

# IDC SOLUTION BRIEF

---

## Assessing the Business Value of Scale-Out Network-Attached Storage Solutions

*Sponsored by EMC*

By Randy Perry and Benjamin S. Woo  
May 2011

---

### Overview

With the growth of Web, cloud, and data-intensive applications, companies are increasingly challenged to find better ways to scale and manage their storage environments without disrupting business operations or adding costly IT staff. One solution gaining popularity is scale-out storage, which supports data growth by adding new storage "nodes" when extra storage capacity is needed. This approach differs from the traditional scale-up architecture, which handles data growth by aggregating lots of individual disk drives behind one or two large controller servers.

With scale-out storage solutions, a software layer allows the multiple nodes to behave like a single, larger file system, known as single global namespace, making it easier to manage data growth without having to hire additional IT specialists. Also, as server prices drop and capabilities and capacities increase, scale-out storage solutions allow enterprises to replace large, costly equipment with a readily sized pool of inexpensive servers.

The scale-out approach can be applied to direct-attached storage and storage area networks, but it is more commonly associated with network-attached storage (NAS) systems, which use servers, disks, and management software to serve files over a network. Scale-out NAS systems become a cluster of NAS nodes, which allow for capacity and performance to scale in combination or independently while maintaining a single system image.

To assess the business value of scale-out NAS solutions, IDC used an extensive research database of several thousands of organizations to quantify the benefits realized by companies implementing the scale-out approach to manage data growth. IDC found that, on average, companies achieved annual benefits of \$8,200 per terabyte from deploying scale-out NAS solutions.

---

## **Business Benefits of Scale-Out Storage**

By combining and configuring many small, inexpensive storage components, scale-out solutions create an aggregate storage pool that can exceed the capacity of a single, traditional storage array. Scale-out solutions can also provide more timely processing and greater resource utilization than scale-up systems because they allow companies to scale storage resources in line with datacenter demands and changing business needs. Also, as nodes are added, the processing power and storage I/O bandwidth increase, along with the storage capacity, so scale-out systems become increasingly faster and more powerful as capacity is added to the storage infrastructure.

Additionally, the simple, inexpensive nodes allow for granular expansion of storage resources as needed without disrupting business operations or risking the overbuying that can occur with scale-up systems. Further, with scale-up solutions, each additional system increases the management burden and the likelihood of having to add IT staff. By avoiding these problems, scale-out solutions can hold down costs and make IT staff more productive. Also, by supporting data growth in a dynamic, nondisruptive manner, scale-out solutions ensure that system capabilities always remain as productive as possible.

A key component of scale-out solutions is the software that enables the nodes to be interconnected and referenced as a single object by storage administrators. With such solutions, administrators are responsible for managing just the data and not the storage hardware. The software also provides all the data services IT professionals have come to expect from their storage solutions, including snapshots, replication, and data protection.

## **Example of a Scale-Out NAS Solution Portfolio**

One company with a highly developed range of scale-out NAS solutions is EMC. Its Isilon OneFS operating system supports three series of platform nodes — S, X, and NL — as well as data replication, data retention, and other software capabilities. With Isilon scale-out storage solutions, companies can scale capacity from 20 terabytes (TB) to 10 petabytes (PB) in a single file system and scale performance to 50 gigabits per second.

The recently announced Isilon S200 platform is intended to speed time to market and maximize return on investment with ultrafast primary storage for mission-critical, highly transactional, and random-access file-based applications. It scales with extremely demanding workflows and applications and delivers high performance in a small footprint. The platform uses enterprise solid state drives (SSDs) to accelerate namespace-intensive metadata operations.

The companion X200 is a versatile midtier platform, designed specifically to accelerate highly concurrent and sequential throughput applications while reducing cost and complexity. The company's highest-density platform, the IQ 108NL, also has the most aggressive cost per TB. Intended for nearline archive applications, it scales from 100TB to 15.5PB in a single pool of storage with global namespace, eliminating the need to manage multiple volumes.

EMC's Isilon OneFS 6.5, its sixth-generation operating system, combines three layers of traditional storage architectures — file system, volume manager, and RAID — into one unified software layer, creating a single intelligent file system that spans all nodes in a cluster. OneFS 6.5 provides integrated support for SSDs within SmartPools, a feature that allows users to combine pools of different system types in the same cluster, all in a single file system, and to automatically migrate data among nodes with different performance and capacity characteristics. As a result, critical SSD resources can be used to improve performance cost-effectively across a wide range of workflows.

---

The company has also enhanced disaster recovery and business continuity with the release of SyncIQ 3.0 data replication software, which provides high-speed replication of millions of files and TB of data over the WAN and LAN for disk-to-disk backup, disaster recovery, and content distribution. Integrated snapshots quickly determine which data has changed between syncs, reducing the time needed to protect the same amount of data.

### **IDC's Methodology for Quantifying the Business Benefits of Scale-Out NAS Solutions**

In assessing the quantifiable benefits of scale-out NAS solutions, IDC first determines improvements in key metrics, such as time spent provisioning, utilization percentage, and backup schedule, and translates them into financial savings. IDC groups the benefits into the three major categories of cost savings — reduced costs, increased IT staff productivity, and improved end-user productivity — and normalizes the results by expressing them in dollar savings per TB of storage. This allows companies with various storage capacities to extrapolate the results to their own organization.

#### ***Cost Reduction***

Cost savings may be achieved through better leveraging of SAN resources and by reducing the number of standalone servers needed. Elimination of old hardware and contractor fees can also lower costs, along with reduced spending on disk storage and services.

#### ***IT Staff Productivity***

For companies to remain competitive, IT managers and their staff must use their time more effectively. Besides reducing costs, gains in IT staff productivity can free up staff to implement new initiatives more rapidly, helping to create a competitive edge. Productivity gains may come from managing capacity more efficiently, reducing provisioning time, or reducing the time spent on backup and recovery operations. IT productivity gains are quantified by multiplying the savings in IT staff time by an annual loaded salary of \$100,000.

#### ***End-User Productivity***

When users are unable to access stored data and applications, their productivity may be severely impaired, so increasing storage availability can generate significant savings. Minimizing business disruptions during software upgrades also helps to increase user productivity. To quantify end-user productivity gains, IDC multiplies the increased availability of the data and applications that end users need to do their jobs by a loaded annual salary of \$67,500, scaled by a productivity factor to account for the fact that they may be able to do other jobs during data and application downtime.

### **Business Value Results for Scale-Out NAS Solutions**

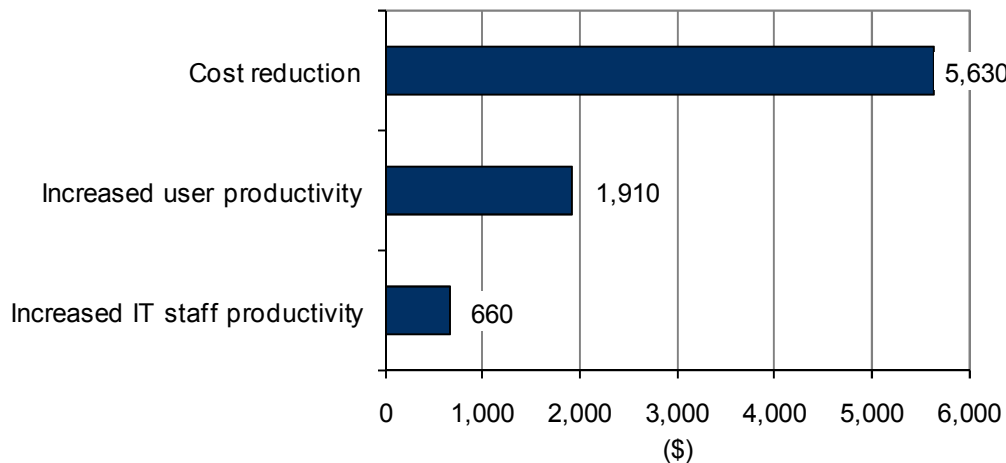
IDC's assessment confirmed that scale-out NAS solutions can improve IT management's ability to provide more timely provisioning, greater resource utilization, improved performance, and higher levels of data and system availability. Performance improvements were significant (see Table 1). On average, companies with scale-out NAS solutions reduced provisioning time by 60% and the time needed for backup by 76%. Utilization increased by an average of 30%, while storage capacity was managed 27% more efficiently.

<b>TABLE 1</b>	
Performance Improvements	
Increased utilization	30%
Reduced provisioning time	60%
Reduced time for backup	76%
More efficient management of storage capacity	27%
Source: IDC, 2011	

On average, companies with scale-out NAS solutions realized annual cost savings of \$5,630 per TB (see Figure 1). Annual savings from improved end-user productivity averaged \$1,910 per TB, while savings from increased IT staff productivity averaged \$660 per TB annually. In total, the companies achieved average annual benefits of \$8,200 per TB from deploying scale-out NAS solutions.

**FIGURE 1**

Average Annual Benefit per TB



Source: IDC, 2011

**Copyright Notice**

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2011 IDC. Reproduction without written permission is completely forbidden.