



# All About CVS & CM

CVS SUITE AND CM SUITE 2009 Build 3644 January 2010

WinCVS



Tortoise



CVSNT



Change  
Manager



Workspace  
Manager



Release  
Manager



Defect  
Tracking



Audit

Optimized for



March Hare Software Ltd



**Legal Notices**

**There are various product or company names used herein that are the trademarks, service marks, or trade names of their respective owners, and March Hare Software Limited makes no claim of ownership to, nor intends to imply an endorsement of, such products or companies by their usage.**

**This document and all information contained herein are the property of March Hare Software Limited, and may not be reproduced, disclosed, revealed, or used in any way without prior written consent of March Hare Software Limited.**

**This document and the information contained herein are subject to confidentiality agreement, violation of which will subject the violator to all remedies and penalties provided by the law.**

**LIMITED WARRANTY.**

**TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, March Hare Software Limited AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES AND CONDITIONS, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT, WITH REGARD TO THIS DOCUMENT, AND ANY ADVICE OR RECOMMENDATION CONTAINED IN THIS DOCUMENT.**

**NO OTHER WARRANTIES.**

**TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL March Hare Software Limited OR ITS SUPPLIERS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR ANY OTHER PECUNIARY LOSS) ARISING OUT OF THE USE OF OR INABILITY TO USE THE FOLLOWING DOCUMENTATION INCLUDING ANY RECOMMENDATION OR ADVICE THERIN, EVEN IF March Hare Software Limited HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN ANY CASE, March Hare Software Limited's ENTIRE LIABILITY UNDER ANY PROVISION OF THIS DOCUMENT INCLUDING ANY RECOMMENDATION OR ADVICE THERIN SHALL BE LIMITED TO THE GREATER OF THE AMOUNT ACTUALLY PAID BY YOU FOR THE DOCUMENT OR £5.00; PROVIDED.**

**© Copyright 2004 - 2010 March Hare Software Ltd**


**march-hare.com**

**sales@march-hare.com**



# Multi Site, Repository Replication and WAN Performance

CVSNT is designed to work well when users are distributed worldwide:



The purpose of this section is to explain in theory how to use CVS well in a distributed environment, and what should be avoided.

## Overview

Before starting a technical discussion about Multi Site version control with CVSNT there are some introductory message that we must clearly advise you of:

- If you have multiple sites relying on CVSNT then you should be supporting the project financially and you should ensure that you have technical support from the authors
- No proprietary software solutions are required
- Define the problems you are trying to solve with metrics

### *Support of CVSNT Server and your Technical Account Manager*

Ensuring the reliability, resilience and integrity of your repository are important when implementing a SCM versioning server. The people who wrote the server control software are the people who can ensure that. If you are unhappy with the performance or reliability of an existing CVS or CVSNT server then the first step is to ensure that you have support from the people who wrote it, and speak to your Technical Account Manager regarding the specific problems you are experiencing – March Hare Software write CVSNT and can assist with resolving problems through CVS Professional Support Services.

### *Software Solutions*

There are unscrupulous individuals and even some organisations that attempt to convince CVS and CVSNT users that they are experts on Multi Site Version Control and based on that expertise they recommend purchase of a proprietary software solution.

No author of CVS or CVSNT has ever written or said that additional proprietary software is required to resolve any multi-site, performance or reliability issue with the CVS or CVSNT server.

Some of the proprietary software that is advertised for enabling Multi Site CVS or CVSNT also claims to provide added security benefits, reports and controls. CVSNT contains sophisticated and reliable security controls that are better integrated and more reliable than any proprietary add-on package, including:

- Access Control Lists
- Authentication
- Deny/Allow clients

No proprietary software is required to resolve multi-site, performance security or reliability issues.

### *Define the Problem*

It is our experience that people ask about Multi Site CVS for the following reasons:

- They are not currently using CVSNT and they expect problems due to previous experience with other SCM tools like ClearCase, Visual SourceSafe or Continuous/CM Synergy.
- They are not currently using CVSNT and they have been approached by a vendor of a proprietary Multi-Site product or have read a review of such a product.
- They are currently using CVS not CVSNT, CVS Suite or CM Suite.
- They are using CVSNT and do not have support and have not received professional advice by the authors on how to configure and use the system, eg: setting up ‘stand by’ servers. The eBook *All About CVS* (a part of CVS Professional or CVS Suite) describes how to set up a standby server and also a suite of all the necessary software that has been tested to work together as well as e-mail based support).
- They are using CVSNT and have a specific performance problem (eg: a ‘build’ process takes too long or a ‘tag’ process takes too long).

If the last point applies to you then you should document the specific problem you are experiencing and contact your March Hare Software Technical Account Manager.

For all other problems described above you should purchase CVS Suite or CM Suite with Software Maintenance and Professional Support from March Hare Software, and optionally the on-site training and installation services. See the section *Learning more about CVSNT* above and specifically the headings *Professional Support* and *Contacting Us*.

As with any IT project – you should define the expected outcomes before you begin the project so that after the project is completed then you can establish the relative success of it. If you require the tag to complete with 10 minutes then write that down, if you require the build to complete in 1 hour then write that down – and then give that information to your Technical Account Manager or pre-sales technical support.

*If there is no proprietary software needed why do I need to pay for support?*

March Hare Software donates all the effort to write CVSNT and publish CVSNT according to the Free Software/Open Source Software GPL and LGPL licenses and does so with no charge to the millions of users. There are many thousands of users each week who begin using this software with only the support they receive from the other users via the newsgroup and wiki.

By encouraging those who require more assistance to pay for support we can ensure that the software continues to be developed and provide the best levels of service to those users who rely on CVSNT for their productivity at work.

### *Why to March Hare require Silver or higher support for Multi Site?*

March Hare Software base our charges on the effort required to resolve problems. It is our experience that if a company has users at multiple sites that their problems are more difficult to resolve. Additional (open source) software may also be provided for customers with multiple sites designed to resolve particular technical problems that only occur in multi-site installations.

Silver level support gives the customer access to telephone support as well as optionally purchasing additional on site services which are very likely to be needed in the first year after installation of a multi-site product.

### *Multi Site authentication*

When working with multiple sites it is important that the authentication scheme is carefully designed. If users are being authenticated by a remote authentication server (eg: a domain controller) then CVSNT client operations may be delayed while that authentication takes place for each command. These types of authentication issues are often not noticeable with operations such as logging in to a PC or connecting network drives, however since CVSNT client operations can be performed many times a minute (especially if using SCCI/Visual Studio Integration) the lag for authentication becomes critical. Also see the sections *Unix and Linux Server Synchronisation to a Reliable Time Source* and *Windows Server Synchronisation to a Reliable Time Source* for other reasons why client authentication may be slow in a multi site installation.

A second multi site issue is that users who are not registered on the central authentication system (such as a domain controller) may require access to your CVSNT server. See the section *CVS only users (no system user with that username)* for details of how to overcome this problem.

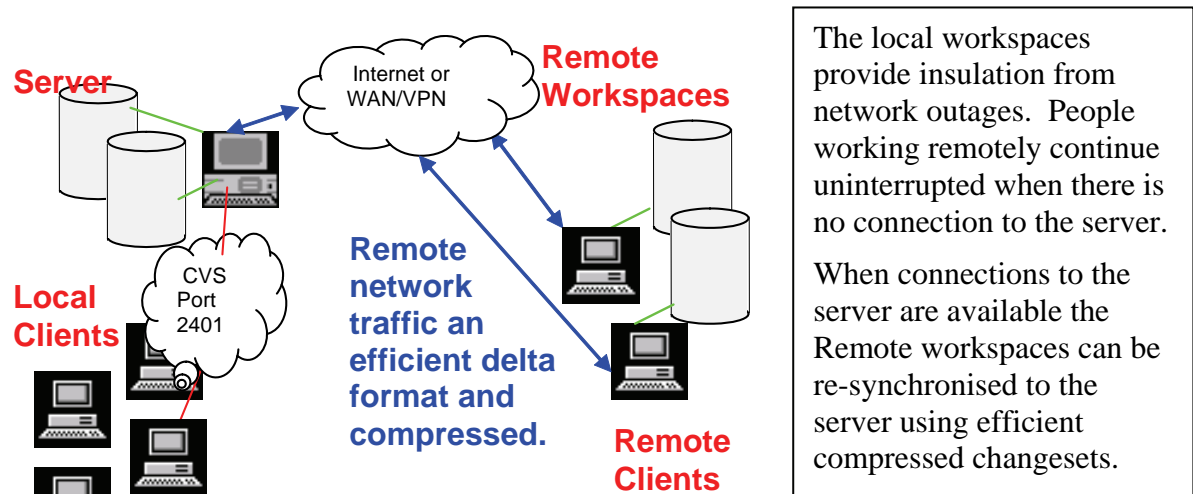
### *If performance should be good without multi-site software, why is my performance poor?*

The most common reason why we see poor performance is because:

- The existing tools are not being used in conjunction with CM and CVSNT best practice.
- There are existing procedures that are not required and are inefficient. They may have been introduced to resolve a problem that there were insufficient skills to resolve any other way or a problem that has been resolved in later releases of the software.
- The expectations are unrealistic (defined as performance expectations that could not be matched with a local workgroup repository). Eg: if people using the software are trying to highlight the advantages of an alternative version control system where the repository is not stored or managed centrally and stored only on their own computer.
- The performance has not been measured and the performance problem overstated. This is occasionally seen with 'outsource' suppliers who are looking for a scapegoat to blame for missed deadlines.
- The performance problems are a direct result of insufficient training of the people who are using the software (the people experiencing the performance problem). This is not common, however can be occasionally seen with 'outsource' suppliers who are using insufficiently trained staff and do not want to bear the cost of training.

## WAN Optimised Workspaces and Procedures

Most people who need to use CVSNT at multiple sites or geographically distributed locations will use this solution. This solution provides high performance for remote sites as insulation from network connection failures.



When implementing WAN Optimised Workspaces it is important that the software and the procedures are correctly configured to reduce the network transmission overheads.

For example: it is typical to ensure that binary objects (eg: compiled objects such as .jar, .obj, .exe files) are not regularly checked out and updated using CVSNT, but instead that a 'build server' is located at each site that can create these for the users at the remote site. Implementing a more complex strategy for binary compiled objects is counter-productive since they still need to be regularly synchronised between the servers and will result in more administration of conflicts.

Another example is to ensure that developers do not delete workspaces. If a developer was working on Trunk but now needs to work on branch-rel-1 then the trunk workspace should be kept and the developer switches to working on the branch-rel-1 workspace.

At the remote site it is typical to maintain several sandboxes (at least one per branch). These can be set to automatically synchronise to the server whenever the server is available, and most importantly set to synchronise before the work day for the remote site commences. This sandbox can be copied to another person's PC at the same remote location if they need a 'clean' copy of the latest sandbox.

Often remote developers are fearful of making changes when there is no connection to the server because they fear that their workspace may be out of date. Configuring the CVSNT server with automatic e-mail notification will advise users at all sites (via a group e-mail address) of what is changing on the server. A more complex repository replication strategy would not provide a more reliable solution since if the link is down the cache will be out of date.



## Multiple Repositories

Multiple servers (repository replication) may be considered for the following reasons:

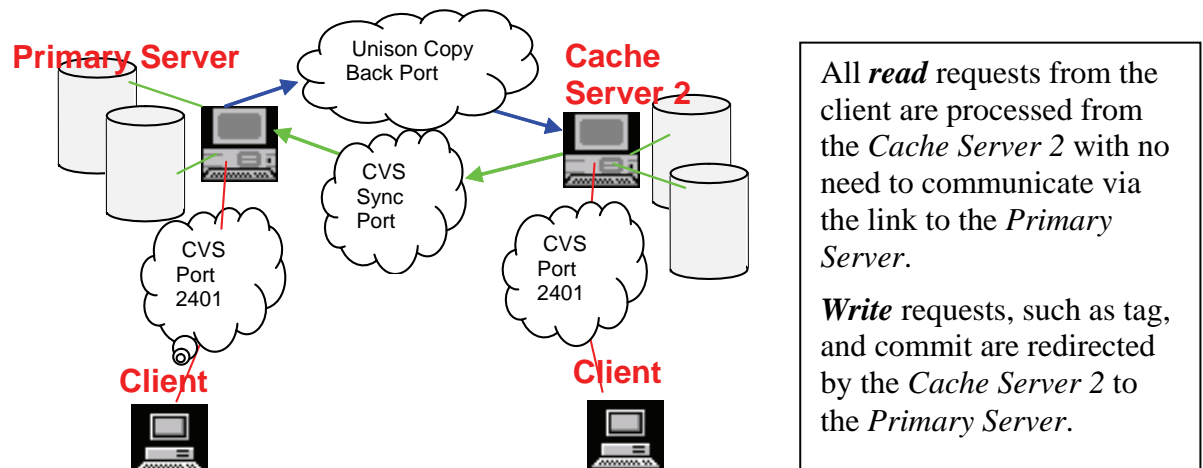
- Large Dual Site Development (Performance / Repository Replication Caching)  
*Typically there are two sites in different timezones so that work can proceed 24 hours a day on a project. It is common for each site to house more than 500 users.*
- Repository Mirrors for Disaster Recovery or Ownership  
*The majority of work is performed at one location (eg: India) however the work is owned by an organisation elsewhere (eg: North America)*
- Multi Site parallel Development - Offshoring  
*Where distinct work is carried out at different sites, eg: In India version 2 is developed but England maintains version 1 and tests version 2.*

## Large Dual Site Development - Repository Replication Cache

CVSNT client server communication was designed for the kind of networks that were available over twenty years ago. The networks typically in use today exceed the performance and bandwidth of those many times over. Therefore it is not necessary to use Repository Replication to obtain good performance from your CVSNT based CM system.

Implementing replication to improve performance generally does not work: if there is a performance problem it is most likely in the underlying infrastructure or CM methodology. If you have less than 50 users at each location then the best solution for you is described in the section *WAN Optimised Workspaces and Procedures*.

In a few rare cases it may be beneficial to have a local (or regional) cache if there are hundreds of users at each site or if policy prevents the introduction of more efficient workflows and procedures.



This technology is available to CVS Suite and CVS Suite x64 customers with Silver (or higher) level support who purchase additional on site consulting for the installation and certification of the replication software. Configuring repository replication requires:

- CVS Suite and Silver (or higher) support for *each* person using the software at *each* site
- CM Design and CVS Administration Training Course
- Installation Consulting for *each* site
- Professional Services (usually 1-3 days per site)

Prices and descriptions for each of these are available on the web store, including a sample package for two sites with up to a total of 100 users:


<http://store.march-hare.com/s.nl/sc.2/category.12/.f>

- 100 x CVS Suite (32 bit edition) with Silver Support [may vary please contact sales]
- 1 x CVS Design and Administration Training Course
- 2 x CVS Installation Consulting for *each* site
- 3 x days on site CVS Professional Services [may vary please contact sales]

### Read Only Mirrors - Repository Replication Disaster Recovery or Ownership

Your business may require copies of the repository to be stored at multiple physical locations in case of a critical failure in infrastructure at the primary location, or to ensure that the ownership of the property is centralised.

Included in CVS Suite is a repository replicator named Unison designed for this purpose.



Replication may be triggered at each commit to the primary server, or at a time based interval. See the section on *Unison Repository Replication* included in the eBook *All About CVS*<sup>4</sup> for more information.

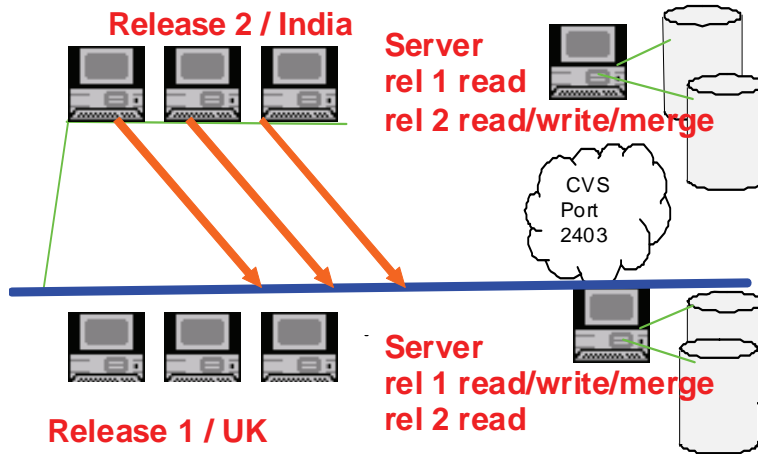
If the primary server becomes unavailable, a typical disaster recovery procedure would be to use the DR DNS server to assign the *Read Only Server 2* the same name as the original *Primary Server* therefore allowing users to recommence working with the least disruption.

---

<sup>4</sup> The eBook *All About CVS* is available with CVS Suite, and CVS Pro level 2 or higher.

## Multi Site Parallel Development

This is required when some of the development of a project has been outsourced to a remote site. In this example each stream of development is clearly performed at a different physical location – though users at all locations need to see the work in progress or initiate merges.



This functionality is only available in CM Suite Multi Site – please discuss your requirements with your Technical Account Manager.

## Other Multi Site Techniques

### *Active-Active Repository Replication (or Quorum based Multi Site)*

This solution is not recommended – do not use the services of any consultant advertising these services. March Hare Software regularly receives support requests from companies that have implemented this software and experience problems that the vendor cannot resolve. Our recommendation is to uninstall this software and purchase professional support from the people who wrote CVS and CVSNT.

A quorum based multi site solution uses a large number of servers (usually a minimum of 5 or 8) which all regularly synchronise with each other. Based on the number of servers that can be reached via the network at any time (the quorum) the ‘authoritative’ server(s) are established and changes sent to servers that are not a part of the quorum are excluded.

This solution is typically offered based on a perception that it improves performance on unreliable networks (eg: India and China), however it does not achieve the desired result. As soon as the nodes in other continents cannot be reached they are excluded from the quorum and any changes there become non-authoritative.

Technically this proprietary software is implemented as a proxy that attempts to intercept the protocol between the client and server – since the authors of the proxy are not members of the active open source CVS and CVSNT communities and the CVSNT clients and servers are written by March Hare Software and open source community members around the world it appears that the proxy intervention regularly causes simple operations to fail.

Due to the large number of servers required there are also very high procurement costs and administrative overhead to this solution.

### *CVS Server Clustering*

This solution is not recommended – typically only organisations hosting projects for many geographically distributed companies require clustering, eg: SourceForge or gnu.org. Some enterprise customers with more than 500 users at a single location may consider this option once initial workflow and procedural issues have been resolved on a phased deployment of smaller stand alone servers.

CVSNT supports clustering ‘out of the box’ – the lock service started on only one node of the cluster co-ordinates the activities of cvs server processes on all nodes of the cluster. The repository is stored on a suitable cluster filesystem<sup>5</sup> on a SAN.

This cluster software is typically combined with a network load balancing switch such as the Barracuda Load Balancer<sup>6</sup> or Cisco Catalyst 4500/4000 Series Switch<sup>7</sup> or with Cluster load balancing software (built into the operating system, eg: Windows 2003 Advanced Edition or Red Hat Linux AS 4).

A cluster provides both redundancy and high performance.

---

<sup>5</sup> Eg: SAN Filesystem from DataPlow, or  
IBM Tivoli SANergy: <http://www-306.ibm.com/software/tivoli/products/sanergy/>

<sup>6</sup> [http://www.barracudanetworks.com/ns/products/balancer\\_overview.php](http://www.barracudanetworks.com/ns/products/balancer_overview.php)

<sup>7</sup>

[http://www.cisco.com/en/US/tech/tk389/tk213/technologies\\_tech\\_note09186a0080094714.shtml#cat4k](http://www.cisco.com/en/US/tech/tk389/tk213/technologies_tech_note09186a0080094714.shtml#cat4k)