



IT Infrastructure Risk Management

From the Experts at
Scale Computing™





Understanding Risk in Modern IT Infrastructure

Infrastructure risk is the likelihood that failures, misconfigurations, security gaps, or capacity constraints will disrupt the availability, performance, or integrity of the systems your business depends on. In modern environments—spanning on-premises, cloud, and edge—risk increases as complexity rises, and small issues can cascade into outages or exposures.

How Scale Computing Helps Manage Infrastructure Risk

When implementing a 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™** edge computing solution can reduce or even eliminate risks that have become common with traditional virtualization infrastructure.

Right-Sizing the Infrastructure

Determining the right amount of compute and storage resources, with room for growth, can be complex. 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 enables us to provide a right-sized recommendation for your current needs and to anticipate future needs. This significantly reduces the risk of under-provisioning or over-provisioning the infrastructure.



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 scaled out quickly and easily at any time, with no downtime for workloads.

Hardware Compatibility and Lifecycle Risk

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 hardware and software components have been tested together to deliver a near-turnkey solution that can be up and running in minutes.

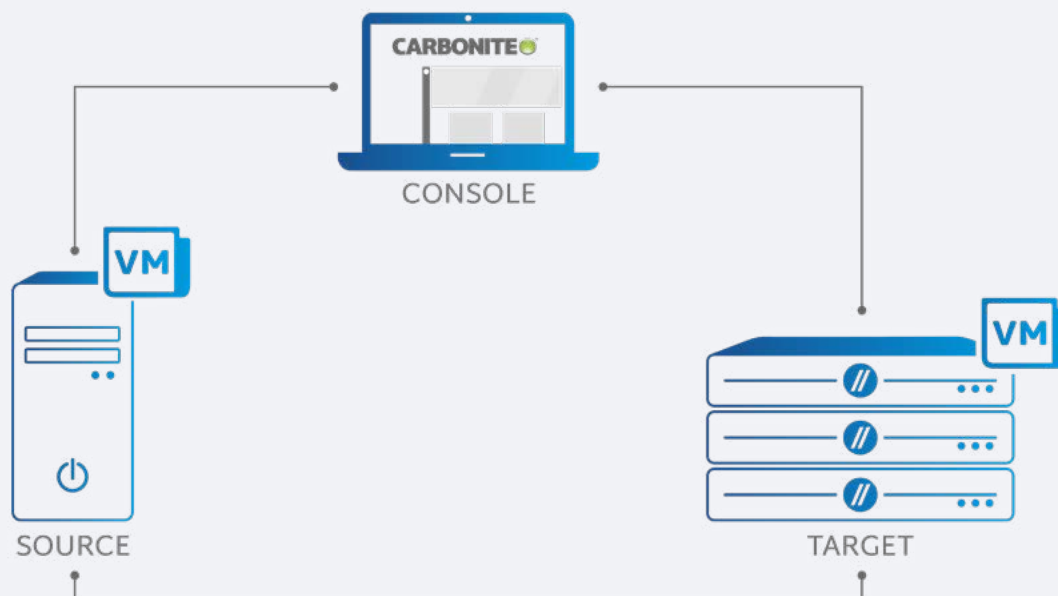
With traditional infrastructure design, it may take your IT staff weeks of implementation and testing to validate the entire infrastructure solution. With Scale Computing, we have done all of that work for you and will support you every step of the way in implementation and migration.

Workload Migration and Platform Transition Risk

When implementing a new virtualization solution, such as SC//Platform™, you need to migrate your existing workloads from your legacy infrastructure. These may be physical workloads or virtual workloads, but either way, you want to avoid both downtime and data loss. We offer options to reduce downtime with migration tools that eliminate data loss risk. For critical workloads where downtime must be minimal, we use SC//Platform Move™, which replicates data from the running workload and takes the workload offline for only a few minutes during the 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 is investigated and the migration can be performed again.

For non-critical workloads that can tolerate some downtime, we tend to use free solutions such as Clonezilla to 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.

Hardware Compatibility



Downtime Risks in Traditional Virtualization

Understanding Risk in Modern IT Infrastructure

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

Unplanned Downtime Risks

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 and Maintenance Windows

Unplanned downtime is the most impactful on business, but even planned downtime is undesirable. Planned downtime for infrastructure is often used to update firmware and hypervisors, with an administrator spending hours on the process. With Scale Computing, these updates are automated and can be performed without any cluster workload downtime. Workloads are automatically moved between cluster nodes without taking any node offline to update it. The process has no manual steps other than initiation. Similarly, adding a new node to a cluster requires no downtime and only a few user steps. Most of the process is automated to improve usability.

Data Protection and Business Continuity Risks


Disaster Recovery and Resilience

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 occurring as often as every 5 minutes, sending only changed data, and compressed and secured via 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, depending on 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 customers who do not have, or do not 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





Operational and Vendor Risks

Vendor Reliability and Lock-In

Scale Computing has built a reputation for its focus on providing solutions. A look at our customer success stories reveals a consistent theme of customer satisfaction driven by our ongoing commitment to customer support.

SC//Platform was designed to provide highly available, scalable compute and storage services while maintaining operational simplicity through intelligent software automation and simplified architecture. Scale Computing tightly controls, reviews, and maintains all third-party and open-source software used within **Scale Computing HyperCore™** virtualization suite; 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 does not open the system to external parties. First, the hypervisor and management tools are included in SC//HyperCore™ and locked behind the software and a built-in firewall. More critically, the entire virtualization layer is completely embedded into the system itself. No "controller" VM or VSA is required to access or manage the cluster.

Unexpected Costs and Licensing Exposure

SC//Platform's primary design strategy is simplicity. It is this simplicity that helps reduce many of the additional costs associated with traditional infrastructure. These costs may include training, consulting, testing, and troubleshooting.

Simply reducing the number of vendors involved in the infrastructure can significantly reduce the runaround you typically encounter with infrastructure supported by multiple vendors. There is no finger-pointing. With Scale Computing, we find solutions and resolve issues as quickly as possible. Many customers underestimate the hidden costs of vendor runaround until they face a serious issue exacerbated 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 completed in under an hour. The infrastructure enables many customers to reduce management hours from days to minutes. When it comes to SC//Platform, it is easier to discuss unexpected savings than unexpected costs.

Summary: A Practical Framework for IT Infrastructure Risk Management.

The risks of moving from one infrastructure to another often stem from legacy approaches that rely on disparate components from multiple vendors, creating complexity, integration gaps, and operational inconsistencies. As organizations modernize, solutions like hyperconvergence are designed to simplify day-to-day management and reduce the points at which failures and misconfigurations occur—[explore how Scale Computing reduces infrastructure risk](#) as part of that shift. 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](#)

CORPORATE HEADQUARTERS

3307 Northland Dr #500 // Austin, TX 78731
P. +1 317-856-9959 // scalecomputing.com

