

EMC ISILON HOME DIRECTORY STORAGE SOLUTIONS FOR NFS AND SMB ENVIRONMENTS

Best-Practices Recommendations for Capacity and Performance

Abstract

This white paper provides technical information to plan and implement an EMC Isilon NAS solution for home directories, including access management, data protection, storage capacity, and cluster performance. This solution offers SMB and NFS-based network access support for centralizing end-user home directory provisioning, management and support.

November 2012

Copyright © 2012 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided "as is." EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

EMC2, EMC, the EMC logo, Isilon, OneFS, SmartCache, SmartConnect, SmartPools, SmartQuotas, SnapshotIQ, and SyncIQ are registered trademarks or trademarks of EMC Corporation in the United States and other countries.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number H11152

Table of Contents

Executive summary	5
Introduction	5
About this guide	6
Intended audience	6
Assumptions	6
Home-directory requirements	6
Audience assumptions	6
Prerequisites	7
Revision history	7
Overview of Isilon features	8
File Sharing on an Isilon storage cluster	8
NFS access and compatibility	8
SMB access and compatibility	9
Authentication providers and access management	9
Active Directory	10
Integration features and functions	10
Lightweight Directory Access Protocol integration	12
Considerations and guidelines for directory authentication	13
Best-practices recommendations for directory authentication	13
Capacity planning and management	14
Home-directory usage profiles	14
Sizing guidance for storage capacity	15
Managing capacity utilization with SmartQuotas	15
SnapshotIQ	16
Capacity-planning considerations	17
Capacity planning best practices	18
Performance and sizing guidelines	18
File-services protocol differences	18
Total vs. active user connections	19
Planning for disk-pool performance and management	19
Disk Pools	19
File pool policies	20
File-pool management	21
Workload types	21
Additional information	21
Considerations and guidelines for file-pool management	21
Best-practices recommendations for file-pool management	22
Planning for network performance and throughput	22
SmartConnect overview	22
Using SmartConnect to manage client connections	23

Configuration options.....	23
Additional information.....	25
SmartConnect considerations and guidelines.....	25
Best practices recommendations for SmartConnect management.....	26
Conclusion.....	26
References	27
About EMC Isilon.....	28
Contact EMC Isilon.....	28

Executive summary

Organizations across industries are challenged by the unprecedented scale of growth of end-user enterprise data. At the same time, IT departments need to ensure that, as the storage footprint of critical end-user grows, they can scale out the underlying storage platform in a way that minimizes the disruption resulting from storage expansions and migrations, while protecting data availability, capacity, and performance.

EMC® Isilon® Scale-Out Storage addresses these challenges by providing a highly scalable storage platform that integrates seamlessly with existing environments—both Windows- and UNIX-based enterprises—while enabling the consolidation of enterprise file services to a single access point. In addition to the ease of scalability and management, an Isilon storage cluster also provides seamless data protection and recovery, simplifying storage management for both end-users and administrators.

The simplicity of implementing, managing, and scaling Isilon storage increases business efficiency and while reducing operational and administrative overhead for IT and business users alike. An Isilon storage cluster addresses the challenges of large-scale file shares and home directories as it consolidates existing file services, eliminates file-server sprawl, improves storage utilization and reduces administrative costs.

The ease of implementing, managing, and growing an Isilon cluster results in high business efficiency with little or no change in administrators' time or efforts as an Isilon storage cluster grows.

- Isilon storage eliminates data migrations—the complex storage management tasks needed to mitigate hot spots, to expand existing capacity, and to balance capacity among multiple file-system or RAID-group volumes.
- Automated, policy-based storage tiering reduces total cost of ownership by moving inactive files to the most cost-effective tier.
- Quota management and thin provisioning are made simple and flexible.

Introduction

This white paper outlines the principles and concepts for deploying an Isilon storage cluster as a file-server/storage repository for end-user home directories in an enterprise environment. It includes architectural explanations of the technologies and features associated with providing both NFS- and SMB-based file services, as well as EMC's technical recommendations and best-practices settings for optimal performance, management and support.

Much of the relevant information for planning, provisioning and supporting end-user home directories on an Isilon storage cluster is available through existing white papers and guides from EMC Isilon. As such, this Guide will attempt to minimize the amount of duplicate content by including only the information relevant to home-directory file services. The References section at the end of this Guide includes links to those other guides and white papers for additional information.

About this guide

While home-directory services are often categorized and treated as simply a subset of general file services, the workflow and performance characteristics often differ significantly from 'file services' as a generalized solution in many cases. This Guide is intended to assist storage and file-services administrators in planning for the use cases and technical recommendations specific to provisioning and supporting end-user home directories on an Isilon storage cluster.

Intended audience

This Guide is intended for experienced system and storage administrators who are familiar with file-services and network storage administration.

Assumptions

The following assumptions are made as part of this Guide.

Home-directory requirements

This Guide is predicated on the assumption that as a network service, home directories have the following characteristics in common that differentiate them from enterprise file services:

- Home directories are typically used for smaller-scale storage of less-critical organizational—or even personal—data, and generally have less-stringent throughput requirements overall, than enterprise file services.
- End-user access to home directories is much more intermittent, with periodic short bursts of traffic from a given user followed by long periods of inactivity.
- IT service-level agreements (SLAs) for home-directory support, particularly in the area of data-restore requests, tend to be less rigorous than enterprise file services.
- Per-user storage capacity is restricted to a standard amount of storage space, with hard enforcement limits preventing users from exceeding their allotted capacity.
- Individual home directories have unique permissions settings that restrict access to the intended user. Additional access is provided for administrative support and backup-services as necessary.

Audience assumptions

This Guide assumes the reader has an understanding and working knowledge of the following:

- NFS and/or SMB storage protocols, as appropriate for the specific organizational requirements
- Isilon scale-out storage architecture and the EMC Isilon OneFS® operating system
- Additional Isilon data protection and management software including Isilon SmartConnect™, SmartPools™, SnapshotIQ™, and SmartQuotas™
- File-system management concepts and practices, including provisioning, permissions, and performance optimization.

- Integration practices for connecting and establishing authentication relationships with Active Directory or other LDAP-based platform(s), as appropriate

While this Guide is intended to provide a consolidated reference point for systems administrators and managers looking to deploy end-user home directories on an Isilon storage cluster, it is not intended to be the authoritative source of information on the technologies and features used to provide and support a file-services platform. Please refer to the referenced documentation in the [References](#) section at the end of this Guide for more information.

Prerequisites

Some of the features recommended and described within this document may require additional per-node licensing from EMC to enable and use. For more information, please contact a representative from EMC or an authorized reseller.

Revision history

Table I: Revision History

Date	Version	Author	Change Summary
September 2012	0.9	Brad Garvey	Initial Document
October 2012	1.0	James Walkenhorst	Expanded content to include performance and sizing guidance, best practices

Overview of Isilon features

The successful administration and support of home-directory services in an enterprise environment require that the management overhead be minimized and automated as much as possible.

To be successful at a large scale, enterprise home-directory services require that the management burden be offloaded from an organization's IT administrative team(s) as much as possible. In most cases, management overhead can be reduced through a combination of environment simplicity and automated management features.

Isilon simplifies the management of petabytes of file data through automated processes that balance client and internal workloads and capacities: rapid and simple provisioning of storage capacity and storage-related services; automatic generation of snapshot and quota services to protect user data and storage capacity utilization.

The unique architecture of an Isilon storage cluster enables the following:

- Multi-petabyte-sized NAS-based home-directory data storage managed by a single administrator
- High business efficiency through automated storage tiering, rapid and simple user-managed data recovery, automatic distribution of user and client connections across the storage cluster, and inherent storage-utilization rates greater than 80 percent.
- High-availability storage, including a high degree of resilience against storage component and network failures, and the elimination of client and data migrations

This section provides an overview of Isilon storage capabilities and their applicability to an enterprise home-directory services environment.

File Sharing on an Isilon storage cluster

The /ifs directory is the root directory for all file-system data on an Isilon storage cluster, and is automatically shared via both NFS and SMB. Additional shares and exports can be created within the /ifs directory tree.

Under the /ifs directory, two additional directories are automatically created when the Isilon storage cluster is initially provisioned: the /ifs/data directory for departmental and organization-level data; and /ifs/home for end-user home directories.

Note: Physical placement of home-directory data can be applied on a per-user basis, independent of the path of user home directories within the file-system tree. More information about the physical location of Isilon file data is provided in the [SmartPools disk pools](#) section of this Guide.

NFS access and compatibility

For Network File System (NFS) clients, OneFS supports NFS v2, NFS v3, and NFS v4 protocols. The Isilon cluster's default NFS mount point, the /ifs directory, enables UNIX and/or Linux clients to remotely mount any subdirectory, including subdirectories created by Windows clients. Linux and UNIX clients can also mount ACL-protected subdirectories created by a storage administrator.

For NFS v3 and NFS v4 clients, OneFS offers kerberized session support for authentication and access management.

SMB access and compatibility

For Windows clients and users, OneFS supports SMB1 and SMB2 access protocols. The default `\ifs` shared folder gives Windows users access to file system resources via the network, including directories and files created by UNIX and Linux clients.

This SMB compatibility feature means that, to a Windows client, an Isilon storage cluster looks and acts just like a Windows file server on the network: it can be accessed via a standard Windows server name or Fully Qualified Domain Name (FQDN), it can integrate with Active Directory for authentication and permissions management, as well as acting as a Distributed File System (DFS) file server node on the network. SMB users see data on an Isilon storage cluster via mapped network drive or via Universal Naming Convention (UNC) address and share name.

More information about Active Directory integration and access control for files and directories is provided in the following section of this Guide.

Authentication providers and access management

Authentication services provide data security by verifying users' identities before granting access to files and directories. Isilon storage supports several methods of user authentication. Based on the results of these administrator-defined authentication policies, an Isilon storage cluster allows or blocks access to stored data. An Isilon storage cluster can be accessed via any of several different application-layer protocols—including SMB, NFS, HTTP, FTP, and SSH—and authentication policies can be planned and implemented across any or all of these protocols. The appropriate data-security and protocol-level access settings should be consistent with the organization's specific requirements and overall security policies.

File- and directory-access permissions are enforced consistently across protocols, regardless of the actual security model being used. A user is granted or denied access to a file when using SMB as with NFS. Traditional UNIX (NFS) permissions are set on the file system by default. By using Windows Explorer or OneFS administrative tools, standard Windows Access Control List (ACL) permissions can be applied to directories and files.

In a multi-authentication-provider scenario, OneFS also supports configuring Windows-based ACLs for local, NIS and LDAP groups, as well as AD groups and users. After a file or directory has been configured with an ACL, however, the previous UNIX-mode bits for that object are no longer enforced.

Isilon supports the following authentication providers:

- Active Directory (AD) services
- Lightweight Directory Access Protocol (LDAP)
- File-based database
- Local security database
- Network Information Service (NIS)

This Guide focuses primarily on integration with AD and LDAP security providers.

Active Directory

Since an Isilon cluster can seamlessly emulate a Windows-based file server, it can also seamlessly integrate with Active Directory to provide home-directory file services as an enterprise platform. In addition to the native SMB functionality that Isilon provides, it also integrates with an organization's existing file- and directory-permissions models, and automated mapping of AD users with their designated home directories on the Isilon storage cluster.

This section reviews Active Directory integration features and EMC best practices for planning and provisioning AD-based home-directory services.

Integration features and functions

When an Isilon storage cluster is joined to an AD domain, a single computer account is created in the domain. This account is used to establish and maintain the trust relationship between the cluster and the domain, and enables the authentication of users from the AD forest and the authorization of object access on the cluster, regardless of which node within the cluster a user or client connects to. Unless specified otherwise, joining an Isilon storage cluster to AD automatically enables AD domain mode integration, in which the cluster leverages existing user and group objects to create and enforce ACLs on files and directories.

Cluster naming

While an Isilon storage cluster is joined to an Active Directory domain using a single name, the resulting trust relationship between the cluster and AD is automatically extended to include all names by which the cluster is available on the network.

As an example, if SmartConnect network pools are used to manage client connections for optimal performance and client load balancing, each SmartConnect pool will require a unique FQDN on the network. Despite the multiple names to which an Isilon storage cluster can respond, only one machine account—the one corresponding to the name configured on the Cluster Identity settings page—is necessary for the cluster within AD.

SmartConnect pools are explained in more detail in the [SmartConnect Overview](#) section of this Guide.

Permissions

ACL permissions, as applied to files and directories shared via SMB, can be created and modified through any of the following tools:

Shared-folder permissions

- Computer Management MMC console from a Windows client
- OneFS File System Explorer WebUI
- Command-line interface (SSH): 'isi smb' utility
- Command line interface (Windows): 'net share' utility

In addition to creating and editing share-level permissions, any of these tools can also be used for creating shares on an Isilon storage cluster.

File and directory permissions

- Windows Explorer
- Command-line interface (SSH): 'chmod' utility
- Command-line interface (Windows): 'icacls' utility

Note: Using the 'chmod' utility to modify ACLs for files and directories is only available when logged in via SSH to an Isilon cluster directly. The ACL editing feature of 'chmod' is unavailable on Linux and UNIX systems that remotely mount a OneFS export.

Integration with UNIX-style permissions

When an administrator uses a Windows client to change the ACL settings of a file, no information is lost because OneFS stores the original ACL and replaces it. The same holds true when an administrator uses a Windows client to change the permissions of a file with mode bits. OneFS maps the mode bits to a synthetic ACL and because the ACL model can capture the full range of POSIX permissions, so no security information is lost. In such cases, OneFS replaces the file's synthetic ACL with an actual ACL that is equivalent to the mode bits.

The situation is different when a 'chmod' or 'chown' command modifies the permissions of a file protected by an ACL. In this circumstance, OneFS must map the permission changes between two disparate security models. To do so, OneFS, in its default setting, merges the ACL with a patch derived from the change in mode bits. In most cases, this preserves the ACL information and minimizes conflicts between expected and actual effective permissions.

OneFS creates a heterogeneous environment with four primary types of network access to files. These access types also apply to directories and other securable system objects:

- UNIX client accessing a file stored on the Isilon cluster over NFS
- Windows client accessing a file stored on the cluster over SMB
- UNIX client connecting by NFS to a file that was stored on the cluster by a Windows client over SMB
- A Windows client connecting by SMB to a file that was stored on the cluster by a UNIX client over NFS

When a UNIX user requests a file protected with POSIX mode bits over NFS, OneFS controls access by using the file's POSIX permissions. The process is similar when a Windows user requests a file with an ACL: The rights in the user's access token are evaluated against the file's ACL under the Windows security model.

When a Windows user attempts to access a UNIX file, a problem occurs because the permissions set with POSIX-mode bits are incompatible with the Windows security model. Similarly, when a UNIX user requests access to a Windows file, the permissions set with an ACL do not work with the UNIX security model.

Because the two models are different and because the Windows model has a richer set of rights, there is no one-to-one mapping between the two types of permissions. As a result, a file's security could be compromised, or a user who expects to gain access to a file could be denied access.

More information on running an Isilon storage cluster in multi-protocol mode is provided in the “EMC Isilon Multiprotocol Data Access with a Unified Security Model” white paper, a link to which is included in the [References](#) section at the end of this Guide.

Creating home directories in Active Directory

An Isilon storage cluster that functions as a Windows file server in an AD forest supports the same method of home-directory creation and end-user association as any other Windows file server.

Since the /ifs/home directory is already present on the cluster, creating an SMB share that presents the /ifs/home directory to SMB clients as a shared folder can be accomplished via any of the tools mentioned in [Shared-folder permissions](#) above. Once this folder has been shared, home directories can be automatically populated as subdirectories of the /ifs/home directory using standard AD methods: PowerShell scripts, batch-mode scripts, or Active Directory Users and Computers (ADUC) using the ‘Profiles’ tab and the %username% variable. Home directory redirection in Active Directory

Organizations with end-user home directories on an Isilon storage cluster can use Group Policy Object (GPO) redirection of the default end-user “My Documents” directory from the local Windows client to their Isilon-based home directory. This redirection facilitates centralized user-document storage and end-user portability across multiple Windows devices.

Home-directory default and custom permissions

If the %username% variable is used in ADUC, the corresponding directories will be created automatically on the Isilon storage cluster under the designated home-directory share, as with any Windows server. While AD will automatically set the home-directory permissions to grant Full Control access to the corresponding user account for that directory, it will also propagate the file-level permissions from the parent directory into the newly-created home directory.

Domains and trusts

In a multi-domain environment, file and directory access between the Isilon storage cluster and any AD domains that are trusted by the cluster’s parent domain can be granted explicitly to users and groups in those trusted domains.

Users and groups in untrusted domains can be granted access to data on an Isilon storage cluster through the Advanced Authentication Settings page of the OneFS administrative console, in which unknown or untrusted domains are mapped to a trusted domain.

Lightweight Directory Access Protocol integration

In addition to the features and integrated security settings enabled by Active Directory, an Isilon storage cluster can be configured to authenticate users and groups against a Lightweight Directory Access Protocol (LDAP) repository to grant or block access to data stored on the cluster.

- The LDAP service of an Isilon storage cluster supports the following features:

- Users, groups, and netgroups
- Customized LDAP attribute mapping
- Simple BIND authentication (both with and without SSL)
- Redundancy and load balancing across all servers with identical directory data
- Encrypted passwords

Configuring and enabling LDAP integration

On an Isilon cluster, the necessary LDAP settings, including the cluster's base distinguished name (base DN), port number, and at least one LDAP server, must be configured on the cluster in order to enable LDAP.

LDAP configuration

The base DN, also referred to as the search base, identifies the record in the directory from which searches initiated by LDAP clients occur. Base DNS may include a common name (cn), locality (l), domain controller (dc), organizational unit (ou), or other components.

Enabling and disabling LDAP integration

Once the LDAP configuration options have been set in OneFS, the LDAP service is automatically enabled. It can be disabled at any time, either explicitly, or by removing all servers and the base DN from the LDAP configuration settings page. Removing these entries will automatically disable the LDAP service.

Creating home directories for LDAP users

As with Active Directory integration, home-directory creation can be created manually, or scripted and provisioned automatically, and mounted via either SMB or NFS protocol connection.

Considerations and guidelines for directory authentication

The following guidelines are included and recommended within this Guide to satisfy common availability, performance, and security requirements:

- The default share permission for an Isilon storage cluster grants the Read share permission to the Everyone group and Full Control to the Domain Admins group. Unless this default setting is changed from the default setting, non-administrative domain users may not be able to write to their own home directories.
- Redirecting the default location of each user's "My Documents" directory to their respective home directory on the Isilon storage cluster is recommended in most instances. The practice of enabling roaming profiles on the Isilon storage cluster should be undertaken with more care and consideration for the overall impact to the organization. While the storage cluster itself can scale to whatever disk-space levels are required of it, roaming profiles will likely increase the total logon time for end users.

Best-practices recommendations for directory authentication

This section contains EMC's overall recommendations for optimizing directory security and management simplicity.

- Prior to any bulk creation operation of home directories in AD, EMC recommends modifying the parent directory permissions to restrict file and directory access to the appropriate administrative and service accounts only, then testing the resulting configuration to ensure compliance on the cluster's designated home-directory repository to ensure compliance with organizational security policies.
- Using ADUC, along with the %username% variable, to manage home-directories usually simplifies the process of creating directories and managing permissions. Using scripts to create new directories, assign directories to AD user objects, and automate permissions settings is more error-prone in general. If a scripted-management process is the preferred approach, EMC recommends testing the resulting configuration to ensure the correct execution of the process as intended.
- If ACL changes to a file or directory are necessary, EMC generally recommends using native Windows tools, such as Windows Explorer or Windows-based command-line utility, to make the change(s). Using UNIX-based tools may introduce conflicting settings that require extensive troubleshooting efforts to resolve.
- If the same home-directory data will be accessed by both NFS and SMB clients, EMC recommends setting the Isilon storage cluster's global permissions policy to balanced mode. If any of the available balanced-mode policies are unsuitable for a particular mixed environment, then manual configuration of the policy is recommended.

Capacity planning and management

At a high level, capacity planning entails scaling an Isilon storage cluster to accommodate the multiple, competing demands of the combined workload(s) that those resources will need to support. In the case of home directories, workload requirements are driven by multiple factors: disk capacity to accommodate the combined data-storage requirements of all targeted users; sufficient disk throughput to support the combined transactional requirements of all users, and enough network bandwidth to provide adequate throughput between users and storage. Capacity planning entails designing a storage configuration that simultaneously meets all these performance objectives, and ensuring that, as the end-user population grows, or as workload profiles change, the Isilon storage cluster is reconfigured to adjust to the new capacity and performance requirements.

Home-directory usage profiles

Most file-services utilization profiles for end-user home directories are likely to include the following workload characteristics:

- User directories are mapped automatically at the time of user login, typically via either login script or persistent user-profile connection.
- Most users store traditional Office files (documents, spreadsheets, presentations, etc.), images, and streaming media on their home directories, rather than high-transaction, high volume data sets.
- Per-user connections in most scenarios are highly intermittent, resulting in short bursts of per-user data transfers on demand, followed by long pauses with no activity between the user and the directory-storage server.

- Enterprise requirements and user expectations for home-directory throughput and performance are subject to different standards than enterprise file services requirements.
- Home-directory data is often retained for very long periods of time without being accessed or modified.
- Home-directory snapshots and backups are often managed under separate capture and retention policies than enterprise file-services data.

The specific performance requirements for hosting end-user home directories on a consolidated storage platform will vary by organization—or even by department within an organization—but user directory utilization patterns will broadly fit the above description, with variations in degree rather than in substance.

Sizing guidance for storage capacity

Determining the overall disk capacity requirement is a relatively straightforward process. The objective of this process is to calculate the amount of disk space necessary to provide sufficient disk capacity to all expected users over the period of time that will be covered by the initial acquisition period, i.e. how long before a capacity expansion adds more disk space to provide home-directory storage. If the Isilon storage cluster undergoes a capacity expansion once per year, then the initial node acquisition needs to provide sufficient disk capacity to last for the entire year.

The following sizing factors are necessary to accurately estimate the amount of disk capacity necessary for home-directory storage on an Isilon storage cluster.

- Number of users with home directories on the cluster
- Expected disk-capacity allocation per user. An accurate assessment of this number may require the inclusion of additional factors, such as snapshot settings (e.g., the size of each SnapshotIQ operation and the rate of change to the home-directory data set), as well as space, archive and retention policies, and quota enforcement settings, as applicable.
- SmartPools protection overhead for /ifs/home (or whatever top-level directory will contain the end-user home directories), e.g. N+2:1, N+1, etc.
- Expected rate of growth for the home-directory data set. This may be the result of adding more users, increasing the per-user disk-capacity allocation, or changing SnapshotIQ policies to require more space.
- Expected performance requirement for the home-directories data set, as determined in the performance calculation process [above](#).

Once this information is known, an Isilon technical consultant can determine the total amount of storage space required for home-directory data, as well as the specific node type and configuration necessary to satisfy the capacity requirement.

Managing capacity utilization with SmartQuotas

To help enterprises maximize the long-term value of their critical business data and drive down storage management cost and complexity, Isilon offers SmartQuotas: a simple, scalable and flexible quota management and provisioning software application that integrates with the EMC Isilon OneFS operating system.

SmartQuotas allows administrators to control and limit storage usage across their organization and provision a single pool of Isilon clustered storage to best meet their unique storage challenges.

A quota by definition is the permissible share or proportional part of a total. Applied to storage requirements, it is the amount of storage capacity that is permissible to a certain entity within the Isilon cluster. At its core, a quota system is a combination of accounting, enforcement and reporting. Accounting refers to tracking data owned by resource entities, such as users, groups and directories. 'Enforcement' refers to setting and forcing limits for certain counts. 'Reporting' refers to the mechanism by which such enforcement can be conveyed to the administrators or users.

To apply SmartQuotas, two types of capacity quotas need to be considered: Accounting Quotas and Enforcement Quotas.

Accounting quotas

Accounting quotas monitor but do not limit disk storage utilization, are useful for auditing, planning, or billing purposes. Using Accounting quotas enables the following capabilities:

- Track the amount of disk space used by various users or groups to bill each entity for only the disk space used.
- Review and analyze reports to help identify storage usage patterns, which can then be used to plan for future storage acquisitions or engage with end users to review their overall usage levels for educational or planning purposes.

Enforcement quotas

There are four types of Enforcement quotas:

- Hard quotas, which cannot be exceeded. Write operations which exceed the quota limit will be blocked.
- Soft quotas, which can be exceeded until a predefined grace period has expired, at which point they are treated the same as hard quotas.
- Advisory quotas, which are for informational purposes only, and can be exceeded with no limits
- None. These are accounting-only quotas.

More information on SmartQuotas, including detailed planning and management information, is available in the "Storage Quota Management and Provisioning with EMC Isilon SmartQuotas" white paper, which is linked in the [References](#) section of this Guide.

SnapshotIQ

To effectively protect a file system that is hundreds of terabytes or petabytes in size requires an extensive use of multiple data availability and data protection technologies. As the demand for storage is continuing to grow exponentially the demand for ways to protect and manage that storage also increases.

Historically, data protection was always synonymous with tape backup. However, over the past decade, several technologies like replication, synchronization and Data

snapshots have become main stream. Snapshots offer rapid, user-driven restores without the need for administrative assistance.

OneFS snapshots are highly scalable and typically take less than one second to create. They create little performance overhead, regardless of the level of file-system activity, the size of the file system, or the size of the directory being copied. Also, only the changed blocks of a file are stored when updating the snapshots, thereby ensuring highly-efficient snapshot storage utilization. User access to the available snapshots is via a /.snapshot hidden directory under each file system directory.

Isilon SnapshotIQ software can also be used to create unlimited snapshots on a cluster. This provides a substantial benefit over the majority of other snapshot implementations because the snapshot intervals can be far more granular and hence offer improved RPO time frames. SnapshotIQ can take read-only, point-in-time copies of any directory or subdirectory within OneFS, providing the following benefits:

- Snapshots are created at the directory-level instead of the volume-level, thereby providing improved granularity.
- There is no requirement for reserved space for snapshots in OneFS. Snapshots can use as much or little of the available file system space as desirable.
- Integration with Windows Volume Shadow Copy Service (VSS) allows end-users on Windows clients running Windows XP and later versions to restore from using the “Previous Versions” tab for the file or directory. This reduces the amount of assistance that IT resources need to provide, by enabling users to recover their own data.
- Snapshots are easily managed using flexible policies and schedules.
- Using Isilon SmartPools software, snapshots can physically reside on a different disk tier than the original data.
- Up to 1,024 snapshots can be created per directory, and there is no hard limit of snapshots at the cluster-level.

More information on SnapshotIQ is available from the EMC white paper, *High Availability & Data Protection with EMC Isilon Scale-Out NAS*, a link to which is included at the end of this Guide.

Capacity-planning considerations

Based on the above performance characteristics for home-directory services planning and management, the following guidelines are included and recommended within this Guide to satisfy capacity requirements for a successful home-directory services platform:

- While an application-driven workload hosted on an Isilon storage cluster may require that the cluster be sized first for performance, and second for capacity, the reduced SLAs typically in effect around home directories often mean that disk space is the primary consideration when planning for an adequate home-directory-services solution that leverages Isilon storage.
- Even with SmartQuotas policies in effect, organizations may have exceptions policies that allow certain users or data types to bypass the standard quota-enforcement restrictions. Policies that are less restrictive may result in the home-directories data set growing faster than upfront planning would have suggested.

- Different users may warrant different quota settings. Rather than a single per-user capacity quota and a single exceptions policy for all users, consider a tiered-quota approach, in which different categories of users (e.g. managers, IT administrators, etc.) are allocated a higher quota than other users.
- Some organizations use a stair-step approach to quota-policy allocation, e.g. a base 10GB policy for most users, then a 25GB policy for the next tier of users, then 50GB, then 100GB, etc. This approach allows administrators to increase user quotas as necessary without removing them entirely for the largest home directories.
- While SmartQuotas may provide a constraining factor on data growth, the use of snapshots on home directories has the opposite effect: more frequent snapshots lead to automatic increases in the amount of space in use, and the length of the retention policy (i.e. the standard lifecycle of a single snapshot) determines how long that disk space will remain in use.
- SmartQuotas capacity-restriction policies can be configured to include or exclude the overhead associated with SnapshotIQ data.
- In addition to raw disk space, Isilon's unique architecture means that capacity planning also needs to apply when considering the number of user connections per node that will be required to provide acceptable performance for home-directory data. These considerations and recommendations are listed in the [SmartConnect considerations and best practices](#) section of this Guide.

Capacity planning best practices

This section contains EMC's overall best-practices recommendations for planning and managing home-directory disk capacity on an Isilon storage cluster.

- EMC recommends configuring a separate accounting quota for the /ifs/home directory (or wherever home directories are provisioned on the cluster) to monitor overall disk-space usage and issue administrative alerts as necessary to avoid running out of space unexpectedly.
- If an organization's SLA with respect to home-directory data is different than the default general file-services SLA, then the snapshot schedule and snapshot-retention settings can be adjusted accordingly to reduce the amount of capacity that snapshot operations will consume on the cluster.

Performance and sizing guidelines

This section reviews the concepts and analysis processes necessary for determining the appropriate size and configuration of an Isilon storage cluster that will be used to host user home directories. It is not intended to be an authoritative source for all organizational and technical requirements.

File-services protocol differences

In terms of planning for per-user overhead on an Isilon storage cluster, the following guidelines typify most use cases. They are not intended as a definitive rule for all environments, so EMC recommends validating the individual requirements and profile characteristics of the specific environment prior to committing to a particular Isilon storage cluster design or configuration.

Generally, network- and disk-throughput rates are more a function of the type of workload and the specific use case in effect on a given storage cluster rather than the network protocol in use. The per-user overhead placed on the other components of an Isilon storage cluster—including CPU, memory, and network bandwidth—is more directly a result of which file-services protocol is in use. File services based on SMB client connections typically require a higher amount of overhead per user than do comparable workloads from NFS clients, particularly NFS v3. The exact footprint per user varies by environment, configuration and workload characteristics.

The differences in per-user overhead between NFS and SMB connection protocols may require changes in the Isilon storage cluster design to ensure that cluster resources are sufficient to meet the overall performance targets for the appropriate protocol.

Total vs. active user connections

Storage workloads for home-directory file services tend to be driven more by the number of active user connections rather than the number of connections overall. Even an organization with over a thousand users with home directories on an Isilon storage cluster may not require a high level of sustained throughput to the cluster if only a hundred or so of those user connections are in active use at any given time.

The critical factor for ensuring that an Isilon storage cluster support an organization's required performance targets, therefore, is a clear understanding of the specific end-user access patterns: how many users an Isilon storage cluster will need to support, what percentage of those connections can be expected to be active at any one moment, and what volume and type of workload those active connections will be carrying.

Once these factors have been successfully quantified, a suitable storage cluster configuration can be determined that will satisfy the appropriate performance requirements for the given workload type and volume.

Planning for disk-pool performance and management

This section provides an overview of SmartPools disk pools and disk-pool policies with respect to planning for optimal storage performance and data protection. It is not intended to be the definitive source of best-practices information on SmartPools.

Disk Pools

A Disk Pool is a logical grouping of disks across multiple nodes within a cluster. Each Disk Pool is a homogenous group of nodes with an Isilon storage cluster. For example, S Series nodes with 300 GB SAS drives and 1,200 GB SSD per node would be in one pool, whereas NL Series with 3 TB SATA Drives would be in another.

A single Isilon storage cluster may consist of multiple disk pools, since OneFS enables the grouping of multiple node types into a single file system. Each node type is optimized for a different capacity-to-performance ratio, so it's common for organizations to leverage all these architectures into one cluster in order to ensure an optimum match between their different data sets and the storage nodes that they reside on.

File pool policies

The placement of a particular file or directory on a particular disk pool is automated by the Isilon storage cluster, based on a series of standard and customized file-pool policies, and occurs entirely transparently to users and clients within the logical file-system hierarchy. Moving a file or directory from one disk pool to another does not alter the file's location within the file-system tree, nor does it require clients to be reconfigured to access the file in a new location.

Policy selection criteria

File-pool policies are created and applied at either the directory or file level on an Isilon storage cluster based, on one or more attributes of the data. Attributes that can be used to determine inclusion or exclusion by a file-pool policy may include:

- File name
- File path
- File type, e.g. extension
- File size
- Create time
- Modify time
- Access time
- Metadata-change time
- User attributes

Policy-driven actions

Storage administrators and business units can leverage these attributes to ensure that a particular data set is optimized for both capacity and performance. In creating a file-pool policy based on one or more of the above criteria, the following actions can be applied by the policy:

- Data location, i.e. disk-pool placement
- Data-layout settings: File-pool policies can be used to change the layout of the data on the underlying disk for optimal write and read performance.
- Read and write optimization settings: The SmartCache feature of OneFS, which can be managed within file-pool policies as well, manages read and write caching settings that are applied to a particular policy's data. For certain workloads, this yields better read and write performance, increasing throughput and decreasing latency.
- Metadata placement: File and directory metadata refers to the attributes listed in the [Policy selection criteria](#) above. Moving the metadata for files and directories from hard disk drives (HDD) to solid-state drives (SSD) on the cluster can yield significant performance improvements, particularly for operations focused primarily on accessing metadata, such as directory browsing or file/data search operations based on any of the associated attributes.

Data-access performance and protection can be adjusted and optimized any or all of these available policy-based settings. High performance data, as identified by one of

the above attributes, can be placed on a pool of S nodes, or its metadata moved to SSD (also known as 'metadata acceleration', or both, to increase throughput rates and decrease latency levels.

Alternately, older data, as identified by last-modified or last-accessed date, can be moved to a pool of NL disks to free up capacity on the S or X disk pools for newer, higher-performance data.

Note: The use of metadata acceleration in conjunction with one or more SmartPools file policies requires that the node(s) on which the targeted data is stored be SSD-equipped. A file-pool policy that places data in an NL-node pool, for example, removes the option to enable metadata acceleration since NL nodes do not include SSD.

File-pool management

This section outlines several guidelines for configuring SmartPools file-pool policies for optimal home-directory performance. As guidelines, they are applicable in most environments and most configurations, but specific tuning settings may vary by organization.

For optimal network throughput:

- X-node vs. NL node throughput
- Protection settings and performance impacts

Workload types

As mentioned in the [Home-Directory requirements](#) section at the beginning of this Guide, home-directory workflows and use cases are broadly similar across multiple organizations. Most home-directory usage follows this pattern, in descending order (i.e. most to least frequent usage):

1. Browsing home-directory data, including listing directory contents and searching for data
2. Read operations
3. Write operations

General utilization may vary by organization, but configuring home-directory capacity and performance on an Isilon storage cluster should generally be based on these use cases. While general file-services performance expectations and use cases may drive an entirely different configuration, these factors should be included in the home-directories planning process.

Additional information

For more information on planning and implementing SmartPools file-pool policies, consult the "Next Generation Storage Tiering with EMC Isilon SmartPools" white paper, a link to which is provided in the [References](#) section of this Guide.

Considerations and guidelines for file-pool management

The following guidelines are recommended to ensure that overall performance objectives for home-directory users are satisfied:

- Care should be taken when planning the file-pool policy configuration, in terms of what policies will be applied and in what sequence. Policies can conflict or override one another if not properly planned, analyzed for overall effect, and implemented on the Isilon storage cluster.
- Enabling SmartCache for home-directory data can improve performance, but can lead to data loss if a node loses power or crashes while uncommitted data is in the write cache.
- Home directories with large amounts of stale data—more than 60 days since last access, for example—can be migrated automatically via file-pool policy to archive storage on an NL pool, if available. OneFS provides a file-pool policy template to simplify the management of this process.

Best-practices recommendations for file-pool management

This section contains EMC's overall best-practices recommendations for planning and managing file pools for optimal home-directory performance on an Isilon storage cluster.

- End users whose home directories take a long time to load in a GUI-based file manager, e.g. Windows Explorer, due to a large number of objects—directories, files, or both—should see significantly improved performance by enabling metadata acceleration on home-directory data.
- Most home-directory data should see improved read and write performance by setting the data-access pattern to Streaming, rather than the default setting of Concurrent.

Planning for network performance and throughput

As the number of users connecting to home directories and shared increases the need to balance these connections across network interfaces becomes increasingly important. This section outlines configuration options for network performance optimization.

SmartConnect overview

SmartConnect is a software module of Isilon's OneFS operating system that optimizes network-throughput performance and availability by enabling intelligent client-connection load-balancing and failover capabilities.

Through a single host name, SmartConnect enables client connection load balancing, as well as dynamic NFS failover and fallback of client connections across storage nodes to provide optimal utilization of the cluster's available network connections. By leveraging an organization's existing DNS infrastructure, SmartConnect provides universal compatibility with all client types, eliminating the need for complicated connection management on the client side. With SmartConnect, in the event of a node or path failure, file-system stability and availability are maintained for NFS clients that support automatic path failover.

To a client system, the cluster appears as a single network element. SmartConnect automatically balances incoming client connections across all available interfaces on the Isilon storage cluster, improving performance on the cluster by distributing the workload evenly across multiple network paths and multiple nodes.

Finally, for an Isilon storage cluster that hosts multiple concurrent workloads in addition to end-user home directories, SmartConnect provides administrators the ability to partition workloads by type across the available node interfaces in a cluster. By maintaining multiple SmartConnect pools, and minimizing the number of pools that overlap on a particular node interface, administrators can maintain sufficient network bandwidth for critical workloads on dedicated interface connections.

Using SmartConnect to manage client connections

SmartConnect is available in two versions:

- The SmartConnect Basic version manages client-connection balancing using a simple round-robin balancing policy. SmartConnect Basic is restricted to static IP addresses, and to only one IP address pool per external network subnet. SmartConnect Basic is included with all versions of OneFS as a standard feature.
- The SmartConnect Advanced version of the module, in addition to a basic round-robin policy, offers balancing policies based on CPU utilization, connection count, or network throughput. It also allows the creation of multiple IP address pools (also known as zones) that can then be mapped to multiple unique DNS hostnames. Finally, it supports NFS failover using one of several standard failover policies. SmartConnect Advanced requires a separate license for all nodes in an Isilon storage cluster.

Configuration options

Using SmartConnect successfully for load balancing across an entire cluster's worth of storage nodes, or across a limited subset of interfaces on those nodes, requires the proper configuration of a number of interdependent components. This section describes those components.

SmartConnect static and dynamic pool comparison

While SmartConnect static pools use the same initial-connection balancing algorithms when queried for the zone name of the static SmartConnect pool, they do not provide path failover in the event of an interface failure on the cluster. Static pools assign a specific IP address to a specific node interface in a pool, and if that node interface goes offline, its SmartConnect static IP address goes offline as well.

Dynamic pools provide seamless failover only for NFS clients. Other connection types, including SMB/CIFS and iSCSI, do not support the failover mechanism that SmartConnect dynamic pools provide. Static pools are recommended for connecting those workloads.

DNS configuration

SmartConnect leverages an organization's existing DNS infrastructure by providing a layer of intelligence within the OneFS software application. The client attempts to connect to the Isilon cluster using a SmartConnect name which appears to the client as the hostname of the cluster. It does so by requesting a lookup for that host name from the environment's DNS server. The resident DNS server will forward the lookup request for the delegated zone to the delegated zone's server of authority, in this case the SmartConnect Service IP (SIP) address on the cluster. If the node providing the SmartConnect service becomes unavailable, the SIP will move to a different node in the pool automatically.

SmartConnect pools can be configured on the cluster so that all pools are resolved via a single DNS delegation (if all DNS servers are connected to the same subnet), or via multiple DNS delegations (for configurations in which multiple DNS servers are used on multiple subnets connecting to the cluster). To mount a particular SmartConnect pool, regardless of connection protocol, simply use the FQDN corresponding to the name that was assigned the pool at the time of its creation.

Note: While the name of an existing SmartConnect pool can be changed at any time, this will break any persistent connections from client workstations the next time they are rebooted.

SmartConnect service subnets

The SmartConnect service subnet is the name of the external network subnet whose SmartConnect service will answer DNS requests on behalf of the IP address pool. A pool can have only one SmartConnect service answering DNS requests, though this subnet can be changed at any time.

Note: If the service subnet option is not configured for given SmartConnect pool, then all incoming DNS requests to the cluster for that particular pool will be ignored.

In most circumstances, where the DNS infrastructure and Isilon storage cluster all connect to one network, a single DNS delegation should be adequate for all SmartConnect pools. To ensure proper name resolution, configure all SmartConnect pools to use the same service subnet and ensure that the subnet's service IP address is provided as the delegated name server address on the organization's DNS servers.

Service IP address

The SmartConnect Service IP is the IP address that receives all incoming DNS requests from outside the cluster. SmartConnect answers these DNS requests for each IP address pool according to the pool's client connection policy.

Connection policy

The connection policy determines how incoming requests are distributed to across members of the address pool. The following options are available when setting the connection policy:

Simple Round Robin: Round Robin works on a rotating basis. As one server IP address is handed out, it moves to the back of the list; the next server IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of servers being used. This insures that an orderly sequence occurs. This is the SmartConnect default state.

CPU Utilization: This method examines average CPU use in each node, and then attempts to distribute the connections to balance the workload evenly across all nodes in the cluster.

Aggregate Throughput: This method relies on an evaluation of the overall average throughput volume, and then connection balancing policies are set based on optimizing this volume.

Connection Count: In this algorithm, the number of open TCP connections is determined and an attempt is made to balance these connections evenly per node.

IP address allocation

When choosing how the method by which IP addresses are assigned to member interfaces for the address it is important to understand the difference between “Static node IPs” vs. “dynamic IPs”.

Static IP address allocation will assign IP addresses when member interfaces are added to the IP pool. As members are added to the pool, this method allocates the next unused IP address from the pool to each new member. Once an IP address is allocated, the pool member keeps the address indefinitely unless:

- The member interface is removed from the network pool
- The member node is removed from the cluster
- The member interface is migrated to another IP address pool

Dynamic IP address allocation ensures that all IP addresses in the IP address pool are assigned to member interfaces, which allows clients to connect to any IP addresses in the pool and be guaranteed a response. If a node or an interface becomes unavailable, its IP addresses are automatically moved to other available member interfaces in the pool.

Failover policy settings

If a node goes offline that has client connections established, the behavior is protocol-specific. If the IP address gets moved off an interface because that interface went down, the TCP connection is reset. NFS will re-establish the connection with the IP on the new interface and retry the last NFS operation.

SMB-protocol based connections, however, are stateful. When an IP is moved to an interface on a different node, the connection is broken because the state is lost. In this case the connection must be reestablished.

Additional information

More information on planning, deploying and supporting a SmartConnect pool for optimal availability and performance can be found in the “SmartConnect: Optimize Scale-out Storage Performance and Availability” white paper. A link to this paper has been included in the [References](#) section of this Guide.

SmartConnect considerations and guidelines

The following guidelines and recommendations are included to assist in leveraging SmartConnect for optimal network connection balancing and performance:

- Connection speed (1Gb/s vs. 10Gb/s) has less impact on overall home-directory performance than the number of active users per node.
- For home-directory solutions utilizing an Isilon storage cluster via NFS protocol, connecting users via SmartConnect dynamic pools yields both the highest available throughput results and the highest level of availability in the event of a connection or path failure between cluster and node.

- SMB-based connections should use SmartConnect static pools to mount the cluster, since the SMB protocol does not support session failover to another interface without re-authenticating the connection before a new session can be established.

Best practices recommendations for SmartConnect management

This section contains EMC's overall best-practices recommendations for planning and managing home-directory disk capacity on an Isilon storage cluster.

- Overall network performance on an Isilon storage cluster is at its highest when the network connection is to the same group of nodes hosting the underlying storage data. In other words, if home directories are stored on an X-node disk pool, then mounting end users to their home directories on the cluster via the X nodes' network interfaces will yield the best performance.
- Idle end-user connections, as described in [Home-directory usage profiles](#) above, have a negligible effect overall on cluster resources. When determining how best to configure an Isilon storage cluster's SmartConnect pools for home-directory load balancing, the determining factor should be the number of active user connections, rather than the number of total connections.
- For an Isilon storage cluster that hosts multiple workloads, EMC recommends a connection-balancing policy based on active connections or CPU utilization levels, rather than the default round-robin policy, which does not factor in a given node's existing workloads assigning new connections to the node.

Conclusion

As an NFS or SMB-based file-services storage provider, EMC's Isilon storage specializes in the type of workloads that a home-directory-services solution requires. In addition to delivering a storage platform that is cost-effective in terms of both capacity and performance, an Isilon storage cluster includes features that simplify both the deployment and ongoing support of an enterprise home-directory infrastructure.

- Isilon offers native integration with a number of centralized end-user directory authentication and security management platforms: Active Directory, Lightweight Directory Access Protocol, and Network Information Service are all supported platforms for centralizing and standardizing on end-user accounts and authentication management.
- Home directories can be created and mapped or mounted automatically for end-users, eliminating the difficulty of creating and managing the directories themselves and the unique permissions sets associated with each individual directory.
- Network bandwidth scales as more capacity is added, unlike traditional storage architectures whose network connections typically hit a fixed limit. The flexibility of SmartConnect enables administrators to partition dedicated network connections for specific workloads and to balance those connections across a pool of available interfaces on the Isilon storage cluster. For NFS connections, SmartConnect provides automatic failover to a good connection in the event of a path failure on any of the nodes.

- Isilon simplifies the management of the underlying storage that hosts end-user home directories. OneFS eliminates the need to create and manage any underlying storage volumes or RAID groups prior to provisioning end-user directory capacity. With Isilon, all capacity is immediately and automatically added to a single file-system hierarchy. Capacity expansion is simply the process of joining additional storage nodes to the cluster and allowing OneFS to expand the file-system space automatically.
- In addition to eliminating the overhead of RAID-group, volume, or individual LUN management, SmartPools from Isilon enable administrators to manage protection and I/O optimization levels at the individual file or directory level, without migrating data, taking data offline, or reconfiguring client computers to remap the data.
- In addition to the customizable and dynamic protections against component failures offered by SmartPools' file-pool policies, a home-directory solution based on Isilon storage can leverage the multiple layers of data protection that Isilon offers. The ability to create regular snapshots of file-system data offered by SnapshotIQ integrates with the Windows Previous Versions feature to let users manage their own file-recovery efforts. SyncIQ enables the large-scale recovery of an entire Isilon storage cluster's data set to an offsite cluster. Native support for NDMP-based backups and ICAP-enabled antivirus scanning provides additional layers of protection against data loss or corruption.

Finally, an Isilon storage cluster offers simplicity of management, even at large scales, to reduce the amount of overhead necessary to provide administrative support. By eliminating the need for RAID, volume, or file-system management, Isilon reduces the total cost of ownership associated with provisioning and managing end-user home directories. At the same time, Isilon improves performance levels, simplifies the protection and recovery of end-user data, and reduces downtime.

References

The following documents provide additional and relevant information. Access to these documents may depend on your login credentials. If you do not have access to a document, contact your EMC representative.

- [EMC Isilon Multi-protocol Data Access with Unified Security Model for SMB and NFS](#)
- [Storage Quota Management and Provisioning with EMC Isilon SmartQuotas](#)
- [Next Generation Storage Tiering with EMC Isilon SmartPools](#)
- [SmartConnect: Optimize Scale-Out Storage Performance and Availability](#)
- [EMC Isilon SnapshotIQ](#)
- [Best Practices for Data Replication with EMC Isilon SyncIQ](#)
- [EMC Isilon OneFS Operating System](#)
- [High Availability & Data Protection with EMC Isilon Scale-Out NAS](#)
- [OneFS 6.5 User Guide](#)

About EMC Isilon

EMC Isilon is the global leader in scale-out NAS. We provide powerful yet simple solutions for enterprises that want to manage their data, not their storage. Isilon products are simple to install, manage and scale, at any size and, unlike traditional enterprise storage, Isilon stays simple no matter how much storage is added, how much performance is required, or how business needs change in the future. We're challenging enterprises to think differently about their storage, because when they do, they'll recognize there's a better, simpler way. Learn what we mean at www.emc.com/isilon.

U.S. Patent Numbers 7,146,524; 7,346,720; 7,386,675. Other patents pending.

Contact EMC Isilon

www.emc.com/isilon

505 1st Avenue South, Seattle, WA 98104

Toll-free: 877-2-ISILON | Phone: +1-206-315-7602

Fax: +1-206-315-7501 | Email: sales@isilon.com