WANdisco Subversion High Availability Administration Guide

© 2009 WANdisco

Revision History

REVISION	DATE
1.0	October 2008
2.0	October 17, 2008

This material is confidential to WANdisco and may not be disclosed in whole or in part to any third party nor used in any manner whatsoever other than for the purposes expressly consented to by WANdisco in writing.

This material is also copyright protected and may not be reproduced, stored in a retrieval system or transmitted in any form or by any means in whole or in part without the express written consent of WANdisco.

WANdisco Subversion High Availability - Copyright © 2009 WANdisco

Contents

© 2009 WANdisco

Contents, cont'd

© 2009 WANdisco

Contents, cont'd

© 2009 WANdisco

Contents, cont'd

© 2009 WANdisco

1 Introduction

Welcome to the WANdisco world of replication. WANdisco Subversion High Availability provides high availability and disaster recovery for Subversion. A WANdisco Subversion High Availability deployment consists of a Failover Agent and multiple Subversion servers, each with an associated WANdisco proxy, functioning as exact replicas of each other on a LAN. It offers Subversion users business continuity. If one Subversion server fails, the Failover Agent transparently directs users to the next available replica.

WANdisco Subversion High Availability guarantees RPO (Recovery Point Objective) equal to 0, i.e., zero data loss, even if the failure occurs in the middle of a transaction. Your High Availability group of Subversion nodes are synchronized at all times: each Subversion database is a functional equivalent of the other Subversion databases. High Availability employs WANdisco replication technology to ensure that all replicas remain synchronized, so that if one of the Subversion servers goes down for some reason, there is zero data loss.

WANdisco provides the Admin Console, a web-based user interface, to administer and monitor the HA group.

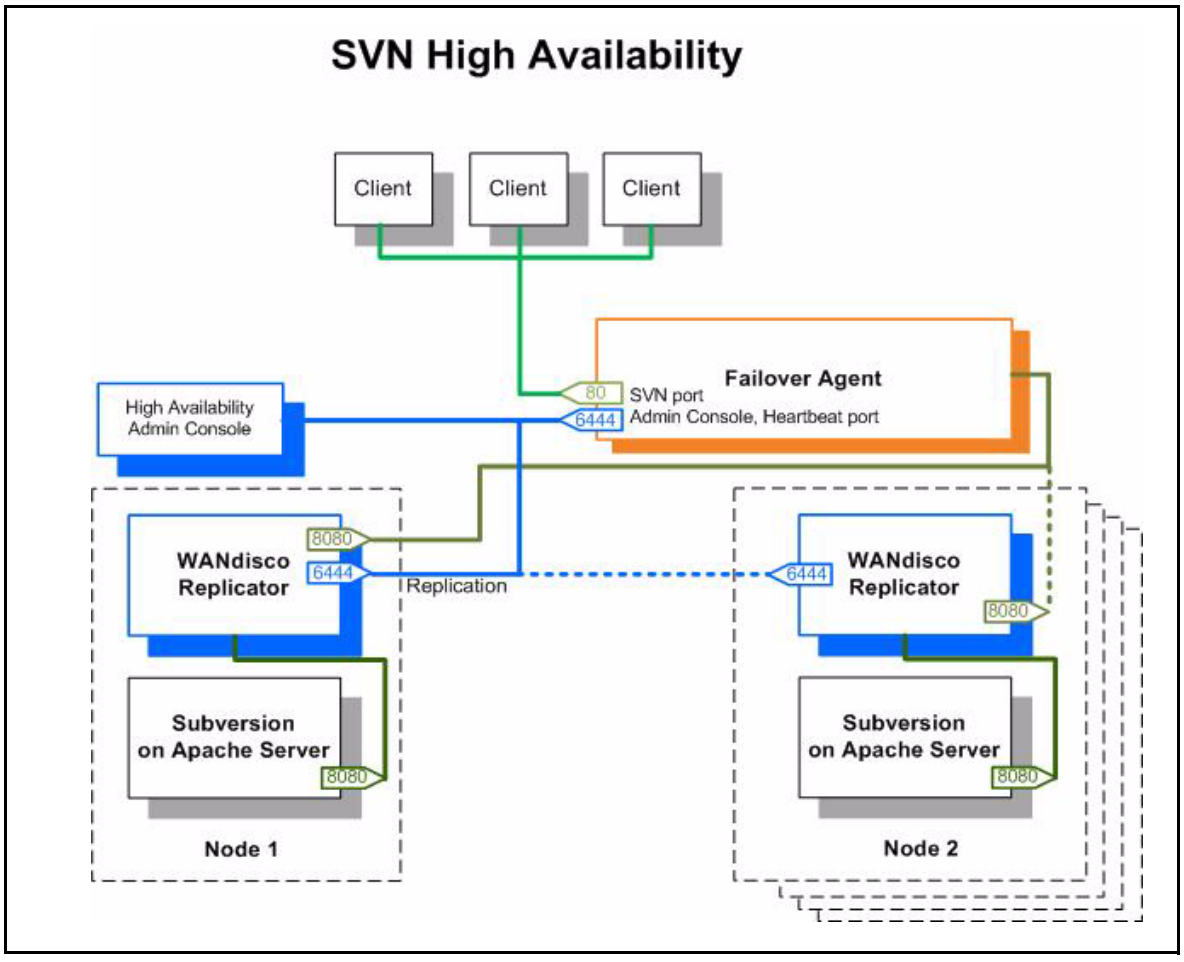

© 2009 WANdisco

A stand-alone Subversion instance is hosted in an Apache web server. Users access the application with a standard web browser and HTTP protocol (via the HTTP Failover Agent port, by default 80). WANdisco High Availability (HA) configures an additional port (by default 6444) for replication communications (known as DConE) using the DConENet protocol. The Admin Console also communicates through this port, using HTTP protocol.

The Subversion user connects to the WANdisco Failover Agent on the standard Subversion port (configurable), such as 80. For the replicators, the Subversion client port is by default 8080, which helps ensure that no Subversion user can bypass the replicator and connect directly to Subversion. The actual Subversion server port can also be 8080, or you can change it to another port number.

1.1 How Replication Works Within an HA Group

The nodes in the HA group continuously communicate any Subversion transactions that users are making. The HA group has to agree to an order for the transactions.

The group establishes transaction ordering through the agreement of the quorum. By default, High Availability has a majority response quorum, which means a majority of the nodes must agree on the transaction order. So if the group has five nodes, three of them must agree on the position of a transaction in the transaction order.

Having a majority quorum also ensures that if one node goes down in an HA group, the other nodes can continue uninterrupted. The HA group catches up the errored node when that node rejoins the group.

An HA group needs at least three nodes to avoid having a single point of failure. In a two node group, one node has the potential to be a single point of failure, which is why WANdisco does not recommend having just two nodes in a group. If you are contemplating a two node group, please read section 2.4, Deployment Checklist.

© 2009 WANdisco

1.2 Failover and the Heartbeat

The Failover Agent continuously sends out a heartbeat, or an “are you alive?” message, to each replica, on the WANdisco port using the DConENet protocol. By default the heartbeat messages are sent at intervals of one second. The replicas respond to this with an “I am alive” message. The Failover Agent expects a response from each replica. If the Failover Agent does not get a response within a configurable number of heartbeats, the Failover Agent marks that node as unresponsive. HA administrators can change the heartbeat interval and missing heartbeat count in the Admin Console.

If the Failover Agent has marked node 1 (the current node) as unresponsive, actual failover to the next node occurs lazily, that is, only when a Subversion client request comes through. This reduces the number of false failover alarms, as a replicator may not respond within the configurable number of heartbeats during a restart, or a node may be restarted before the next client request, eliminating the need for failover.

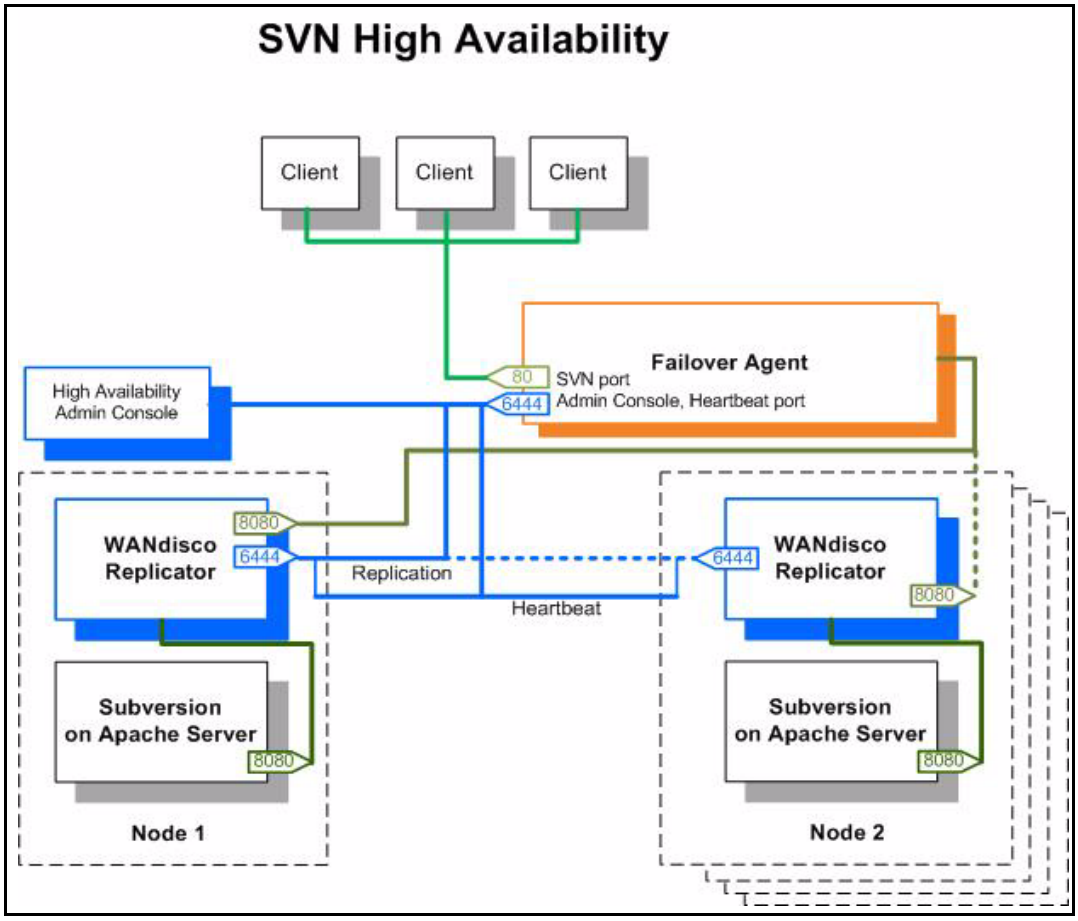

© 2009 WANdisco

1.3 Failover Sequence

The Failover Agent passes Subversion transaction data to only one node. That node then replicates the transactions to the entire group, so that every node in the High Availability group is a functional replica of each other.

1.3.1 The Current Node and the Designated Node

The node that communicates with the Failover Agent is called the “current” node. In normal operation, the current node is also the “designated” node, the node the Failover Agent expects to talk to. If the designated node should fail, and failover occurs, the current node would become node two, however node one is still the designated node. When node one comes back online, the Failover Agent resumes talking to node one (which again becomes the current node as well as the designated node).

1.3.2 Priority Order

During HA installation, you determine the failover order for the group. For example, node 1 is first in the failover order (priority one), node 2 is the second in the failover order (priority two), and so on.

The first HA node, as the current node, receives all Subversion transactions through the Failover Agent. If that node fails, the Failover Agent immediately begins communicating directly with node 2. If node 2 fails, then the Failover Agent immediately begins communicating with node 3, and so on. The order is referred to as the priority order, meaning the order the Failover Agent takes during a failover.

If node 1 comes back up, the Failover Agent immediately returns to communicating with that node.

You can see the failover priority order by looking at the value for the `PriorityOrder` element in the prefs.xml file (shown in **bold** in the following example). In the example, you see that this node is the current node, the first node in the priority order.

```
        ...
        <ServerProxy>
           <ListenerIP>192.168.1.184</ListenerIP>
           <ListenerPort>2401</ListenerPort>
           <PriorityOrder>1</PriorityOrder>
        </ServerProxy>
...
```


© 2009 WANdisco

1.4 The Distinguished Node

Each HA group has a distinguished node (user-assigned at installation). If there are an even number of nodes, the distinguished node acts as a tie breaker. For example, in a four node group, say two nodes want transaction X to be committed first, while the other two nodes want transaction Z to be committed first. The distinguished node's vote carries a heavier weight, and its transactions always take precedence.

1.5 Replication Example

Here is an overview of what occurs when a write transaction is received by the current active primary replicator in the HA group.

Step 1 The originating client sends the transaction to Subversion, passing through the Failover Agent.

Step 2 The Failover Agent relays the information to node 1 (the current active primary).

Step 3 Node 1 replicates the transaction throughout the HA group.

Step 4 Transaction data is successfully received by the quorum (a majority of the nodes). The quorum agrees and assigns it a Global Sequence Number (GSN).

Step 5 High Availability passes the transaction data to Subversion.

Step 6 Subversion processes the transaction.

Step 7 High Availability waits for Subversion to complete the transaction. High Availability only marks the transaction complete when Subversion returns a completion status.

If for some reason node 1 goes down during this process, the Failover Agent communicates with the next node in the failover order. Through replication, all nodes have a record of all transactions, and so are immediately available to act as the current active primary, should the need arise.

1.6 Replication and Site or Network Failures

High Availability defaults to majority quorum, which ensures continuos operation of the replication group, as long as a majority of the nodes are up. That is, if there are five nodes in an HA group, and three go down, then replication is suspended until at least one more node rejoins the group (three being a majority of the nodes).

© 2009 WANdisco

1.6.1 Site Failures

Say you have a five node High Availability group, spread across three continents. One of the sites goes down. If that site was the current active primary, then the Failover Agent begins communicating (fails over) to the next node in the priority order. Replication continues uninterrupted.

If that site was not the current active primary, then replication continues uninterrupted and the Failover Agent does not need to take any action. Replication continues unless a majority of the sites go down.

As soon as a site comes back up, the replication group catches up the site on its missing transactions, so that all sites are again synchronized.

1.6.2 Network Failures

If a network link goes down and outside connectivity is completely lost, then replication halts and SVN users hang if and only if they are doing a write operation (tag, commit, etc.) that requires a quorum. Read operations (like up, co, log, etc.) continue to work with the stale data.

When connectivity is restored or the errored server comes back online, the local replicator syncs up with the replication group automatically. First, the local replicator consults its local recovery journal (similar to a database redo log), and then, if necessary, attempts recovery from any of the quorum sites.

The recovery infrastructure and details of WANdisco fault-tolerance can be found at http://www.wandisco.com/pdf/dcone-whitepaper.pdf.

1.6.3 Failover Agent Failures

A watchdog runs on the Failover Agent, so if the Failover Agent crashes, the watchdog immediately restarts it. If the machine should crash, service is unavailable until you reboot and restart the Failover Agent.

If you do not want to wait for the Failover Agent machine to be restored, you could run the failover agent on a hardware cluster. The Veritas Cluster Server is an example of a commercial solution. See http://www.symantec.com/business/products/overview.jsp?pcid=pcat_business_cont&pvid=20_1.

Linux-HA is an example of an open-source solution. See http://www.linux-ha.org/.

© 2009 WANdisco

1.7 Establishing a Replication Baseline

When you deploy High Availability, you must ensure that all the replicas start out in sync, meaning that all of them are identical. Once High Availability is deployed, WANdisco's replication technology ensures they remain in sync.

You start with one SVN database, referred to as a replication baseline. To create this baseline, follow the procedure in 5.2, Establishing a Baseline for Replication.

1.8 Failover Group of Two Replicators

WANdisco strongly discourages using High Availability with just two nodes. Having only two replicators in the HA group forces changes to the failover mechanism. Whereas the objective of a larger failover group is continuous operation, the main objective of a two-replicator based deployment becomes dealing with a single replicator node failure.

In a two node deployment, High Availability defaults to singleton quorum, and the distinguished node is the second node.

© 2009 WANdisco

For a two-node deployment, in addition to the standard procedure chapter, see Chapter 6, Procedures for Two-Node HA Groups.

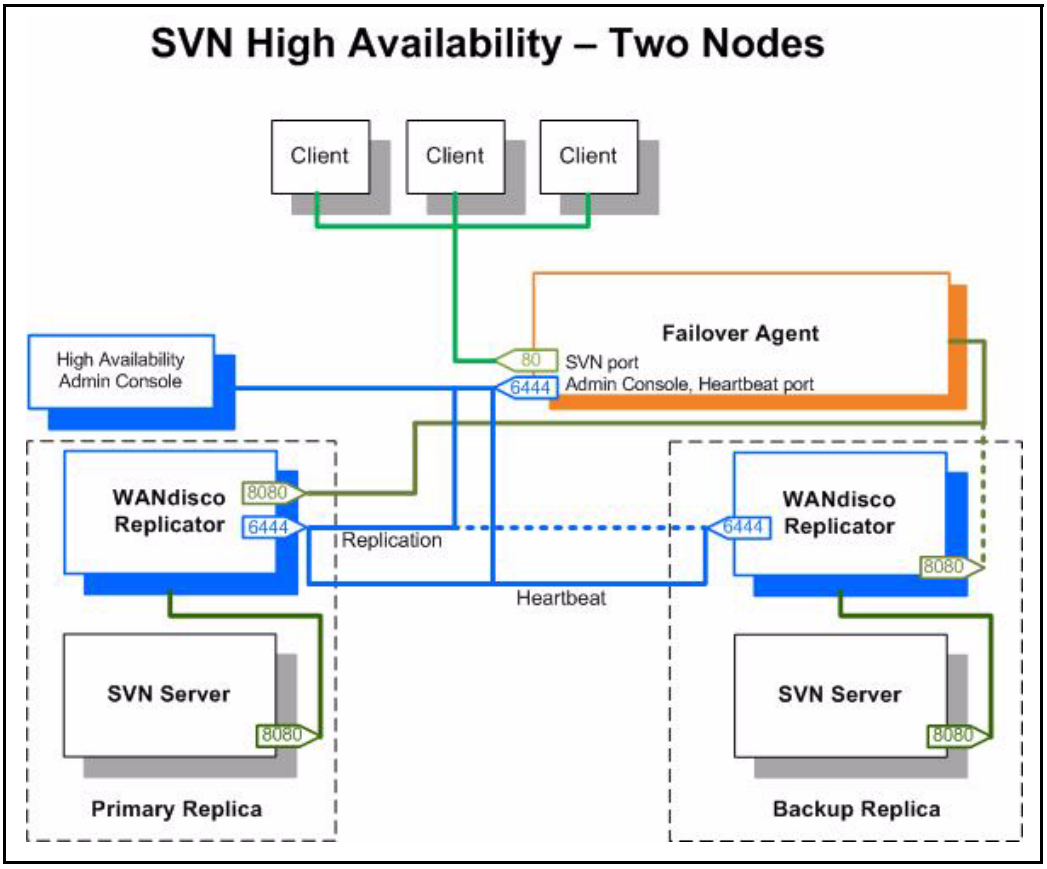

1.8.1 What Happens If the First (Designated) Node Fails

If the designated node fails, the Failover Agent engages the second node, setting the “failed to backup” flag. In order to have the designated node rejoin the HA group, you must perform some steps in the Admin Console, where a wizard guides you.

1.8.2 What Happens If the Second Node Fails

If the second node fails, users see no interruption, since the SVN transactions are engaging the first node anyway. The Admin Console displays the current status. Since the second node is the distinguished node, its failure forces the first node to go into unilateral mode, setting the unilateral flag. Returning the distinguished node to the replication group is described in 6.2, Recovering from Backup Node Failure.

© 2009 WANdisco

1.9 Terms

You should familiarize yourself with these terms.

TERM	DEFINITION
replica	an SVN instance that is an exact equivalent or copy of another SVN instance. In WANdisco's High Availability product, a replica is also called a site.
replicator	It is the intermediary that acts as an application proxy/gateway between the Failover Agent and a given SCM server. Each *Replica* has an associated *Replicator*. It coordinates with other peer replicators to ensure that all replicas of the SCM repositories stay in sync with each other.
Failover Agent	It is the intermediary that acts as an application proxy/gateway between the SVN client and the replicators. The Failover Agent keeps track of which replicas are available, and proxies the SVN client's request to one of them.
replication group	a collection of replicators that work together to keep replicas of a SVN repository in sync.
current node	The node the Failover Agent is communicating with (which then replicates them throughout the HA group).
designated node	The node with priority order 1 (as listed in the prefs.xml file). This is the node the Failover Agent expects to communicate with.
one copy equivalence	all replicas are functionally equivalent copies of each other
GUID	Globally Unique Identifier. WANdisco Subversion High Availability assigns each node a GUID on installation. The nodes identify each other by their GUIDs.
site or node	a server on which is installed a replicator and a replica. The sites comprise the replication group.
distinguished node	The distinguished node acts as a tie breaker for a majority quorum when there are an even number of nodes, making the final decision on replicator operations.
Quorum	A set of nodes that can reach agreement without participation from all nodes. In the case of an even number of nodes, the distinguished node settles a conflict. Quorum is defined in the prefs.xml file. HA by default has Majority Response quorum.
prefs.xml	The preferences files contain information on the replication group. Each site contains all preference files for the entire replication group. The files are specific to each site. The preference files are located in `svn-ha\config`.
SCM Repository	Software Configuration Management repository like SVN
SCM Server	A network server that provides remote access to an SCM Repository

© 2009 WANdisco

TERM	DEFINITION
installDir	this is the installation directory for WANdisco High Availability, for each node and the Failover Agent.
DConE	WANdisco's Distributed Coordinated Engine, the software engine underlying replication.
Install node	The node where you run the High Availability install program.
heartbeat	mechanism the Failover Agent uses to monitor availability of all HA nodes
heartbeat interval	the interval, in seconds, the Failover Agent waits to send an "are you alive?" message to all HA nodes
heartbeat connection timeout	The time the Failover Agent waits before assuming the non-responding node is unavailable. If the current active primary is unavailable, this triggers failover, and the Failover Agent begins communicating with the next node in the priority order.

© 2009 WANdisco

2 Recommended Deployment Practices

2.1 Administrative Pre-requisites

This guide is intended for a Subversion administrator or a user who is reasonably comfortable with:

- Setting up a Subversion based repository
- Configuring inetd/xinetd service on Unix/Cygwin or Windows service
- Installing Perl and required Perl modules
- Installing Java
- Unix or Windows system administration

If you don't meet the above pre-requisites, you may want to contact your Subversion administrator or request that WANdisco perform a professional install for you.

2.2 Physical Environment

WANdisco strongly recommends that you follow these guidelines to ensure the successful installation and use of High Availability:

- a dedicated server for the Failover Agent
- running servers for each node in the HA group, pre-configured with
  - the same operating system
  - the same version of Subversion server
  - a command line zip/unzip utility
  - Java (see Appendix A - Installing Java and Perl)
  - Perl (see Appendix A - Installing Java and Perl)
  - browser with network access to all servers
- e-mail from WANdisco containing the tar file link and the attached production licence key file

2.2.1 Firewalls and Virus Scanners

2.2.1.1 Firewalls

You must determine if your HA group sits inside a firewall or outside of one. If the HA group is inside a firewall, the HA group ports are untouched by the firewall and you need take no action.

However, if any part of your HA group sits outside a firewall, you must configure the firewall to not block or filter the port numbers you specify during installation.

© 2009 WANdisco

2.2.1.2 Virus Scanners

If you have a virus scanner running on your network, you must configure it to not filter traffic on the ports you specify during installation.

2.3 Subversion User Password Management

WANdisco recommends you allow High Availability to manage Subversion user passwords across the HA group. The Admin Console offers an easy user interface for administrators.

If you do not allow WANdisco to manage Subversion user passwords, you must ensure that each Subversion user is entered also in High Availability.

During the installation process, you are asked if you want WANdisco to handle the Subversion user passwords.

© 2009 WANdisco

2.4 Deployment Checklist

You may be familiar with this checklist from an evaluation copy of Subversion HA. It is included here as reference.

System Setup ❖ All sites must share the same operating system	
Supported Operating Systems	Fedora (32 or 64 bit): 6, 7, 8, 9 Red Hat Linux Enterprise Server (32 or 64 bit): 4, 5.2 Sun Solaris (32 or 64 bit): 9, 10 Linux: Linux kernel 2.6 or higher CentOS-4 Windows Server, (32 or 64 bit) 2003 Please read this for more information on NPTL. **Note:** VMware has a tendency to become unresponsive due to memory paging issues even without WANdisco present. Extra tuning may be needed to ensure optimal performance.
Subversion Server Version	1.3 and above. If you are using Subversion 1.5.4, use version 1.3.0 of Apache Portable Runtime.
Subversion Client Version	compatible with local Subversion servers
System Memory	Ensure RAM and swapping containers are at least three or four times the largest SVN file you have. Recommended: 1 GB RAM; 2 GB swapping container
Disk Space	• SVN: depends on the number of projects and issues • HA Transaction Journal: Recommended - equivalent of seven days of changes
File Descriptor limit	Ensure hard and soft limits are set to 64000 or higher. Check with the `ulimit` or `limit` command.
Journaling File System	Replicator logs should be on a journaling file system, for example, ext3 on Linux or VXFS from Veritas. Notes: NTFS is not a journaling file system: ext4 is a journaling file system, however WANdisco does not support its use because of its deferred writes.
Maximum User Process Limit	At least three times the number of SVN users.
Java	Install JDK 1.5.0. **Note:** There should not be any spaces or control characters in the path where Java is installed. For example, c:\Program Files\java does not work with WANdisco as a JAVA install directory. See Appendix A - Installing Java and Perl.
Perl	Install version 5.6.1 or later. See Appendix A - Installing Java and Perl.

© 2009 WANdisco

Network Setup	
Reserved Port (i.e. 6444, another port for synchronizing)	SVN HA needs a dedicated port for DConENet (replication protocol) and HTTP (for Admin Console). WANdisco also recomends having a port available in case you have to copy (rsync) the repository from one site to another. If your network has a firewall, notify the firewall of the port numbers.
Firewall or virus scanner	Notify the firewall and any virus scanners of the SVN HA port numbers.
VPN	Set up IP sec tunnel. Ensure WAN connectivity. (optional)
Persistent Connection Keep Alive	Ensure VPN doesn't reset persistent connections for WANdisco, or else ensure there are no RST bugs
DNS Setup	Best to use IP address for WANdisco related hostnames, or else ensure DNS availability
WNAdisco Setup	
Agreement Threads	Tune based on number of concurrent Subversion writers
Reader/Writer Network IO Thread Pool	Tune based on Subversion client connection rate, file transfer rate
ConnectionKeepAlive timeout	Tune inactivity timeout for persistent DConeNet/DFTP connections based on VPN/WAN router set up
Message Queue Max Thread Pool Size	Tune based on SVN write concurrency
Maximum connections per IO thread	Tune if active SVN user population is large (greater than 100)
Disk space for recovery journal	Provision large disk for `logs/tmp`, at least number of commits within a two to four hour window
Reader/Writer Network IO Thread Pool	Tune based on client connection rate, file transfer rate
ConnectionKeepAlive timeout	Tune inactivity timeout for persistent DConeNet/DFTP connections based on network delays
HeartBeatInterval	Tune based on WAN latency, ensure interval is >> max ping time between Failover agent machine and replicator
MissingHeartBeat Count	Default is 4, if you see too many false alarms (spurious failover events) increase
Admin Email Address	To generate email notifications from the failover agent. Requires `/usr/sbin/sendmail`.
Notify all users that they must flush their client cache.	

© 2009 WANdisco

Apache 2 Setup (for http:// access)	
Apache version	All sites have the same version, 2.2.3 and higher.
Apache modules version	All sites have the same version of mod_dav and mod_svn_dav WANdisco does not support the mod_deflate.c for SVN_DAV
Apache KeepAlive	WANdisco recommends setting KeepAlive to on
Apache KeepAliveTimeout	WANdisco recommends setting the KeepAliveTimeout to at least 30 seconds
Location URI	All sites' apache conf files have same location URI for SVN repository access
Require valid user for write methods	Ensure that all WebDAV methods require authentication for SVN-DAV protocol
Repositories are in sync	Before starting WANdisco, ensure that repositories at all sites are in sync
User Name/Password Data-base	Must be the same at each replica
Using port 80 for Replicator	Std Port 80 avoids confusion, change default Apache port if using 80
Apache server port	Non-standard Apache server port to avoid conflict with replicator port?
File Permissions in svnroot	See this article (http://www.reallylongword.org/articles/svn/)

© 2009 WANdisco

3 Installation

Before you start the installation, make sure you have met the requirements listed in the first paragraph of Chapter 2, Recommended Deployment Practices.

Current Subversion users must back up their current repository to use it as a baseline for replication. See 5.2, Establishing a Baseline for Replication.

3.1 First Time Installation

You can run the installation program from any machine. The output is zip files, which you then install on the High Availability nodes and the Failover Agent.

NOTE:

This procedure is for an HA group of three or more nodes. For a two node group, see 3.1.12, First Time Installation for Two-Node Group.

Step 1 Save the svnha.tar.gz file.

Step 2 For Windows, create a directory, and unzip the file in that directory. For other platforms, untar the file. The uncompressed file produces a directory, `svn-ha-installer`.

Step 3 Copy the license evaluation key file to the `config` folder in the `svn-ha-installer` folder.

Step 4 At the command prompt (or editor), go to `svn-ha-installer/bin`.

Step 5 Type

```
perl setup
```

The last line of text returned contains a WANdisco URL.

Step 6 Copy the URL returned in the last step and paste it into a browser. The High Availability Welcome page appears.

Step 7 Click **Continue** at the Welcome text.

Step 8 Read the User Licence Agreement. You must agree to the terms to continue.

© 2009 WANdisco

3.1.1 Configuring the High Availability Group

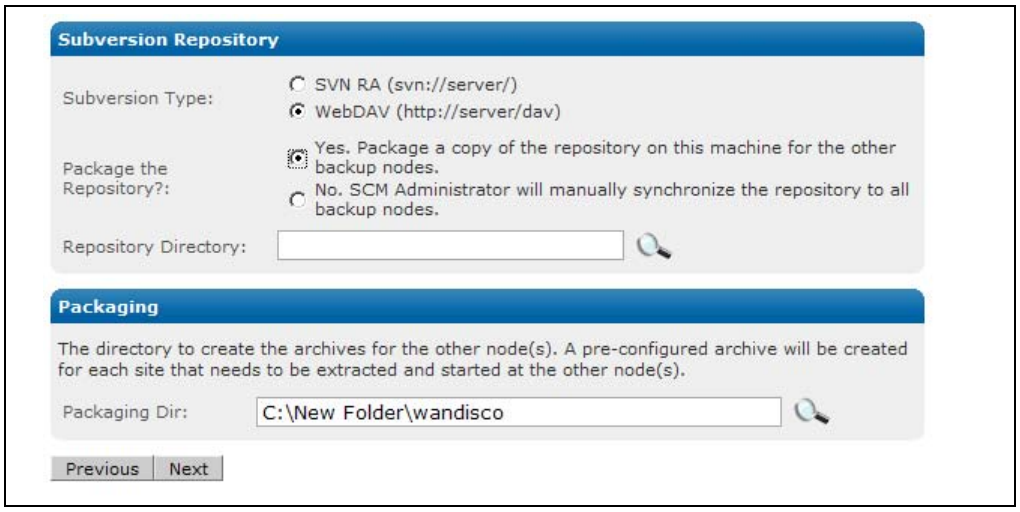

Step 9 Select the Subversion type.

Step 10 The installer can copy your Subversion repository and distribute it during the installation process, or an administrator can manually synchronize the repository to all the sites. If you check the Yes radio button, you are asked to browse to the repository's location.

Step 11 Select a packaging directory.

Step 12 Click **Next**.

© 2009 WANdisco

3.1.2 Configuring the Failover Agent

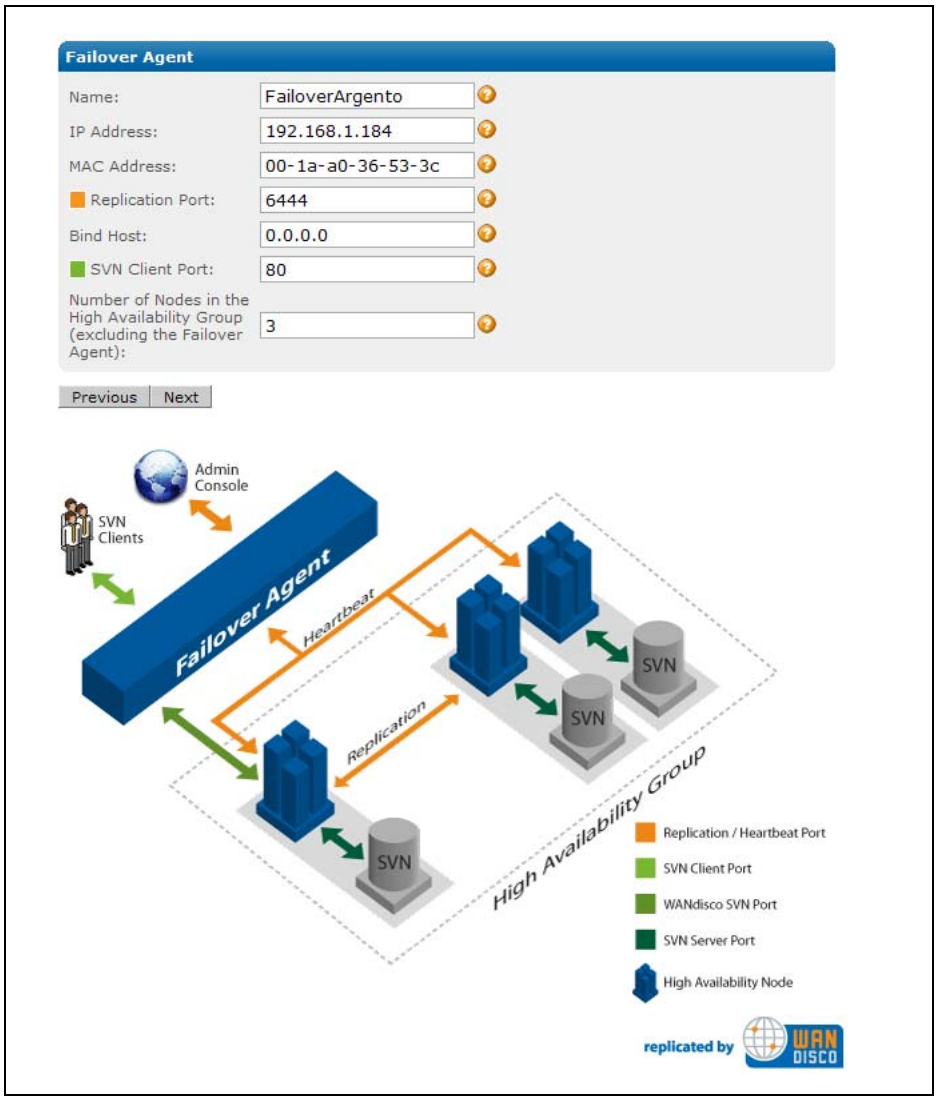

Step 13 In the Name filed, specify a name for the Failover Agent. The filed accepts alphanumeric characters, no spaces are allowed.

Step 14 Enter the IP and MAC address. To find these addresses, go to that server's command prompt. For Unix, type

```
ifconfig
```

For Windows, type

```
ipconfig /all
```

Make sure the IP address is one you know is valid. The MAC address is the physical address.

© 2009 WANdisco

For example, on a Windows machine, the output would look like:

```
Physical Address. . . . . . . . . : 00-1A-A0-36-53-3C
Dhcp Enabled. . . . . . . . . . . : Yes
Autoconfiguration Enabled . . . . : Yes
IP Address. . . . . . . . . . . . : 192.168.1.124
Subnet Mask . . . . . . . . . . . : 255.224.255.0
Default Gateway . . . . . . . . . : 192.124.1.1
DHCP Server . . . . . . . . . . . : 192.122.1.1
DNS Servers . . . . . . . . . . . : 192.128.1.50
                                   64.405.172.26
```

Step 15 Replication Port is the port that WANdisco uses for its communications. (Protocols are replication, HTTP and heartbeat.)

Step 16 The Bind Host is the IP address/host name that WANdisco listens on for incoming connections. Unless there are multiple network cards on the Failover Agent, this is typically set to 0.0.0.0.

Step 17 The SVN Port is for Subversion client communication.

Step 18 Specify the number of nodes in the High Availability group, excluding the Failover Agent.

NOTE:

WANdisco recommends using more than two nodes in the HA group. A two-node group is vulnerable to a single point of failure if one node goes down.

If you did enter 2 in step 18 for a two-node replication group, then proceed to section 3.1.12, First Time Installation for Two-Node Group.

© 2009 WANdisco

3.1.3 Configuring the First Node

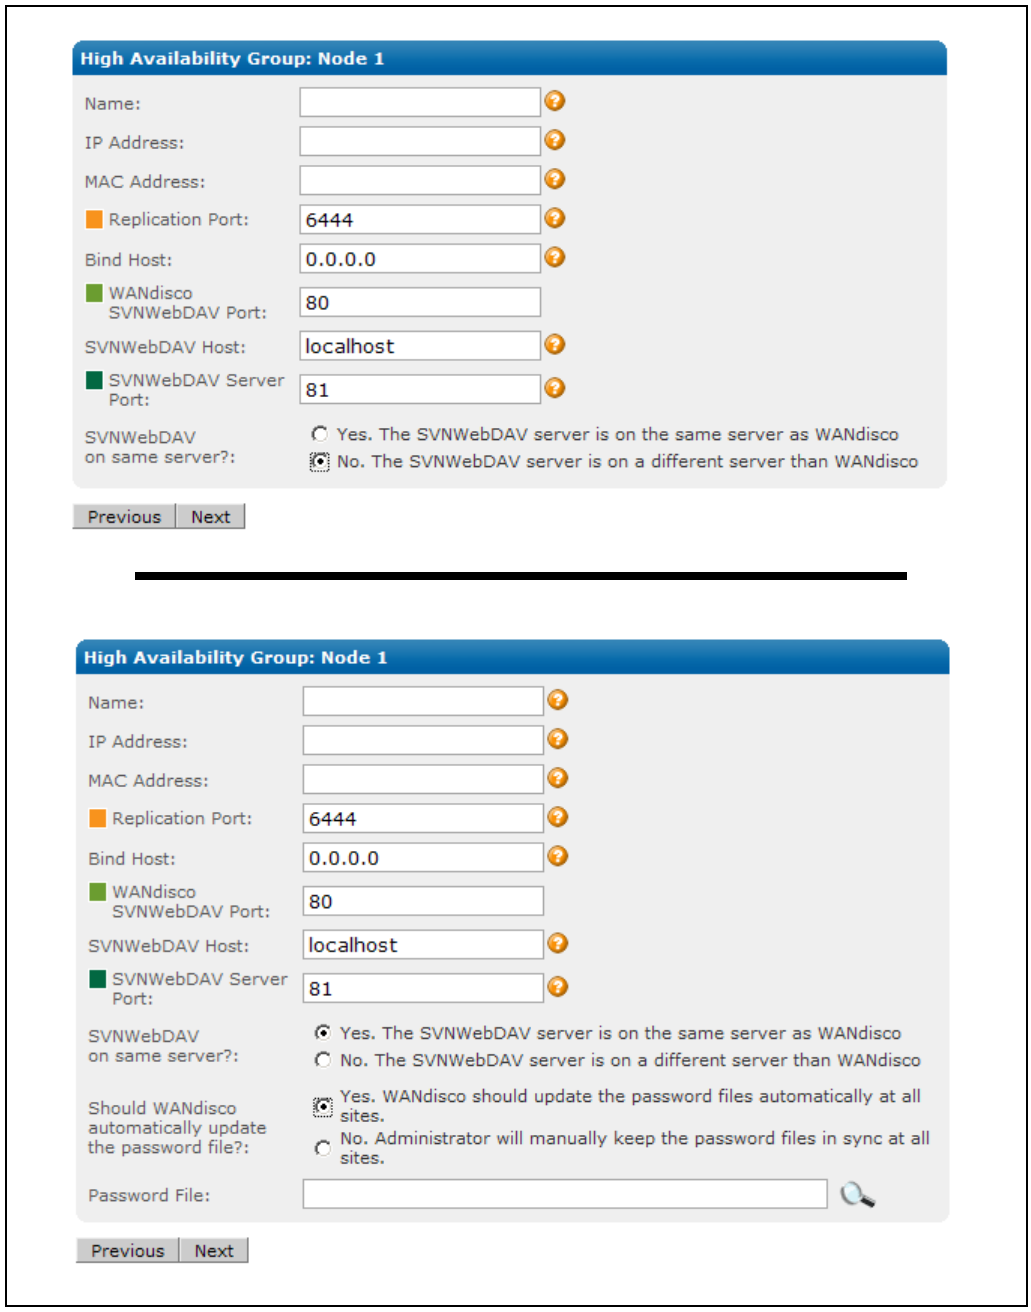

NOTE:

The priority order for failover is based on the order you specify here. For a full discussion on priority order, see section 1.5, Replication Example.

Step 19 Fill in the Name, IP Address, and MAC Address fields.

© 2009 WANdisco

Step 20 For SVN Port, use a different value than the Failover Agent's SVN port number.

Step 21 For the SVN server port, WANdisco recommends you use a different port than the SVN port.

Step 22 WANdisco needs to know if it is sharing a server with Subversion or not: answer yes or no to the *SVN on same server?* question.

Step 23 WANdisco needs to know if it is updating the Subversion password file, or whether you (or a Subversion administrator) is going to synchronize the password files: answer yes or no to the *Should WANdisco automatically update the password file?* question.

Step 24 If you answered yes to the previous question, enter the Subversion password file path. Use the magnifying glass to browse.

3.1.4 Configuring Subsequent Nodes

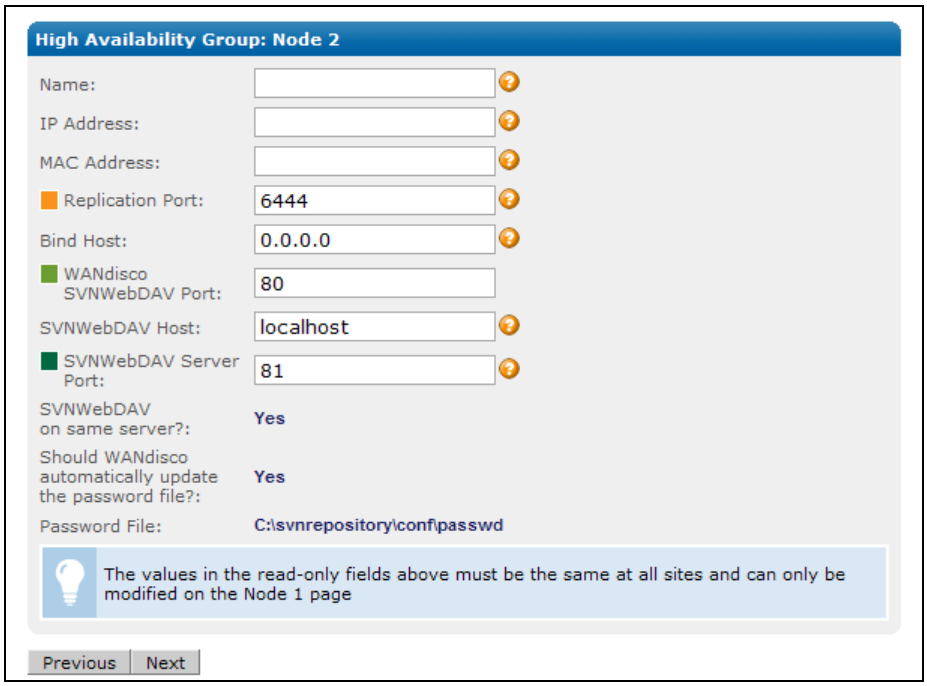

Step 25 Fill in the values. WANdisco recommends you use the same port numbers as you did for the previous node, as it can eliminate confusion.

Step 26 Since all the nodes must be configured identically, the install program uses the response from previous pages for the questions *SVN on same server?* and *Should WANdisco automatically update the password file?*

Step 27 Fill out a page for each node in your HA group.

© 2009 WANdisco

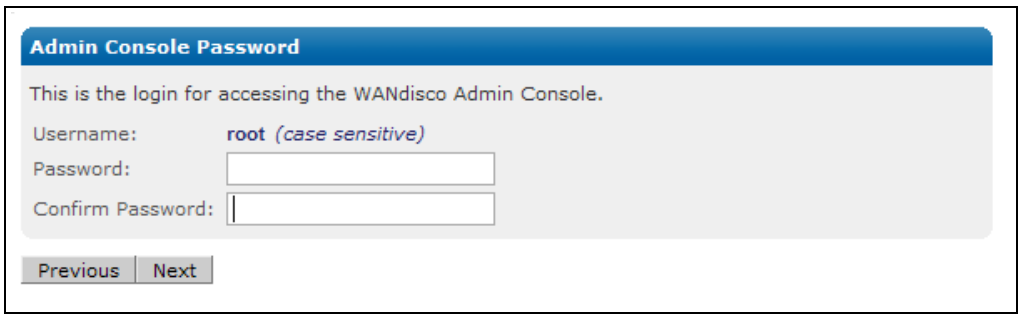

Step 28 Enter and confirm a password for WANdisco's username.

Step 29 Click **Next**.

© 2009 WANdisco

3.1.5 Looking Over the Summary Page

After you have defined each node in your cluster, you are presented with the Summary page.

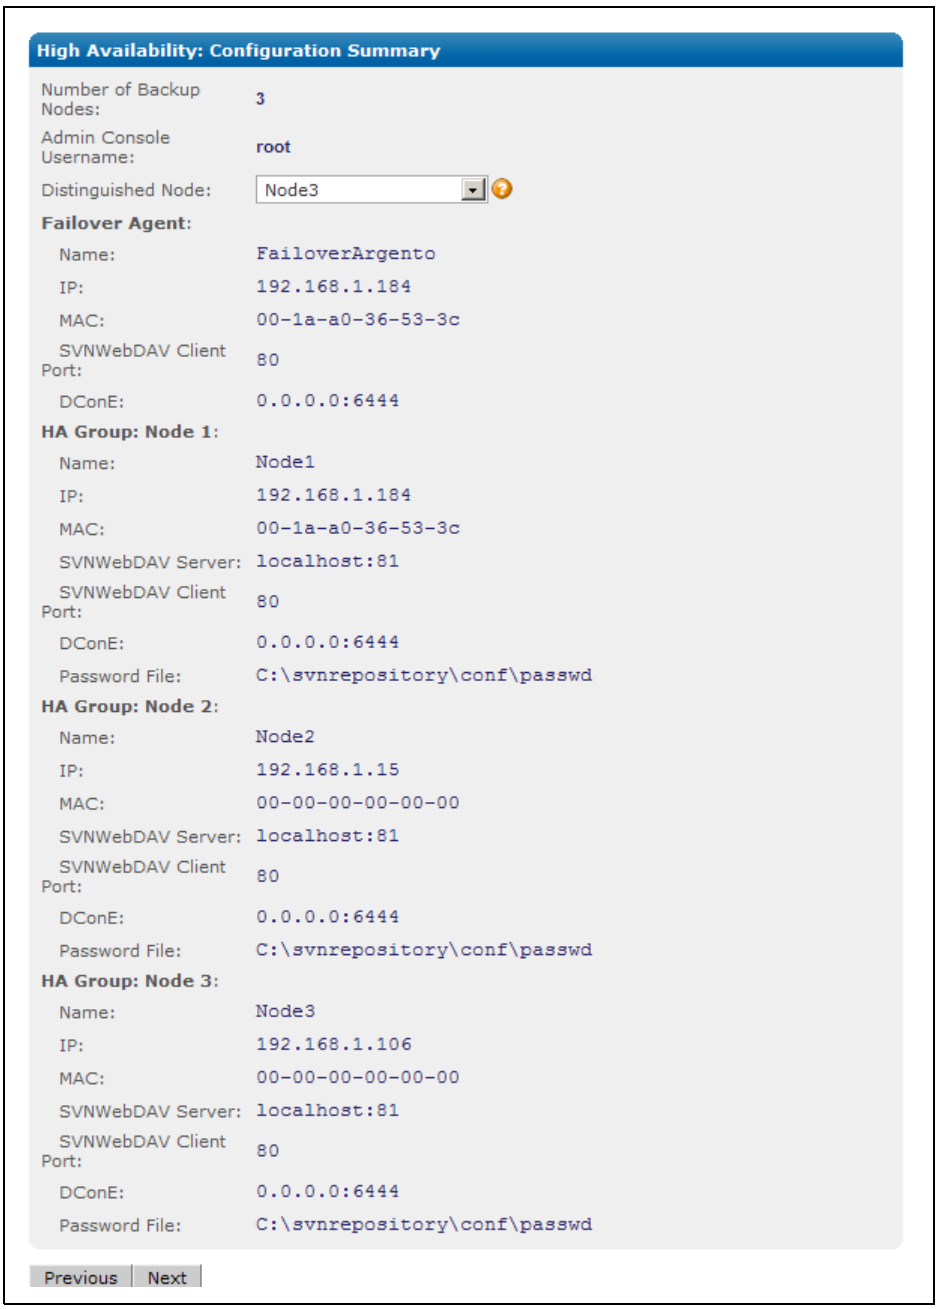

© 2009 WANdisco

3.1.5.1 Establishing a Distinguished Node

In the **Distinguished Node** field, select the node to be the HA group's distinguished node. The pull-down list includes all of the group's nodes. WANdisco recommends you choose the most reliable server in the group as the distinguished node.

NOTE:

> If you have a two-node group, the backup node is always the distinguished node.

WANdisco recommends you print the Summary page for handy reference.

Step 30 After verifying all settings are appropriate, check each checkbox.

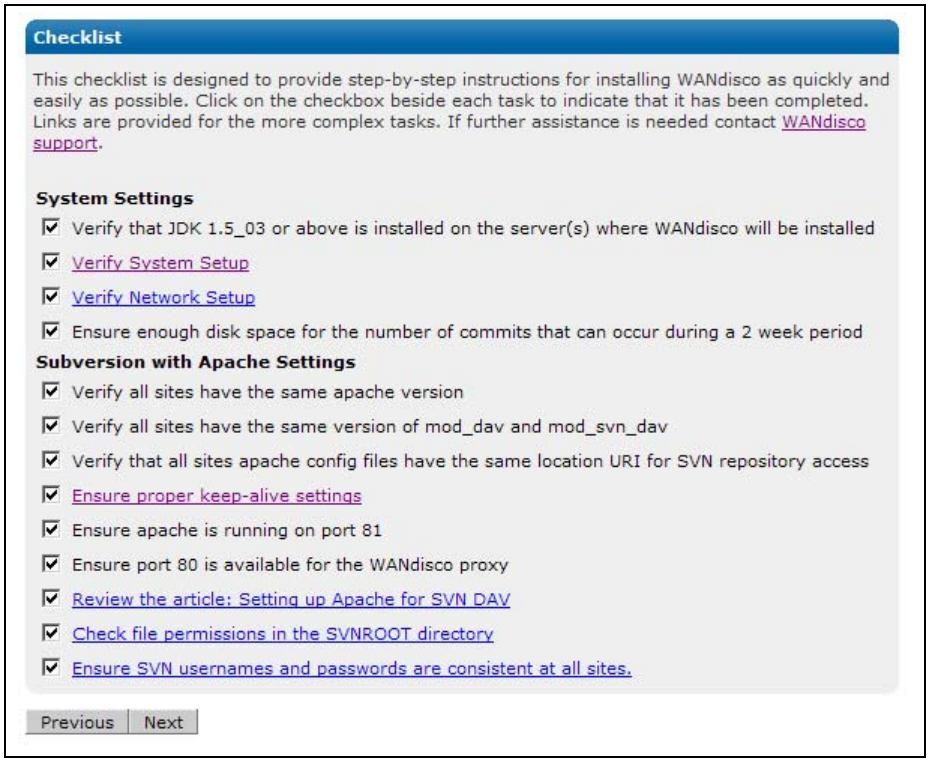

Step 31 Click **Next**.

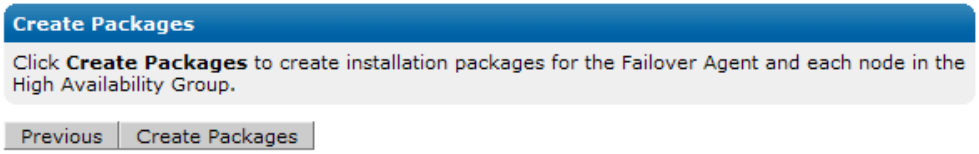

Step 32 Click **Create Packages**.

© 2009 WANdisco

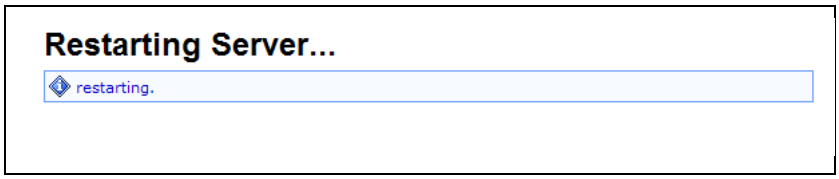

3.1.6 Package Creation

The installer creates a package for the Failover Agent and each node in the group.

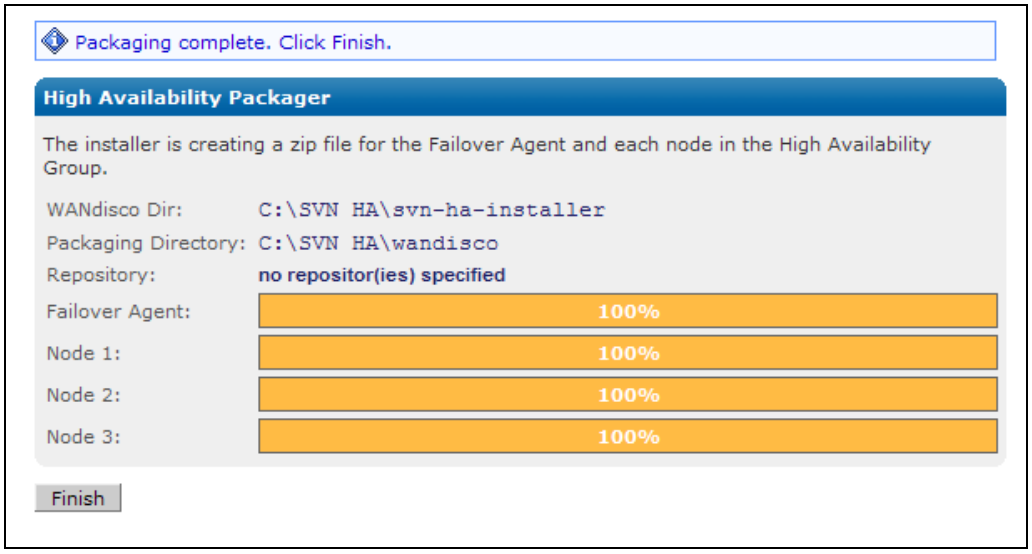

Step 33 Click **Finish**.

© 2009 WANdisco

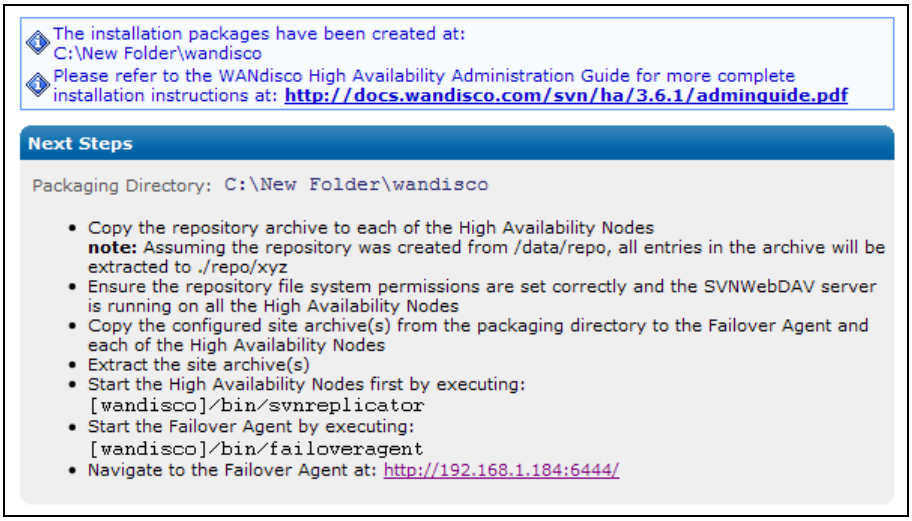

The install program places zip files for the Failover Agent and the nodes in the `wandisco` directory.

Step 34 You can follow the instructions on the installation page, or continue with this procedure. However, make sure you start the Failover Agent last.

3.1.7 Installing the Nodes and the Failover Agent

Step 35 Copy the *<node name>.zip* files to a new directory their respective sites.

Step 36 Unzip the files for each node and for the Failover Agent. This creates a directory, `svn-ha`, where all the uncompressed files are.

Step 37 For Unix, go to the *<node name>/bin* directory, and type

```
chmod +x *
```

Each node and the Failover Agent contains these High Availability directories:

DIRECTORY	CONTENTS
bin	Contains scripts like `failoveragent`, `shutdown`
config	Contains the [replicator]/config/prefs.xml file used to configure HA.
lib	Contains the `jar` files and DLLs that are required to run the product.
docs	Contains the administration guide in PDF format.
logs	Contains the pid file, log files and other temporary files. WANdisco HA's log file is named `FailoverAgent?prefs.log.0` and `svnProxyServer-prefs.log`.
systemdb	Contains the system database with its transaction journal. **Warning**: Deleting or modifying files from systemdb will likely corrupt your installation.

© 2009 WANdisco

3.1.8 Handling Subversion Data

Step 38 If you had WANdisco include a copy of the repository in the installation package, it is extracted to a folder with the same name as the original repository.

If you are manually synchronizing the Subversion repositories, see 5.2, Establishing a Baseline for Replication.

3.1.9 Validating the Installation

It is a good idea to take the time to validate the installation. That includes ensuring:

- all the databases are identical
- all database user privileges are identical at all nodes

3.1.10 Starting the Nodes and the Failover Agent

Step 39 At each node, go to `/svn-ha/bin` and type

```
svnreplicator
```

Step 40 At the Failover Agent, go to `svn-failover` and type

```
failoveragent
```

High Availability is now up and ready for use.

3.1.11 Post Installation Configuration

There are a few issues you may want to address pretty quickly after installation. They are described here, but any action you take is through the Admin Console.

3.1.11.1 Tuning the Heartbeat Frequency

The WANdisco Failover Agent by default uses a heartbeat frequency of one heartbeat every second. This is appropriate for a LAN deployment.

If you are deploying High Availability over a WAN, then you may want to change the heartbeat frequency based on LAN latencies. The heartbeat interval should be two to three times the expected LAN latency between the Failover Agent and the replicator site. This ensures that the Failover Agent can distinguish between missing heartbeats and a slow network link.

© 2009 WANdisco

You can also adjust the number of missing heartbeats that dictates when the WANdisco Failover Agent marks a replicator node as unresponsive, triggering failover. The default is four.

For a complete discussion of failover and the heartbeat, see section 1.2, Failover and the Heartbeat. To change the default values for the heartbeat, see the heartbeat definition in Chapter 4, Using the Admin Console.

3.1.11.2 Configuring Node Start Up Commands

The WANdisco Failover Agent is capable of starting the WANdisco replicator nodes from the Admin Console directly, provided you enter the startup commands in the Admin Console. The WANdisco Failover Agent also can use `ssh` to launch WANdisco replicators on remote machines.

To configure start up commands, see the Start Command definition in Chapter 4, Using the Admin Console.

3.1.11.3 Email Alerts for Failover Events

The WANdisco Failover Agent can generate email alerts whenever it detects an event related to failover. Examples of such events are:

- unable to transition to unilateral mode
- when transitioned to unilateral mode
- when unilateral mode starts
- when failed to backup
- when the primary is not available when the failover agent starts

To set up email alerts, see the Config definition in Chapter 4, Using the Admin Console.

NOTE:

Email notifications are not supported on the Windows platform.

© 2009 WANdisco

3.1.12 First Time Installation for Two-Node Group

This section is for users who have selected a two-node group on the Failover Agent page.

Although WANdisco does not recommend it, you can choose a two-node group. A two-node High Availability group forces changes to the failover logic as well as the quorum. For a two-node group, the default quorum is singleton response, and the backup node is the distinguished node.

There is a separate section for procedures for two-node groups. See Chapter 6, Procedures for Two-Node HA Groups.

3.1.12.1 Configuring the Primary Node

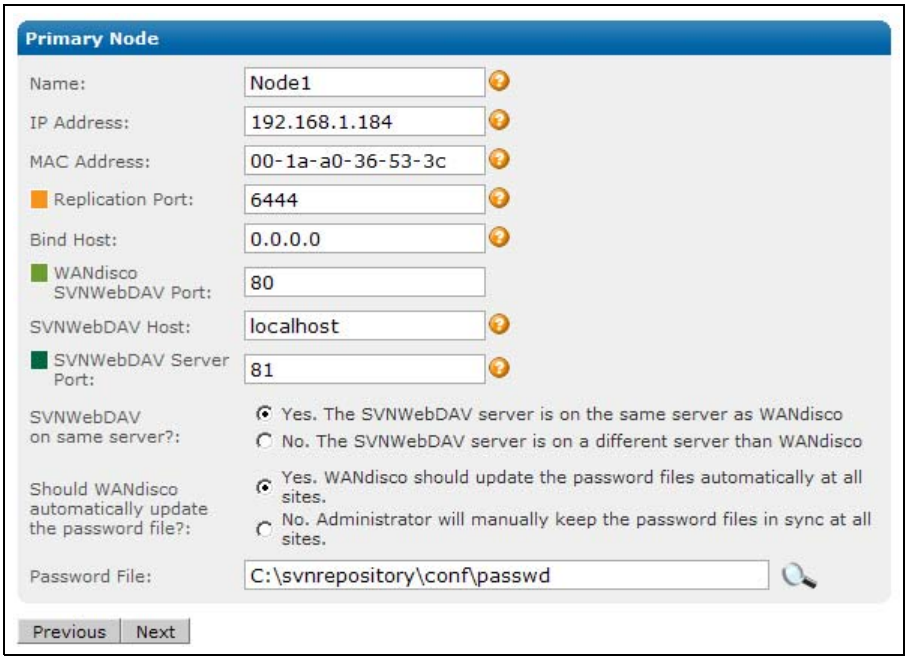

NOTE:

The priority order for failover for a two-node group is always the same: the primary node is priority 1, the backup node is priority 2.

Step 1 Fill in the fields.

Step 2 For SVN Server Port, use a different value than the Failover Agent's SVN port number.

Step 3 For the SVN server port, WANdisco recommends you use a different port than the SVN port.

Step 4 WANdisco needs to know if it is sharing a server with Subversion or not: answer yes or no to the *SVN on same server?* question.

© 2009 WANdisco

Step 5 WANdisco needs to know if it is updating the Subversion password file, or whether you (or a Subversion administrator) is going to synchronize the password files: answer yes or no to the *Should WANdisco automatically update the password file?* question.

Step 6 If you answered yes to the previous question, enter the Subversion password file path. Use the magnifying glass to browse.

3.1.12.2 Configuring the Backup Node

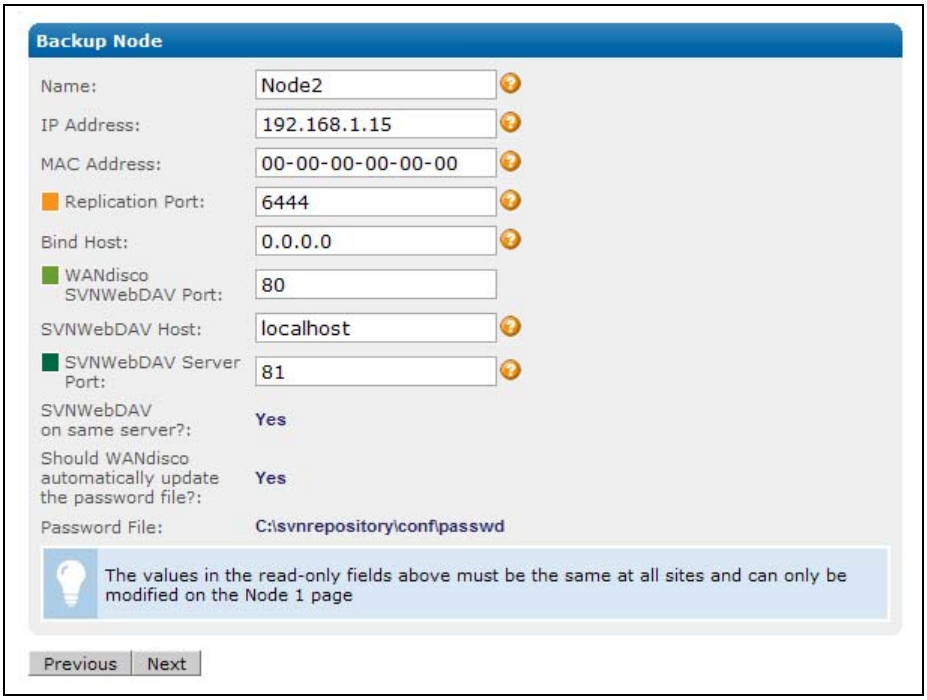

Step 7 Fill in the values. WANdisco recommends you use the same port numbers as you did for the previous node, as it can eliminate confusion.

Step 8 Since all the nodes must be configured identically, the install program uses the response from previous pages for the questions *SVN on same server?* and *Should WANdisco automatically update the password file?*

Step 9 Click **Next**.

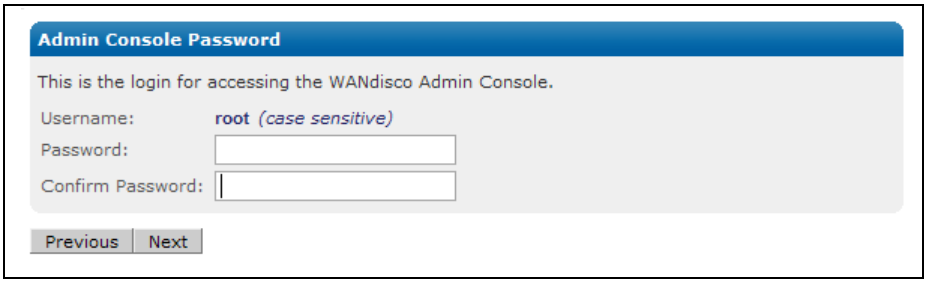

© 2009 WANdisco

Step 10 Enter and confirm a password for WANdisco's username.

Step 11 Click **Next**.

Now you can complete the installation by proceeding with section 3.1.5, Looking Over the Summary Page.

3.2 Installing Upgrades

NOTES:

This procedure involves taking Subversion offline. Please follow your company procedures about notifying Subversion users of down time.

You must either export the existing database or manually synchronize to the new databases.

3.2.1 Disconnect Subversion Users

Step 1 After notifying Subversion users of the downtime, navigate to the Dashboard.

Step 2 Watch for **Total Transactions Pending** to reach **0**. The **Total Transactions Pending** must be 0. If you do not wait for the **0**, the databases will be out of synch.

Step 3 When the count reaches 0, shut down the Failover Agent.

Step 4 Stop all the HA nodes.

3.2.2 Back Up or Export Subversion Data

Step 5 Backup or export the Subversion repository.

3.2.3 Back Up WANdisco Data

Step 6 Go to the HA Node page. Click **Export Settings**.

© 2009 WANdisco

3.2.4 Preparing the Install Node

Step 7 On the original install node, zip these directories:

```
config
```

Step 8 Delete this directory:

```
systemdb
```

Step 9 In the `svn-ha-installer/config` directory, delete these directories:

```
membership
security
passwd
```

3.2.5 Preparing Subsequent Nodes

Step 10 On each node, zip the `svn-ha` or `svn-failover` directory.

Step 11 On each node, delete the `svn-ha` or `svn-failover` directory.

3.2.6 WANdisco-supplied .jar and License Key Files

Step 12 Save the svnha.tar.gz file.

Step 13 Verify the md5 checksum. For Unix, type

```
md5sum
```

Step 14 Copy the jar file to the `svn-ha-installer/lib` directory.

Step 15 Unzip or untar the file.

Step 16 Copy the licence evaluation key file to the **config** folder in the **svn-ha** folder.

3.2.7 Running the Install Program

Step 17 Run this command on the original install site. The install program automatically populates all previous configuration information.

```
svn-ha-installer/bin/setup
```


© 2009 WANdisco

The Setup page appears.

Step 18 Enter the number of nodes in the replication group.

Step 19 Click **Next** to continue.

Step 20 Make changes as needed.

Step 21 Verify the configuration on the Summary page. Click **Next**.

Step 22 Wait while the wizard creates packages for all the nodes. Click **Next**.

Step 23 Complete packaging by clicking **Finish**.

3.2.8 Installing and Running the HA Group

Step 24 Copy the `<node name>.zip` to each site.

Step 25 For all nodes, unzip the `<node name>.zip` file.

Step 26 In the `/bin` directory, type

```
svnreplicator
```

3.2.9 Importing WANdisco Data

Step 27 Import WANdisco data. Go to the HA Node page on any node, and click **Import Settings**. Locate the file you exported in step 6. You only have to import it once, WANdisco replicates it on the other nodes.

3.2.10 Validating the Installation

It is a good idea to take the time to validate the installation. That includes ensuring:

- all the databases are identical
- all database user privileges are identical at all nodes

3.2.11 Installing and Running the Failover Agent

Step 28 Copy the `svn-failover.zip` to each site.

© 2009 WANdisco

Step 29 Unzip the `svn-failover.zip` file.

Step 30 In the `/bin` directory, type

`failoveragent`

The new installation is complete.

© 2009 WANdisco

4 Using the Admin Console

The Admin Console is a simple interface that allows you to monitor and perform simple tasks for your HA group and the Failover Agent. The Admin Console for the Failover Agent shows the other HA nodes onscreen, so you may want to have the Failover Agent's Admin Console running.

4.1 Starting the Admin Console

To start the Admin Console for the Failover Agent or any node, in a browser's address bar, type

http://<IP address>:<replication port number>

The Admin Console's Home page appears.

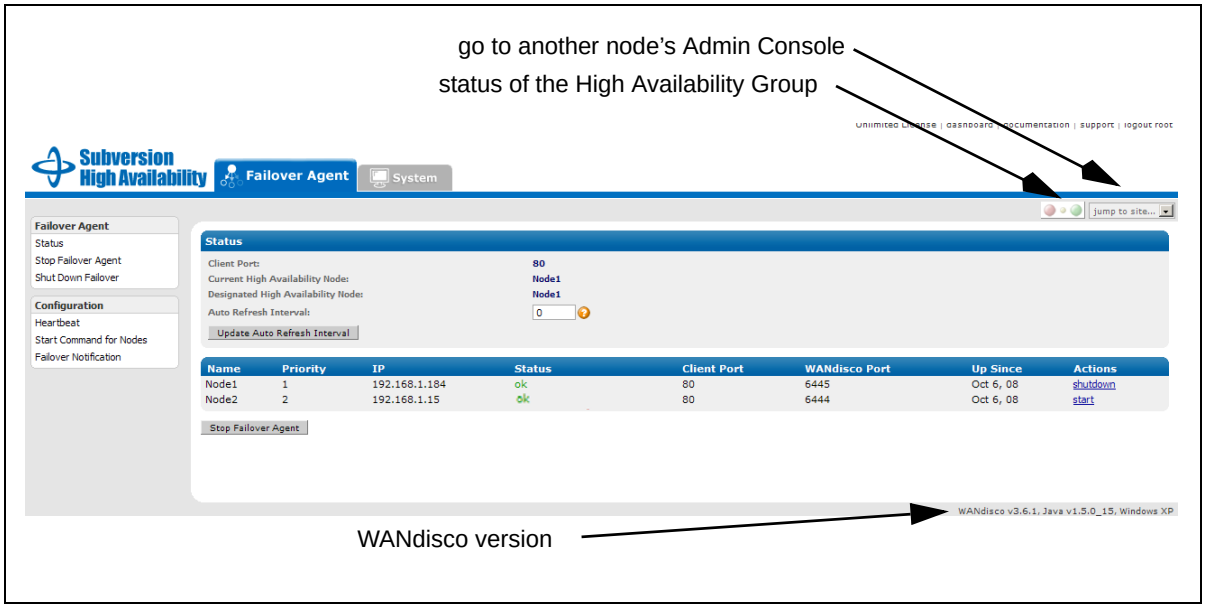

© 2009 WANdisco

4.2 Failover Agent Pages

The Failover Agent's Admin Console has two pages: The Failover Agent page and the System page.

4.2.1 Failover Agent Page

The Failover Agent page displays the status of the other nodes in the HA group, gives information on the current active primary, and offers several commands in the left menu.

The status of the High Availability group displays in the upper right. One of three circles pulses, giving the current status: green signifies all nodes are up, yellow signifies at least one node is not responding, and red signifies that all nodes are not responding.

The Status section names the Current Active Primary (the node listening to the Subversion client through the Failover Agent) and the Designated Primary (priority level 1 in the prefs.xml file). These are the same node unless failover has occurred. The Auto Refresh Interval offers you the ability to automatically update the page. The default is 0.

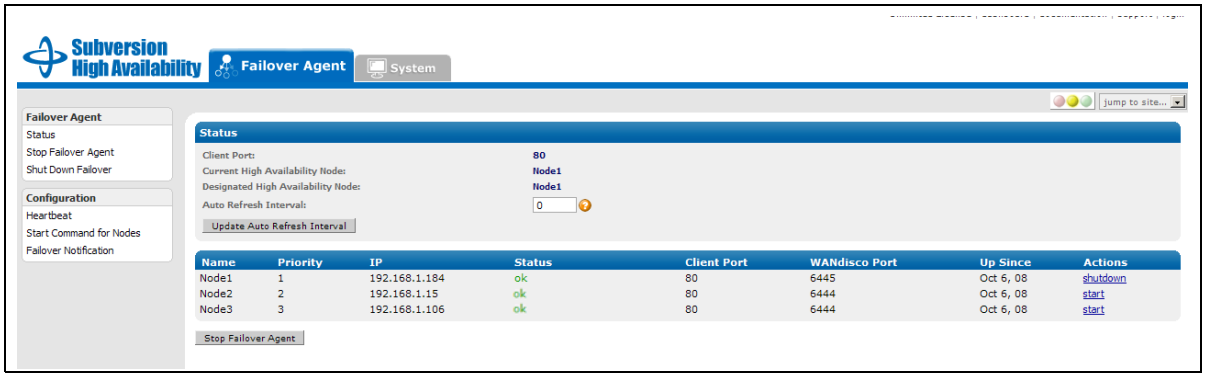

Each node is listed by name, priority order, IP address, current status (**ok** and **not responding**), client port (in use only by the current active primary), replication port, date of last start up, and action (**shutdown** or **start**).

NOTE:

If you shut down the current node, you trigger a failover. Shutting down any other node does not result in a failover.

A High Availability group of three or more nodes has a majority response quorum. If you shut down a majority of nodes in your HA group, replication stops, and can only continue when a majority of the HA nodes are running.

© 2009 WANdisco

4.2.1.1 Left Side Menu

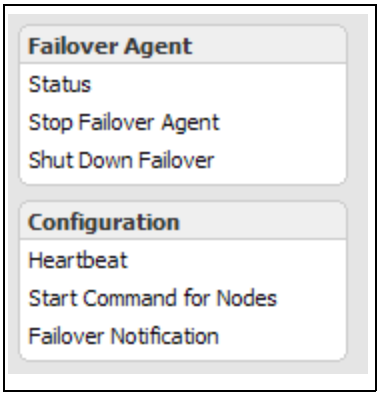

High Availability

Status	This command displays the HA group's status in the page's main area.
Stop Failover Agent	This stops Subversion access to Subversion clients. You must confirm your action.
Shutdown Failover	This shuts down the Failover Agent.

Config

Heartbeat	For a complete discussion on the heartbeat, see 1.2, Failover and the Heartbeat.
Missing Heartbeat Count	The number of missing heartbeat responses the FA gets before marking a node as unavailable.
Interval	The time between queries from the HA to the nodes. You can specify a different number for each node.
Connection Timeout	The time to wait before assuming the query has failed. You can specify a different number for each node.
Save All	Save changes for all nodes.
Start Command	Enter in commands to start the nodes.
SSH Command	Enter syntax for the SSH command.
Start Command	Enter syntax to start the node. For Windows, in the Start Command field, enter `perl -x -S <pathname>\svn-ha\bin\svn-replicator`
Save All	Save changes for all nodes.
Notifications	The Failover Agent can send emails when: a) the current active primary dies, triggering failover; b) priority 1 node comes back online; c) in a two node group, the second node dies, triggering backup exclusion protocol.
Admin Email	Enter in the address to receive Failover Agent emails.

Primary

Start Primary	Start the current acting primary (if the commands are programmed with the Start command on this page).

© 2009 WANdisco

4.2.2 The System Page

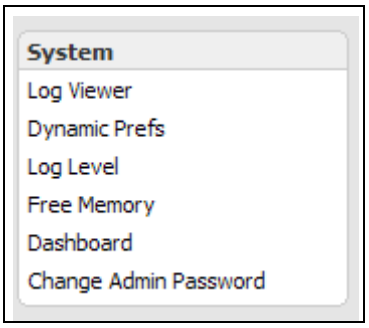

The System page offers several commands and utilities for the High Availability group.

4.2.2.1 Left Side Menu

System

Log Viewer	
FailoverAgent-prefs.log	WANdisco's log file for the Failover Agent.
svnProxyServer-prefs.log	The log file for replication.
Show Log	Click this to display the log in the Dashboard.
Dynamic Prefs	These fields are used internally and for customer support.
Log Level	WANdisco uses one log, and the default level is **info**. The levels vary from **severe**, where you get only the most severe warnings, to **finest**, which logs every action.
Free Memory	This command frees the memory (GB stands for garbage collection) for the current node. The command occurs when you click on this menu selection. The display shows information on the command that was just performed.
max mem used by JVM	the maximum memory that JVM can use on the current node
free memory before GC	the amount of free memory before you ran this command
free memory after GC	the amount of free memory after you ran this command
memory freed	the total amount of memory freed at the command's completion
Dashboard	offers another way to get to the Dashboard. The Dashboard is discussed in detail in section 4.2.3, The Dashboard.
Change Admin Password	You can change the WANdisco admin password here.

© 2009 WANdisco

4.2.3 The Dashboard

There are two ways to get to the Dashboard from the Failover Agent pages.

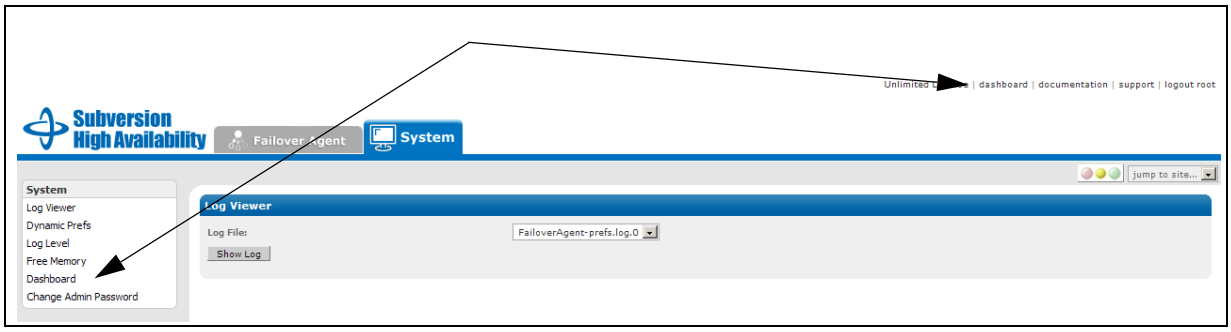

The Dashboard shows each node's transactions. As soon as High Availability receives a Subversion transaction request, that transaction joins the replication group's queue. (There is one queue for the replication group.) High Availability keeps track of which node it emanated from, but otherwise, the transaction joins the queue for the replication group.

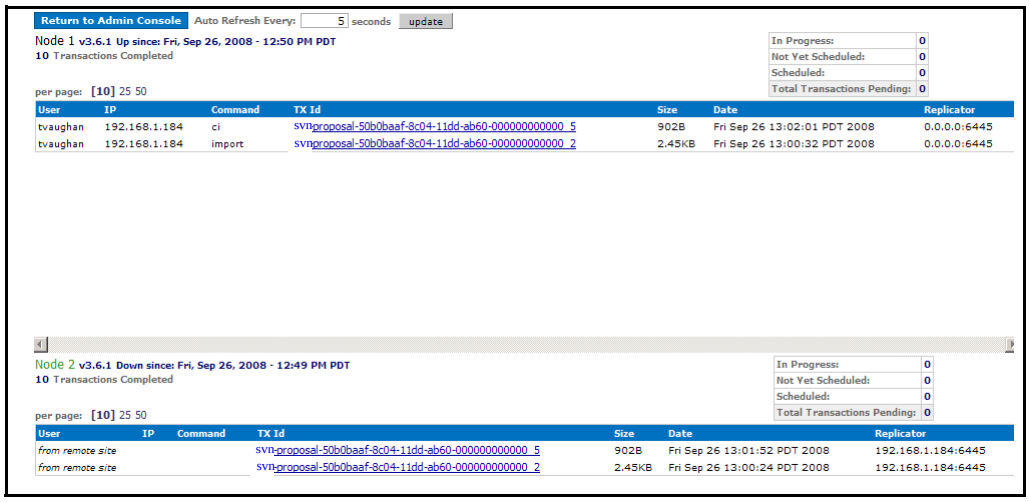

© 2009 WANdisco

4.2.3.1 Pending Transactions

There are three statuses of replicated transactions before they are committed to Subversion.

- **Not Yet Scheduled** - these transactions are in the queue
- **Scheduled** - these transactions are approved by the quorum and are waiting to be executed
- **In Progress** - these transactions are in the process of being completed

The **Total Transactions Pending** lists the total number of transactions in all statuses.

4.2.3.2 Completed Transactions

Once a transaction is completed, the Dashboard lists it in the transaction list, and that transaction is no longer considered in the **In Progress** count. The completed transaction joins the count in the **Transaction Completed** display. This display keeps count of transactions since replication began.

4.2.3.3 Refreshing the Dashboard Display

The Dashboard by default does not refresh. As High Availability is completing many transactions, a display that refreshes often would be very confusing to read. While the order of completed transactions is the same at all nodes, the real time of when a transaction is posted may vary from node to node. To refresh the Dashboard, click **Update**. You can also set the Dashboard to refresh automatically by entering a number in seconds in the field.

NOTE:

The value you enter in the Auto Refresh box is replicated throughout the HA Group and the Failover Agent. Whenever you change this value, at any node, that value becomes the Auto Refresh value for the HA group.

© 2009 WANdisco

4.3 HA Group Pages

4.3.1 The Node Page

The HA Node tab offers the status of this node.

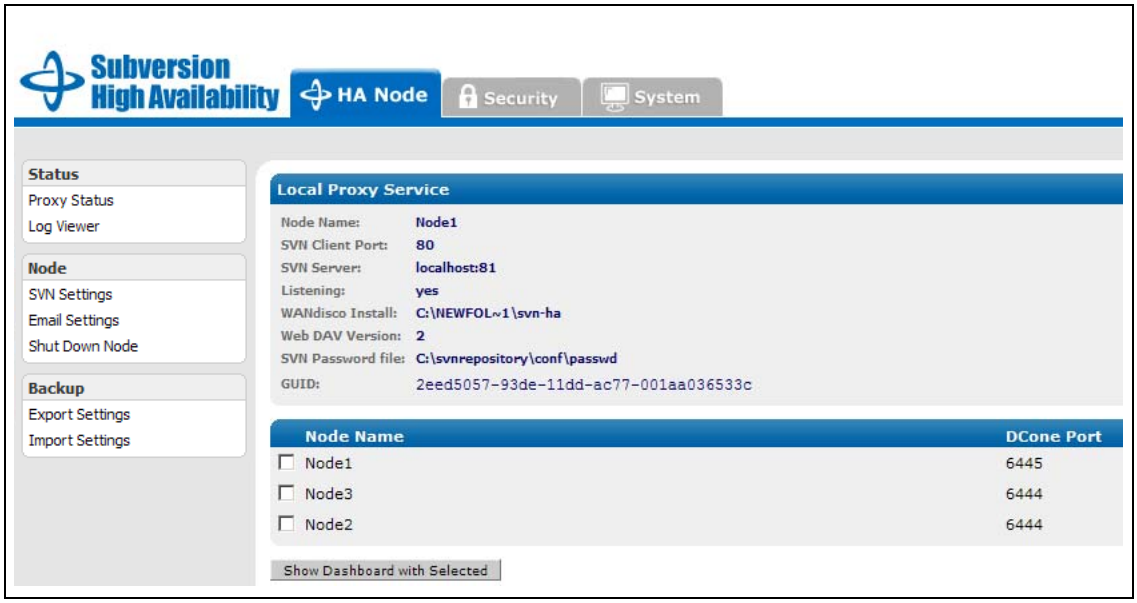

Local Proxy Service

Node Name	This node's name.
SVN Client Port	Set during installation. This node receives Subversion client communications through the Failover Agent on this port.
SVN Server	This identifies which Subversion server this node communicates with, and which port number.
Listening	This node is listening on the Subversion Client Port. Only the current active primary is listening; all other nodes are not listening to the Subversion client.
SVN Type	The type of Subversion specified upon WANdisco installation.
WANdisco install	The directory where WANdisco is installed.
Web DAV Version	This specifies the version of Web DAV.
SVN Password file	Specified upon WANdisco installation: either WANdisco manages the Subversion passwd file or
GUID	This node's GUID (Globally Unique IDentifier). Also listed in prefs.xml file.

Node Name The HA group nodes are listed.

Show Dashboard with Selected

This is another quick way to see the Dashboard. Use the checkboxes to include HA nodes on the Dashboard display.

© 2009 WANdisco

4.3.1.1 Left Side Menu

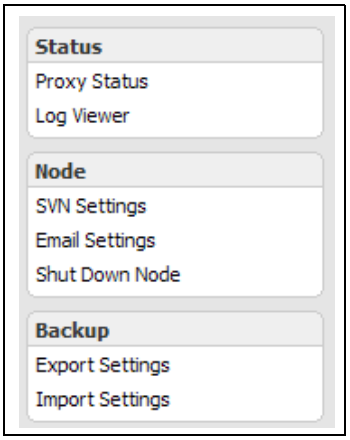

Status

Proxy Status	Described in the previous section.
Log Viewer	Select from two logs to view in the Dashboard.

Node

SVN Settings

SVN Executable	Identify the path to the Subversion executables. For remote password management for Unix.
Username	For remote password management for Unix.
Password	For remote password management for Unix.
SVNROOT (s)	For remote password management for Unix.
Default User	For remote password management for Unix.
Temp Directory	For remote password management for Unix.

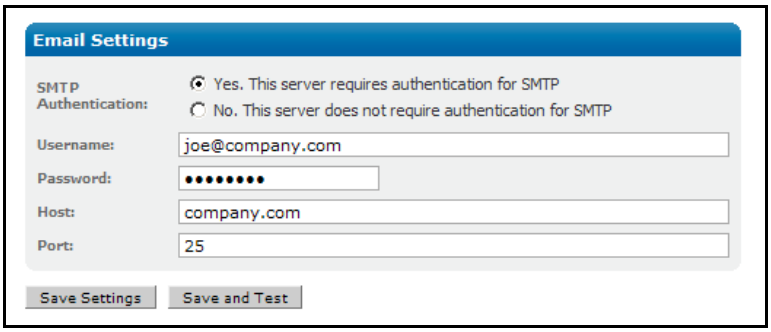

Email Settings	Set the email settings for the watchdog contact. You have to set the email in watchdog also. See 5.16, About Watchdog Mode.
SMTP Authentication	Select whether you need SMTP Authorization.
Username	For authentication, enter a valid username.
Password	For authentication: enter a valid password.
Host	Enter the email host.
Port	Enter the emal port number.
Save Settings	Save any changes.

© 2009 WANdisco

Save and Test Save any changes, and run an email test. Selecting this displays the test page.

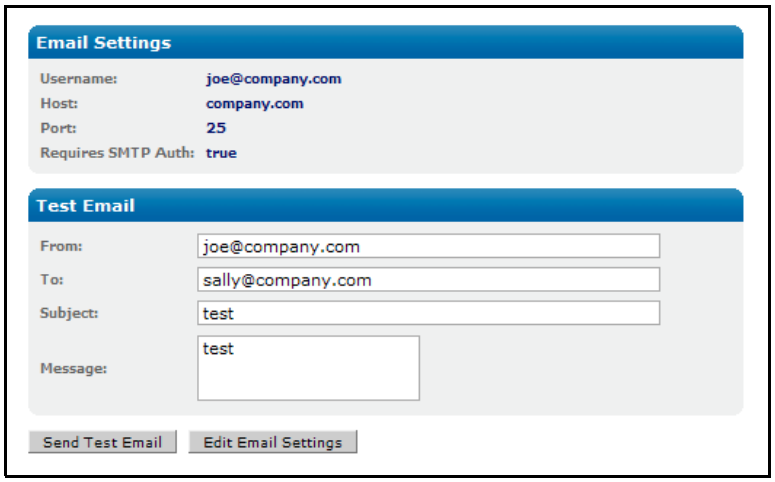

Test Email

From	Enter the sender address.
To	Enter in another address to send a test email.
Subject	Name a subject.
Message	Enter test text.
Send Test Email	Click to send the test email.
Edit Email Settings	Click to edit the email settings.

Shut Down Node — This command shuts down this node.

NOTE:

If you shut down the current node, you trigger a failover. Shutting down any other node does not result in a failover.

A High Availability group of three or more nodes has a majority response quorum. If you shut down a majority of nodes in your HA group, replication stops, and can only continue when a majority of the HA nodes are running.

Backup

Export Settings	Export WANdisco settings with this command.
Import Settings	Import WANdisco settings with this command.

© 2009 WANdisco

4.3.2 The Security Tab

The Security tab offers password control for users across the HA group.

Most users, during installation, elect to have WANdisco control the Subversion password files. You can verify this on the HA Node page. If the field **SVN password file** lists a path, WANdisco is controlling the Subversion password file. If the message `not managed by WANdisco` displays, you, the WANdisco administrator, must control and synchronize the Subversion password files at each node.

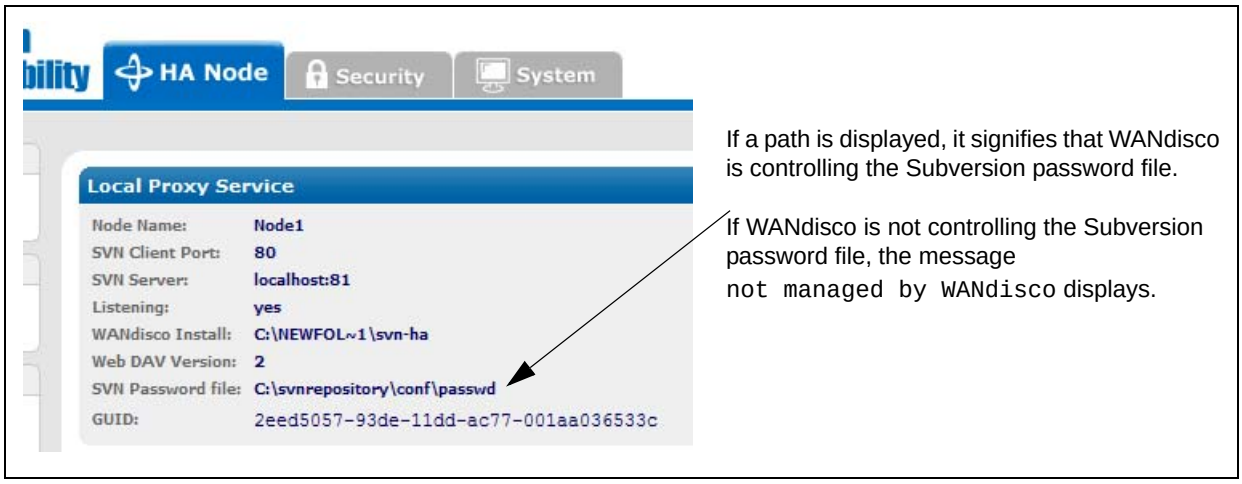

What Happens When WANdisco Controls the SVN Password File	What Happens When WANdisco Does Not Control SVN Password File
Any changes entered in the Security tab are replicated throughout the HA group.	Any changes entered in the Security tab are replicated throughout the HA group.
Changes are immediate and are reflected in the Subversion password file, at each node.	The WANdisco administrator must update each node's Subversion password file to reflect any changes in WANdisco.

© 2009 WANdisco

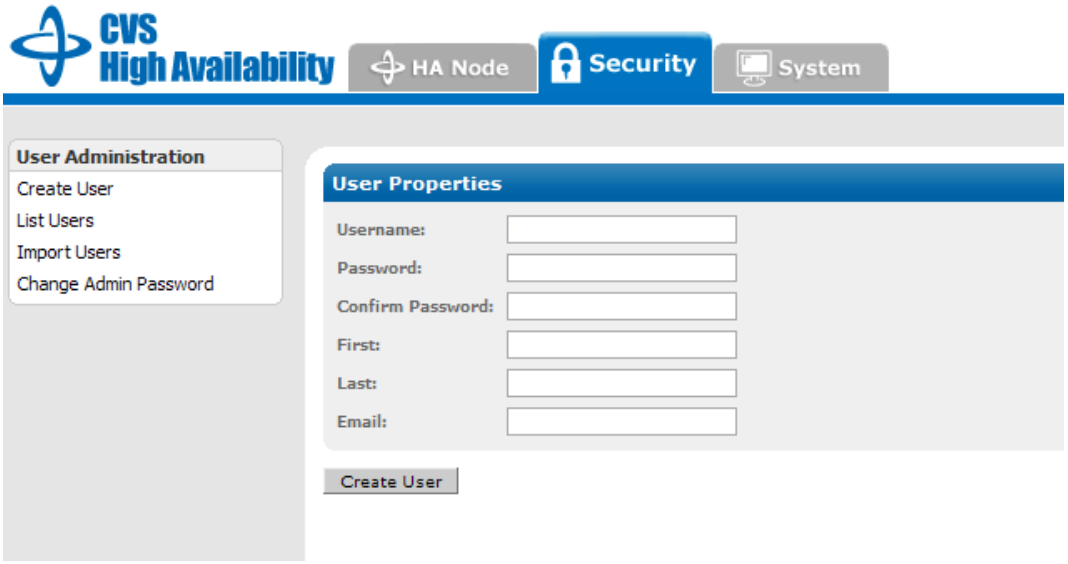

User Administration

Create User	Create Subversion users here.
Username	Enter the username.
First	Enter the user's first name.
Last	Enter the user's last name.
Email	Enter the user's email address.
Create User	Click when you have entered all information.
List Users	This command lists all the current users. You can order the users by userid, first name, last name or email, by clicking on the title in the title bar.
Import Users	You can import a text file of users, of the format `username,lastname,firstname,email`. Do not use spaces.
Change Admin Password	Change the Admin password here.

4.3.3 The System Page

Most of the commands and options on this page are described in 4.2.2, The System Page for the Failover Agent. There is one new command:

Transaction Status	You can check the status of a specific transaction number. You can select another node's transaction, in which case you are switched to that node's Admin Console.

© 2009 WANdisco

4.3.4 The Dashboard

This page offers the same commands and options as the Failover System page. See 4.2.3, The Dashboard.

By default, the Dashboard displays all nodes. If you only want certain nodes displayed, check just those nodes and click **Show Dashboard with Selected**.

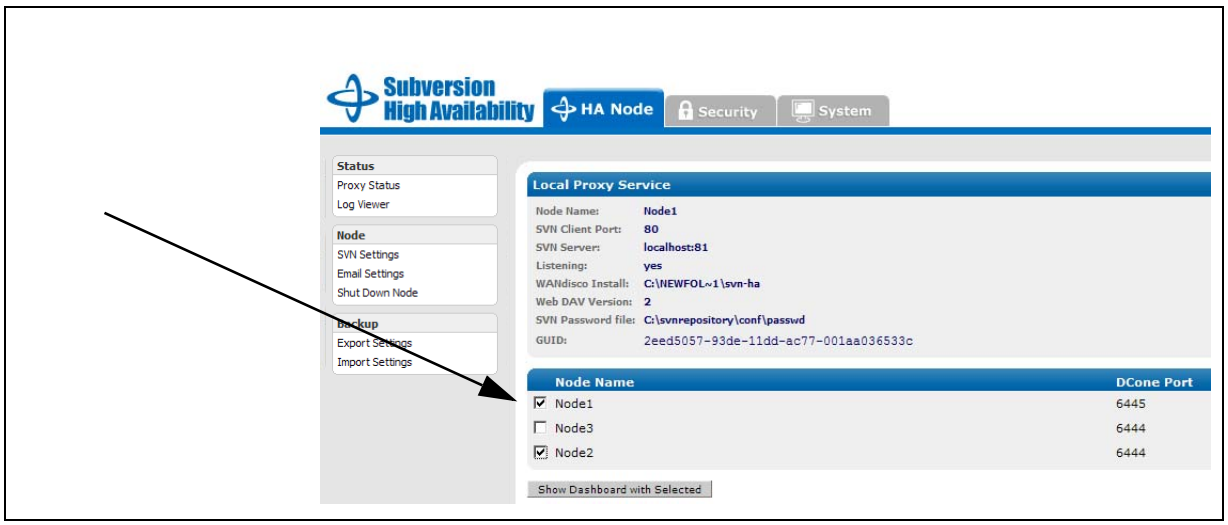

© 2009 WANdisco

5 Procedures

5.1 Preventing Subversion Users From Making Transactions

NOTE:

This procedure involves taking the High Availability offline. Please follow your company procedures about notifying Subversion users of down time.

Step 1 After notifying your users of the downtime, navigate to the Dashboard.

Step 2 Stop the Failover Agent. In the Admin Console, in the Failover Agent tab, click **Stop Failover Agent**.

Step 3 Watch for **Total Transactions Pending** to reach 0. The **Total Transactions Pending** must be 0. If you do not wait for the 0, you jeopardize the synchronization of the databases.

Step 4 To start the Failover Agent, click **Start Failover Agent**.

5.2 Establishing a Baseline for Replication

Before starting WANdisco, you should ensure that all sites start with an identical copy of the repository (the svnroots) - identical in all respects, except as noted below.

Depending on the size of your repository and available bandwidth to the remote sites, you can decide whether to copy or sync the repository over the network or ship a copy of the repository on a physical medium (for example, a CD, DVD or hard disk). Select the method that works best for your situation.

If you already have an older copy of the repository at the remote sites, for example, if, prior to deploying WANdisco, you were using a master-slave replication solution such as svnup, choose the **Synchronize** procedure.

5.2.1 Copying the Subversion Database

Otherwise, start by estimating how long it may take you to copy the repository over the network by determining the size of your repository and the bandwidth available to the remote sites. If you conclude that it takes too long, you will want to ship the repository to the remote sites on a physical medium.

© 2009 WANdisco

Step 1 Determine the size of the repository. From a Unix command prompt, `cd` to your repository and type

```
du -s
```

This reports the size of your repository in kilobytes.

Step 2 Determine the network bandwidth. Copy a reasonable-sized file (say 100 megabytes) to the remote site using any means available (example, `scp` or `ftp`). Time the copy.

Step 3 Estimate how long the copy will take. Using the information gathered above, you can estimate how long it can take you to copy the repository to the remote sites over the network. For example, if copying a 100 megabyte file over the network took 10 minutes, copying your 5 gigabyte repository may take about 500 minutes (8 hours and 20 minutes).

5.2.2 Synchronizing with an older remote copy.

You can use `rsync` to sync up an older remote copy with your master copy. For example, from the machine with the master copy of myRepository, type

```
rsync -rvlHt /path/to/myRepository remoteHost:/path/to
```

Note that the final element, `myRepository`, is not specified in the remotehost's path. For further information, consult the `rsync` man pages.

5.2.3 Copying Over the Network

Use this procedure if:

- You do not have an older remote copy; i.e., you are copying the entire repository over.
- Your repository is small enough.
- You have enough network bandwidth to copy the repository to the remote sites in reasonable time.

Step 1 Ensure that the repository is not in use. For example, type

```
/sbin/service httpd stop
```

Step 2 Package the master copy of your repository.

Step 3 Copy the package to the remote host.

© 2009 WANdisco

Step 4 Log in to your remote host and unpackage the repository. For example, on the server with the master copy, type

```
cd /path/to
tar pzcf myRepository.tgz myRepository
ssh remoteHost mkdir -p /path/to
scp myRepository.tgz remoteHost:/path/to/
log into remoteHost
cd /path/to
tar pzxf myRepository.tgz
```

Step 1 Start the SCM server. Type

```
/sbin/service httpd start
```

5.2.4 Shipping on a Physical Medium

If copying over the network may take too long, you can ship the repository to the remote destination on a physical medium, such as a CD, DVD or hard disk. Note that you do not have to wait for the baseline to be available at all sites before using WANdisco. Instead, you can follow the procedure below.

Using WANdisco before the baseline is available at all sites

Step 1 Deploy WANdisco as usual, but do not start the WANdisco server at the sites where the baseline is not yet available.

Step 2 When choosing a quorum, ensure that the sites where WANdisco can be started are sufficient to form a quorum. The simplest way to do this is to choose the Singleton Quorum policy, and choose the site that has the master copy of the repository as the distinguished node.

Safe Differences

The only things that can safely differ in the baselines across your sites are post-commit triggers. For example, if you generate email notifications from a post-commit trigger, it is a good idea to do that at only one site to avoid generating duplicate email notifications.

Common Pitfalls

It is important that the repositories are identical in all respects except as noted above. A common mistake when the desired baseline is an empty repository is to `create` a new empty repository at each site. Instead, you should `create` the repository at one site, and copy the empty repository to other sites.

© 2009 WANdisco

5.3 Finding the Last Committed Transaction

Even though committed transactions are always in the same order for each node, the timing of the commits usually varies from node to node. So unless there are no Subversion users logged in, you probably are going to have variations per node for committed transactions.

Go to any node's Dashboard. Type

```
http://<IP address>:<HA port number>/dashboard2
```

You see all the nodes on the Dashboard to compare the listed transactions.

5.4 Adding a Node to the HA Group

You can add a node to an existing HA group. Notify WANdisco of the IP address of the new node or nodes, and WANdisco sends you a new installation file and license key file.

Follow the instructions in 3.2, Installing Upgrades. The instructions are identical; you are just adding more nodes to the group.

5.5 Deleting a Node from the HA Group

You can delete a node or nodes from an existing HA group. Follow the instructions in 3.4, Installing Upgrades, but just eliminate the node from the group.

5.6 Changing a prefs.xml File

The prefs.xml files for nodes are located in `svn-ha/config`, and for the Failover Agent in `svn-failover agent/config`. Each file contains all preference information for the nodes in the group.

If you make changes that affect more than one node, you must change each node's specific file. But if your change affects just one node, you can change just that node's prefs.xml file.

© 2009 WANdisco

5.7 Performing a Synchronized Stop

NOTE:

This procedure involves taking Subversion offline. Please follow your company guidelines on notifying Subversion users of downtime.

Step 1 Bring up the Admin Console for each site.

Step 2 Stop the Failover Agent. On the Failover Agent page, click on **Stop Failover**.

Step 3 Watch the count for Pending Transactions go to 0. View all the nodes on the Dashboard.

NOTE:

You must wait for the count to reach 0 for all nodes. If you shut down a node before the count is 0, the nodes are no longer in sync and must be re synchronized.

Step 4 When the Pending Transaction count is 0 for all nodes, stop all the nodes. Go to the Failover Agent page and click on **Shut Down Node** for each node.

Step 5 If desired, shut down the Failover Agent. Click **Shut Down Failover Agent**.

Step 6 (Optional) Manually verify that the repositories are in sync. Make sure that there are no lingering lock files in the Subversion repository.

Lingering SVN locks can happen for a variety of reasons. This discussion is quite illuminative, and includes a script for fixing the problem: http://lists.gnu.org/archive/html/info-svn/2005-06/msg00211.html.

© 2009 WANdisco

5.8 Performing a Clean Restart After a Synchronized Stop

WARNING:

This procedure involves running the `reset` script. You must always reset all the nodes in the group.

Step 1 At each node's command line, run the `reset` script in `svn-replicator/bin`. (For Windows, type `perl reset`.)The command resets the system database, and the transaction count resets to 0.

Step 2 Start up all the nodes. At each node's prompt, type

```
svnreplicator
```

Step 3 Start up the Failover Agent. At the Failover Agent's prompt, type

```
failoveragent
```

5.9 Verifying That the Replicator is Working

There are two ways you can check. You can make a minor change in SVN on one client, wait a minute, and go to another client to ensure the change is reflected.

Another way to check if High Availability is replicating, is to verify there are commit transactions posted to the log file `svn-ha/logs/svnProxyServer-prefs.log`.

```
INFO:  [listen-1] Listening on port : 0.0.0.0/0.0.0.0:6445
1219077847375 org.nirala.communication.transport.DConeNet.AsyncConnector
makeConnection
INFO:  [main] Connection request to Node Id = c66b6db9-6a50-11dd-8675-
001aa036534c, host = 192.168.1.15, port = 6666, timed out in 500ms
1219077847875 org.nirala.communication.transport.DConeNet.AsyncConnector
makeConnection
INFO:  [main] Connection request to DFTPEndpoint - Node Id =
192.168.1.156666, host = 192.168.1.15, port = 6666, timed out in 500ms
1219077848578 org.nirala.admin.DiskMon start
INFO:  [main] Diskmon is monitoring C:\Thursday\svn-ha\systemdb every
15min
1219077849000 org.nirala.communication.transport.DConeNet.ListenReactor
setupListener
INFO:  [listen-1] Listening on port : 0.0.0.0/0.0.0.0:2403
1219077849000 org.nirala.communication.transport.svnproxy.ProxyServer
onStartedProxyListen
INFO:  [main] SVN Proxy listener is now turned ON at port :2403
```


© 2009 WANdisco

```
1219077853765 org.nirala.communication.transport.DConeNet.Listen-
Stage$TCPStopListening onStop
INFO:  [listen-1] Host: 0.0.0.0, Port: 2403 Stopped Listening.
1219077853765 org.nirala.communication.transport.svnproxy.ProxyServer
onStopProxyListener
INFO:  [p-queue-1] SVN Proxy listener is now turned OFF at port :2403
1219077872328 org.nirala.communication.transport.DConeNet.ListenReactor
setupListener
INFO:  [listen-1] Listening on port : 0.0.0.0/0.0.0.0:2403
1219077872328 org.nirala.communication.transport.svnproxy.ProxyServer
onStartedProxyListen
INFO:  [mqueue-1] SVN Proxy listener is now turned ON at port :2403
```

5.10 Installing a .jar File Patch

Follow these steps to install a `svn-replicator.jar` patch to an existing High Availability installation. You are going to copy the same jar file to the Failover Agent's and each node's `lib` directory.

NOTE:

Always read the `readme` file and release notes first.

Step 1 Download the `svn-replicator.jar` file.

Step 2 Verify the md5 checksum.

Step 3 At the Failover Agent and all the nodes, move the existing jar file to a back up directory. (All jar files in the `/lib` directories are in the WANdisco `CLASSPATH`.)

Step 4 Perform a synchronized stop. See 5.7, Performing a Synchronized Stop.

Step 5 Copy the new jar file to each node's `lib` directory.

Step 6 Restart the group. See 5.8, Performing a Clean Restart After a Synchronized Stop.

Step 7 Confirm the upgrade by checking the dashboard for the newer version, and check the log file under `svn-ha/logs` for the start header with the new version.

© 2009 WANdisco

5.11 Setting Replicator to Start Up on System Boot

You can easily have the replicator start up on system boot. To start up on boot, edit the `init.d` scripts. You must make modifications based on your operating system. Make this change at the Failover Agent and at each node, and edit the script accordingly.

For instance, here is an `/etc/init.d/svnreplicator` script for Gentoo Linux.

```
#!/sbin/runscript
#
# Gentoo Linux dist compatible rc script for
# starting/stoppping svnreplicator
#
# Copyright WANdisco
#

REP_HOME="/home/admin0/svn-replicator"
REP_OPTS="-wdog -email admins@example.com"
export JAVA_HOME="/export/share/apps/jdks/1.5.0"
USER="admin0"

pidfile="my.pid"

depend() {
need net
}


checkconfig() {
if [ ! -f $REP_HOME/bin/svnreplicator ]; then
eerror "No ${REP_HOME}/bin/svnreplicator present"
return 1
fi
prog=$"svnreplicator"
}

start() {
checkconfig || return 1
ebegin "Starting $prog:"
ulimit -S -c 0 >/dev/null 2>&1
ulimit -n 65000 >/dev/null 2>&1
RETVAL=0
start-stop-daemon --start --quiet -u ${USER} --chuid ${USER} --exec
${REP_HOME}/bin/svnreplicator -- ${REP_OPTS}
RETVAL=$?

if [ "$RETVAL" -gt 0 ]; then
eend $RETVAL "Failed to bring up svnreplicator"
return $RETVAL
fi
eend $RETVAL
}
```


© 2009 WANdisco

```
stop() {
checkconfig || return 1
ebegin "Shutting down $prog:"
su ${USER} -c \"${REP_HOME}/bin/shutdown\" >/dev/null 2>&1
start-stop-daemon --stop --quiet -u ${USER} --pidfile ${REP_HOME}/logs/
${pidfile}
RETVAL=$?
if [ "$RETVAL" -gt 0 ]; then
eend $RETVAL "Failed to shutdown svnreplicator"
return $RETVAL
fi
eend $RETVAL
}
```

5.12 Setting the Replicator Up as a Windows Service

To set the replicator to run as a Windows service, perform the following command at the command prompt:

```
sc create SVN-Replicator binpath= C:\perl\bin\perl.exe –x –S C:\svn-rep-
licator\bin\svnreplicator start= auto
```

Substitute the path for Perl in your environment, and give a different path to the SVN replicator perl script, depending on where that was installed. You may want to also set `type= share`. The MicroSoft knowledge base article (http://support.microsoft.com/kb/251192) indicates that that is the default, but the `sc.exe` help for create indicates that `type= own` is the default. Note that there is a space between the equals sign, `=`, and the parameter's value.

The Services Control Panel indicates that the service has not started, because our Perl script is currently not exiting because the watchdog is running to restart the replicator. This is actually fine, because the Perl script really takes over.

5.13 Verifying System Integrity

WANdisco recommends performing periodic incremental diffs to ensure all the remote SVN repositories are in sync. This could be done periodically using a cron job, for example. But do remember that because of the way WANdisco replication works, it is possible for two SVN repositories to appear to be different while updates are happening. In other words, if you take a snapshot while updates are being applied, the snapshots could differ from each other. This does not indicate a failure of WANdisco replication.

So, when making snapshots for comparing, please make sure repositories are not being updated. The easiest way is to:

© 2009 WANdisco

- disable client access to the replicator using Admin Console
- wait briefly for the replication group to quiesce
- Do a local SVN checkout without the RCS keywords expansion, using the -kk option. This will help avoid spurious conflicts.
- Bring the snapshot to the machine where you would be doing the diff, perhaps using rsync.

5.14 Changing SVN Port on Unix Flavor

You can change the port in the `httpd.conf` configuration file in the Apache server. Please see the `Listen` directive, discussed in the article at http://httpd.apache.org/docs/2.2/bind.html.

Here is a snippet of an `httpd.conf` file:

```
#
# Listen: Allows you to bind Apache to specific IP addresses and/or
# ports, instead of the default. See also the <VirtualHost>
# directive.
#
# Change this to Listen on specific IP addresses as shown below to
# prevent Apache from glomming onto all bound IP addresses (0.0.0.0)
#
#Listen 12.34.56.78:80
Listen 8080
```

With this configuration, Apache server listens on port 8080 instead of default port 80.

5.15 Using SVN Triggers for Sending E-mails

Many administrators like to set up SVN backend triggers that fire whenever a SVN user commits a set of file changes. With a single/master SVN server setup, e-mails can be initiated once when the `post-commit` trigger fires.

However, with the addition of WANdisco replicator, unless some safeguards are put in place, all your SVN replicas may fire the `post-commit` trigger. This could potentially cause multiple e-mail notifications. Most likely, developers do not want several e-mails for the same transaction.

The easiest way to remedy this is to designate any one node as the “e-mail hub.” Just enable the `post-commit` trigger to fire from a single site within the replication group. Alternatively, you could use the time of day to fire the e-mail alerts from a specific site. For example, you could modify the `post-commit` trigger to send e-mails from India during 9:00 a.m. to 5:00 p.m. IST, and from the US during 9:00 a.m. to 5:00 p.m. PST.

© 2009 WANdisco

It is allowable to have asymmetry in the e-mail triggers, but make sure not to disable the `pre-commit` trigger on any node. That may cause a SVN commit transaction to abort at some sites but commit at other sites. The `pre-commit` trigger behavior at each site should be deterministic and should not cause the replicas to go out of sync.

When sending e-mail, it is important to set up the e-mail configuration to avoid long blockages or delays. Many times, an administrator uses the default SMTP settings on the SVN host. These settings by default try to use the organization domain specific e-mail server to send e-mails (by looking up the MX records corresponding to the organization's domain).

The organization-wide SMTP server may be located on a remote WAN, or it may have throttling policies for e-mails originating from the same IP address to cut down on spam. This can cause it to block or reject e-mails, which may in turn cause scripts (like the `post-commit` script) to hang or terminate. To avoid such problems with e-mail triggers, WANdisco recommends that you set up a local e-mail hub or a local SMTP agent/server. The local SMTP server should preferably be on the same host as the SVN server. It should be set up to forward/relay e-mails to the organization-wide SMTP server. This ensures the e-mail triggers are a lot faster and just need to enqueue the e-mails to the local SMTP server.

5.16 About Watchdog Mode

By default, WANdisco starts in watchdog mode. Whenever the replicator goes down, the watchdog mode restarts it. In watchdog mode, the replication process automatically disassociates from the terminal and becomes a daemon process, so you should not try running it in the background (with &).

NOTE:

> Watchdog mode is not supported in Windows, but it is in Windows Cygwin.

You can turn off watchdog by typing `-nowdog`.

If WANdisco is unable to start up, for example if it terminates several times in quick secession, watchdog starts WANdisco in read-only mode.

```
$ ./bin/svnreplicator -h
Usage: svnreplicator [-v] [-verbose] [-nowdog] [-pause time]
[-email email-address]
```

-v	Print the svnreplicator version
-verbose	Verbose, console messages go to STDOUT/STDERR instead of logs/console.txt
-nowdog	Turn off watchdog mode. WANdisco will not restart automatically if it terminates. Use this option for testing.
-pause	Time in seconds that the watchdog pauses for, before restarting service. Defaults to 0 seconds.

© 2009 WANdisco

-email	Specify an email address to send an alert to, whenever the Watchdog restarts or shuts down WANdisco. WANdisco generates an email per local replicator activity. Set up the email account for each site with the Email Settings, described in Chapter 4, Using the Admin Console.

Use the `-email` option to generate email alerts whenever WANdisco restarts. For instance:

```
$ svn-replicator/bin/svnreplicator -pause 5 -email “admin@blueand-
gold.com, scmuser@blueandgold.com”
```

In order to have WANdisco Subversion Replicator automatically started on system reboots, see 5.11, Setting Replicator to Start Up on System Boot.

© 2009 WANdisco

6 Procedures for Two-Node HA Groups

6.1 Recovering from Primary Node Failure

When the primary node fails, the Failover Agent sends transactions to the second node. This sets a failover flag. Use the onscreen wizard to clear the flag. WANdisco ensures the two repositories are in sync.

6.2 Recovering from Backup Node Failure

When the backup node fails, there are a few more steps to complete.

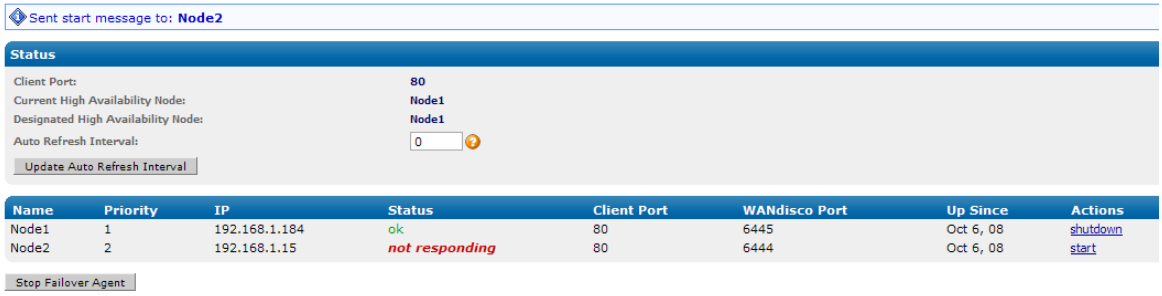

Upon an SVN commit, you see this error message about going into Unilateral mode.

Click **Click here to restore the High Availability Group**.

© 2009 WANdisco

Step 1 Click **Stop Failover Agent**.

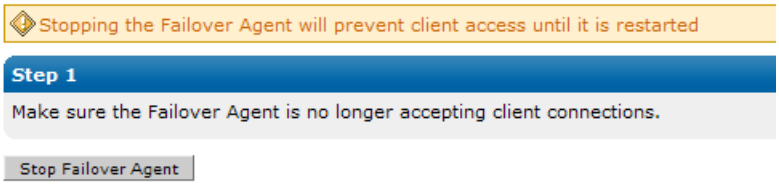

Step 2 Make sure there are no pending transactions on the dashboard, and reset both nodes. Reset the nodes by typing in the `/bin` directory

```
reset
```

On Windows, type

```
perl reset
```

When you have done this, click **Yes**.

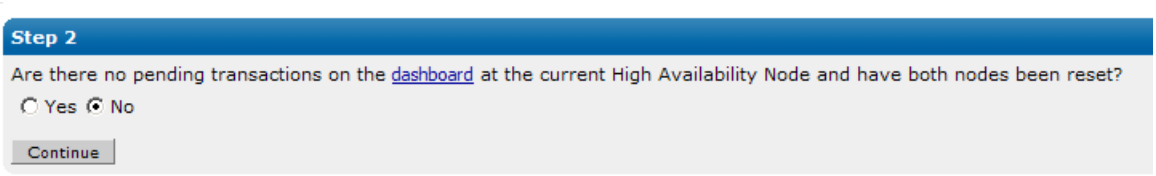

Step 3 Make sure the repositories are in sync. Click **Continue**.

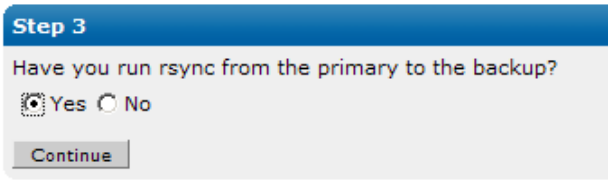

Step 4 Start the nodes. In the /bin directory, type

```
svnreplicator
```

Start the Failover Agent. Type

```
failoveragent
```


© 2009 WANdisco

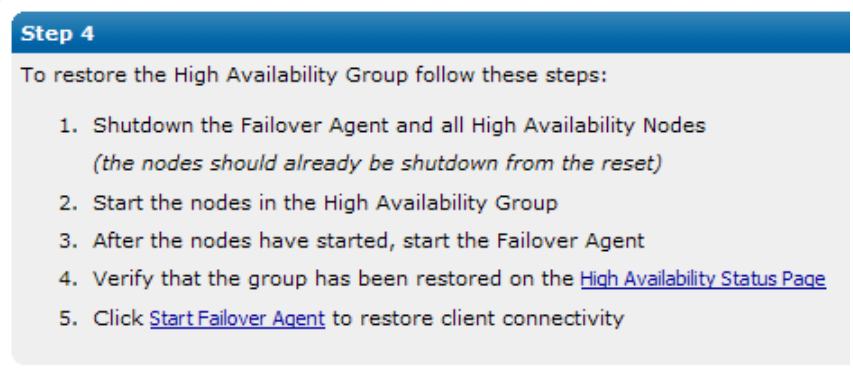

Your HA group is restored.

© 2009 WANdisco

7 Troubleshooting

7.1 How Do I Get WANdisco Support?

Before opening a ticket or submitting a new issue, always search the Knowledge Base on *http://www.wandisco.com/php/support_login.php*.

If you want to open a ticket, you can do so at that URL.

7.1.1 How Do I Run the Talkback Script?

When you do contact WANdisco with a problem, the first thing WANdisco support asks for is the talkback file. Run the file by typing, for Unix

```
svn-replicator/bin/talkback
```

For Windows, type

```
svn-replicator\bin\perl talkback
```

Type in the pathname to SVNROOT when prompted. The output looks like this:

```
Please open a ticket by visiting http://support.wandisco.com and upload
the /talkback-<machine name>.zip, with a description of the issue.

Note: do not email the talkback files, only attach them via the web
ticket user interface.
```

The zip file is located at the root directory. Do not email the .zip file, just attach it to an issue at *http://www.support.wandisco.com*.

© 2009 WANdisco

7.2 General Subversion High Availability

7.2.1 Connection Request Timeout Messages

Sometimes in the WANdisco logs, you see connection request timeout information messages logged. These are informational messages and should be ignored unless it is guaranteed that the connection can be established in xxx milli-seconds and happens often.

In normal operation of WANdisco, two connections are established between each of the replicated machines, a DConE connection and a DFTP connection. These two connections were established when High Availability started and are used when required. A keep-alive signal is sent on the DConE port periodically. There is no traffic on DFTP until a file transfer.

It takes between 300 to 400 milli-seconds to establish a network connection even on a slow Wide Area Network (WAN). By default, High Availability waits for 500 milli-seconds before giving up that a connection cannot be established to a peer machine and prints this informational message. What if the establishments of connection always take 501 milli-seconds. In this case, a connection is never established. To solve this problem, the timeout value is adjusted in 10% increments of the last timeout, starting at 500 milli-seconds, to a maximum of 10 seconds for each timed out connection. Upon establishment of a successful connection, this timeout value is used for subsequent connection establishment unless an adjustment is required for failed attempts.Replication Processes

7.3 Error Messages

7.3.1 Missing License Key File

Subversion High Availability depends on a license key file being present in the `svn-failover/config` and `svn-ha/config` directories for each node. Please get a valid license from WANdisco and copy the file to the config directory. The Replicator does not start without the license file.

7.3.2 I’m Getting a SEVERE Exception

I’m getting a SEVERE exception, and replicator is aborting the SVN transaction and shutting down.

If you get a message in the logs/svnproxy*.log file similar to

© 2009 WANdisco

```
SEVERE: [reader-1] Encountered a system error from svn server : E svn
[commit aborted]: could not find desired version 1.16 in /development/
software/svnroot/libcobra/Build,v
```

it means the replicator has detected an out of sync condition. Remember the replicator continuously monitors your repository for any out of sync issues. If it detects this has occurred, it triggers an automatic shutdown to prevent further corruption.

This could happen if some one accidently committed directly to SVN, bypassing the replicator, and ramped up the version in one site without giving the replicator any chance of replicating. This can be easily resolved by following the reset procedure (See 7.4.1, I Directly Committed to SVN, How Do I Rsync?).

Follow all precaution to avoid bypassing the replicator:

Step 1 Ensure only svnreplicator host/IP address is allowed to connect to svnpserver.

Step 2 Protect direct logins to svn server box from end user.

Step 3 When administering, ensure a valid `SVNROOT` that points to the replicator.

7.4 Oops!

7.4.1 I Directly Committed to SVN, How Do I Rsync?

If you bypassed the replicator, you can reset the replicator state with these steps:

Step 1 Shut down all replicators.

Step 2 Reset each replicator. For Unix, type:

```
$ svn-ha/bin/reset
```

For Windows, type

```
svn-ha\bin\perl reset
```

Step 3 If this happened during an initial setup/evaluation stage, delete the old project in SVN and create a new one.

If this happened on a production repository, just re-sync all the repositories to the same state/data.

© 2009 WANdisco

Step 4 Restart all the replicators.

NOTE:

It is very important that you take all precautions to avoid directly checking in or committing to the backend SVN repository.

7.4.2 I Pressed Ctrl-C During an SVN Command!

If you were executing a read command (a command that does not modify the SVN repository), you do not have to do anything.

If you were executing a write command, update your sandbox after the replicator has applied the command to the repository.

In addition, if you were adding files to the repository (either svn import or svn add, followed by svn commit), wait until you update your sandbox before you continue to use it.

© 2009 WANdisco

8 Frequently Asked Questions

8.1 Why Are So Many Java Processes Running?

On older versions of Linux, every thread is listed as a process by the `ps` command. This does not affect the operation of High Availability. WANdisco does not support the older versions of Linux.

8.2 How Do I Deal with Failover Agent Failure?

If the failover agent fails, the watchdog script immediately restarts it. If the machine crashes, service is unavailable until the machine is rebooted and the failover agent is restarted.

You could also run the failover agent on a hardware cluster. The Veritas Cluster Server is an example of a commercial solution. See http://www.symantec.com/business/products/overview.jsp?pcid=pcat_business_cont&pvid=20_1.

Linux-HA is an example of an open-source solution. See http://www.linux-ha.org/.

8.3 Can I Store Logs or Content on NFS?

NFS (Network File System) allows files and directories to be accessed remotely over a network using NFS clients. NFS clients are typically built into the operation system kernel these days. However, some operations, like renaming a file, are not guaranteed to be atomic over NFS. Here is a snippet from the rename function's `man` page on Linux, for example:

```
BUGS
On NFS filesystems, you can not assume that if the operation failed the
file was not renamed. If the server does the rename operation and then
crashes, the retransmitted RPC which will be processed when the server is
up again causes a failure. The application is expected to deal with this.
See link(2) for a similar problem.
```

Code management systems such as SVN make heavy use of the `rename` operation to modify the underlying databases. Independent of WANdisco, it is a risky practice to store SVN database content on NFS. The code management community at large recommends not using NFS for storing repositories.

WANdisco High Availability is bundled with a built-in transactional journal and an object database. These are by default stored in the `svn-ha/systemdb` and `svn-ha/config` directories. These directories should not be mounted on an NFS drive. The replicator itself may be installed on an NFS drive but the systemdb and config directories should be on direct storage (non-NFS

© 2009 WANdisco

options like RAID, SCSI, SAN, etc). Replicator's transactional integrity can be compromised if writes to an NFS server are lost due to a potential NFS client cache crash after the NFS server has indicated IO completion.

8.4 Why is Set Up Configuring IP Addresses as 0.0.0.0?

The address 0.0.0.0 is a special IP address, treated as a wild-card IP address. In other words, on a machine with multiple NICs (Network Interface Cards), it binds to all interfaces. The advantages of using wild-card IP address include:

- It avoids binding to a fixed IP address. If the host's IP address changes, (for example, the subnet changes, or the machine is moved to a different location) you don't have to change the wild-card IP in the prefs.xml file to the new IP address.
- There is wider bandwidth to TCP clients. Now TCP clients can connect to any NIC, because High Availability is listening on multiple NICs.

The disadvantage to using the wild-card IP is that it gives coarser access control at the IP address level, as all address are being listened to at the specified port.

You can always switch from the wildcard IP address to a fixed, static IP address or a DNS hostname, though for the most part, WANdisco recommends you stick with wild-card addressing.

8.5 Should I Worry About Time Changes?

Time changes have no effect on the operation of High Availability.

8.6 Does WANdisco Support Dynamic DNS?

Yes, WANdisco supports dynamic DNS, but strongly discourages its use.

If a hostname is specified during the setup process, WANdisco requires that it should be able to connect to a valid DNS and resolve the hostname to valid IP address upon startup. If the hostname cannot be resolved to an IP (either by not being able to connect to DNS, or no entry is found at the given hostname), WANdisco dies gracefully. This has never been a problem during production and with static IPs.

However, if dynamic DNS support is required, please modify the `prefs.xml` file at each site and set `UseDynamicDNS` to `true` in `DConeNet` element.

```
<Preferences>
  ...
  <DConeNet>
```


© 2009 WANdisco

```
    ...
    <UseDynamicDNS>true</UseDynamicDNS>
</DConeNet>
```

In addition, the following Java security properties should be set to different Time-to-live (TTL).

```
networkaddress.cache.ttl
networkaddress.cache.negative.ttl
```

Please read InetAddress Caching for more details.

8.7 WANdisco Authentication

Authentication is the process of determining whether someone or something is, in fact, who or what it is declared to be. Authorization is the process of giving someone permission to do or have something.

The Apache user-names and passwords should match at all sites. The WANdisco Subversion replicator's license manager requires a valid user-name inside the HTTP authorization header to be passed for all DAV commands, except `OPTIONS` and `PROPFIND`. In other words, anonymous access to Apache is not allowed to enforce license requirements, unless you have an unlimited or an evaluation license. With an unlimited or evaluation license, you are not required to register the user. This typically means ensuring a `Require` valid-user line is specified in the Apache SVN DAV configuration files in the `/etc/httpd/conf/httpd.conf` and `/etc/httpd/conf/conf.d/*` directories. When using Basic Authentication, it is the end user or administrator's responsibility to keep Apache authentication databases in sync across all sites.

8.8 Apache 2.2 with SVNDAV on Windows

Step 1 Install Apache2.2.

Step 2 Install svn-win32-1.4.4 for Apache 2.2. Make sure it's SVN for Apache 2.2.

Step 3 Copy `svn-win32-1.4.4/bin/intl3_svn.dll` to `apache/bin`.

Step 4 Copy `svn-win32-1.4.4/bin/libdb44.dll` to `apache/bin`.

Step 5 Copy `svn-win32-1.4.4/mod_authz_svn.so` to `apache/modules`.

Step 6 Copy `svn-win32-1.4.4/mod_dav_svn.so` to `apache/modules`.

Step 7 Uncomment these lines in `apache/conf/httpd.conf`:

© 2009 WANdisco

```
LoadModule dav_module modules/mod_dav.so
LoadModule dav_fs_module modules/mod_dav_fs.so
```

Step 8 Add these lines to `apache/conf/httpd.conf`:

```
LoadModule dav_svn_module modules/mod_dav_svn.so
LoadModule authz_svn_module modules/mod_authz_svn.so

<Location /myDavLocation>
DAV svn
SVNPath C:\repo
SVNAutoversioning on
AuthType Basic
AuthName "SVN Repo"
AuthUserFile C:\repo\dav-auth
Require valid-user
</Location>
```

Step 9 Check that the users have been added to the `C:\repo\dav-auth` file. To add new users or change passwords, use `apache/bin/htpasswd.exe`.

Step 10 Restart Apache.

Step 11 Point a web browser to: *http://server:port/myDavLocation*.

8.9 Setting Up Apache for SVN-DAV

Sometimes, it is required that Apache server be running on port 80 and server multiple locations in addition to SubVersion. This can be accomplished with proxy enabled.

This section outlines the steps to set up Apache for such configuration and also for setting up SVN-DAV for HTTPS access. The reader should be familiar with the Apache set up and basic WANdisco replicator set up. The assumption is that Apache is running on UNIX, though the same steps apply for the Windows platform. These points are assumed:

- Apache server is running on port 80
- Apache web-dav module is running on port 8181
- WANdisco SVN replicator is configured to listen on port 8080, and it forwards the requests to apache web-dav module on port 8181
- WANdisco DConE port is listening on port 6444
- Apache SSL is running on port 443 to handle HTTPS requests (if this is set up, the stunnel package is not required)
- All the processes are running on the same machine
- Apache is compiled with `mod-proxy` and `mod-ssl` modules
- The SVN URL is */svnrepos* (either as a parent-path or path)

Run the `<svn-replicator>/bin/setup` utility and specify the following ports for each replicator:

© 2009 WANdisco

- DConE port 6444
- Proxy port 8080
- Apache web-dav port 8181 on localhost

In the httpd.conf file, specify the following parameters:

```
    #

# Define apache port and pass anything that matches location /svnrepos to
WANdisco SVN Replicator

#

    NameVirtualHost *:80
<VirtualHost *:80>
    ProxyPass /svnrepos http://127.0.0.1:8080/svnrepos
    ProxyPass /!svn http://127.0.0.1:8080/svnrepos/!svn
    ProxyPassReverse  /svnrepos http://127.0.0.1:8080/svnrepos
    ProxyPassReverse  /!svn http://127.0.0.1:8080/svnrepos/!svn
    RequestHeader edit Destination ^https: http: early
</VirtualHost>

  Listen 443

<VirtualHost *:443>
    ProxyPass /svnrepos http://127.0.0.1:8080/svnrepos
    ProxyPass /!svn http://127.0.0.1:8080/svnrepos/!svn
    ProxyPassReverse  /svnrepos http://127.0.0.1:8080/svnrepos
    ProxyPassReverse  /!svn http://127.0.0.1:8080/svnrepos/!svn
    RequestHeader edit Destination ^https: http: early
    SSLEngine on
    SSLCertificateFile /etc/httpd/conf/ssl.crt/server.crt
    SSLCertificateKeyFile /etc/httpd/conf/ssl.key/server.key
    SSLCACertificateFile /etc/httpd/conf/ssl.crt/ca-bundle.crt
</VirtualHost>

Listen 8181
    NameVirtualHost *:8181
<VirtualHost *:8181>
    KeepAlive On
<Location /svnrepos>
    AllowOverride None
    Order allow,deny
    Allow from 127.0.0.1
    DAV svn
    SVNParentPath /tmp/dav
    AuthType Basic
    AuthName wandisco
    AuthUserFile /etc/httpd/conf/htpasswd
    Require valid-user
```


© 2009 WANdisco

```
</Location>
</VirtualHost>
```

8.10 Encryption Around WANdisco Protocol

Here are the details about any ECCN classifications you may have applied for and been granted from the U.S. Government for export (due to the encryption capabilities in the client for DAV over SSL).

WANdisco Subversion MultiSite distribution does not actually perform any encryption or decryption of the DAV traffic. We rely on Apache to decrypt the SSL traffic, and then use a proxypass definition within the Apache configuration to redirect the unencrypted request to the WANdisco replicator.

The WANdisco replicators do not directly encrypt communication between sites. Instead, many customers may use something like a persistent VPN connection for the replicator-to-replicator traffic over an encrypted connection, but our code actually does no encryption.

Lastly, the WANdisco replicator simply sits as a proxy on the SVN server itself (the host running Apache and ModDAV), so there is no client component that we provide that would be sending any traffic to the SVN server.

8.11 How Do I Restrict Direct Access to My Repository?

If you would like to prevent users from directly accessing your Subversion repository, use the Subversion `Location` directive as suggested below. You allow only specific IP addresses to access the repository.

This assumes that WANdisco and Apache server are running on the same machine.

From the example shown in 8.9, Setting Up Apache for SVN-DAV:

```
<Location /svnrepos>
AllowOverride None
Order allow,deny
Allow from 127.0.0.1
DAV svn
SVNParentPath /tmp/dav
AuthType Basic
AuthName wandisco
AuthUserFile /etc/httpd/conf/htpasswd
Require valid-user
</Location>
```


© 2009 WANdisco

8.12 About WANdisco Log Files

WANdisco uses Java logging. See http://java.sun.com/j2se/1.4.2/docs/guide/util/logging/overview.html for a discussion on Java logging. Make any changes to the `svn-replicator/config/log.properties` file.

WANdisco places the log files in `svn-replicator/logs`. The current file is always `SVNProxyServer-prefs.log.0`, and the files are rotated out and eventually garbage collected. For rotation schedule, see the `svn-replicator/config/log.properties` file.

© 2009 WANdisco

Appendix A - Installing Java and Perl

You should have already installed Java and Perl at all the sites in your replication group for your trial evaluation. However, any new site you add to the replication group needs Java and Perl installed as well.

Installing Java

Step 1 Install JDK 1.5 and define the `JAVA_HOME` environment variable to point to the directory where the JDK is installed. You can download JDK 1.5 from the URL below.

```
http://java.sun.com/javase/downloads/index_jdk5.jsp
```

Step 2 Add `$JAVA_HOME/bin` to the path and ensure that no other java (JDK or JRE) is on the path.

```
$ which java
/usr/bin/java

$export JAVA_HOME="/usr"
```

or

```
$which java
/export/share/apps/jdk/1.5.0/bin/java

$export JAVA_HOME="/export/share/apps/jdk/1.5.0"
```

Step 3 Ensure the full JDK is installed, not just the JRE. This can be confirmed by running `java -server-version`. If it generates a **not found** error, repeat Steps 1 and 2.

If you find package management problems or conflicts with the JDK version you are downloading (for example, rpm download for Linux), you may want to use the self-extracting download file instead of the rpm (on Linux) package. The self-extracting download easily installs in any directory without any dependency checks.

© 2009 WANdisco

Installing Perl

Step 1 On UNIX or Cygwin, install perl version 5.6 or greater and ensure that the perl executable is on the system path.

Step 2 On Windows, install ActivePerl version 5.8 or greater and ensure that the perl executable is on the system path. You can download the MSI installer for ActivePerl from the URL below.

`http://activestate.com/Products/Download/Download.plex?id=ActivePerl`.

© 2009 WANdisco