IT Infrastructure Risk Management

How Scale Computing Helps Manage Risk in Infrastructure

Introduction

When implementing new IT infrastructure there are always risks. These risks include under-provisioning or over-provisioning, hardware incompatibility, software incompatibility, network issues and outages, migration issues, downtime, disaster recovery, vendor reliability, and unexpected costs. These risks can be inflated when ripping and replacing an entire infrastructure, but that doesn't have to be the case. Hyperconverged infrastructure solutions like Scale Computing Platform can reduce or even eliminate risks that have become common with traditional virtualization infrastructure.

Right-Sizing the Infrastructure

Sizing up the right amount of compute and storage resources with room for growth can be a complex process. Scale Computing simplifies this process in two ways. First, system engineers assist your administrators in using an infrastructure risk assessment and sizing tool to gather system usage and performance data from your existing environment. This information allows us to provide the right-sized recommendation for your current needs and make recommendations on future needs. This significantly reduces the risk of under-provisioning or over-provisioning the infrastructure.

VM Count → IOPs → Storage → RAM

Secondly, a Scale Computing cluster can be scaled out very quickly and easily with any appliance configuration. Nodes can be mixed and matched within clusters to scale out both performance and capacity. When it is time to add more infrastructure resources, you can add only what you need rather than being locked into more of the same nodes you started with. You no longer need to over-provision for years of growth when you implement the initial solution. The infrastructure can be quickly and easily scaled out at any time, with no downtime to workloads.

Hardware Compatibility

Unlike traditional virtualization infrastructure architectures that force you to integrate separate components like servers, storage, and virtualization from different vendors, Scale Computing has integrated and pre-validated the hardware before delivering an appliance. With Scale Computing, you get a single vendor supporting the infrastructure, including the hypervisor. All the hardware and software components have been tested together to provide a near turnkey solution that can be up and running in minutes.

With traditional infrastructure design, it might take weeks of implementation and testing by your IT staff to validate the whole infrastructure solution. With Scale Computing, we have done all of that work for you and will support you every step of the way in implementation and migration.

Workload Migration

When you are implementing a new virtualization solution like SC//Platform, you need to migrate your existing workloads to it from your old infrastructure. These may be physical workloads or virtual workloads, but either way, you want to avoid both downtime and data loss. We provide options to reduce downtime using migration tools that eliminate the risk of data loss.

For critical workloads where downtime must be minimal, we use SC//Platform Move (based on Carbonite Move™), which replicates the data from the running workload and only takes the workload offline for a few short minutes during migration cutover. As with the other solutions, there is no risk of data loss, and downtime is minimal. In both cases, if the migration fails for any reason, the original workload can be brought back online and continue running until the failure can be investigated, and migration can be performed again.

For non-critical workloads that can handle some downtime, we tend to use free solutions like Clonezilla that copy the workload in an offline state. There is no risk of data loss with this type of tool; however, the workload must be offline for the duration of the migration. The only real risk here is that downtime will be longer than anticipated.
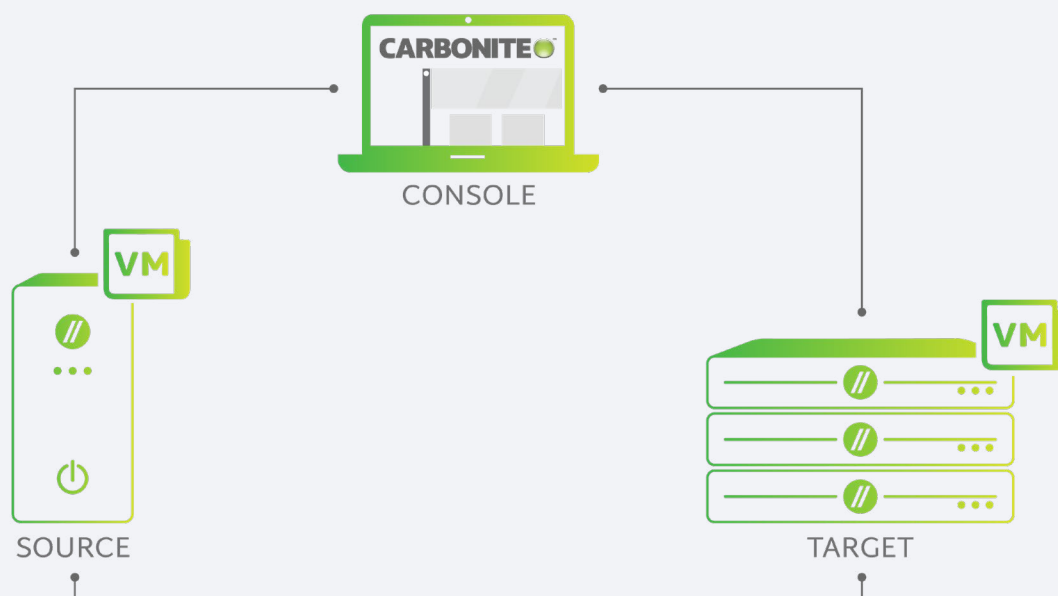
Unplanned and Planned Downtime

Downtime can be extremely costly to organizations, and as business becomes 24/7/365, it can be critical to avoid. Scale Computing has built-in high availability into every aspect of the infrastructure to help customers avoid downtime.

Unplanned Downtime

Beginning with some hardware best practices, such as providing redundant components in the hardware build, we are able to achieve impressive levels of fault tolerance in our clustering with wide striping of data across the entire cluster and high availability of VMs between cluster nodes. If a node fails, VMs are automatically failed over to other nodes in the cluster. Additionally, our built-in disaster recovery options, including failover and failback minimize downtime even for site disasters and failures.
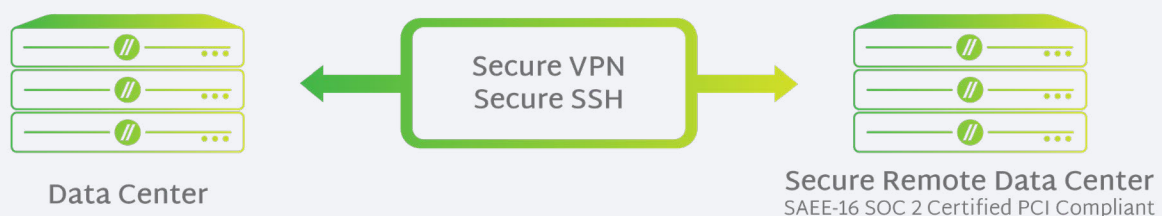
Planned Downtime

Unplanned downtime is the most impactful to business, but even planned downtime is undesirable. Planned downtime for infrastructure is often used for updating firmware and hypervisors, with an administrator taking hours to perform this process. With Scale Computing, these types of updates are automated and can be done without any workload downtime within a cluster. Workloads are automatically moved around between cluster nodes without being taken offline to update each node. The process has no manual steps other than initiation. Similarly, adding a new node to a cluster requires no downtime and minimal user steps to add the cluster. Most of the process is automated for ease of use.

Disaster Recovery

Implementing disaster recovery is often yet another vendor solution that must be integrated and tested for compatibility. Scale Computing has built disaster recovery into SC//Platform and also provides disaster recovery as a service (DRaaS). The built-in capabilities include continuous replication, failover, failback, and recovery.

When combined with the ScaleCare Support services, disaster recovery planning is documented within a runbook to ensure critical VMs are up and running quickly in the event of a disaster. With replication that can occur as often as every 5 minutes, sends only data that has changed, and is compressed and secured by SSH encryption, VMs can be protected between clusters or appliances across any distance. Replication is configured on a per-VM basis, so you can protect some or all of your VMs, or whatever fits your DR needs. In the event of failure or disaster, VMs can be failed over to the remote cluster or appliance within minutes. When the primary site is recovered, VMs and data can be restored and failed back, also with only a minute of downtime.

For our customers who do not have or want to host a DR site of their own, they can use our DRaaS option to replicate VMs directly to our secure hosted facility. The same built-in capabilities simply direct VM replication to the DRaaS facility. Whatever DR strategy you use, our ScaleCare engineers will always be on hand to help you through your disaster to get you back up and running as quickly as possible.

**Secure VPN
Secure SSH**

Data Center

Secure Remote Data Center
SAEE-16 SOC 2 Certified PCI Compliant

Vendor Reliability

Scale Computing has built a reputation for our focus on providing solutions. A look at our customer success stories will reveal a theme of customer satisfaction based on our ongoing commitment to customer support and success.

SC//Platform was designed to provide highly available and scalable compute and storage services while maintaining operational simplicity through highly intelligent software automation and architecture simplification. Scale Computing tightly controls, reviews, and maintains all third-party and open-source software used within Scale Computing HyperCore; common vulnerabilities and exposures (CVEs) are monitored and patched as needed at the source-code level by Scale Computing employees (with no dependencies on outside third parties), and no root or privileged access is granted to end-users or other outside representatives.

Scale Computing has complete ownership and control over SC//Platform's design, the components included, and the updates applied to our products. Trusted Scale Computing engineers manage all software - not unreliable third-party entities or outsourced engineering teams. No root or privileged access is available to general users or outside vendors.

Some hyperconverged solutions leave hooks to plug in your own hypervisor and related management tools. This can be a complex and dangerous combination, especially concerning security management.

SC//Platform avoids opening the system to outside parties. First, the hypervisor and management tools are included in SC//HyperCore and locked behind the software and a built-in firewall. Second, and more critical, the entire virtualization layer is completely embedded into the system itself. There is no "controller" VM or VSA needed to access or manage the cluster.

Recognized as a Leader in Hyperconverged Infrastructure and Edge Computing

Here's a sample of our recent recognition:

- Gartner Peer Insights, Customer's Choice for Hyperconverged Infrastructure, April 2022 (second consecutive year) » Read Press Release
- CRN®, a brand of The Channel Company, recognized Scale Computing:
    - 2022 Data Center 50 List » Read Press Release
    - 2021 Tech Innovator Award » Read Press Release
    - 2021 Edge Computing List » Read Press Release
- GigaOm named Scale Computing a leader in two February 2022 Radar Reports - Hyperconverged Infrastructure Solutions: Edge Deployments and Hyperconverged Infrastructure Solutions: Small-to-Medium Business » Read Press Release

Unexpected Costs

SC//Platform's principal design strategy is based on simplicity. It is this simplicity that helps reduce many of the extra costs of traditional infrastructures. These costs may come in the form of training, consulting, testing, and troubleshooting.

Simply reducing the number of vendors involved in the infrastructure can significantly reduce the runaround you would typically encounter in an infrastructure supported by several different vendors. There is no finger-pointing. With Scale Computing, it is finding solutions and resolving the issue as quickly as possible. Many customers underestimate the hidden cost of vendor runaround until they have a serious issue that is made worse by vendor finger-pointing.

SC//Platform is so easy to deploy and manage that we do not require any training for our users. We walk them through the process of racking, stacking, and configuring a cluster, which can be done as quickly as under an hour. The infrastructure allows many customers to see a reduction of management hours from days to minutes. When it comes to SC//Platform, it would be easier to talk about unexpected savings than unexpected costs.

Summary

The risks of moving from one infrastructure to another are largely based on outdated practices of building infrastructure with disparate components from different vendors. These ideas are simply out of date with new, easy-to-use solutions like hyperconvergence and SC//Platform. We are at the forefront of hyperconvergence, eliminating the complexity that creates the risks our customers work so hard to avoid.

Additional Resources

» DR Strategies with Scale Computing White Paper
» How SC//Platform Lowers the Total Cost of Infrastructure White Paper
» Implementing Business Resilience Best Practices White Paper