



Tiger Serve 2.7.2 Release Notes

<i>What's New</i>	2
<i>Fixed Known Issues in Version 2.7.2</i>	2
<i>Upgrading to Version 2.7.2</i>	3
<i>Best Practices</i>	5
<i>New Known Issues</i>	12
<i>Unresolved Known Issues</i>	13

This document provides release information for Tiger Serve version 2.7.2. It discusses new features in this release as well as best practices, fixed and known issues.

What's New

New Kernels Supported on Red Hat Enterprise Linux/ CentOS Linux 6.x

Version 2.7.2 adds support for the following kernels on Red Hat Enterprise Linux/CentOS Linux 6.x Tiger Client computers:

- 2.6.32-431.3.1.el6
- 2.6.32-573.7.1.el6
- 2.6.32-642.1.1.el6

Fixed Known Issues in Version 2.7.2

Improved Performance with Small Files on Mac OS X Clients

Version 2.7.2 provides a fix, which visibly improves the performance with small files of Mac OS X client computers.

Improved Performance on Mac OS X Fast LAN Clients

Version 2.7.2 provides a fix that visibly improves the performance of Mac OS X Tiger Client computers accessing the shared storage over 10GbE connection.

Emptying the Trash of a Mac OS X Tiger Client

With version 2.7.2 Active Directory domain users on Mac OS X Tiger Client computer can normally delete a file moved to the Trash of the client computer.

Disconnecting from a Shared Volume While Browsing it in Terminal on Mac OS X/Linux

With version 2.7.2 Mac OS X and Linux Tiger Clients can normally disconnect from the storage server whether or not they are browsing its shared volumes in the Terminal.

Improved Stability with Files with Long Names

Version 2.7.2 provides a fix, which improves the system stability, when processing files with longer names (above 240 characters).

Restarting the Add-on Services

Version 2.7.2 resolves a problem that could prevent the starting of some add-on services after you had restarted them in the web interface of the appliance.

Faster Firmware Updates

Version 2.7.2 resolves a problem that could slow down the firmware update on the appliance.

Improved Playback with Autodesk Flame on Linux Clients

Version 2.7.2 provides improved support for Autodesk Flame playback from Linux Tiger Clients.

Upgrading to Version 2.7.2

Important: *All computers on your Tiger Serve network must run the same version of the software.*

To upgrade your Tiger Serve network to version 2.7.2 you must first install the update on the appliance and after that download and install the updated Tiger Client software on each Tiger Client computer. The procedure for upgrading your Tiger Serve appliance differs depending on whether both server nodes are online or just one.

The firmware update doesn't keep any settings except the IP addresses and name of the storage server. On an appliance with two server nodes, both of which are online, to preserve the settings configuration you can back it up before uploading the firmware update and then restore it after upgrading to this new version. On an appliance with just one server node, or with two server nodes, but one of them currently offline, you cannot back up the settings configuration and will have to manually re-configure all settings after upgrading to this new version.

Once you upgrade the appliance, the web interface provides the updated list with Tiger Client software and you can upgrade all client computers. On Mac OS X and Linux, you should first uninstall the previous version of the Tiger Client software and after that perform a clean installation. On Windows, you can upgrade the Tiger Client software without uninstalling the previous version, following the steps below.

Note: *Version 2.7.2 is compatible with Project Store 3.5.1 and later. If you want to benefit from the features available in Project Store 3.7.1, you need to upgrade it on your storage server, following the steps described in the Project Store 3.7.1 release notes: <http://www.tiger-technology.com/projectstore-3-7-1-release-notes/>*

To upgrade a Tiger Serve appliance with 2 server nodes to version 2.7.2:

1. In the left pane of the web interface, click System and then Maintenance.
2. In Back Up System Settings, click Back Up.
3. Provide the password for the web interface and click Continue.
A standard dialog opens that allows you to browse for a location where to save the settings backup file.
4. Select the location where to save the file and click Save.
5. In the Maintenance page, click Enter Maintenance mode and then confirm that you want to enter Maintenance mode.
Important: *All connected Tiger Clients will be automatically disconnected from the shared volumes and any file operation going on at the moment will be canceled.*
6. Go to HA Nodes and in the tile of each server node, click Manage Node.
The node view of each server node opens in a new tab/window of your web browser.
7. In the node view of the first server node, go to System | Maintenance.
8. In Firmware Update, click Browse/Choose file and then browse for and select the file containing the 2.7.2 update to upload.
9. Click Continue to confirm that you want to perform firmware update.
10. In the node view of the second server node, go to System | Maintenance.
11. In Firmware Update, click Browse/Choose file and then browse for and select the file containing the 2.7.2 update to upload.
12. Click Continue to confirm that you want to perform firmware update.
When the firmware update finishes, both server nodes automatically restart.
13. In the left pane of the web interface, click System and then Maintenance.
14. In Restore System Settings, click Restore.
15. Provide the password for the web interface and click Continue.
16. In the dialog that opens, browse to and select the file containing the configuration backup, then click Open.
Tip: *By default, the file with the backed up configuration is saved with the following name: [storage server name]_[time]_[date]_backup.json*
17. Tiger Serve imports the selected backup file and then automatically restarts.
Once the appliance restarts, all clients can connect to the web interface and download the updated version of the Tiger Client software for their operating system.

To upgrade a Tiger Serve appliance with one server node to version 2.7.2:

1. In the left pane of the web interface, click System and then Maintenance.
2. In the Maintenance page, click Enter Maintenance mode and then confirm that you want to enter Maintenance mode.

Important: *All connected Tiger Clients will be automatically disconnected from the shared volumes and any file operation going on at the moment will be canceled.*

3. In Firmware Update, click Browse/Choose file and then browse for and select the file containing the 2.7.2 update to upload.
4. Click Continue to confirm that you want to perform firmware update.

When the firmware update finishes, the appliance automatically restarts.

Once the appliance restarts, all clients can connect to the web interface and download the updated version of the Tiger Client software for their operating system.

To upgrade Tiger Client to version 2.7.2:

Important: *On Mac OS X and Linux, you should first uninstall the previous version of the Tiger Client software, following the steps described in the Tiger Serve Administration Guide. On Windows, you can upgrade the Tiger Client software without uninstalling the previous version.*

1. Access the web interface of the storage server.
2. Find the Tiger Client installation for your operating system and click Download.
3. Once the installation downloads to your computer, run it, following the installation steps described in the Tiger Serve Administration Guide.

Best Practices

IP Addresses Limitations

When assigning IP addresses to the network ports of Tiger Serve, make sure that you avoid the following IP addresses as they are either link-local addresses or are in use by Tiger Serve services:

- Link-local addresses in the range 169.254.0.0 - 169.254.255.255
- IPMI addresses in the range 1.1.1.0 - 1.1.1.255

Avoid to Back Up the Settings Configuration in Maintenance Mode

Tiger Serve cannot ensure the consistency of the settings configuration backup, when you save it while in Maintenance mode. To avoid possible problems, it is advisable to always back up and restore the settings configuration when the appliance is not running in Maintenance mode.

Granting Full Control to the System Account

In Windows, the built-in System account, used by the operating system and by services, by default has Full Control to all files on an NTFS volume. To provide for the normal operation of Tiger Serve, it is advisable not to change the Full Control permissions of this account.

For details about displaying the System account when setting permissions and granting Full Control to it, refer to your Windows documentation.

Detaching Shared Volumes from Tiger Serve for Use on Another Computer

To protect volumes from data corruption that could occur when multiple computers have access to them simultaneously, Tiger Serve marks the volumes it manages. To be able to mount a volume without Tiger Serve, this volume needs to be unmarked. Tiger Serve automatically unmarks all volumes with Offline status. If you intend to physically disconnect a volume for use on another computer, you must first change the status of this volume to Offline. If you fail to set the volume to Offline, you need to contact Tiger Technology support for assistance on unmarking the volume.

Assigning IP Addresses to Network Ports

The appliance must have an IP address for each network port through which it is connected to Tiger Clients (directly or via an Ethernet switch) - the Public port and all available 10GbE ports. The IP address you assign to each network port must be on the same LAN as the Tiger Client that will communicate with the appliance through this port.

IP Addresses Limitations

When assigning IP addresses to the network ports of Tiger Serve, make sure that you avoid the following IP addresses as they are either link-local addresses or are in use by Tiger Serve services:

- Link-local addresses in the range 169.254.0.0 - 169.254.255.255
- IPMI addresses in the range 1.1.1.0 - 1.1.1.255

IP Addresses of 10GbE Ports

When a Fast LAN Tiger Client can reach the appliance both through the Public port and the 10GbE port, it is sometimes possible a request from this Tiger Client to be redirected through the 1GbE port instead of the fast 10GbE port. The problem can be observed with Tiger Serve appliances with two server nodes and is reproducible when the node you have added to your storage servers list tries to redirect you to the other server node, which is currently playing the role of a primary node. To avoid the problem, it is highly advisable to set the IP addresses of the 10GbE ports of both the appliance and the Tiger Clients to be on a separate LAN or on a subnet of the LAN used by the Public port.

For example, if the Public port uses IP address 172.16.1.1 (with subnet mask 255.255.255.0), set the IP address of the first 10GbE port to 172.16.2.1 (with subnet mask 255.255.255.0) and change the IP address of the Tiger Client 10GbE port to be on the same subnet. For the next 10GbE port of the appliance, set IP address 172.16.3.1 (with subnet mask 255.255.255.0) and respectively change the IP address of the 10GbE port of a Tiger Client to be on the same subnet.

Another important thing is to add Tiger Serve to the storage servers list on each Fast LAN Tiger Client using the IP address set for a 10GbE port and not the Public port.

Synchronization of IP Settings Between Server Nodes

To guarantee the failover between the two server nodes in the appliance, the network ports with identical names on both nodes (the Public ports, for example) must use consecutive IP addresses and identical settings for subnet, router and DNS server. In case you cannot meet these requirements, the failover between server nodes cannot take place. When you configure the IP settings of one node, as long as the other node is also online, Tiger Serve attempts to automatically synchronize the IP settings and assigns the same settings for subnet, router and DNS server and a consecutive IP address of the port with the same name on the other server node. If the consecutive IP address on the network is already taken, Tiger Serve will assign a random IP address to the identical port of the other server node and failover between the two nodes will not be possible.

When one of the nodes is offline while you change the IP configuration settings of the other node, you must synchronize the IP settings of the two nodes manually, in order not to obstruct failover between server nodes. You can do this in the Node view of the respective server node.

Migrating Replicated Data

Migrating Replicated Data from One Replication Target to Another

While the Replication & Tiering mechanism of Tiger Serve doesn't restrict you about changing the replication target, to avoid problems, you should set a new replication target with caution and after careful planning. If you need to set a new replication target to which files to be copied, you should keep in mind that all already replicated files will be inaccessible from the new replication target. Additionally, offline files (replicated files that have a copy only on the replication target) cannot be retrieved unless they have a counterpart on the new replication target.

When changing the replication target of your shared storage, it is recommended to migrate all already replicated data from the previous replication target to the new one. The safest workflow to do this is to let the Replication & Tiering mechanism replicate it again on the new replication target. For the purpose, you have to make sure that:

- all already replicated files are available on the shared storage (note that lack of free space on your shared storage may hamper the process).

- the status of all already replicated files is not “replicated” but “normal” as the Replication and Tiering mechanism will skip files on the shared storage that are with “replicated” status and will not copy them to the new replication target.

You can achieve this following these steps:

1. To ensure that each replicated file has an offline counterpart on the shared volume(s), in the drop-down box of the Volume Browser, select a volume with enabled replication and then click Rescan Volume.
2. Repeat the above step for each volume on which replication is enabled.
3. In the left pane of the Volume Browser, select the root of the volume and in the right pane select a folder and click Restore.

Note: *You have to manually select all offline files in the root of the volume and click Restore.*

4. Repeat the above step for each folder on each volume with enabled replication in order to restore all replicated files to their original locations.
5. In the Replication & Tiering page of the web UI, specify a new replication target.
6. In the Replication Options dialog, select “Do not replicate” and leave the Exceptions list empty.
7. In the left pane of the Volume Browser, select the root of the volume and in the right pane select a folder and click Make Offline, to revert all files from “Replicated” status to “Normal” status.

Note: *You have to manually select all restored file in the root of the volume and click Make Offline to revert their status to “Normal”.*

Important: *Clicking Make Offline transfers the status of replicated files to “normal” only if the Replication & Tiering mechanism detects that the selected files don’t have copies on the replication target. If you issue the command for a file with “replicated” status, which actually has a copy on the replication target, the command will replace the file on the shared storage with a stub file pointing to its copy on the replication target.*

8. Repeat the above step for each folder on each volume with enabled replication.
9. Set replication policy parameters for the new replication target to let the Replication & Tiering mechanism to copy all files.

Migrating Replicated Data to a New Volume

When you need to reformat a volume from which data has been replicated or to replace it with a volume with bigger capacity, instead of manually moving data between volumes in order to restore data on the new volume, you can use the Replication & Tiering mechanism to automate the process. Following the steps below, you will be able to restore all already replicated data from the replication target to the new volume. Data will be restored in the form of stub files, which you can make online in the Volume Browser or on demand (by accessing them) in your file browser.

To ensure that all needed data is copied to the replication target, you may have to manually replicate files or whole folders in the Volume Browser of your appliance:

1. In the web interface of your appliance, go to System | Replication & Tiering and click Browse.
2. In the Volume drop-down box, select the old volume.

Note: *If you are migrating Project Store data, in the Volume Browser select the "Projects" check box to display all project folders.*

3. In the right pane, select a file or folder with Normal status and click Replicate.

Note: *In the Volume Browser a folder is with Normal status if it contains even just one file, which is not yet replicated.*

When all files/folders you want to migrate to the new volume are with replicated status in the Volume Browser, you are ready to proceed with the migration steps below.

Important: *If you're using a Private volume as replication target, before you reformat or detach the volume from which data has been replicated, you must get its GUID as well as the GUID of the volume, on which you will restore data. For the purpose you can access the web interface of your appliance through port 9125, typing the following in the address bar of a web browser:*

[IP address of the storage server]:9125/volumes

The storage server lists information about each shared volume. Search the list for the label of the respective volume and copy its GUID from the "guid" field of the list.

To restore replicated data on a new volume:

1. If you are migrating Project Store data, enable support for Project Store on the new volume.
2. In the Replication & Tiering page, enable replication on the new volume, by setting exactly the same replication target as for the old volume (the same WOS policy, private volume, folder on the SMB/CIFS share, BlackPearl bucket).
3. (for replication on a Private volume only) Disable replication (select "No Replication") and share the private volume, then in your file browser go to the .tt_rt folder in the root of the volume and rename the folder with the GUID of the old volume with the GUID of the new volume, then enable replication again using the same settings as in the step above.
4. In the Replication & Tiering page, click Browse to open the Volume Browser.
5. In the Volume drop-down box, select the new volume and then click Rescan Volume.

The Replication & Tiering module creates a stub file for each replicated file on the replication target, which it cannot find on the volume i.e. all files that have been replicated from the old volume.

6. (for migrated Project Store data only) Convert all restored files from offline to replicated: in the Volume Browser, select the new volume in the drop-down box and select the "Projects" check box, then one by one select each project and click Restore.

7. (Project Store only) To let Project Store detect the restored projects, reset the Project Store database:
 - (Project Store SE) Click Settings and go to Project Management, then click Reset DB.
 - (Project Store) In the administrative interface, go to Project Management and click Reset DB.
8. (Project Store only) In the administrative interface, go to Project Management, select the check box of each project in the list and click Rescan, to regenerate project information such as project size, users, application type, number of files, etc. as well as to queue proxy media generation.

Restarting A Server Node While in Maintenance Mode

When you enter Maintenance mode, Tiger Clients are automatically disconnected from both server nodes of your Tiger Serve appliance, until you exit Maintenance mode in the web UI or restart the appliance. It is advisable not to restart just one server node (in node view) as it will exit Maintenance mode and will attempt to supervise the volumes managed by Tiger Serve, while the other node is in Maintenance mode and performs a maintenance operation.

Setting Up the SAN Policy on Windows Tiger Clients

By default, on computers running Windows Vista and above all shared disks visible to the computer are delivered as offline and read-only. You can bring these disks to online mode in Disk Management, still, there's a possibility on the next boot of the system the disks to be offline again. To change this behaviour, you should make sure that the SAN policy of the operating system is set to "Online All". To view and set the SAN policy of your computer you can use diskpart.exe command-line utility that is stored in the system drive of your computer.

To view and change the SAN policy of your computer:

1. Run command prompt as administrator.
2. In command prompt, start diskpart utility by typing:
diskpart
3. To view the current SAN policy, type this followed by return:
SAN
The utility displays one of the following:
 - Offline Shared - all newly discovered disks that do not reside on a shared bus are brought online and made read-write.
 - Offline - all newly discovered disks remain offline and read-only.
 - Online All - all newly discovered disks are brought online and made read-write.
4. Do one of the following:

- If the SAN policy of your computer is “Online All”, exit diskpart utility and command prompt.
 - If the SAN policy of your computer is not “Online All”, proceed with the steps below.
5. Run command prompt as administrator.
 6. In command prompt, start diskpart utility by typing:
`diskpart`
 7. Execute the following:
`san policy=OnlineAll`

Using Identical Time and Date Settings

It is advisable to use identical settings for time and date on all computers on your Tiger Serve network (the storage server and all client computers). Using different time and date settings may hamper communication with the storage server and the Traffic and Events reports may display incorrect information.

Disabling Mac OS X Time Machine for Shared Volumes

It is highly advisable to exclude all volumes shared by Tiger Serve from the list of locations Time Machine is set to back up on your Mac OS X. Keeping this feature turned on for shared volumes may hamper performance.

To exclude a shared volume from Time Machine's backup:

1. In Time Machine preferences, click the Options button.
2. In the sheet with locations, select a shared Tiger Serve volume and click the "-" button.
3. Repeat the above step for each shared Tiger Serve volume.
4. Click Save.

Calculating Folder Size (Mac OS X)

Use the Calculate Sizes option, available in List view of the Finder, to calculate sizes per folder and per window. Activate this option only for folders on your local disks. Do not enable it for folders on the shared storage volumes, as this will temporarily slow down performance. If you enable this option on a shared volume, network workstations that see the volume will be unable to work on it.

For details about managing the “Calculate Sizes” option, refer to your Mac OS X® documentation.

File Systems Behaviour and Files Compatibility

The three operating systems (Windows, Mac OS X and Linux) on which you can install the Tiger Client software differ in the way they manage files. On Windows, files can have many streams, in

Tiger Serve 2.7.2 Release Notes

contrast to Mac OS X (only 2 streams - data fork and resource fork), and Linux (only data stream). In addition, files created by Mac OS X also have extended attributes (such as the Finder Info, for example) that are not available on the other two platforms.

These peculiarities of the file systems lead to partial loss of information when transferring a file created on one platform through a machine running on another platform.

For instance, when on Mac OS X you copy a file with 50 streams created by Windows, all file streams' information is lost, except the information of the data stream. The result will be the same when a Linux computer copies a 50-stream Windows file. A Mac OS X file copied by Linux will also lose its resource fork's info.

Streams' information is not lost when copying Mac OS X or Linux files by a Windows computer, but in this case a Mac OS X file will lose the extended attributes information.

There is another problem with files sharing common file format. For example, a movie file created with Final Cut Pro on Mac OS X is playable by QuickTime. If a Windows workstation copies this file, it will not recognize it as QuickTime-playable until adding the *.mov extension to it.

See also "Final Cut Pro Project Names on the Shared Volumes" on page 12.

Avoiding Long File Names

As the volumes Tiger Serve uses are NTFS formatted, each file operation on them is subject to the limitations of the file system. By default, Windows usually limits the length of file names to 260 characters. Still, as the full path to the file is also included in the character count, it is possible to receive errors when attempting to create or copy a file with long file name. That is why it is advisable to avoid long file names in order to ensure that file operations on the shared storage are normally processed.

Final Cut Pro Project Names on the Shared Volumes

It is advisable that the Final Cut Pro projects you create on the shared volumes have names that do not exceed 31 symbols.

New Known Issues

Duplicated Files on Spectra Logic BlackPearl

When the Replication & Tiering module is set up to replicate on a Spectra Logic BlackPearl appliance, it is possible duplication of files to occur on both the replication tier and the shared volume(s). As the BlackPearl appliance cannot tell that an already replicated file is moved to another directory on the shared volume, it perceives it as a new file to replicate and keeps both its copy in its previous location and the new one. Similarly, should you attempt to restore a replicated

file that has been moved, the mechanism will restore two copies of the file - one in its original location and another in the location to which it has been moved.

File Moved to Another Volume in the Pool After Modifying It

If you use an application that creates a temporary file when modifying an existing file on the volume pool, after you save the file it might be moved to another volume in the pool (the one on which the distribution policy in the pool creates new files).

Unresolved Known Issues

Support for Hard Links Limitations

The support for hard links on volumes shared by a Tiger Serve storage server is provided with the following exceptions:

- due to incompatibility between NTFS and HFS+ hard link attributes, the Finder may not show hard links on shared volumes properly. For example, if a file has two hard links in the same directory, when you browse this directory in the Finder only one of the hard links will be shown;
- you cannot use hard links on pooled volumes (smart storage pooling);

Support for Fast User Switching Known Issues

The Fast User Switching option allows switching between users on a single machine without quitting applications and logging out. It lets users leave their applications running and their files open when another user needs to log in.

While operations on the shared volumes are not interrupted when you switch users on a Tiger Client computer, the Tiger Client tray application/menulet may not be present when a new user logs on, thus preventing that user from controlling the available Tiger Client commands. A workaround to this problem is to force display the Tiger Client tray icon/menulet by restarting it.

Note: *As users on Linux computers use the command-line interface to manage the Tiger Client driver, you can use fast user switching without problems.*

To restart the Tiger Client tray application on a Windows computer:

1. Run the Task Manager as administrator.
2. In the Processes tab, right-click smct.exe and select End Process.
3. In the Run command (press Windows logo key+R to display it), execute this
C:\Program Files\Tiger Technology\Tbox\smct.exe

Note: Where C: is the drive, on which the Tiger Client driver is installed.

To restart the Tiger Client menulet application on Mac OS X:

1. Start the Activity Monitor (go to Applications | Utilities | Activity Monitor).
2. In the Disk tab, in the Process name column select smct.
3. Click End Process in the Activity Monitor task bar.
4. Provide an administrator's user name and password.
5. On the drive, on which the operating system is installed, go to /Library/Application Support/embedded/smct and double-click it.

Teaming Network Ports Known Issues

After enabling/disabling the setting for teaming of network ports, the web interface may display error, whereas the setting is actually applied. Additionally, on a Tiger Serve appliance with 2 server nodes, the default IP address assigned to the team is identical on both server nodes. Note that manually assigning an IP address to the team in cluster view, synchronizes the setting on both server nodes if they are online and there's no need to assign it in node view.

Restoring Deleted Project Store Project from the Replication Tier

When you delete a replicated project in the Project Store web interface, Project Store automatically moves it to a temporary Trash folder in the projects depot. Should the Replication & Tiering happen to replicate the Trash folder before Project Store deletes it, the deleted project is automatically moved to a Trash folder on your replication tier.

Note: *Currently the replication on Spectra Logic BlackPearl doesn't support paths synchronization when renaming the containing folder and the deleted project will be present in both the projects depot and the Trash folder on the replication tier. Thus, you will be able to restore the deleted folder directly from the projects depot.*

To restore the project in the projects depot on your main storage from the Trash folder on a Private volume, SMB/CIFS network share or a DDN Web Object Scaler as a replication tier, you should:

- (SMB/CIFS network share) Manually move the project from the Trash folder to the Projects folder and then rescan the Volume in order to re-generate the deleted files as stub files, which you can then restore.
- (Private volume) Disable Replication & Tiering on the Private volume and share it to Tiger Clients, then unhide the .tt_rt folder in the root of the volume and manually move the deleted project from the Trash folder to the Projects folder. Make the volume Private again and enable Replication & Tiering on it, then rescan the shared volume in the Volume Browser in order to re-generate the deleted files as stub files, which you can then restore.

- (DDN Web Object Scaler) Rescan the volume and regenerate the deleted data as stub files and then contact Tiger Technology support for assistance on moving the data from the hidden Trash folder to the Projects folder.

Regardless of the replication target you use, the final step to allow Project Store to detect and let you work with the project in Project Store is to restart the add-on services in the Maintenance page of Tiger Serve.

Only Automatic Defragmentation of Volumes Available

Currently, you cannot perform manual defragmentation of your volumes. To keep your volumes defragmented it is advisable to enable the automatic defragmentation.

Renaming a Volume

You may have to rename a volume twice in the web interface of Tiger Serve in order to permanently change its name.

Mac OS X Tiger Clients Mounting NTFS Stripe Partitions over the LAN

Currently, Mac OS X Tiger Clients can mount just the first partition of a striped NTFS volume over the Fibre Channel and all other partitions are mounted over the Ethernet.

Label of SAN Disks Missing in Windows Explorer of Windows Tiger Clients

On some Windows Tiger Clients, the label of shared volumes mounted over the Fibre Channel may be missing in Windows Explorer.

To attempt to resolve the problem restart the Tiger Client computer or disconnect and then reconnect to Tiger Serve:

1. Click the Tiger Client tray application and then click Connect/Disconnect.
2. In the list of all storage servers, click the Tiger Serve appliance to disconnect from it.
3. Again click Connect/Disconnect and click the Tiger Serve appliance in the list to connect to it.

Specifying Name and Preferred Mount Location of a Volume Pool

You cannot rename a smart storage pool or set preferred mount location on Windows Tiger Clients.

Volume Disappears from Web UI during Check and Repair Procedure

While you perform the Check & Repair procedure on a Tiger Serve volume, it is not listed in the web UI. Once the procedure finishes, you can manage the volume in the web interface again.

Exporting a SMB Share in Domain Environment

To ensure that Tiger Serve can export a SMB share to computers without the Tiger Client software, when your appliance is deployed in an Active Directory domain, the domain account “Everyone” must have at least Read Only access to the root of the volume from which the share is exported.

Mounting a Storage Pool on Tiger Client in Domain Environment

In case you have deployed your Tiger Serve in Active Directory domain and smart storage pooling is enabled, Windows Tiger Clients may fail to mount the storage pool unless the group ‘Everyone’ has Read access to the virtual volume.

Incorrect Connection Status on Tiger Clients Running Mac OS X

The Tiger Client software on Mac OS X may sometimes fail to correctly display the connection status to the Tiger Serve appliance. For example, the Settings dialog of the Tiger Client menulet may display that the Tiger Serve is online, when it is actually offline. Another example is when the Tiger Client loses its Fibre Channel connection to the appliance and mounts the volume(s) over the LAN, while the volume icons in the Finder and on the desktop are not of volumes mounted over the Ethernet. Additionally, in this case Tiger Clients running Mac OS X may fail to disconnect from the volume(s).

Allowing smct.app and vpd During Tiger Client Installation on Mac OS X

During installation of Tiger Client on Mac OS X, even if you specify that smct.app and the vpd daemon must be added as firewall exceptions, you should manually allow them in the firewall database of your computer. Note that as long as the firewall of your computer is turned on, you will have to add them to the firewall database after each reboot of the computer.

Connecting to a Newly Added Secondary Node from Tiger Clients

If you have added a Tiger Serve appliance on a Tiger Client’s list of storage servers, while only one of its server nodes has been online, when you turn on the other server node, the Tiger Client

computer will not be able to automatically connect to it (in case of failover, for example). To detect both server nodes, you will have to remove the appliance from the storage servers' list and then add it again.

Re-adding a Tiger Serve Appliance to the Storage Servers List

In case you have added a Tiger Serve appliance to your storage servers list when only one of its nodes has been online, when the other server node is also online and an IP address is assigned to it, you may have to remove the appliance from the list and add it anew, in order to ensure that in case of failover you can connect to it.

Connect to Tiger Serve from Mac OS X Without DNS Server

If there isn't a DNS server on the network through which a Mac OS X Tiger Client communicates with Tiger Serve, the Tiger Client computer will fail to connect to the storage server and mount its volumes.

A workaround to the problem is to add an entry about each node of the storage server you want to connect to in the hosts file on the Mac OS X Tiger Client:

1. In Terminal, type:
`sudo nano /etc/hosts`
2. At the end of the hosts file, add a new line for each storage server you want to connect to, which contains the following:
`[IP address of the storage server] [name of the storage server]`
For example, if your Tiger Serve has two server nodes that use respectively 172.16.1.1 and 172.16.1.2 as IP addresses and the name of the appliance in the web UI is Tiger Serve, in the hosts file enter the following lines:
`172.16.1.1 Tiger Serve1`
`172.16.1.2 Tiger Serve2`
3. Press CTRL+X, then Y and then Enter to exit the nano editor and save the changes.
4. Restart the Tiger Client.

Applying New Network Environment Settings

It is possible when changing the settings for network environment (domain or workgroup) of your appliance after clicking the Apply button in the Settings page an error to be displayed, when in fact the changes have been applied. It is advisable to refresh the web UI of Tiger Serve in order to verify whether the new settings have been applied.

Renaming the Appliance

Should you rename your Tiger Serve appliance after it has been added to an Active Directory domain, Tiger Clients may fail to browse the volumes it shares even if they re-add it to their storage servers list. It is advisable before renaming the appliance, to set the network environment to workgroup and manually delete the records about the appliance from the domain controller. It is advisable to do this each time you switch the network environment of your Tiger Serve from domain to workgroup.

Moving Replicated or Offline Files/Folders

If you move a replicated or offline file/folder from its original location to a folder that is not subject to replication on the shared storage, this folder will also appear on the replication target, but will contain just the files/folders that are subject to replication. For example, if you move an offline file “Report” from its original folder “April” to folder “Drafts” that is not subject to replication and contains other files as well, the replication and tiering mechanism of Tiger Serve will create on the replication target a folder “Drafts”, but containing just the file “Report”.

Importing Offline Dynamic Disks in Maintenance Mode

Offline dynamic disks cannot be imported when your Tiger Serve appliance operates in Maintenance mode. To import offline disks, you have to exit Maintenance mode.

Incorrect LAN Traffic Statistics

If you monitor the LAN traffic statistics of Tiger Serve by accessing its web interface through the secondary node (the server node that is in standby mode), the graph will display no LAN traffic. To view the actual statistics, open the web interface using the IP address of the current primary node.

Errors When Entering/Exiting Maintenance Mode

When entering/exiting Maintenance mode, the web UI of Tiger Serve may display errors, when in fact the operation has succeeded. To verify that you have entered/exited Maintenance mode, refresh the browser. If on exiting Maintenance mode the error persists, restart the storage server.

Executing Operation on the Tiger Serve Volume from Terminal

When Tiger Serve is deployed in an Active Directory domain, you won't be able to execute certain operations on the Tiger Serve volume(s) from the Terminal of a Mac OS X client computer even if you have logged on as a user that has Full Control. These security mapping issues are not valid for operations performed in the file browser.

Mirrored/Spanned Volumes Mounted Over LAN on Mac OS X and Linux

Currently mirrored/spanned SAN volumes can be mounted on Mac OS X/Linux SAN Tiger Clients only over the LAN. Such volumes are mounted normally on Windows SAN Tiger Clients.

No Support for Spotlight Indexing

Currently, Spotlight Indexing is not supported on Mac OS X Tiger Clients. To index and search through shared volumes from Tiger Clients running Mac OS X, use Spotlight alternatives.

No Support for Indexing Service on Shared Volumes

Currently, Tiger Serve does not support the Windows indexing service on shared volumes mounted on Tiger Clients.

No Support for Sparse Files

Tiger Serve does not provide support for sparse files on SAN volumes.

No Authorization Required When Removing Tiger Serve from the Domain

If you decide to switch your Tiger Serve from domain to workgroup environment, the domain controller does not require that you authenticate yourself as a user authorized to remove computers from the domain.

