

Nov. 20, 2013  
Jeremy Li

## Big Data ([Hadoop](#)) with Scale-out Storage (NAS)

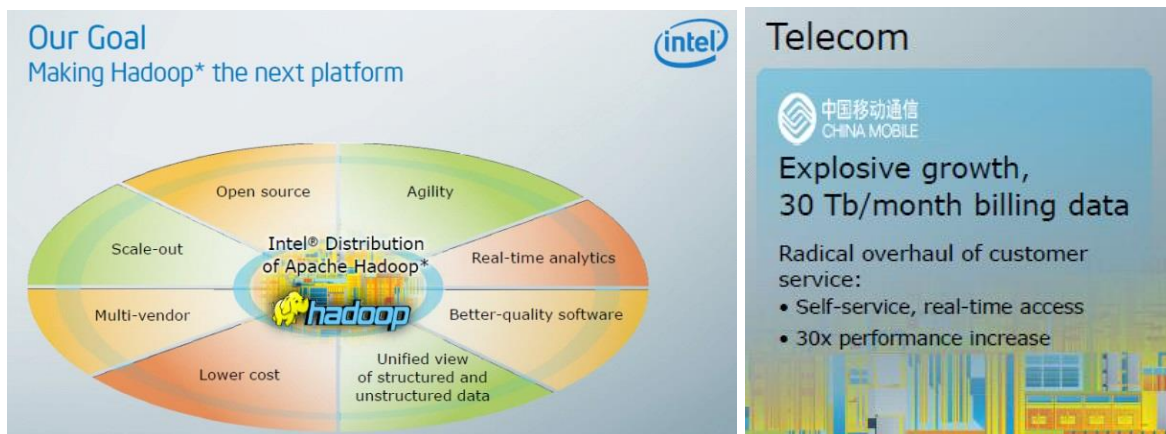
Isilon was acquired by EMC in late 2010 and became EMC Isilon Storage Division. Isilon scale-out NAS solutions comprise an appliance that includes central processing units (CPUs), Random Access Memory (RAM), Power Supplies, Infinite Band (IB), 10/1 Gigabit ports/connections, and its 7<sup>th</sup> generation operating system, known as [OneFS](#) that puts three components of File System, Volume Manager and Data Protection into one File System innovatively.

Isilon has more than 4,000 customers globally with growth rate at 64% year over year in 2012, generated product revenue run-rate at \$1 billion and shipped 85PB storage in the 4<sup>th</sup> quarter, 2012.

Hadoop is pioneered by Google's [MapReduce](#), a programming model and framework where an application is broken down into numerous small parts, usually written in Java. Now, it is based on open-source data storage and processing API with massively scalable and automatically parallelizable capability. The core parts of a Hadoop distribution consists of HDFS storage, MapReduce API and Other Libraries.

[Apache Hadoop ecosystem](#) consists of many components, two of the important components are storage and HDFS (Hadoop Distributed File System) that provides high-throughput access and comprises of NameNode, Secondary NameNode and DataNodes.

[Apache Hadoop](#) is an [open-source software framework](#) for storing and large scale processing of data-sets on clusters of [commodity hardware](#) (Source: Wikipedia). Intel and other leading IT vendors are making Hadoop the next platform, as illustrated in Figure 1. [Click here](#) for a comparison of RDBMS and Hadoop.



**Figure 1 Intel is making Hadoop the next platform**

**Note:** INTEL and CHINA MOBLE, the world's largest Teleco with around 700 million users and the No. 1 brand name in China from 2010 to 2013, are working together and have achieved the savings of \$30 million every month via Hadoop technology, making Hadoop the next platform.

Isilon is very good for digital content and other unstructured data and semi-structured data used by a variety of industries. It may also be a good candidate in High Performance Computing (HPC) in [a few special scenarios](#), even though it is never listed at [2013-annual-hpcwire-readers-choice-awards](#) and EMC only wants to concentrate its energy on two pillars: 1) Storage and 2) Big Data.

**Note:** The true HPC requires accessing and computing multiple databases simultaneously, while Isilon can only access the database one at a time because Isilon is a cluster NFS with a point-to-point connection, one at a time, while HPC (e.g., DDN) is a traditional parallel computing architecture that is much faster than Isilon's HPC.

Intel is doing a good job to bring Big Data (Hadoop) into the HPC. A traditional ["Share-Nothing" Hadoop](#) must retain at least three (3) redundant data copies that must be copied to the Hadoop cluster before analysis can begin. It is very difficult to maintain data consistency when a file changes on a primary storage.

Isilon created a new idea with ["Share-Everything" Hadoop](#). An analysis can begin with the first VM (Virtual Machine) using already existing data without needing to rely on and wait for the completion of three identical copies of identical data by accessing the original data via NAS and HDFS (Hadoop Distributed File System) protocols in addition to being able to use unused processing and RAM available in a VMware environment. Therefore, the replication is not required. Hadoop is becoming "computing only" cluster under Isilon and looks like having hours-needed for actual copies without copies with a pointer technique!

Therefore, Isilon might be a better choice for Hadoop due to its distributed computing architecture – [parallel processing](#) with rapid data transfer rates among nodes and still having the uptime in case of a node failure (N+1, N+2, N+3 and N+4) by OneFS scale-out operating system (OS) and its unique technique by using a pointer to generate three identical copies (a compliance for a Hadoop) on a scale-out storage in Big Data Analytics scenarios – a massive storage savings only Isilon can deliver at this time because:

- At least three (3) identical data must be copied to the Hadoop cluster before analysis can begin
- Simple install, manage and scale without going through a forklift upgrade
- Delivers powerful yet simple solutions for enterprises that want to manage their data, not their storage.
- Overcome the [traditional "Share-Nothing" Hadoop](#) with [innovative "Share-Everything" Hadoop](#)
- [Independent Scaling](#), instead of Dependent Scaling

**Note:** When a cluster becomes large, a massive parallel processing capability is formed.

Below are highlights of EMC Isilon:

- Single File System with [Simplicity](#)
- [High Performance Computing](#)
- [Automated Tiering](#)
- [Easy Growth](#)
- [Linear Scalability](#)
- [HDFS is a Native Integrated Protocol](#)

**Note:** Common Internet File Systems (CIFS), also known as Server Message Block (SMB) for Windows file sharing, is widely being used by Microsoft Windows applications, while Network File System (NFS) for Unix and Linux file sharing, is widely being used by UNIX and Linux applications. NFS is a more efficient protocol than CIFS that must go through TCP/IP handshaking process. As a result, many large enterprises use NFS, instead of CIFS in a NAS environment, instead of using Windows on top of Linux OS such as VMware ESXi for NAS purpose to get rid of additional layers.

IDC MarketScape Names [EMC Isilon a Leader in Scale-Out File-Based Storage Market](#)

Additional Information from B. Scott Cassell, EMC<sup>2</sup>| Isilon Storage Division

- Hadoop is traditionally not compatible with either SAN or NAS. It is only compatible with DAS (Direct Attached Storage) utilizing HDFS with the ONE exception of Isilon who has developed the ability to transform HDFS into a protocol like NFS or CIFS/SMB. With Isilon, Hadoop is now compatible with NAS including NFS and CIFS/SMB, as well as FTP, HTTP, etc.
- ONLY Isilon lets Hadoop users separate the storage component from the computer component. Meaning with Isilon, you can add just compute to MapReduce jobs faster without adding any additional storage, or just add storage to hold larger data sets without adding any additional compute components
- In a traditional Hadoop scenario, adding storage will require installing a server and mounting DAS hard drives to those servers (in triplicate) and the servers needs to match the resources of the other DataNodes already in existence. The DataNodes node resources are used for MapReducing and the MapReducing Nodes are used for DataNodes storage. You bloat one if you only need the other
- ONLY Isilon lets Hadoop users start MapReducing immediately without the need to duplicate and move the duplicated data to the target data nodes
- ONLY Isilon lets Hadoop users perform analytics-in-place because there is no need to duplicate and ship data. Analytics can be performed on the one copy of actual original data in real time without preparatory steps.

## High Performance Computing (HPC)

Isilon's architecture is also good for HPC. The more appliances are added, the more scale-up for storage and scale-out for performance achieved because each additional new node brings more compute, memory, disk capacity and network connectivity. This allows performance to scale in line with capacity - [Linear Scalability](#) ([Click here](#) for a comparison of Isilon Scaling vs. RAID Scaling). For example, finding deposits of natural resources (oil, natural gas, precious metals, etc.) for the major mining companies must require HPC in order to achieve the result.

IBM is well-positioned to participate in the rapidly growing business analytics market because of its participation in the high-performance computing (HPC) market as a server, storage and software provider, and its technical expertise in basic research and vertical markets (**Source:** Gartner)

EMC does not have server and networking business and relies on a partnership with other vendors such as Cisco and Brocade. As a result, many big enterprises choose to do business with IBM that can offer the entire solution from single vendor.

As you can see, Hadoop is big and comprises many components. EMC Isilon is an excellent choice in storage, but the ecosystem of Hadoop needs more components to achieve Big Data Analytics.

It is worth mentioning: 1) only a few vendors who are able to offer HPC for storage support - EMC Isilon, IBM Scale Out Network Attached Storage (SONAS) based on [IBM DCS3700](#), for which NetApp is the OEM, [SGI](#), [DataDirect Networks \(DDN\)](#) and etc.; 2) most local government agencies do not need HPC capability except for agencies that require a video surveillance system [e.g., the Los Angeles County Sheriff's Department (LASD) is deploying an expanded [video surveillance and data storage system](#).] or a massive analytics system.

Click on the link [2013-annual-hpcwire-readers-choice-awards](#) for more.

**Note:** The topic of HPC is beyond the scope of this writing.

## Deduplication and Compression

[In-line dedupe](#) method increases efficiency of global deduplication since immediate processing of data stream coming from several sources allows eliminating duplicates on the fly, while [post-process deduplication](#) requires an identical storage to hold the data first prior to starting the dedupe process, and is mostly used in the backup applications, virtual tape libraries and the like. Click [Post-process deduplication vs In-line deduplication](#) for details.

While providing the appropriate level of IO performance is an important task for a storage platform, providing efficient usable utilization is also very important and affects

the total cost of ownership. Employing modern storage optimization technologies such as Deduplication and Compression helps reduce the cost of storage by allowing more data to be stored per GB of usable capacity. Isilon just added a few new features in OneFS 7.1 including post-process [SmartDedupe](#) along with Audit, Data Encryption, Hadoop 2.0 and Software Defined Network in Oct., 2013. Although still lacking native data compression and does not offer a data compression, Isilon claims that it can still achieve high useable capacity after taking into consideration base raw-to-usable efficiencies enhanced by deduplication, compression, and/or both.

EMC VNX and NetApp Data ONTAP offer both post-process deduplication and compression technologies, while IBM Storwise only offers real-time compression. The post-process deduplication or compression technologies are not good as a real-time deduplication or compression because either of them occurs only after the data has been landed to storage ([click here](#) for details) in addition to the fact detailed above. To address the post-process dedupe inefficient issue, EMC just released a new inline dedupe on its All Flash Array in Nov., 2013 that put the post-process dedupe in its secondary class because EMC must realize the benefit of inline dedupe or [IBM real-time compression's advantage](#) since IBM Storwise real-time compression has an ability to compress data in real time prior to landing the data to the storage. Therefore, it truly reduces the initial storage acquisition cost by not purchasing a full planned raw storage capacity size at its initial acquisition time, while both EMC and NetApp recommend sizing initial storage installations based on containing data at its original occupied size. Gartner rated IBM real-time compression that was one of the crowning achievements is [a potential game changer](#).

Below is a quote from a white paper, titled "[Compares IBM Real-time Compression for Storwise V7000 vs EMC and NetApp compression for Block Storage](#)" that also illustrates the benefit of the Real-time Compression Technology.

"For this reason, both EMC and NetApp recommend sizing compressed devices for data at its original size. If followed, this practice may eliminate the benefits of deploying compression." IBM V7000 is able to compress data up to 80% in real time and reduce its initial storage capacity size at its acquisition time."

A quote from B. Scott Cassell "Because some data is highly deduplicatable and/or compressible while other data is not, it is important to understand the dedupe and compress potential of the actual data being stored. The different brands each apply their own proprietary deduplication and compression methodologies so a significant consideration should be to assess the Net Usable storage after raw-to-usable calculations, space reclamation from deduplication, and space reclamation from compression calculations have been made."

**Note:**

- EMC Isilon estimates [35% space reclamation](#) from its dedupe algorithm
- NetApp estimates [23% space reclamation](#) from its dedupe algorithm

## Storage Vendors' 50% Less Storage Guarantee

Due to either in-line (real-time) or post process compression or dedupe, leading storage vendors have made the pledge:


- IBM offers a solution, titled “**Deploy Real-time Compression and Save**” - If you can't store 50% more data, [IBM will make up the difference!](#)
- NetApp guarantees its customers who will use 50% less storage for virtual environments or [NetApp will make up the difference!](#)
- HP guarantees that you'll reduce your storage capacity requirements by 50% or more or [HP will make up the difference!](#)
- EMC VNX Unified Storage Delivers 25% More Storage Efficiency — Guaranteed or [EMC will make up the difference!](#)

## Usable Storage Capacity

Any storage system will have a different usable capacity, even though the raw storage capacity is identical.

- [EMC Isilon identifies 80% usable from raw capacity](#)
- [NetApp identifies 55% usable from raw](#)
- [IBM benchmark identifies 50% usable from raw](#)

Below is the snapshot of a few vendors' usable capacity, snapshot usage, SATA Checksum overhead, File System Reserve and Parity.



Source: Nimble

Net Usable relative to raw - factoring in deduplication and compression:

- EMC Isilon implies a total Net Usable of 104% of raw

- (35% of 80% usable reclaimed = 80% + 24% = 108% Net Usable)
- NetApp implies a total Net Usable of 91.3%
  - (66% of 55% usable reclaimed = 55% + 36.3% = 91.3% Net Usable)
- IBM implies a total Net Usable of 90%
  - (80% of 50% usable reclaimed = 50% + 40% = 90% Net Usable)

Combined Dedupe and Compression:

- EMC Isilon offers deduplication but not compression and estimates [35% space reclamation](#)
- NetApp estimates [66% for combined dedupe and compression](#)
- IBM estimates [80% space reclamation](#) from its [compression with Easy Tier](#) and [both are used together](#) to achieve the highest TPC-C bench mark

## Replication

Three main replication issues are: latency, quality, and quantity.

1. Latency is associated with distance and will increase (e.g., from 5 ms to 100 ms) as distance increases from Local Area Network (LAN) to Wide Area Network (WAN) or from East and West coast. Latency is static or cannot be changed and will decrease the amount of data that can be sent across the network. However, leveraging the protocols and applications behavior can put more data into the replication pipe, thus, increasing the application performance.
2. Quality is associated with any issues on either the network or wide area network (WAN). When MPLS and VPNs are deployed, the quality (reliability) of the network is reduced due to dropped packets. As a result, the throughput of replication will be reduced dramatically, even though the amount of bandwidth is still available.
3. Quantity is associated with network bandwidth. A larger the bandwidth (more cost), more data can be replicated.

Isilon claims its SyncIQ is a fast flexible asynchronous replication for disaster recovery protection. Since EMC Isilon lacks the compression and relies on post-process deduplication, a third party Wan Optimization Controller (WOC) solution such as Silver-Peak or Riverbed might be able to accelerate its replication by 10X or 20X.

**Note:** Silver-Peak or Riverbed WOC solution will not help IBM V7000 or Nimble Storage because both storage arrays use a real-time compression. On the other hand, Silver-peak and Riverbed solution can help NetApp or EMC VNX to achieve faster replication because both vendors' replication is based on post-process replication.

However, in Hadoop scenario, Isilon will take advantage of copying 100TB data from one place to another over a 10GB link (also can be considered as a different kind of replication) simply because Isilon uses a pointer, not an actual copy of data, as illustrated in the [link](#) here.

## Isilon 140-node Cluster

Isilon is able to put 140 nodes into one NFS cluster to achieve a massive scale-out NAS.

The newly updated [All SPEC SFS2008 Results](#) published by SPEC, on Dec. 9, 2013 shows an interesting observation with a maximum to maximum comparison between two top scale-out NAS vendors in the world:

288TB/1728 disks with 72 network ports from NetApp vs. 864TB/3360 disks with 140 network ports from EMC Isilon.

Below are excerpts from the [blogs](#):

NetApp exhibits traditional storage system behavior – latency is very low initially and gradually gets higher the more the box is pushed, as one would expect. Isilon on the other hand starts out slow and gets faster as more metadata gets cached, until the controllers run out of steam (SPEC SFS is very heavy in NAS metadata ops, and should not be compared to heavy-duty block benchmarks like SPC-1).

This is one of the reasons an Isilon cluster is *not* really applicable for low-latency DB-type apps, or low-latency VMs. It is a great architecture designed to provide high sequential speeds for large files over NAS protocols, and is *not* a general-purpose storage system. Kudos to the Isilon guys for even getting the great SPEC result in the first place, given that this isn't what the box is designed to do (the extreme Isilon configuration needed to run the benchmark is testament to that). The better application for Isilon would be capacity-optimized configs (which is what the system is designed for to begin with).

Since EMC acquired Isilon in late December, 2010, Isilon has achieved an explosive growth in the scale-out NAS market due to no fork-lift upgrade, easy installation and EMC reputation.

**Isilon's response:** To be continued...

## Hadoop Implementation via “Agile” Approach to “Metrics”

Due to a very complicated Hadoop ecosystem, it is highly recommended to follow the [SCRUM](#) that is a framework for managing the development and deployment of complex products, in order to implement the entire ecosystem correctly. Agile that follows the principle of “Inspect and Adapt” and advocates team empowerment uses Scrum.

### **A few tips to use SCRUM for achieving your goal by Michael Vincent:**

- Don't be tempted to change Scrum



- Scrum exposes inefficiency
- Fix the problem
- Don't shoot the messenger
- Don't reward a wrong person
- Scrum exposes need for change

**Source:** 1) [Scrum Fundamentals Do It Right](#) by Michael Vincent

2) [Case Study of a Difficult Federal Government Scrum Project: FBI Sentinel](#)

**Note:** The FBI abandoned the VCF project in 2005 after spending \$170 million. The project went live on July 1, 2012 after spending 600 million in 12 years by switching to the Scrum.

3) <http://agile2013.agilealliance.org>

**Note:** The topic of SCRUM or Agile is beyond the scope of this writing.

### Gartner Review

Below are some additional expert reviews of the various storage products by Gartner:

- [Market Share Analysis: Network-Attached Storage and Unified Storage, Worldwide, 2012](#)
- [Critical Capabilities for Scale-Out File System Storage](#)

Vendor/Product Name	Dell Fluid FS	EMC Isilon	Hitachi NAS (HNAS) Platform	HP StoreAll Storage	IBM SONAS	NetApp Clustered Data Ontap	Nexenta NexentaStor	Quantum StorNext	Red Hat Storage Server
Product Viability	Fair	Outstanding	Good	Good	Good	Excellent	Good	Good	Fair

**Table 1 Product Viability Assessment**

### Cautions – Information below May Not be a Fact

EMC Isilon may not be the best choice for everyone, although EMC Isilon seems to be an enterprise class scale-out NAS detailed above. It is important to verify features and information about the various products in order to discern facts from “FUD” (fear, uncertainty, and doubt).

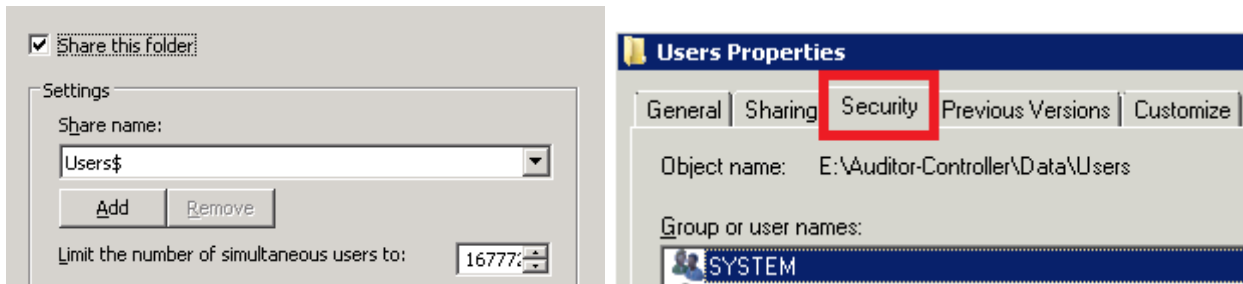
Below are some details about Isilon features or lack of features which were verified with the manufacturer during the research of this evaluation. Note: The topics may not apply to all environments:

## 1. Hiding Shared Folders

Isilon does not provide for the ability to hide shares and mounts from users who are looking for something to hack into the storage; All folders are all visible on the network. In many environments from which a security concern will be a top priority, any organization will consider to implement a system that must be extremely difficult for rogue users to compromise it. Therefore, often, the root of the file share folder(s) and many other critical shares (e.g., sharefoldername\$) will be hidden from the network for security reasons as a first layer of defense. Again and again, the security breaches come from within the enterprises. That might be one of the reasons why Federal Government implements NetApp storage from which NetApp earns storage market share at 60% as of today.

**Note:** [Leading Market Research Firm Names NetApp Data ONTAP as #1 Branded Storage Operating System](#)

Please note that hiding a share folder does not mean a folder cannot be accessed via network. The folder permission via “Security Tab” in Figure 2 must still be applied in order to prevent an unauthorized user to access a folder. Hiding a share folder that also hides the physical path to that folder is just an additional security enhancement to prevent a rogue user from browsing the network to act as reconnaissance. [Click here](#) to view misuse by insiders in Public Sector at 44% (Source: Forrester Resear).



**Figure 2 A Hidden Share Folder**

**Security – AD ACLS**

**Isilon’s response:** This is not accurate. OneFS does in fact support the folder concealment based on permissions. The feature is called Access-based Enumeration and it allows users to see only files and folders that they have access to on a file server. This feature is not enabled by default. The 3-step process below can accomplish the task:

- `isi smb shares modify <share>`
- `--access-based-enumeration {yes | no}`
- Specifies to only enumerate files and folders that the requesting user has access to

**Below is a direct response from B. Scott note to Jeremy:**

I am a bit concerned by the notion that if one cannot see it, it must be secure. Simply inserting a \$ into the share name does not prevent access to the share. If the permissions are set to allow access, the share can still be accessed via UNC pathing. Isilon does not believe in security by obscurity. Access and presentation of shares are defined specifically by the explicit permissions set for the folder and the user. If a user is supposed to access the folder, set the permissions accordingly. If the user is not supposed to access the folder, set the permissions accordingly including share view (enumeration) and folder listing

in AD ACLS. I have seen many problems where things were believed to be secure because it was “out of sight, out of mind” only to find out that the permissions were set incorrectly which allowed undesired access via UNC.

**Final comments on this subject from Jeremy:**

Any organization should consider changing the Isilon’s default setting – disabling the feature for hiding a folder feature - in order to enable the Access-based Enumeration to hide a shared folder for security reasons, as illustrated in a 3-step process detailed above. Author believes that most Isilon customers may not hide a shared folder in its 4,000 installed customers because the feature is turned off by default.

Obviously, Isilon does not support Microsoft’s method to hide a share name in an easier way by simply adding a sign “\$” at the end of the share name, as shown in Figure 2 above.

**2. Scale-up While Scale-out is not Needed**

Scaling should be non-disruptive and independent. All scaling operations should be independent, storage efficiency and transparent to system operation and free of disruption.

- Every time customers want to add capacity on Isilon while the IOPS or additional performance is not needed, they must also pay for CPU, memory, power supplies, IB+NICs and additional network ports because Isilon does not have the ability to simply add another disk shelf, like EMC VNX, HP 3PAR 7400, Dell Compellent SC8000 and etc. This will cause many unnecessary components (CPU, Memory, Power Supplies, IB and Gigabit Ports) to be added into the storage systems when a scale-up is only needed – a storage efficiency issue.
- When computing performance or network access to the storage must be increased; Isilon scaling out approach will accomplish the task by adding additional nodes, but the storage capacity is also increased, even though it is not needed

With a unified storage (e.g., EMC VNX2) or SAN (e.g., HP 3PAR 7400), storage can be scaled to hundreds of TB or above PB by just adding disk shelves. Performance (IPOS) can also be enhanced through the addition of SSDs that can support active data, or upgrade the controller within the same chassis.

As a result, Isilon Scale-up NAS appliances may not be a good choice for those enterprises that do not need a Scale-out NAS and only need a scale-up NAS. Therefore, one unified storage system can address both block-access and file-access storage with more efficient way to increase the storage with more **GREEN** approach. The NetApp unified storage approach is based on the following diagram:

High-end FC	SAN:	Primary
Midrange FC	FCP/FCoE	Secondary
Midrange SAS	iSCSI	Backup
Low-end SATA	NAS:	Compliance
Virtualization	CIFS/NFS	HA & DR

However, there is an exception for large enterprises that need a massive volume to support data. Also, for ease of use, simply adding additional node might be easier than the upgrade. This is a tradeoff.

Author agrees with B. Scott Cassell [whose opinion](#) makes sense.

**Isilon's response:** Correct and intentional so. Isilon is a Scale-Out technology and by design, incorporates the resources necessary to achieve linear scalability into each node. Frankly out customers who recognize this design value and capability select Isilon in large part because of this very attribute. Not including those resources have consequences. RAID based controllers hit a throughput maximum "glass ceiling" because those limited CPU/RAM/Network resources are handling the disk only shelves. Adding more disks further task the limited handling resources and at some point start to create a diminishing performance effects. Isilon intentionally couples CPU/RAM/Network handling capabilities to each node so adding more disks bring along with it the CPU, RAM, and network access capabilities necessary to fully leverage the added performance and capacity. It is this unique capability which makes Isilon the technology of choice for many companies in just about every business sector including the most demanding US Government department that require almost boundless scalability.

### 3. A Non-disruptive Update vs. Disruptive Update

Isilon does not support a non-disruptive update (e.g., a major update from version 7.01 to 7.10 or 7.1.20). A reboot of the cluster is required, even though it may only take about 1 minute or so. A larger the cluster, the longer the rebooting time might be required. If a cluster has 20 nodes, more than 20 minutes outage will occur at least. According to EMC, a non-disruptive operation (NDO) is in its roadmap in the future.

However, Isilon can support a minor update (e.g., from 7.0.1 to 7.0.2 or 7.0.20) as a non-disruptive "rolling update". Each node is updated in sequence allowing the cluster to stay in full operation. Isilon lacks the Quality of Services (QoS) and cannot guarantee performance levels and cannot provide.

**Isilon's response:** "Correct, As of OneFS 7.1, Isilon is designed to 5 nines high availability (99.999%) and minor updates (e.g. from 7.0.1 to 7.0.2) are applied in a rolling update fashion thus each node is updated in sequence allowing for a non-disruptive update on an in production cluster. Isilon does not yet support non-disruptive updates for major code version changes (e.g., a major update from version 7.0.1 to 7.1.0). When a major update is elected and applied, a reboot of the cluster is required and typically takes less than 5 minutes to apply to the entire cluster. Some configurations may take slightly longer than 5 minutes. A non-disruptive major and minor update feature is on the roadmap for 2015."

### 4. Five Nine's (99.999%) vs. Six Nine's (99.9999%)

Due to a non-disruptive update (NDU) not supported by Isilon at this time, EMC has rated Isilon and VNX at 99.999% high availability while EMC VMAX is rated at 99.9999% high availability. Refer to Figure 3 for details.

The VNX or VNX2 is an N by M architecture that relies on dual storage processors for block, while file-based data mover (DMs) sits on the top of the block-level storage. The File systems capacity in a single VNX is limited by the amount of data each data mover can manage. Each DM capacity in VNX2 (the sum of all file systems on a single DM) is 256TB. 512TB/DM will be introduced in the future. The VNX 2 can:

- a) Support 1,792TBs of file systems

- b) Each file system is 16TBs.
- c) The VNX8000 can scale to 6,000 TBs of capacity in total.

The Isilon system delivers availability via node redundancy and node-level protection, while VNX architecture mentioned above uses drive-level protection to deliver 5 nines from which it can support a non-disruptive update.

### What is the meaning of the 9's?

- Disaster Avoidance - 99.9999% Uptime (Six-Nines – only allows for 32 seconds or less of downtime per year!)
- Mission Critical - 99.999% Uptime (Five-Nines - only allows for 5 minutes and 15 seconds or less of downtime per year!)
- Business Critical - 99.90% Uptime (Three-Nines - only allows for 8 hours and 46 minutes or less of downtime per year!)
- Business Important - 99% Uptime (Two-Nines - only allows for 3 days 15 hours and 40 minutes or less of downtime per year!)

**Isilon's response:** "Isilon is designed to 5 nines high availability (99.999%) as of OneFS 7.1. B. Scott Cassell asserted that elective maintenance for major software updates does not count against Isilon's uptime. "Some customers are still running on OneFS versions for several years which is 3 major revs old, and have never experienced an interruption in service."

## 5. Quality of Services (QoS)

Isilon lacks the Quality of Services (QoS) and cannot guarantee performance levels and cannot provide "prioritization of CPU or network resources" while others can.

**Isilon's response:** This is not accurate. OneFS specifically provides "SmartPools" which is performance dedicated tiering to prioritize CPU, disk, and network resources to specific workloads to ensure Quality of Service to the intended workloads. OneFS has a feature called "SmartConnect Zones" to further ensure just those devices that should connect to the dedicated performance tier are allowed to.

## 6. Large files vs. Small files

Isilon is good for large files such as Media files (e.g., videos), but small files are clearly not Isilon's sweet spot due to its large block size at 128KB vs. 4KB from NetApp Fabric-Attached Storage (FAS) storage.

**Isilon's response:** This is not accurate. Isilon is optimized for files sizes for file sizes that are 128KB and larger. Most media files are in the GB and TB range in size and while yes Isilon specifically performs exceedingly well in handling those very large files, most regular documents and other files are larger than 128KB as well.

Isilon handles smaller files than 128KB fine, the difference is the storage is not as efficient at storing the file and corresponding parity data as it is with files 128K and larger. Remember, Isilon provides 80% useable space. Most of the other platforms can only provide 60% or so. If the context is referring to 4k blocks, I would agree with the assertion. 4K blocks is the block based storage unit size. If we are talking

files, which Isilon is by design a file based file system not a block based file system, there are profound benefits to the file based architecture.

## 7. Lost Writes is Not Supported

Isilon cannot protect any data from “Lost Writes” (aka “misdirected writes” or “torn pages”). Only VMAX can achieve this protection.

**Isilon’s response:** This is not accurate. Isilon uses NVRAM Journaling to ensure any write packets that are delivered to the cluster are kept in redundant Non-Volatile RAM until it has been committed to disk and validated. This may have been an issue years ago, but was resolved some time ago. Remember Isilon is a file based storage system not a block based platform.

## 8. Highly Skilled Storage Administrators are Required to Manage Isilon

Below are the excerpts from [GARTNER MAGIC QUADRANT FOR GENERAL PURPOSE DISK ARRAYS-2013](#)

### Cautions

- The VNX system is a dual-controller design that lacks the automatic controller load balancing and scalability available with some competitive products.
- Isilon, which is a true scale-out NAS solution, still lacks native data compression and/or deduplication that may increase its cost as a big data solution.
- The overhead and limitations of various value-add features, such as thin provisioning and autotiering (Fully Automated Storage Tiering for Virtual Pools [FAST VP]) and architectural holdovers, require storage administrators to have high skill levels to ensure optimal performance and the use of new features.

**Isilon’s response:** This is inaccurate. Isilon is by far and according to the leading experts in the industry, the easiest scale-out platform to manage

- Gartner rates Isilon the best in manageability by a significant margin. [Click here](#) to find more comparison information.

It is worth of mentioning that managing EMC VMAX storage also requires storage administrators to have high skill levels in order to create and manage storage to ensure the best performance.

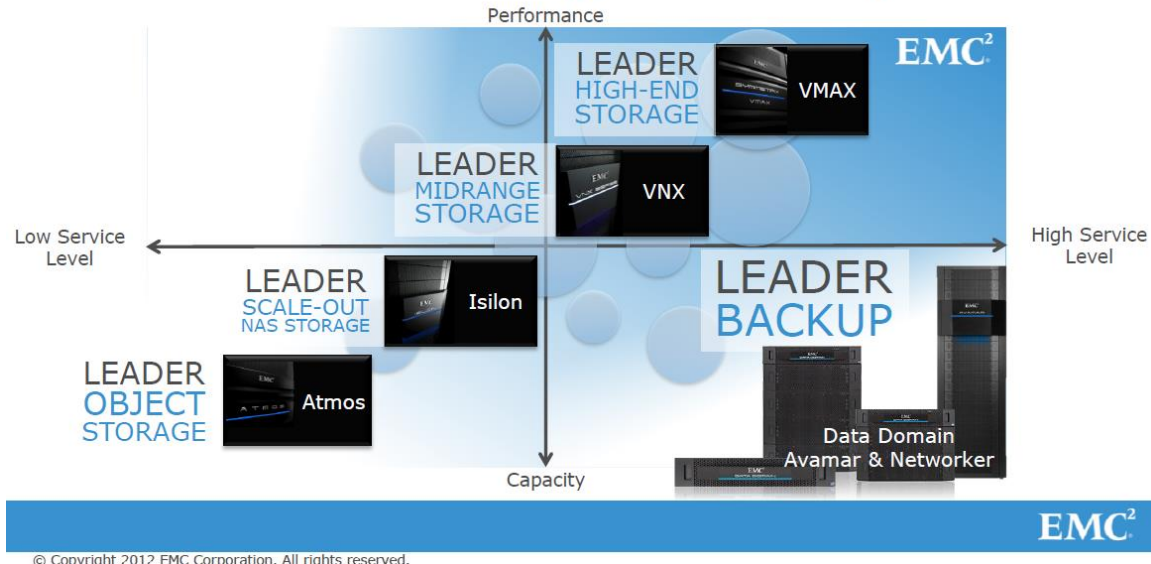
Although Isilon claims it has five 9’s (99.999%) high availability, EMC places Isilon as a leader in the midrange storage for scale-out NAS storage, and VNX and VMAX in a high-end storage environment, as illustrated in Figure 3 that was displayed on a gigantic screen at the EMC Forum, August, 2013 in Long Beach, Calif.

With flash innovation, [click here](#) to view EMC’s position on the Flash technology range from low service level to high service level. As you can see, the EMC Isilon service level is behind the EMC VNX service level due to Isilon does not support a major Firmware upgrade without any downtime (Non-Disruptive Operation).

**The reason is simple:**

Neither Hadoop technology nor a scale-out NAS solution, including NAS only storage (no application can be run from NAS) must be used by an enterprise while a unified storage such as EMC VNX can do both in one box to achieve two birds with one stone.

## Best Of Breed Storage & Backup



**Figure 3 EMC's Vision at EMC Forum 2013**

For very large enterprises who have massive data growth challenges like Facebook and others, EMC Isilon is an excellent choice because every box is inserted into its cluster, higher IOPS and more storage will be achieved with additional parallel processing capability (scale-out NAS) at the cost of additional central processing units (CPUs), Random Access Memory (RAM), Power Supplies, IB and 10/1 Gigabit ports and connections.

[Click here](#) to view the position among VMAX, VNX and Isilon from B. Scott Cassell.

### Challenges

1. Although EMC Isilon is an excellent choice for very large enterprises that have massive data growth challenges like Facebook and others that requires a large scale-out NAS, it will not be a good fit for enterprises who do not need a large Scale-out NAS. For example, one public agency with 8,000 employees only has 3,000 IPOS all the time and needs a large storage due to data growth rapidly. Under this scenario, [EMC VNX2](#), [HP 3PAR 7400](#) or NetApp FAS Series might be a better choice since they all can meet the needs of Scale-up by just adding more disk drives (spindles) into a shelf, instead of adding more CPU, Memory (RAM), Power Supplies, IB+NICs and additional ports detailed above under the section titled "Caution".
2. Microsoft Storage Server based on Windows Server 2012 R2 has dramatically changed the NAS storage landscape for file-based access business:

- Additional enhanced new [server message block \(SMB 3.0\)](#) supports by providing high-performance (higher IOPS) and low-latency access to virtualized data sets with scale-out file server capability due to reduced file protocol chatter
- A Windows 2012 R2 built-in native (RAW) NFS feature that eliminates an external NFS gateway for supporting NFS in a block-access SAN in addition to supporting multicore processors
- Both Native NFS and SMB 3.0 implementations have very good performance - the IT industry benchmark tests show a significant performance increase and surpassed the Isilon's CIFS performance in a significant margin due to the SMB 3.0 features not being supported by Isilon at this writing
- With new SMB 3.0 features built into the Windows Server 2012, Microsoft might drive a wholesale shift away from SANs' complexities and to a simpler, NAS-like storage, such as Microsoft Storage Server HP, Dell and other storage vendors are betting on without spending a lot of money on R&D on NAS development (EMC spent [686.00 million in one quarter](#) in R&D as of Sept. 30, 2013)
- Scale-out file serving with large clustering capability:
  - Up to 64 nodes
  - Up to 8,000 VMs
  - Up to 1TB DRAM/VM
  - Up to 64TB/VM

**Note:** In order to utilize the SMB 3.0 high performance features, a Windows 8 client is required unless Microsoft decides to patch Windows 7 to support SMB 3.0 in the future.

3. NetApp Clustered Data ONTAP adds a horizontal scale-out capability to FAS and V Series systems (which are different from traditional enterprise) by just adding disk drives without sacrificing the unnecessary costs of additional central processing units (CPUs), Random Access Memory (RAM), Power Supplies, Infinite Band (IB) and 10/1 Gigabit ports and connections, from which, EMC Isilon must absorb the additional costs in order to increase the scale-up storage capacity, if an enterprise does not need to have additional IOPS or scale-out NAS needs.

With a scale-up method, Isilon's approach is more expensive than NetApp's scale-up approach detailed above. That's another reason NetApp scale-out NAS is rated No. 2 just behind the EMC Isilon scale-out NAS because it can do both scale-up and scale-out within the single storage system or its own systems.

4. A Few HPC vendors are entering into and providing the Hadoop solutions:

- IBM is well-positioned to participate in the rapidly growing business analytics market because of its participation in the high-performance computing (HPC) market as a server, storage and software provider, and its technical expertise in basic research and vertical markets (**Source:** Gartner).
- [DDN hScaler](#) (DDN received unprecedented Six 2013 HPCwire Readers' and Editors' Choice Awards) enables organizations to scale both storage and compute independently, allowing organizations to implement a solution without costly overprovisioning that fits both workload and budget. The World's First Enterprise Apache™ Hadoop® Appliance (hScaler) is able to



reduce the deployment of Hadoop from a typical 6 to 9 months (just bring it to the floor) to just a few weeks. According to DDN, it solved the following problems:

- Complexity
- Inflexible Configuration
- Multiple Management Interfaces
- Support Headaches
- High TCO

5. EMC does not have server and networking business and relies on a partnership with other vendors such as Cisco and Brocade. As a result, many big enterprises choose to do business with IBM or HP that can offer the entire solution from single vendor.

6. EMC focuses on Storage and Big Data (Hadoop), not on Mobile and HPC according to EMC Forum 2013, while Intel does a good job to bring Hadoop into the HPC field. As a result, Isilon will face more competition from HPC vendors for Hadoop.

### Conclusion

EMC Isilon is an excellent choice for very large enterprises that have massive data growth challenges like Facebook and others that requires a large scale-out NAS. It also should be considered as one of the top choices for Big Data Storage due to [its unique pointer technique](#) detailed above.

According to IDC, 49.8% increase occurred in NAS (file-based storage platform) installation. Isilon is able to address and manage unstructured and semi-structured data explosion with Scale-out NAS appliance.

Generally speaking, EMC Isilon is expensive and has its vertical markets such as entertainment industry, enterprises that needs a massive NAS Scale-out storage and etc., but EMC has its reputation to provide an excellent presales and postsales support.

### Quotes from B. Scott Cassell:

“Most companies use applications that utilize two categories of file types; “structured” and “unstructured”. Structured files are typically files that contain databases, email system data, active virtual machine environment images, etc. Unstructured files are those which are more familiar such as Word documents, text files, spreadsheets, pictures, video files, computer logs, etc. It is estimated that [80% of all data stored is composed of “unstructured” data](#). Take a look at [IDC analyst report](#) for details.

EMC Isilon is a Scale-out NAS (Network Attached Storage) solution which is optimized to efficiently handle small, medium, and large datasets of unstructured files, shares, and application workloads that depend on linear scalability which should ideally not disrupt how

applications use storage as storage is expanded. Isilon's design supports structured data, although only on the small to medium scale. For example it is not recommended to run more than 500 VMs on a high performance Isilon cluster. For small environments that have a limited amount structured data and significantly more unstructured data like the current 80 unstructured /20 structured ratio suggests, Isilon provides the majority of the desired features and performance.

SAN technology is in many cases a better fit for the structured data scenarios and for application that always update the same file very rapidly. NAS is a better fit for providing storage access to applications that produce files as an output which are occasionally updated such as documents and logs rather constantly updated like an Exchange Server Information Store. Isilon is often found in environments where the unstructured file growth is putting capacity or performance upgrade pressure on SAN or unified storage system. When unstructured file storage exceeds 50TB Isilon provides a mid-price option to dedicate SAN storage to structured data and NAS to unstructured.

Planning what type of data is being stored will be an important exercise in determining which technology or technologies are the best fit for the enterprise."

However, Isilon Scale-up NAS appliances may not be a good choice for those enterprises who do not need a Scale-out NAS and only need one unified storage system that can address both block-access and file-access with more efficient way to increase the storage with more **GREEN** approach in addition to avoid paying for CPUs, Memory (RAM), Power Supplies, Network Switches (IB Ports) that connect nodes into a cluster and Gigabit Network Switches that connect cluster to the network of users and applications. With a unified storage, enterprises will not require two separate investments: one for NAS and the other for SAN.

Isilon also has its simplicity in its scale-out NAS except for its Dedupe. According to [GARTNER MAGIC QUADRANT FOR GENERAL PURPOSE DISK ARRAYS-2013](#), Isilon requires storage administrators to have high skill levels to ensure optimal performance and the use of new features (e.g., dedupe).

Always try to keep your system simple that can be easily manageable. Albert Einstein said "Simpler is better, but not simple." A Stacey graphics below illustrates how a system is associated with technology and requirements.


According to Gartner and other analyst reports:

- 80%-85% of IT budgets are spent just "Keeping the Lights On" (KTLO)
- Backlogs are compounding annually at a rate of 10%-20%
- 42% of the initial cost of an application is spent year after year to maintain it (e.g., Oracle Licensing or Microsoft Enterprise Agreements or Software Assurance)

In summary, whether your organization requires a scale-out NAS or a unified storage, it all depends on an environment by considering total cost ownership (TCO) and return on investment (ROI) carefully!

## The Stacey Graph - Complexity

- Simple**  
Everything is known
- Complicated**  
More is known than unknown
- Complex**  
More is unknown than known
- Chaotic**  
Very little is known



Source: Ralph Stacey, University of Hertfordshire

Looking for or evaluating Microsoft Windows Storage Server 2012 R2 (WSS) is always a good idea, especially for small business markets (SBMs) or mid-range enterprises that consist of approximately 85% business in the U.S. because WSS has a built-in native (RAW) NFS feature that eliminates needing an external NFS gateway and introduces new [server message block \(SMB 3.0\)](#) file server features that increase the CIFS performance dramatically and add SMB 3.0 scale-out file server capability with 64 cluster nodes (8,000 VMs per cluster) and ReFS ([Resilient File System](#)) capability from which the newer technology might turn storage into a system that gives an enterprise a competitive IT advantage.

Another import factor is that any customers will not need to wait or depend on a third party's firmware or OS update prior to being able to leverage Microsoft's new and many enhanced features (e.g., the new native NFS or SMB 3.0 features) to maximize their investment.

Whenever possible, always reduce or eliminate a third party dependency to simplify your environment. Today, many enterprises still cannot support Microsoft Windows Server 2012 after Microsoft released its OS on September 4, 2012 while still paying for the Enterprise Agreement (EA) or Software Assurance (SA) year-after-year. This is because the Virtualization brings an added layer of complexity that leads to IT team spending lots of time analyzing data by using various monitoring tools, and trying to figure out how to fix application performance problems or disk bottleneck while either having no time to upgrade the hypervisor up to date or waiting for a hypervisor vendor to fix the various compatibility issues.

Often, today's solutions will not address tomorrow's problems. The added costs associated with labor always challenge any organization to grow its virtualization and cloud environments.

Finally, buying a unified storage or SAN array might be able to address all your storage needs in SMBs or mid-range enterprises in order to leverage your investment and achieve simplicity that in turn reduces the OPEX etc., while Isilon is an excellent scale-out NAS product to address a vertical market or a large enterprise file-based storage demand, including Hadoop etc.

## Recommended Reading

1. [Isilon Data Availability Protection Whitepaper](#)
2. [GARTNER MAGIC QUADRANT FOR GENERAL PURPOSE DISK ARRAYS-2013](#)
3. [DCIG 2013 Private Cloud Storage Array Buyer's Guide.](#)
4. [State of Texas Moves More Than 100,000 State Employees to Microsoft Cloud](#)

The State of Texas is moving more than 100,000 employees onto Office 365 at a cost of about \$3.50 per user, per month, making it the largest statewide deployment of email and collaboration services in the U.S.

5. [How New York City is going to Consolidate 50 Data Centers from 40 City Agencies into One Location](#) (Source: InformationWeek)

## Acknowledgement

Thank IT Summit and B. Scott Cassell, Principal Solutions Architect, EMC Isilon Storage Division, for hosting the event and delivering an outstanding Hadoop presentation at the IT Summit 2013 in Long Beach, Oct. 23, 2013, respectively.

Thanks for [SoCal Code Camp LA 2013](#) at USC, Nov. 9, 2013 – A Scrum topic.

Thanks EMC for allowing me to use a few graphics in my notes for clarification purpose.