance or that can be managed effectively by IT generalists rather than experts can significantly reduce those costs.

Right-Sizing

Finding the right-sized infrastructure solutions becomes more challenging as you target smaller deployments but still requires some of the same features as the data center. High availability, scalability, self-healing automation, and remote management features are not always available for small infrastructure for factors you might want to deploy in a remote site. Is it worth trading off some of these features where their absence may end up costing more later or is it better to overpay by over-provisioning with a larger infrastructure footprint than is needed just to get those features? Hopefully you will not have to choose.

Summary

As discussed, there are many costs to calculate and consider when determining TCO for an infrastructure solution. These costs vary from organization to organization so there is no way to know the exact TCO without digging into the specific processes and practices that drive these costs for each organization. However, with SC//Platform, the TCO is sure to be lowered versus traditional SAN/NAS+server+hypervisor architecture because it is simply easier to use and maintain.

No other solution today can deliver the same level of simplicity combined with offering high availability, disaster recovery, near turn-key deployment, seamless scalability, and non-disruptive rolling updates like Scale Computing Platform. This is why SC//Platform is the leading hyperconverged infrastructure solution for small and midsize organizations. It's time to stop investing in the technical debt of yesterday's infrastructure and continuing the high-cost practices of maintaining that outdated technology.

Learn more about the features that help reduce infrastructure costs in these additional documents or by visiting the Scale Computing website. www.scalecomputing.com



Corporate Headquarters
525 S. Meridian Street - 3E
Indianapolis, IN 46225
P. +1 317-856-9959
scalecomputing.com

EMEA B.V.
Europalaan 28-D
5232BC Den Bosch
The Netherlands
+1 877-722-5359

