



UNIFIED • CLUSTERED • AFFORDABLE • SCALABLE



lower **cost**
increased **control**
more **convenience**

- **Mix & match** node densities with full usability
- **Isolate** intra-node traffic on private network
- **Keep running** even with full node failure
- **True performance** scalability as you add nodes



Scale Computing – Intelligent Clustered Storage (ICS)

Generation 3.0 clustered storage solutions unify SAN and NAS, simultaneously providing iSCSI, NFS and CIFS, with no single point of failure. Based on Scale's Intelligent Clustered Storage™ (ICS) technology, Scale's storage portfolio aims to reduce cost, increase control, and make storage management more convenient. This is what Scale calls the 3 C's (Cost, Control and Convenience). The ICS software allows Scale to leverage off-the-shelf hardware, passing the cost savings onto customers.

► **Unified SAN/NAS: File and Block Storage in One Solution**

Scale supports simultaneous protocols including iSCSI, CIFS, NFS, enabling the end user to utilize the right tool for the job. Since different use cases call for different protocols, one can optimize the storage infrastructure all the way down to the connectivity method. This also allows for the elimination of file servers. By moving the overhead of file level storage from an intermediate server, users can eliminate several bottlenecks and greatly increasing performance. All CIFS and NFS requests are handled by all nodes in a Scale cluster while eliminating single points of failure.

► **Controller-less Architecture: No Single Point of Failure**

All nodes are active and any individual node can take over for any other node throughout the entire cluster, eliminating single points of failure without adding cost or complexity. Scale's architecture allows for failure of up to 49 percent of the nodes in the cluster without losing data access and maintaining full data redundancy.

► **Unlimited Capacity Scalability: Mix and Match Nodes of Different Capacities**

Utilize the full capacity of denser drive nodes as they become available. The need to maintain consistent drive types is eliminated. The Scale clustered file system separates the drives and nodes from Scale software allowing drives of different densities to work in the same cluster and for full access to the entire capacity of the drives. This ensures that customers are never in the position of changes in drive density outdating their investment.

► **Enhanced Performance**

Scale's implementation was designed from the ground up as a full peer-to-peer cluster with a very high performance filesystem. All node-to-node management, communication, data transfer and replication are done out-of-band on a private network. This allows for far greater expansion of the cluster. The filesystem can scale to 8,192 nodes.

► **True Storage Virtualization**

Scale's ICS creates a virtualization layer between the customer's data and the physical spindles that support it. In doing this, ICS removes the need to transplant the dataset from one storage platform to another – you simply add nodes to the cluster and unplug older nodes – data is automatically and seamlessly migrated to the additional storage.

► **True Pay as You Grow Architecture: Don't Pay In Advance for What You May Need Later**

Expanding a Scale cluster is as easy as adding another node and clicking a few GUI buttons. The cluster does the rest. Eliminate forecasting and buy capacity and performance only when you need it.

► **Global Namespace**

Managing islands of storage is time consuming, confusing and expensive. Scale's solution grows from a few TBs to several PBs in a single global namespace. The result is meeting the needs of the most storage intensive applications and achieving simplicity in managing the overall storage solution.

► **Future-Proof Architecture: Your Investment is Protected**

Mix and match nodes to use full capacity of higher capacity drives; retire and replace old nodes easily without the need for data migration; add new protocols when released and customers aren't locked into a single protocol —This is all accomplished with Scale's Protocol Abstraction Layer (PAL) and Clustered File System. The PAL separates the protocols from the hardware allowing Scale to add additional protocols at any time without impacting the customer's prior investments. Scale will be able to add future protocols, even ones that are yet to be developed.

- **Non-disruptive upgrades**
- **Never requires a forklift upgrade**

► **Unlimited Performance Scalability: As Fast as You Need to Go**

Each node adds throughput, IOPS, processing power, and cache. Customers are able to "dial-in" the level of performance required to achieve the throughput needed for each application. Customers are not limited to the throughput of a few host controllers.

► **Multi-way Asynchronous Replication**

One-to-many or many-to-many replication allows for more flexible DR strategies. Great for satellite offices to replicate back to HQ, and HQ to replicate to a DR site. Production sites can replicate to other production sites without the need to have a dedicated DR location. This reduces the impact to budgets and achieves the goal of business continuity.

► **Immediate, Redirect-Based Snapshots**

Copy on write snapshots gives instant protection with low overhead.

No need to reserve a percentage of capacity for Snapshots, allowing for a greater utilization percentage of the available storage.

► **No Storage Traffic on the User Network: Greater Performance and a More Logical Approach**

Results include a 2 – 3x reduction in the number of packets on the wire in the data path and far greater efficiency of use of available data path bandwidth, resulting in 30 to 150 percent greater performance in both IOPS and raw throughput to a Scale cluster vs. competitive offerings. Additionally, this improved efficiency allows for the use of SATA drives in areas previously requiring the extra expense and capacity loss of SAS.

Also, storage performance is not impacted by network level events on the production network, as the nodes are able to freely communicate read requests, replicate IO's and manage traffic independently of the public LAN, as well as a further improve the available bandwidth for new IO on the public LAN data-paths.

Scale utilizes commodity hardware. The power of Scale is in the software. This allows for feature rich functionality and low cost due to commodity hardware. Costs associated with open architecture x64 based approaches are two to four times less expensive than proprietary hardware approaches, resulting in dramatic cost savings to the end user. Scale does not customize the hardware or silicon in anyway, keeping the cost low in order for customers to get more for their budget dollars.

Want a Webinar?

<http://scalecomputing.com/webinar/>

Need Sales or Technical Help Now?

<http://scalecomputing.com/varportal/territories/>

Who is Scale?

The Scale Computing team has a long history of working together across multiple successful organizations over the last two decades. This includes companies such as Radiate, Corvigo, Tumbleweed Communications, Volt Capital, and SolarX. This dynamic lays the foundation for a strong sense of teamwork that permeates the entire organization. Funded by Benchmark Capital, and the 21st Century Fund of Indiana, among others, Scale has hit growth milestones twice as fast as EqualLogic and three times as fast as Lefthand Networks. Based in Indianapolis, IN, Scale Computing was named one of Forbes' Most Promising Companies in 2009 and has expanded internationally to Canada, Japan and Europe.



877-SCALE-59 • www.scalecomputing.com