



Caché Integration with a Network Appliance Filer

Version 2007.1

04 June 2007

Caché Integration with a Network Appliance Filer
Caché Version 2007.1 04 June 2007
Copyright © 2007 InterSystems Corporation
All rights reserved.

This book was assembled and formatted in Adobe Page Description Format (PDF) using tools and information from the following sources: Sun Microsystems, RenderX, Inc., Adobe Systems, and the World Wide Web Consortium at www.w3c.org. The primary document development tools were special-purpose XML-processing applications built by InterSystems using Caché and Java.



Caché WEBLINK, Distributed Cache Protocol, M/SQL, N/NET, and M/PACT are registered trademarks of InterSystems Corporation.



InterSystems Jalapeño Technology, Enterprise Cache Protocol, ECP, and InterSystems Zen are trademarks of InterSystems Corporation.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Memorial Drive, Cambridge, MA 02142, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

InterSystems Worldwide Customer Support

Tel: +1 617 621-0700
Fax: +1 617 374-9391
Email: support@InterSystems.com

Table of Contents

Caché Integration with a Network Appliance Filer	1
1 Purpose and Scope	1
2 Assumptions	2
3 Integration on Windows	2
3.1 Infrastructure Requirements	2
3.2 Configuration Recommendations	3
4 Integration on UNIX-Based Platforms	5
4.1 Infrastructure Requirements	5
4.2 Configuration Recommendations	6
5 NVRAM Reporting	7
6 Support	8
List of Tables	
Recommended NetApp Configuration Options	3

Caché Integration with a Network Appliance Filer

This article is an overview of how to use Caché with network-attached storage (NAS) on the Network Appliance filer. It includes the following sections:

- [Purpose and Scope](#)
- [Assumptions](#)
- [Integration on Windows 2000](#)
- [Integration on UNIX-Based Platforms](#)
- [NVRAM Reporting](#)
- [Support](#)

1 Purpose and Scope

Caché, the post-relational database from [InterSystems](#), has the features professional developers need to quickly create Web and client/server applications. Caché benefits developers by giving them their choice of development tools, programming languages, and methods of data access. Caché benefits transaction processing applications by providing outstanding performance, massive scalability, real-time data analytics, and robust reliability. All these capabilities are tied together in an easy-to-use rapid development environment.

Network Appliance, Inc., (NetApp) manufactures high-performance multiprotocol file servers (or filers) and caching appliances for UNIX (NFS), Windows (CIFS), and Web (HTTP) environments. Since OpenVMS works on DECNET and not IP, there is no OpenVMS support.

The Network Appliance technical report, “[Multiprotocol Data Access: NFS, CIFS, and HTTP](#)”, describes the NetApp multiprotocol filer architecture and explores the implications of multiprotocol filing for system administrators and end users.

This document describes the steps necessary to implement Caché on Windows and UNIX-based operating systems with database data and log files located on a NetApp filer running Data ONTAP(tm) version 5.3.6x and higher.

2 Assumptions

This paper assumes you are familiar with [Caché](#) and the operation of NetApp filers. It also assumes you are familiar with the operation of your specific operating system. These configurations were tested using Caché version 5.0, running on Windows 2000 (SP4) and various UNIX-based platforms listed in the UNIX section. The filer was running Data ONTAP version 5.3.6R2.

Within the context of Caché, everything you do with a database on a filer is the same as what you would do to a database stored on a local disk. To create a new database on the filer, create it in the NetApp mount directory.

3 Integration on Windows

Integration of a NetApp filer with Caché on Windows has certain [infrastructure requirements](#) and [configuration recommendations](#) to provide optimal performance and reliability.

3.1 Infrastructure Requirements

The following are the infrastructure *requirements* to integrate Caché with a NetApp filer on Windows.

Caché Machine:

- Caché is running on Windows.
- Equipment satisfies the system requirements for running Caché.
- Hardware is on the Windows NT Hardware Compatibility List.
- Caché and the filer participate in a Windows domain.

NetApp Filer:

- Filer is running Data ONTAP version 5.3.6x or later.
- Filer is running the CIFS (Common Internet File System) protocol.

Network:

- Network connection exists between the Caché machine and the filer.
- A Gigabit Ethernet network is recommended to reduce latency and to improve performance.

Note: Make sure there are enough disks (HDs), to allow striping, partitioning, and load balancing algorithms to work properly.

3.2 Configuration Recommendations

Recommended configuration settings and integration techniques used to implement a successful Caché/NetApp Filer configuration are described in the following sections:

- [Filer Configuration](#)
- [Windows Registry Settings](#)
- [Filer CIFS Share](#)
- [Windows and Caché Configuration](#)

For better performance and a more reliable system environment there are several recommended options and configurations. Most of these options are set on both the filer and on the Windows platform. They are described in the following table.

Recommended NetApp Configuration Options

Option	Description	On Filer	On Windows
MaxMpxCt	The maximum of outstanding requests the NT client can have against the filer.	Yes	Yes
TCPWindow	The maximum transfer size for data across the network.	Yes	Yes
minra	Minimal filer read ahead.	Yes	No
no_atime_update	Setting this option can improve CPU usage if accurate times are not important.	Yes	No
security	The security style is the method the filer uses to determine whether a user has access to a file. CIFS qtree should use <code>ntfs</code> security style.	Yes	No

3.2.1 Filer Configuration

On the NetApp filer, set the following options.

CIFS options:

- `options cifs.max_mpx 50` — Must match `MaxMpxCt` (default is 50, other possible values are 126, 253, and 1124). Look at the performance monitor Redirector-Current Commands statistic. If this is constantly running at the current value, increase the value.
- `options cifs.tcp_window_size 64240` — Set this value to 64240 (0xFAF0).

These are recommended settings. See the [na_options](#) manual page for more information.

General options:

- `vol options minra on` — When this option is on, the filer performs minimal read-ahead on the volume. By default, this option is off, causing the filer to perform very aggressive read-ahead on the volume.
- `vol options no_atime_update on` — When this option is on, it prevents the update of the access time on an inode when a file is read. This option is useful for volumes with extremely high read traffic, since it prevents writes to the inode file for the volume from contending with reads from other files. Use this option when you know in advance that the correct access time for inodes is not needed for files on that volume.

These are recommended settings. See the [na_vol](#) manual page for more information.

Qtrees:

`qtree security qtreename ntfs` — The `qtree security` command changes the security style for files and directories. When using the `ntfs` security style for CIFS requests, Windows NT permissions determine user access.

This is the recommended setting. See the [na_qtree](#) manual page for more information.

3.2.2 Windows Registry Settings

On Windows, InterSystems recommends you make the following modifications:

- Increase TCP Windows sizes to 64KB.
- Increase `MaxMpxCt`; must match `cifs.max_mpx` value.
- Increase negotiated CIFS buffer sizing.

Set the following registry keys by running **REGEDT32.EXE**.

Go to the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services` key and make the described entries:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
LanmanServer\parameters\MaxMpxCt
Datatype: DWORD
Value: To match the setting above for cifs.max_mpx
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
Tcpip\Parameters\TcpWindowSize
Datatype: DWORD
Value: 64240 (0xFAF0)
----> Global setting for all interfaces <---
```

```

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
\Tcpip\Parameters\Interfaces\interface\TcpWindowSize
Datatype: DWORD
Value: 64240 (0xFAF0)
  Adapter Setting:
    ----> setting per interface <----

```

3.2.3 Filer CIFS Share

Create a share on the filer with security set to “Everyone Full Control”. For information about how to create a share on a filer, consult the Data ONTAP Software Systems Administration guide that accompanied your system.

3.2.4 Windows and Caché Configuration

The following configuration procedures ensure a proper setup of Caché with a NetApp filer:

- NetApp is properly joined to the Windows domain.
- Users authenticate against a Windows domain server to access shares on the filer.
- The NetApp directory is mounted (use the **Map Network Drive** command from the **Tools** menu of Windows Explorer) with the same username and password as your Windows login.

Note: A useful test to check if the connection is configured correctly is to *remotely* telnet to Caché and open a file on the drive mapped to NetApp. For example, input **Open**
"F:\B.TXT" : "NWS", where F: is your mapped drive and B.TXT is an arbitrary file. If you return to the Caché prompt, the connection is configured properly; if the session hangs indefinitely, it is not.

4 Integration on UNIX-Based Platforms

Integration of a NetApp filer with Caché on UNIX-based platforms has certain [infrastructure requirements](#) and [configuration recommendations](#) to provide optimal performance and reliability.

4.1 Infrastructure Requirements

The following are the infrastructure *requirements* to integrate Caché with a NetApp filer on UNIX-based platforms:

Caché Machine:

- Equipment satisfies the system requirements for running Caché.
- Caché is running on a supported UNIX-based platform. See the [Supported Server Platforms](#) section of the *Caché Supported Platforms* guide.

NetApp Filer:

- Filer is running Data ONTAP version 5.3.6x or later.
- Filer is running the NFS (Network File System) protocol.
- NFS share is exported and mounted writable by the client.

Network:

- Network connection exists between the Caché machine and the filer.
- A Gigabit Ethernet network is recommended to reduce latency and to improve performance.

TCP Protocol:

Where possible, specify TCP protocol, for example: `proto=tcp`.

Note: Make sure there are enough disks (HDs), to allow striping, partitioning, and load balancing algorithms to work properly.

4.2 Configuration Recommendations

Recommended configuration settings and integration techniques used to implement a successful Caché/NetApp Filer configuration on UNIX-based platforms are described in this section.

Mount options:

- *For all mounts:*

```
NFS Version 3, hard, intr, Readsize= 32768, Writesize= 32768, nolock
```

Detailed syntax depends on operating system.

- Example options for *Linux*:

```
rsize=32768, wsize=32768, intr, rw, nolock, nfsvers=3, suid, hard, rw
```

- Example options for *Solaris*:

```
rsize=32768, wsize=32768, intr, rw, vers=3, suid, hard
```

Security style:

Security style is the method the filer uses to determine whether a user has access to a file. Use the `unix` security style for the NFS `qtree`. When the `unix` security style is used, the user's UID and GID, and the UNIX-style permission bits of the file or directory determine user access. This is the recommended setting. See the [na_qtree](#) manual page for more information.

The following example sets the security style of the root volume:

```
qtree security / unix
```

Linux Export Entries

On NFS-mounted file systems on Linux, a file created by an SUID:SGID executable has different, non-UNIX standard, owners than on standard file systems. The SGID bit on the executable fails to take effect, while the SUID bit succeeds in setting the owner of the file to the owner of the executable. The solution is to give root access to all hosts in the `etc/exports` file. For example:

```
/vol/vol1/scratch00 -root=host1:host2:host3,sec=sys
```

See the [na_exports](#) manual page for more information.

The preceding options are recommended for high-speed connections.

Contact Network Appliance for a list of recommended kernel patches for your operating system.

5 NVRAM Reporting

Network Appliance recommends enabling the feature of Data ONTAP (NetApp's operating system software) which supports special error processing when connected to a database server. To do this, enter the following command from the filer's console or a telnet session:

```
vol options vol0 nvfail on
```

This causes the filer to issue appropriate error messages in the `\\FILER1\C$\etc\messages` file in the case of system failure that might affect the Caché database. The administrator learns of these errors either by examining the message logs or by the auto-support email notification feature of the filer.

In particular, this option enables some additional status checking when the filer goes through its initialization sequence at boot time to verify that the NVRAM (non-volatile random access memory) is in a valid state. This is the case for both a clean (normal) shutdown or a dirty (crash, power failure, etc.) outage. Only a failure of the NVRAM card itself should cause it to become invalid. If the contents of NVRAM are found to be invalid, an error message is sent to the system console and to the filer log file. At that point, the Caché system administrator ensures that the state of the database is correct and valid.

Competitive database products may recommend using the optional rename feature. However, to ensure the proper functioning of Caché and the handling of its database, journal, and log files, InterSystems recommends using the Caché built-in "Freeze on Error" feature instead. From the **[Home]** > **[Configuration]** > **[Journal Settings]** page of the System Management Portal, click **Yes** in the **Freeze on error** box and click **Save** to update the setting.

6 Support

InterSystems and Network Appliance support the specific configurations described in this document. If you experience any problems implementing the techniques described in this technical paper, please contact the [InterSystems Worldwide Response Center](#) (WRC) or a [Network Appliance Technical Marketing Engineer](#).