

Release Notes

Product: StorNext FX 1.3

SNFX Client Operating Systems:

SGI IRIX 6.5.26, 27, and 28 (MIPS 64-bit)

Solaris 9 and 10 (SPARC 64-bit)

Red Hat Enterprise Linux Advanced Server 3.0,

Updates 4, 5, and 6 (32-bit Intel and AMD; 64-bit for Itanium, Opteron, and EM64T)

Red Hat Enterprise Linux Advanced Server 4.0, Updates 1 and 2

(32-bit Intel and AMD; 64-bit Itanium, Opteron, and EM64T)

Red Hat Linux 9.0 (32-bit)

SUSE Linux Enterprise Server 9, Service Pack 2 (SLES);

(32-bit Intel and AMD, 64-bit Itanium, Opteron, and EM64T)

IBM AIX 5.2, 5.3 (Power Architecture 64-bit)

HP-UX 11i v2 (PA-RISC and 64-bit Itanium)

Windows 2000, Server 2003, and XP (32-bit Intel and AMD)

Date: April, 2006

Contents

Page	Topic
2	Purpose of this Release
2	New Features
3	Operating System Level Requirements
3	Operating System Level Requirements
4	Certified System Components
4	System Requirements
7	Configuring LDAP
8	Resolved Issues
8	Resolved Issues
12	Known Issues
21	Documentation

Purpose of this Release

StorNext FX (SNFX) 1.3 includes several new features and enhancements that extend its capabilities. These release notes describe these new features and enhancements, as well as currently known issues and issues that were resolved for this release. These notes also provide drive and auxiliary firmware compatibility information.

SNFX is a client only version software that is licensed to run in Apple managed environments and used with an Apple metadata controller. It is fully interoperable and compatible with Apple Xsan.

The first version of SNFX was released with Apple Xsan 1.0. Subsequent releases of SNFX and Apple Xsan followed this compatibility:

- SNFX 2.5 interoperable and compatible with Xsan 1.1
- SNFX 2.6 interoperable and compatible with Xsan 1.2
- SNFX 1.3 interoperable and compatible with Xsan 1.3

With the release of SNFX 1.3, the version numbering is changed to more closely align with the Apple Xsan release versioning. This was done to help minimize confusion of compatible SNFX releases with Apple Xsan releases and to maintain consistency with the Apple Xsan product. The product has not changed, only the release numbering.

Visit www.adic.com and www.adic.com/adicSupportShell.jsp?Page=supportBulletins for additional information and updates about SNFX.

New Features

New features and functionality for SNFX 1.3 include:

- **2TB LUN:** 2TB LUN support enables customers to configure LUNS up to and greater than 2TB. This allows for fewer volumes to store large amounts of data.
- **Lightweight Directory Access Protocol (LDAP):** Simplifies configuration and administration of mixed environments. LDAP (Active Directory) centralizes the ACL translation necessary for Windows and UNIX systems to access a shared data set. All Windows servers are pointed to a single LDAP system for centralized and simplified management of file permissions.

Special Configuration Requirements

For SNFX 1.3, the File System client RAM requirement has changed to a minimum of 512 MB.

Operating System Level Requirements

The following table lists the required operating system levels required to successfully operate SNFX 1.3.

SNFX 1.3 Requirements

Operating System Platform	Operating System Levels
AIX	5.2, 5.3 (Power Architecture 64-bit)
HP-UX	11i v2 (PA-RISC and 64-bit Itanium)
IRIX	6.5.26, 27, and 28 (MIPS 64-bit)
Red Hat Linux	Red Hat Enterprise Linux Advanced Server 3.0 - Updates 4, 5, and 6 (32-bit Intel and AMD; 64-bit Itanium, Opteron, and EM64T) Required kernels: U4 - 2.4.21-27.EL or 2.4.21-27.Elsmg U5 - 2.4.21-32.0.1.EL or 2.4.21-32.0.1.ELsmg U6 - 2.4.21-37.EL Both the kernel and kernel source RPM packages must be installed. Install all tools necessary to build a kernel module (including compilers) on all Linux clients and server systems.
	Red Hat Enterprise Linux Advanced Server 4.0 - Updates 1 and 2 (32-bit Intel and AMD; 64-bit Itanium, Opteron, and EM64T) Required kernels: Update 1 - 2.6.9-11 Update 2 - 2.6.9-22
	Red Hat Linux 9 client only (32-bit Intel and AMD) Required kernel: 2.4.20-31.9.legacy
Solaris	Solaris 9 and 10 (SPARC 64-bit) NOTE: Prior to installing StorNext on a Solaris 10 machine, you must install the Solaris 10 Recommended Patch Cluster dated March 10, 2006 or later. To gain support for LUNs greater than 2TB on Solaris10, the following patches are required: <ul style="list-style-type: none"> • 118822-23 (or greater) Kernel Patch • 118996-03 (or greater) Format Patch • 119374-07 (or greater) SD and SSD Patch • 120998-01 (or greater) SD Headers Patch
SuSE Linux	SUSE Linux Enterprise Server 9, Service Pack 2 (SLES); (32-bit Intel and AMD, 64-bit Itanium, Opteron, and EM64T); kernel 2.6.5-7.191-default. Both the kernel and kernel source RPM packages must be installed. Install all tools necessary to build a kernel module (including compilers) on all Linux clients and server systems.
Windows 2000	Service Pack 4 (32-bit Intel and AMD)

Operating System Platform	Operating System Levels
Windows Server 2003	Service Pack 1 (32-bit Intel and AMD)
Windows XP	Service Pack 2 (32-bit Intel and AMD)

Certified System Components

This table lists certified system components that support SNFX 1.3.

Component	Description
Browsers	Netscape 7.x Microsoft Internet Explorer 5.5 and later Mozilla 1.4 and later FireFox 1.0 and later
HBA	Emulex LP8000, LP850, LP9000, LP90002 QLogic: QLA2200, QLA2310, QLA2340, QLA2342
NFS	Versions 3
Multi-pathing Failover Software	<ul style="list-style-type: none"> • Red Hat Enterprise Linux AS 3.0, Update 3 HP SecurePath - version: 3.0c • Windows 2000 HP SecurePath - version: 4.0c-7 • Windows 2003 EMC PowerPath - version: 3.0.6 <p>For information on supported HBA drivers, refer to the RAID vendor's documentation.</p>

System Requirements



Note

The following requirements assume that SNFX is the only application running on your system.

To successfully install SNFX 1.3, the following system requirements must be met:

- [SNFX 1.3 Requirements](#)
- [Windows Memory Requirements](#) on page 5

SNFX 1.3 Requirements

For each SNFX mounted file system, these requirements must be met.

System/Component	Requirement
AIX	SNFX clients: A minimum of 512 MB of RAM is required.
IRIX	SNFX clients: A minimum of 512 MB of RAM is required.

System/Component	Requirement
HP-UX	SNFX clients: A minimum of 1 GB of RAM is required.
Solaris	SNFX clients: A minimum of 512 MB of RAM is required.
Linux	SNFX clients: A minimum of 512 MB of RAM is required.
Windows 2000 Windows Server 2003 Windows XP	SNFX clients: A minimum of 512 MB of RAM is required.
Network LAN using TCP/IP (all clients and servers must be interconnected)	For the SNFX metadata traffic, ADIC requires that a separate, dedicated, switched Ethernet LAN be used.
SAN	SNFX clients: An FC-HBA or equivalent SAN communication device where the storage is visible and accessible to multiple SAN clients.
	SNFX does not support multiple hosts connected through an FC hub device because the resulting propagation of Loop Initialization Protocol resets can cause data corruption.
Client Hard Disk	SNFX requires 200 MB of hard disk space for binaries, documentation, configuration, and log files.
Disk Drives	SNFX only supports the file system when it is running on FC-3 SCSI drives.

Windows Memory Requirements

The SNFX 1.3 release has a number of performance enhancements that enable it to better react to changing customer load. These enhancements come with additional memory requirements. Because the Windows operating system has unusual memory resource limitations it is sometimes necessary to adjust StorNext memory tuning settings to provide optimal operation.

When running multiple file systems in the Windows environment the StorNext memory tuning parameters must be adjusted or the machine will run out of memory. This can be seen is by bringing up task manager and watching the **Non-paged** tag in the **Kernel Memory** pane in the lower right hand corner. Microsoft warns that the maximum amount of non-paged memory consumed cannot exceed 256MB. However, the actual amount varies depending on your configuration. For example, on a machine with 256MB of memory ADIC testing has found that non-paged memory must not exceed 96MB.

Non-paged memory is the most critical resource limitation but not the only one. For example, Microsoft warns that paged memory consumption must not exceed 470MB. Similarly, the actual amount also varies depending on configuration.

Memory exhaustion can be observed in the following ways:

- commands failing
- messages in the system log about insufficient memory
- `fsmppm` process mysteriously dying
- repeated FSM reconnect attempts
- messages in the application log and **cvlog.txt** file about socket failures with the status code (10555) which is ENOBUFS
- messages in the system log about `ConvertExtent failed with error 20`

The solution is to adjust a few parameters on the **Cache Parameters** tab in the SNFX control panel (`cvntclnt`). These parameters control how much memory is consumed by the directory cache, the buffer cache, and the local file cache.

As always, an understanding of the customer's workload aid in determining the correct values. Tuning is not an exact science, and requires some trial-and-error to come up with values that work best in the customer's environment.



CAUTION

The settings in the Cache Parameters tab are relevant to the selected file system only. When running multiple file systems it is necessary to adjust the Cache Parameters settings for EACH file system. Also, the total amount of memory consumed is the SUM of the Cache Parameters settings for ALL file systems.

Settings and Parameters

- **Directory Cache Size** - The first setting to consider is the **Directory Cache Size**. The default is 10 (MB). If you do not have large directories, or do not perform lots of directory scans, this number can be reduced to 1 or 2MB. The impact will be slightly slower directory lookups in directories that are frequently accessed. Also, in the **Mount Options** tab, you should set the **Paged DirCache** option to allocate the specified memory from paged pool instead of the default non-paged pool.
- **Buffer Cache NonPaged Pool Usage** - The next parameter is the **Buffer Cache NonPaged Pool Usage**; the value is in percent (%) and represents the percentage of available non-paged pool that the buffer cache will consume. By default, this value is 75%. This should typically be set to 25 or at most 50 for two or more file systems to avoid over-consumption of the non-paged pool. The minimum value is 10 and the maximum value is 90. Associated with this setting are the **Data Buffer Cache Minimum** and **Data Buffer Cache Maximum** settings. These settings specify the minimum and maximum amount of paged plus non-paged memory (in megabytes) consumed for buffer cache for the selected file system.
- **Watermarks** - The following parameters control how many file structures are cached on the client; they are. These are controlled by the **Meta-data Cache Low Water Mark**, the **Meta-data Cache High Water Mark** and the **Meta-data Cache Max Water Mark**. Each file structure is represented internally by a data structure called the `cvnode`. The `cvnode` represents all the states about a file or directory. The more `cvnodes` that there are encached on the client, the fewer trips the client has to make over the wire to contact the FSM.

Each `cvnode` is approximately 1462 bytes in size and is allocated from the non-paged pool. The `cvnode` cache is periodically purged so that unused entries are freed. The decision to purge the cache is made based on the Low, High, and Max water mark values. The Low default is 1024, the High default is 3072, and the Max default is 4096.

These values should be adjusted so that the cache does not bloat and consume more memory than it should. These values are highly dependent on the customer's workload and access patterns. Values of 512 for the High water mark will cause the `cvnode` cache to be purged when more than 512 entries are present. The cache will be purged until the low water mark is reached, for example 128. The Max water mark is for situations where memory is very tight. The normal purge algorithm takes access time into account when determining a candidate to evict from the cache; in tight memory situations (when there are more than the maximum entries in the cache), these constraints are relaxed so that memory can be released. A value of 1024 in a tight memory situation should work.

Configuring LDAP

The following information describes how to configure the new StorNext LDAP feature in addition to outlining recent changes to Windows configuration tools.

Using LDAP

SNFX 1.3 introduces support for Light Directory Access Protocol, or LDAP (RFC 2307). This feature allows customers to use Active Directory/LDAP for mapping Windows User ID's (SIDs) to UNIX User ID/Group ID's.

Changes to "Nobody" mapping

As with previous releases, if a Windows user cannot be mapped to a Unix ID, the user is mapped to `Nobody`. SNFX 1.3 allows administrators change the value of `Nobody` by using the file system configuration parameters:

```
UnixNobodyUidOnWindows 60003
```

```
UnixNobodyGidOnWindows 60004
```

These parameters are located in the file system configuration file on the server and are manually modified by the Xsan Administrator GUI.

Changes to UNIX File & Directory Modes

When a file or directory is created on Windows, the UNIX modes are now controlled by the following file system configuration parameters:

```
UnixDirectoryCreationModeOnWindowsDefault 0755
```

```
UnixFileCreationModeOnWindowsDefault 0644
```

In previous releases StorNext used per user mode masks. SNFX 1.3 allows one set of values for all users of each file system.



Note

The default values allow more open access to Windows-created files from UNIX systems than in previous versions. Administrators can manually change these values in the file system configuration file on the server or use the Windows or Web GUI.

LDAP Refresh Timeout

Due to the changes in the Windows Active Directory user mappings, services for UNIX can take up to 10 minutes to be propagated to StorNext clients.

User ID Mapping Precedence

If multiple mappings are found for a given Windows user, the following precedence takes place:

- **NIS/PCNFSD** - If mapping exists
- **Fabricated ID's** - If configured "on"
- **LDAP/RFC 2307** - If defined in Active Directory
- **Nobody** - If no other mapping found

Navigating to the **Windows File System Control Parameters pane > Authentication** tab gives the administrator the ability to disable the NIS/PCNFSD mapping on a client by client basis by checking LDAP.

Domain Server Requirements

The domain server must be at least Windows 2000 or above to support RFC 2307 mapping.

Client Requirements

Mapping is done by searching the domain server that the Windows machine is a member of. The user must be logged in under the user account owned by that domain server. If the user logs in under an account of the client machine or another domain, the user will most likely not have a mapping.

Failover Xsan to/from non-Xsan

There is a known deficiency when failing over from a non-Xsan meta-data server (FSM) to Xsan server and visa versa. On failover (possibly requiring a Windows client reboot), the SID to UID mapping will change if the fabrication configurations are different. Therefore, when Xsan is involved, explicitly set the configuration file parameter `UnixIdFabricationOnWindows` to be the same value on all potential meta-data servers.

Changes to Windows Configuration Tools

- **Fields in Configuration Files** - The Windows configuration tools add more fields to the `.cfg` file. The **Configuration Administrator** now includes all fields when the `.cfg` file is saved.
- **The `use strict` Feature Has Been Removed** - When using the **File System Control Parameters**, note that the `use strict` capability has been removed from the **Authentication** tab.

Resolved Issues

The following tables list resolved issues in this release of StorNext.

File System Issues

The following CRs are specific to the StorNext File System.

Operating System	CR Number	SR Number	Description
AIX	33420	n/a	AIX path failover broken because device always opens first device.
	34450	n/a	AIX 5.2 clients crash simultaneously when the customers Java-based application is run on a CVFS-mounted file system.
	34471	305956	The <code>mount -t cvfs</code> command issued twice under AIX mounts the file system twice. Other file system return busy.
	37503	n/a	AIX does not correctly enumerate HBAs.
	39018	n/a	On AIX, move across file systems fails.

Operating System	CR Number	SR Number	Description
IRIX	38875	310052	On IRIX, Mmap fails for files larger 2 GB.
	51662	n/a	POSIX sparse file support defaults to <code>yes</code> in StorNext 2.7. NOTE: This change can have a negative performance impact. The previous behavior can be achieved with an explicit deactivation of <code>posix-sparse-file</code> support.
Linux	24245	221686	On Linux, license key MAC address enumeration fails with multiple NIC's.
	34919	n/a	Linux <code>ioctl</code> call must use <code>copyin/copyout</code> .
	35317	n/a	<code>ld</code> command hanging SNFS file systems.
Solaris	36567	332658	Stack overflow caused client kernel panic on Solaris.
	51662	n/a	POSIX sparse file support defaults to <code>yes</code> in StorNext 2.7. NOTE: This change can have a negative performance impact. The previous behavior can be achieved with an explicit deactivation of <code>posix-sparse-file</code> support.
Windows	31240	270050	The <code>xcopy /d</code> command should only copy source files newer than the destination files, but it copies all files into SNFS.
	33512	299640 280228	On Windows Explorer, the Properties panel hangs and triggers retrieve of a migrated file.
	33558	299636	On Windows 2003, Portmapper exits during system startup.
	33998	n/a	Failed Windows IRP cancellation not handled and led to BSOD.
	34114	n/a	Windows BSOD: <code>ASSERT(f_spinirq == PASSIVE_LEVEL)</code> .
	34326	n/a	<code>cvfsck</code> <code>ASSERT</code> failed " <code>mapaddr != NULL</code> " file <code>cvfssubrs.c</code>
	34375	n/a	BSOD in <code>PurgeCvnode</code> path.
	34685	n/a	<code>Drivelmage7</code> : Unsupported <code>IOCTLS</code> (Windows only).
	34765	n/a	Windows Explorer claims Windows Recycle Bin is corrupted when removing a directory.
	35651	n/a	Windows client cannot access file system after upgrade from SNFS 2.3.2B40 to 2.4.1B59 (GA).
	36689	331056	Windows BSOD <code>ASSERT</code> failed: <code>f_rwlck->rw_state != RW_IDLE</code> .
	37572	n/a	<code>dircache</code> pool config button gets reset in the Windows GUI.
	41454	n/a	In prior releases of StorNext (2.2 through 2.5), the <code>buffercachemin</code> mount option had no effect. With StorNext 2.6 and later, the option will now cause the file system to pre-allocate blocks in the buffer cache at mount time. Also, the default value for this option has been changed from 16 MB to 1 MB. See the <code>mount_cvfs</code> man-page or Windows help file for more information.
	43199 50946	n/a	<code>WindowsSecurity</code> configuration default setting will be set to <code>Yes</code> in a future release.

Operating System	CR Number	SR Number	Description
All	18144	179324	Pathing method definition not included in c:\cvfs\help\cvfs_config.htm file.
	33490	296206	Client panicked after running cvfs stop.
	33544	n/a	Installation procedures needs a 64-bit operating system check.
	33601	n/a	Windows rename failed with full pathname greater than 128 characters.
	33726	n/a	Bad value passed to "cvdb -R" causes panic.
	33747	288626	SNFS - failure to start fsmpm hangs.
	33834	n/a	cvntclnt fails to map drives.
	34163	n/a	Man page missing MaxMBPerClientReserve.
	34264	311750	Windows files erroneously created with UNIX execute bits.
	34666	n/a	UNIX access denied for all files created by unmapped Windows users.
	34835	n/a	Pinnacle Liquid Blue applications fails on SNFS.
	35022	n/a	Asynchronous I/O fails EOF test.
	35159	n/a	Support asynchronous NFS exports in SNFS 2.5.
	35345	n/a	2.5 trickle writes cause severe fsm cpu spike - order N squared algorithm.
	35350	n/a	dm_get_fileattr returns errno 28 when move issued on remote client, 2 if local.
	35395	n/a	MPI benchmark failure.
	35682	n/a	rcp cannot truncate a file with 444 permissions.
	35961	n/a	Pinnacle Liquid Blue fails certain operations.
	36413	n/a	Permissions on cifs share reset after boot.
	36572	n/a	Move fails when a symbolic link is the target file.
	36762	n/a	Cannot send flush intent message still exists on fast machines.
	36854	340454	NFS hang due to deadlock.
	36884	n/a	dircache not configurable on NT and needs to use paged pool.
	37049	341648	Concurrent C++ programs client opening same file on same SNFS fail.
	37178	n/a	DMA versus buffered I/O imbalance.
	37291	n/a	Unexpected data read from sparse file holes.
	37465	n/a	StorageTek BladeStor is just another LSI-RAID disguise.
	37574	n/a	perfmon fsd read counters never active.

Operating System	CR Number	SR Number	Description
All	50103	n/a	In previous versions of StorNext, the modification time of a file was not updated when a file size was explicitly changed using the <code>truncate(2)</code> system call on UNIX or the <code>_chsize()</code> function on Windows. With StorNext 2.7, this has been corrected and setting the size of a file will cause the modification time of the file to be set to the current time.
	51688	n/a	With StorNext 2.7, the output of the <code>cvlabel -l</code> command has changed. LUNs with valid StorNext labels are now displayed as type SNFS , rather than CVFS (as reported by previous releases).
	53713	n/a	<ul style="list-style-type: none"> • Space consumed by open unlinked files on UNIX clients is now counted against quotas. • Quota update after file removal is now delayed until the space deallocation is fully completed.
	51687	n/a	With StorNext 2.7, the output of the <code>cvlabel -ls</code> command has changed. Previous releases of StorNext may not have found the longest serial number available for a given LUN, leading to serial numbers appearing not to be unique among multiple LUNs on the same RAID controller. This problem has been corrected, which may lead to the reported serial number for a given LUN being different between StorNext 2.6 and StorNext 2.7.

Known Issues

The following tables list known issues in this release of StorNext.

File System Issues

The following CRs are specific to the StorNext File System.

Operating System	CR Number	SR Number	Description	Workaround
AIX	51306	n/a	Using the <code>umount</code> command can cause applications to fail.	Make sure that all applications are stopped before invoking <code>umount</code> .
	52300	n/a	Due to a potential bug in AIX (PMR 32512,033), applications using memory-mapped I/O may hang and become unkillable when unmapping a shared memory segment. In turn, this prevents an SNFS file system from being unmounted.	n/a
	54345	n/a	Unsupported XMEM function; error implies no space, but there is plenty of space.	n/a
	33986	n/a	<p>During a StorNext upgrade, if you encounter a DSM failure that resembles the following output:</p> <pre>SNFS UPGRADE 2.6.1(48) Status 1) Upgrade 2.6.1(48) perl Complete 2) Upgrade 2.6.1(48) GUI Complete 3) Upgrade 2.6.1(48) apache Complete 4) Upgrade 2.6.1(48) DSM Failed Check /tmp/SNFS.install.log log file for errors. Do you wish to continue (Yes or <No>)? No</pre>	<p>Reboot the machine and repeat the upgrade process.</p> <p>NOTE: This may be caused by an inability to unmount the file system. Check the log for details.</p>
IRIX	34763	308348	The <code>cvdb -K</code> command does not work with IRIX.	The <code>-K</code> parameter is not supported at this time.

HP-UX	38359	n/a	FC cable pull does not initiate failover.	Use the SecurePath Fibre Channel driver.
	38892	n/a	Cache coherency between nodes with applications that use mmap can potentially see data corruption.	As of the StorNext 2.5 release date, HP-UX 11iv2 applications using mmap(2) may lose some data under some circumstances. If a file is mmaped for PROT_WRITE and another thread or process uses mmap(2) and then munmap(2) on the file or performs non-mmap I/O, for example, read(2)/write(2), on the same file, some of the mmap writer's data may be lost. With StorNext, this lost data may also occur if the file is modified on another node. For more information, contact HP - problem # (JAGaf50701).
	39987	n/a	Spinlock: locker forgot to unlock.	Disable the StorNext Quotas feature in the file system configuration file, see cvfs_config(4).
Linux	33536	298100	StorNext Linux clients acting as Network File System (NFS) servers do not provide proper cluster filelocks to NFS clients.	n/a
Solaris	17263	n/a	Recursive chmod command fails at a certain depth on Solaris 9 machines.	Solaris chmod, chown, and chgrp depend on "." and "." being the first two entries returned by the readdir command. Compile the GNU not UNIX (GNU) versions of the fileutils package; GNU chmod, chown, and chgrp do not depend on directory entry ordering.
	35223	n/a	LUNs have a maximum limit of 850 GB.	To maximize support across all platforms, set LUNs to be no greater than 850 GB when you are using Solaris 9 without the Big LUN patch.

Solaris	51308	n/a	Solaris NFS clients get I/O errors when using the <code>cp -p</code> command.	Do not use the <code>cp -p</code> command which attempts to copy ACLS not supported by StorNext.
	52120	n/a	Newly labeled disks are not usable in SNFS on Solaris 9 or 10 until system reboot.	Reboot the system.
Windows	14254	n/a	In a Windows/UNIX environment, using path names longer than 220 characters may cause procedures to fail under all versions of Microsoft Windows.	Verify that path names are not longer than 220 characters.
	33837	n/a	Readonly Directories GUI function to be eliminated in a future StorNext release.	n/a
	34699	n/a	File system does not automatically mount after a reboot due to <code>GetHostByName</code> failure.	Do one of the following: <ul style="list-style-type: none"> • Add the server name to the <code>etc\hosts</code> file on the Windows client. - or - • Type the IP address of the machine in the <code>fsnameservers</code> tab instead of the machine name.
	37431	n/a	Need to document usage of Windows 2000 client kit for Windows 2003.	On Windows 2000, 2003, or XP, use the Windows 2000 client kit.
	38979 43192 50943	n/a	On a Windows client, <code>unpack</code> fails with <code>convert to lower-case</code> enabled.	Select the following settings in the Mount Options control panel: <ul style="list-style-type: none"> • File Name Case Insensitive • Preserve Case <p>In a future StorNext release, the following options may be obsoleted:</p> <ul style="list-style-type: none"> • Case Sensitive • Convert to Lower Case • Convert to Upper Case
	40054	n/a	Group Quotas usage is not being reflected on Windows systems.	The StorNext Quotas feature only supports User quotas on Windows systems. Group quotas are not supported. Both user and group quotas are supported on StorNext UNIX/Linux clients.

All	34784	n/a	When performing a <code>cvcp</code> command from a Windows client against a file with international characters, a "cannot find file" error is returned.	n/a
	34903	n/a	Breaking an Ethernet connection causes a disconnect that invalidates locks (file locking).	n/a
	37589	n/a	The <code>cvadmin</code> command does not list all file system servers. This problem can be identified by the following <code>syslog</code> message: ERR NSS: Establish Coordinator failed GetHostByName ...	Manually add entries for each host listed in <code>fsnameservers</code> file to the <code>/etc/hosts</code> file.
	38785	n/a	There is a limitation (OS-specific) on the number of directories deep.	If a customer believes it has reached this limit (the maximum number of directories), then contact ATAC for assistance.
	39805	n/a	StorNext Quotas are not supported in a mixed (UNIX and Windows) environment.	If using StorNext Quotas, ensure that the StorNext environment is homogeneous - machines are either all UNIX machines or all Windows machines.
	39821	n/a	<code>CvApi_AllocSpace</code> may allocate less space than requested.	Use either the <code>fstat(2)</code> system call or <code>CvApi_CvFstat()</code> to determine the amount of space actually allocated.
	39917 43513	n/a	<code>CvApi_PunchHole</code> API function succeeds without performing the requested deallocation when a file is concurrently open on another StorNext client node.	Close the file on the other node and reinvoke <code>CvApi_PunchHole</code> .

All	42761	n/a	<p>Unix clients can create symbolic links in StorNext file systems that are usable on Windows clients if they meet the following criteria:</p> <ul style="list-style-type: none"> • The link is relative (it does not start with a '/' character) • The link is valid (it points to file or directory that is visible on the Windows client) <p>In particular, symbolic links to files can be used to read or write the target file, and symbolic links to directories can be used as components of pathnames.</p> <p>Other operations, such as deleting, renaming or examining properties, may fail or have unexpected results on Windows clients. It is strongly recommended that these operations be limited to Unix clients.</p>	Delete or rename symbolic links only on UNIX clients.
	48542	n/a	<p>In previous StorNext releases, an error encountered when processing the <code>cvpaths</code> file would result in the entire file being ignored; the system would then revert to its usual disk scanning rules.</p> <p>In StorNext 2.7, errors encountered in the <code>cvpaths</code> file are logged, but processing of the file continues. As many entries as can be successfully processed are used.</p>	To force the default disk scanning rules, remove the <code>cvpaths</code> file from the StorNext <code>config</code> directory.
	52929	n/a	The <code>cvlabel -ls</code> command reports incorrect disk serial numbers on Xsan systems.	n/a

All	53141	n/a	An error message may be generated if you have a file system configured with one or more LUNs that are "short-labeled", that is, LUNs that have a raw capacity greater than 4294967295 sectors, but have been labeled with a reduced capacity of 4294967295 sectors (the maximum number of sectors supported in a LUN by StorNext 2.6) and one or more files are trespassing on the restricted area of the LUN.	<p>1 On the metadata controller node, run the <code>cvfsck -t <fsname></code> command to list the names of files or directories allocated in the restricted area ('trespassing' files).</p> <p>2 On a client machine -- which should be same type of system (same OS) as the client machine that created the data originally -- move each of the listed files to non-restricted areas of the file system:</p> <ul style="list-style-type: none"> • For each trespassing file, use the <code>snfsdefrag</code> command to create a new copy of the file. For example: for a file called <code>/stornext/snfs1/xxx</code>, type: <pre>snfsdefrag -m 0 /stornext/snfs1/xxx</pre> • If any trespassing directories are listed, please contact ADIC technical support for instructions to correct the problem.
	54392	497050	When a fullstop of the file system services is initiated, the <code>/etc/init.d/cvfs fullstop</code> script executes a <code>fuser -km \$FSNAME</code> , that kills all processes with open inodes in the file system.	n/a

Limitations

This table lists the limitations that have been discovered in this release of SNFX.

Operating System/ Component Affected	Description
IRIX	<p>SNFS uses 64-bit inode numbers. 32-bit applications (programs compiled without 64-bit) support may experience problems working with files with inode numbers greater than 2147483648. 32-bit applications may also encounter problems when dealing with files greater than 2TB in size. This issue is not unique to StorNext, but may be encountered by some legacy applications when working with SNFS on IRIX.</p> <p>StorNext only supports SGI 64-bit versions of IRIX on operating systems that use the SGI version of the QLogic QLA2200 or QLA2310 FC-HBAs.</p> <p>Use the <code>uname -aR</code> command to determine the running version of IRIX.</p> <ul style="list-style-type: none"> If you are using IRIX level 6.5.26, you must install the maintenance stream of StorNext or SNFS. <p>On many versions of IRIX, the root crontab contains the following entry which is used to remove old application crash dumps and temporary mail files:</p> <pre>find / -local -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{} ' ;'</pre> <p>If StorNext file systems are mounted, they will be traversed by this <code>find</code> command which can have a dramatic impact on the performance of other applications currently using these file systems. To prevent the traversal of StorNext file systems, modify the <code>find</code> command so it reads:</p> <pre>find / -local '(' -type d -fstype cvfs -prune ')' -o -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{} ' ;'</pre>
Linux	<p>On many versions of Linux, the cron system runs a nightly script called <code>slocate.cron</code> that is used to build a database used by the <code>slocate</code> command. If StorNext file systems are mounted, they will be traversed by this cron job which can have a dramatic impact on the performance of other applications currently using these file systems. To prevent cron from traversing StorNext file systems, two files need to be updated. Perform these steps:</p> <ol style="list-style-type: none"> Modify the <code>updatedb</code> command in the <code>/etc/cron.daily/slocate.cron</code> file to read: <pre>/usr/bin/updatedb -f "cvfs,nfs,smbfs,ncpfs,proc,devpts" -e "/tmp,/var/tmp,/usr/tmp,afs,/net"</pre> <p>NOTE: "cvfs" has been added to the exclude list.</p> Add <code>cvfs</code> to the <code>PRUNEFS</code> definition in the <code>/etc/updatedb.conf</code> file. For example: <pre>PRUNEFS="cvfs devpts NFS nfs afs proc smbfs autofs auto iso9660"</pre>

Operating System/ Component Affected	Description
Solaris	<p>On Solaris, by default, the <code>/usr/lib/fs/nfs/nfsfind</code> script is run nightly by the cron daemon. This script contains a <code>find</code> command that traverses any local file system that is exported (shared) via NFS. If the running of this script is interfering with the performance of processes accessing StorNext file systems, it can be modified to skip them. To do so, add the following line to the <code>find</code> command in the script:</p> <pre>'(' -type d -fstype cvfs -prune ')' -o</pre> <p>The final <code>find</code> command should look like this:</p> <pre>find \$dir '(' -type d -fstype cvfs -prune ')' -o -type f -name .nfs* -mtime +7 -mount -exec rm -f {} \;</pre>
Windows	<p>Windows-based SNFS clients do not support symbolic links that point to a file system outside of SNFS. Symbolic links are created either by a UNIX/Linux client or by a Windows client (and are referred to as junctions).</p> <p>If you are using the StorNext client software with Windows 2000, Windows Server 2003, or Windows XP, turn off the Recycle Bin in the StorNext file systems mapped on the Windows machine, so the file systems will work properly.</p> <ol style="list-style-type: none"> 1 On the Windows client machine, right-click the Recycle Bin icon on your desktop and click Properties. 2 Click Global. 3 Click Configure drives independently. 4 Click Local Disk onto which you have mapped StorNext. 5 Select the Do not move files to the Recycle Bin. Remove files immediately when deleted check box. 6 Click Apply and OK.
Windows	<p>Virus-checking software, such as McAfee Virus Scan, can severely degrade SNFS performance. If you have anti-virus software running on your Windows 2000/Windows Server 2003/Windows XP machines, ADIC recommends you configure the software so that it does NOT check SNFS.</p>
All	<p>Hot re-zoning of SAN fabrics is not supported.</p> <p>StorNext is incompatible with CXFS.</p> <p>StorNext File System is incompatible with third-party portmappers.</p>

Operating System/ Component Affected	Description
All	<p>With SNFX 1.3, a change has been made to the way that the Reserved Extents performance feature affects free space reporting. In the previous release, SNFS would reserve a certain amount of disk space which would cause applications to receive an <code>out of space</code> error before the disk capacity reached 100%.</p> <p>In the current release, this reserved space is treated as allocated space. This allows applications to perform allocations until the file system is nearly full.</p> <p>NOTE: Due to allocation rounding, applications may still receive a premature <code>out of space</code> error, but only when there are only a few megabytes of space remaining. In the worst case, the error will be returned when the reported remaining space is:</p> <pre>(InodeExpandMax * #-of-data-stripe-groups)</pre> <p>One side effect of this change is that after creating a new file system, <code>df</code> will show that space has been used, even though no user data has been allocated. This amount of reserved space is determined by the following formula:</p> <pre>reserved space megabytes = (MaxConnections - 1) * MaxMBPerClientReserve * #-of-data-stripe-groups</pre> <p>For example:</p> <p>1 TB file system, 4 data stripe groups, MaxConnections=25, MaxMBPerClientReserve=100</p> <p>After <code>cvmkfs</code> has been used to create the file system, <code>df</code> reports:</p> <pre>File system 1K-blocks Used Available Use% Mounted on snfs1 1073741824 9830400 1063911424 1% /snfs1</pre> <p>While not recommended, the Reserved Extents feature can be disabled by applying the following setting to the Globals section of the FSM configuration file:</p> <pre>MaxMBPerClientReserve 0</pre> <p>This will cause the file system to not reserve space, but will reduce I/O performance in some cases.</p>

Documentation

These documents are currently available for StorNext products:

Document Number	Document Title
6-00360-06 Rev A	<i>StorNext Storage Manager and File System Installation Guide</i>
6-00905-04 Rev A	<i>StorNext File System Installation Guide</i>
6-00362-04 Rev A	<i>StorNext System Administrator Guide</i>
6-00361-09 Rev A	<i>StorNext Storage Manager Quick Reference Booklet</i>
6-00361-10 Rev A	<i>StorNext File System Quick Reference Booklet</i>

