



Avid[®] Interplay[®] Engine Failover Guide

Legal Notices

Product specifications are subject to change without notice and do not represent a commitment on the part of Avid Technology, Inc.

This product is subject to the terms and conditions of a software license agreement provided with the software. The product may only be used in accordance with the license agreement.

Avid products or portions thereof are protected by one or more of the following United States Patents: 5,267,351; 5,309,528; 5,355,450; 5,396,594; 5,440,348; 5,467,288; 5,513,375; 5,528,310; 5,557,423; 5,577,190; 5,584,006; 5,640,601; 5,644,364; 5,654,737; 5,724,605; 5,726,717; 5,745,637; 5,752,029; 5,754,851; 5,799,150; 5,812,216; 5,828,678; 5,842,014; 5,852,435; 5,986,584; 5,999,406; 6,038,573; 6,069,668; 6,141,007; 6,211,869; 6,532,043; 6,546,190; 6,596,031; 6,747,705; 6,763,523; 6,766,357; 6,847,373; 7,081,900; 7,403,561; 7,433,519; 7,671,871; 7,684,096; D352,278; D372,478; D373,778; D392,267; D392,268; D392,269; D395,291; D396,853; D398,912. Other patents are pending.

Avid products or portions thereof are protected by one or more of the following European Patents: 0506870; 0635188; 0674414; 0752174; 1111910; 1629675. Other patents are pending.

This document is protected under copyright law. An authorized licensee of [product name] may reproduce this publication for the licensee's own use in learning how to use the software. This document may not be reproduced or distributed, in whole or in part, for commercial purposes, such as selling copies of this document or providing support or educational services to others. This document is supplied as a guide for [product name]. Reasonable care has been taken in preparing the information it contains. However, this document may contain omissions, technical inaccuracies, or typographical errors. Avid Technology, Inc. does not accept responsibility of any kind for customers' losses due to the use of this document. Product specifications are subject to change without notice.

Copyright © 2010 Avid Technology, Inc. and its licensors. All rights reserved.

The following disclaimer is required by Sam Leffler and Silicon Graphics, Inc. for the use of their TIFF library:

Copyright © 1988–1997 Sam Leffler
Copyright © 1991–1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software [i.e., the TIFF library] and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the names of Sam Leffler and Silicon Graphics may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

The following disclaimer is required by the Independent JPEG Group:

This software is based in part on the work of the Independent JPEG Group.

This Software may contain components licensed under the following conditions:

Copyright (c) 1989 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Copyright (C) 1989, 1991 by Jef Poskanzer.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. This software is provided "as is" without express or implied warranty.

Copyright 1995, Trinity College Computing Center. Written by David Chappell.

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. This software is provided "as is" without express or implied warranty.

Copyright 1996 Daniel Dardailler.

Permission to use, copy, modify, distribute, and sell this software for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Daniel Dardailler not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Daniel Dardailler makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Modifications Copyright 1999 Matt Koss, under the same license as above.

Copyright (c) 1991 by AT&T.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

THIS SOFTWARE IS BEING PROVIDED "AS IS", WITHOUT ANY EXPRESS OR IMPLIED WARRANTY. IN PARTICULAR, NEITHER THE AUTHOR NOR AT&T MAKES ANY REPRESENTATION OR WARRANTY OF ANY KIND CONCERNING THE MERCHANTABILITY OF THIS SOFTWARE OR ITS FITNESS FOR ANY PARTICULAR PURPOSE.

This product includes software developed by the University of California, Berkeley and its contributors.

The following disclaimer is required by Nexidia Inc.:

© 2006 Nexidia. All rights reserved.

Manufactured under license from the Georgia Tech Research Corporation, U.S.A. Patent Pending.

The following disclaimer is required by Paradigm Matrix:

Portions of this software licensed from Paradigm Matrix.

The following disclaimer is required by Ray Sauers Associates, Inc.:

"Install-It" is licensed from Ray Sauers Associates, Inc. End-User is prohibited from taking any action to derive a source code equivalent of "Install-It," including by reverse assembly or reverse compilation, Ray Sauers Associates, Inc. shall in no event be liable for any damages resulting from reseller's failure to perform reseller's obligation; or any damages arising from use or operation of reseller's products or the software; or any other damages, including but not limited to, incidental, direct, indirect, special or consequential Damages including lost profits, or damages resulting from loss of use or inability to use reseller's products or the software for any reason including copyright or patent infringement, or lost data, even if Ray Sauers Associates has been advised, knew or should have known of the possibility of such damages.

The following disclaimer is required by Videomedia, Inc.:

"Videomedia, Inc. makes no warranties whatsoever, either express or implied, regarding this product, including warranties with respect to its merchantability or its fitness for any particular purpose."

"This software contains V-LAN ver. 3.0 Command Protocols which communicate with V-LAN ver. 3.0 products developed by Videomedia, Inc. and V-LAN ver. 3.0 compatible products developed by third parties under license from Videomedia, Inc. Use of this software will allow "frame accurate" editing control of applicable videotape recorder decks, videodisc recorders/players and the like."

The following disclaimer is required by Altura Software, Inc. for the use of its Mac2Win software and Sample Source Code:

©1993–1998 Altura Software, Inc.

The following disclaimer is required by 3Prong.com Inc.:

Certain waveform and vector monitoring capabilities are provided under a license from 3Prong.com Inc.

The following disclaimer is required by Interplay Entertainment Corp.:

The "Interplay" name is used with the permission of Interplay Entertainment Corp., which bears no responsibility for Avid products.

This product includes portions of the Alloy Look & Feel software from Incors GmbH.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org>).

© DevelopMentor

This product may include the JCifs library, for which the following notice applies:

JCifs © Copyright 2004, The JCIFS Project, is licensed under LGPL (<http://jcifs.samba.org>). See the LGPL.txt file in the Third Party Software directory on the installation CD.

Avid Interplay contains components licensed from LavanTech. These components may only be used as part of and in connection with Avid Interplay.

Attn. Government User(s). Restricted Rights Legend

U.S. GOVERNMENT RESTRICTED RIGHTS. This Software and its documentation are "commercial computer software" or "commercial computer software documentation." In the event that such Software or documentation is acquired by or on behalf of a unit or agency of the U.S. Government, all rights with respect to this Software and documentation are subject to the terms of the License Agreement, pursuant to FAR §12.212(a) and/or DFARS §227.7202-1(a), as applicable.

Trademarks

003, 192 Digital I/O, 192 I/O, 96 I/O, 96i I/O, Adrenaline, AirSpeed, ALEX, Alienbrain, AME, AniMatte, Archive, Archive II, Assistant Station, AudioPages, AudioStation, AutoLoop, AutoSync, Avid Active, Avid Advanced Response, Avid DNA, Avid DNxcel, Avid DNxHD, Avid DS Assist Station, Avid Liquid, Avid Media Engine, Avid Media Processor, Avid MEDIAArray, Avid Mojo, Avid Remote Response, Avid Unity, Avid Unity ISIS, Avid VideoRAID, AvidRAID, AvidShare, AVIDstripe, AVX, Axiom, Beat Detective, Beauty Without The Bandwidth, Beyond Reality, BF Essentials, Bomb Factory, Boom, Bruno, C|24, CaptureManager, ChromaCurve, ChromaWheel, Cineractive Engine, Cineractive Player, Cineractive Viewer, Color Conductor, Command|24, Command|8, Conectiv, Control|24, Cosmonaut Voice, CountDown, d2, d3, DAE, Dazzle, Dazzle Digital Video Creator, D-Command, D-Control, Deko, DekoCast, D-Fi, D-fx, Digi 003, DigiBase, DigiDelivery, Digidesign, Digidesign Audio Engine, Digidesign Development Partners, Digidesign Intelligent Noise Reduction, Digidesign TDM Bus, DigiLink, DigiMeter, DigiPanner, DigiProNet, DigiRack, DigiSerial, DigiSnake, DigiSystem, Digital Choreography, Digital Nonlinear Accelerator, DigiTest, DigiTranslator, DigiWear, DINR, DNxchange, DPP-1, D-Show, DSP Manager, DS-StorageCalc, DV Toolkit, DVD Complete, D-Verb, Eleven, EM, EveryPhase, Expander, ExpertRender, Fader Pack, Fairchild, FastBreak, Fast Track, Film Cutter, FilmScribe, Flexevent, FluidMotion, Frame Chase, FXDeko, HD Core, HD Process, HDPack, Home-to-Hollywood, HYBRID, HyperControl, HyperSPACE, HyperSPACE HDCAM, iKnowledge, Image Independence, Impact, Improv, iNEWS, iNEWS Assign, iNEWS ControlAir, Instantwrite, Instinct, Intelligent Content Management, Intelligent Digital Actor Technology, IntelliRender, Intelli-Sat, Intelli-sat Broadcasting Recording Manager, InterFX, Interplay, inTONE, Intraframe, iS Expander, ISIS, IsoSync, iS9, iS18, iS23, iS36, ISIS, IsoSync, KeyRig, KeyStudio, LaunchPad, LeaderPlus, LFX, Lightning, Link & Sync, ListSync, LKT-200, Lo-Fi, Luna, MachineControl, Magic Mask, Make Anything Hollywood, make manage move | media, Marquee, MassivePack, Massive Pack Pro, M-Audio, M-Audio Micro, Maxim, Mbox, Media Composer, MediaDock, MediaDock Shuttle, MediaFlow, MediaLog, MediaMatch, MediaMix, Media Reader, Media Recorder, MEDIAArray, MediaServer, MediaShare, MetaFuze, MetaSync, MicroTrack, MIDI I/O, Midiman, Mix Rack, MixLab, Moviebox, Moviestar, MultiShell, NaturalMatch, NewsCutter, NewsView, Nitris, NL3D, NLP, Nova, NRV-10 interFX, NSDOS, NSWIN, Octane, OMF, OMF Interchange, OMM, OnDVD, Open Media Framework, Open Media Management, Ozone, Ozonic, Painterly Effects, Palladium, Personal Q, PET, Pinnacle, Pinnacle DistanTV, Pinnacle GenieBox, Pinnacle HomeMusic, Pinnacle MediaSuite, Pinnacle Mobile Media, Pinnacle Scorefitter, Pinnacle Studio, Pinnacle Studio MovieBoard, Pinnacle Systems, Pinnacle VideoSpin, Podcast Factory, PowerSwap, PRE, ProControl, ProEncode, Profiler, Pro Tools|HD, Pro Tools LE, Pro Tools M-Powered, Pro Transfer, Pro Tools, QuickPunch, QuietDrive, Realtime Motion Synthesis, Recti-Fi, Reel Tape Delay, Reel Tape Flanger, Reel Tape Saturation, Reprise, Res Rocket Surfer, Reso, RetroLoop, Reverb One, ReVibe, Revolution, rS9, rS18, RTAS, Salesview, Sci-Fi, Scorch, Scorefitter, ScriptSync, SecureProductionEnvironment, Serv|LT, Serv|GT, Session, Shape-to-Shape, ShuttleCase, Sibelius, SIDON, SimulPlay, SimulRecord, Slightly Rude Compressor, Smack!, Soft SampleCell, Soft-Clip Limiter, Solaris, SoundReplacer, SPACE, SPACEShift, SpectraGraph, SpectraMatte, SteadyGlide, Streamfactory, Streamgenie, StreamRAID, Strike, Structure, Studiophile, SubCap, Sundance Digital, Sundance, SurroundScope, Symphony, SYNC HD, Synchronic, SynchroScope, SYNC I/O, Syntax, TDM FlexCable, TechFlix, Tel-Ray, Thunder, Titansync, Titan, TL Aggro, TL AutoPan, TL Drum Rehab, TL Everyphase, TL Fauxlder, TL In Tune, TL MasterMeter, TL Metro, TL Space, TL Utilities, tools for storytellers, Torq, Torq Xponent, Transfuser, Transit, TransJammer, Trigger Finger, Trillium Lane Labs, TruTouch, UnityRAID, Vari-Fi, Velvet, Video the Web Way, VideoRAID, VideoSPACE, VideoSpin, VTEM, Work-N-Play, Xdeck, X-Form, Xmon, XPAND!, Xponent, X-Session, and X-Session Pro are either registered trademarks or trademarks of Avid Technology, Inc. in the United States and/or other countries.

Adobe and Photoshop are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Apple and Macintosh are trademarks of Apple Computer, Inc., registered in the U.S. and other countries. Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. All other trademarks contained herein are the property of their respective owners.

Avid Interplay Engine Failover Guide • 0130-07643-02 Rev H • November 2010 • Created 11/10/10 • This document is distributed by Avid in online (electronic) form only, and is not available for purchase in printed form.

Contents

	Using This Guide	9
	Symbols and Conventions	9
	If You Need Help	10
	Viewing Help and Documentation on the Interplay Portal	11
	Avid Training Services	12
Chapter 1	Automatic Server Failover Introduction	13
	Server Failover Overview	13
	How Server Failover Works	14
	Server Failover Configurations	15
	Server Failover Requirements	19
	Installing the Failover Hardware Components	20
	SR2400 Slot Locations	21
	SR2500 Slot Locations (for Infortrend A16F-R221)	22
	SR2500 Slot Locations (for Infortrend A16F-R2431)	23
	Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-221	24
	Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-R2431	27
	Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R221	29
	Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R2431	32
	Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R221	34
	Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R2431	37
	Clustering Terminology	39

Chapter 2	Creating a Microsoft Failover Cluster	41
	Server Failover Installation Overview	41
	Before You Begin the Server Failover Installation	42
	List of IP Addresses and Network Names	43
	Preparing the Server for the Cluster Service	47
	Setting the QLogic HBA Link Speed	47
	Increasing the Boot Delay	49
	Setting the ATTO Link Speed	50
	Removing Unnecessary Windows Components	51
	Renaming the Local Area Network Interface on Each Node	52
	Configuring the Private Network Adapter on Each Node	55
	Configuring the Binding Order Networks on Each Node	58
	Configuring the Public Network Adapter on Each Node	60
	Joining Both Servers to the Active Directory Domain	60
	Configuring the Cluster Shared-Storage RAID Disks on Each Node	60
	Configuring the Cluster Service	61
	Configuring the Cluster Service on the First Node	62
	Validating the Cluster Service on the First Node	67
	Configuring the Cluster Service on the Second Node	67
	Configuring Rules for the Cluster Networks	70
	Prioritizing the Heartbeat Adapter	71
	After Setting Up the Cluster	72
	Verifying the Quorum Disk	73
	Setting the Startup Times on Each Node	73
	Testing the Cluster Installation	75
	Installing the Distributed Transaction Coordinator	76
	Creating a Resource Group for the Distributed Transaction Coordinator	77
	Assigning an IP Address to the MSDTC Group	78
	Assigning a Network Name to the MSDTC Group	79
	Creating a Physical Resource for the MSDTC Group	80
	Assigning Distributed Transaction Coordinator Resource to the MSDTC Group	81
	Bringing the MSDTC Online	82

Chapter 3	Installing the Interplay Engine for a Failover Cluster	83
	Disabling Any Web Servers	83
	Installing the Interplay Engine on the First Node	84
	Preparation for Installing on the First Node	84
	Starting the Installation and Accepting the License Agreement	85
	Installing the Interplay Engine Using Custom Mode	85
	Specifying Cluster Mode During a Custom Installation	86
	Specifying the Interplay Engine Details	87
	Specifying the Interplay Engine Service Name	89
	Specifying the Destination Location	90
	Specifying the Default Database Folder	90
	Specifying the Share Name	91
	Specifying the Configuration Server	92
	Specifying the Server User	94
	Specifying the Server Cache	95
	Enabling Email Notifications	96
	Installing the Interplay Engine for a Custom Installation on the First Node	98
	Bringing the Disk Resource Online	99
	Installing the Interplay Engine on the Second Node	102
	Bringing the Interplay Engine Online	103
	Installing a Permanent License	103
	Testing the Complete Installation	104
	Updating a Clustered Installation (Rolling Upgrade)	106
	Uninstalling the Interplay Engine on a Clustered System	107
Chapter 4	Automatic Server Failover Tips and Rules	109
	Index	113

Using This Guide

Congratulations on the purchase of your Avid® Interplay™, a powerful system for managing media in a shared storage environment.

This guide is intended for all Avid Interplay administrators who are responsible for installing, configuring, and maintaining an Avid Interplay Engine with the Automatic Server Failover module integrated.



The documentation describes the features and hardware of all models. Therefore, your system might not contain certain features and hardware that are covered in the documentation.

Symbols and Conventions

Avid documentation uses the following symbols and conventions:

Symbol or Convention Meaning or Action



A note provides important related information, reminders, recommendations, and strong suggestions.



A caution means that a specific action you take could cause harm to your computer or cause you to lose data.



A warning describes an action that could cause you physical harm. Follow the guidelines in this document or on the unit itself when handling electrical equipment.



This symbol indicates menu commands (and subcommands) in the order you select them. For example, File > Import means to open the File menu and then select the Import command.



This symbol indicates a single-step procedure. Multiple arrows in a list indicate that you perform one of the actions listed.

(Windows), (Windows only), (Macintosh), or (Macintosh only)

This text indicates that the information applies only to the specified operating system, either Windows or Macintosh OS X.

Symbol or Convention	Meaning or Action
Bold font	Bold font is primarily used in task instructions to identify user interface items and keyboard sequences.
<i>Italic font</i>	Italic font is used to emphasize certain words and to indicate variables.
Courier Bold font	Courier Bold font identifies text that you type.
Ctrl+key or mouse action	Press and hold the first key while you press the last key or perform the mouse action. For example, Command+Option+C or Ctrl+drag.

If You Need Help

If you are having trouble using your Avid product:

1. Retry the action, carefully following the instructions given for that task in this guide. It is especially important to check each step of your workflow.
2. Check the latest information that might have become available after the documentation was published:
 - If the latest information for your Avid product is provided as printed release notes, they ship with your application and are also available online.
If the latest information for your Avid product is provided as a ReadMe file, it is supplied on your Avid installation CD or DVD as a PDF document (README_*product*.pdf) and is also available online.

You should always check online for the most up-to-date release notes or ReadMe because the online version is updated whenever new information becomes available. To view these online versions, select ReadMe from the Help menu, or visit the Knowledge Base at www.avid.com/readme.
3. Check the documentation that came with your Avid application or your hardware for maintenance or hardware-related issues.
4. Visit the online Knowledge Base at www.avid.com/onlinesupport. Online services are available 24 hours per day, 7 days per week. Search this online Knowledge Base to find answers, to view error messages, to access troubleshooting tips, to download updates, and to read or join online message-board discussions.

Viewing Help and Documentation on the Interplay Portal

You can quickly access the Interplay Help, PDF versions of the Interplay guides, and useful external links by viewing the Interplay User Information Center on the Interplay Portal. The Interplay Portal is a web site that runs on the Interplay Engine.

You can access the Interplay User Information Center through a browser from any system in the Interplay environment. You can also access it through the Help menu in Interplay Access and the Interplay Administrator.

The Interplay Help combines information from all Interplay guides in one Help system. It includes a combined index and a full-featured search. From the Interplay Portal, you can run the Help in a browser or download a compiled (.chm) version for use on other systems, such as a laptop.

To open the Interplay User Information Center through a browser:

1. Type the following line in a web browser:

```
http://Interplay_Engine_name
```

For *Interplay_Engine_name* substitute the name of the computer running the Interplay Engine software. For example, the following line opens the portal web page on a system named docwg:

```
http://docwg
```

2. Click the “Avid Interplay Documentation” link to access the User Information Center web page.

To open the Interplay User Information Center from Interplay Access or the Interplay Administrator:

- ▶ Select Help > Documentation Website on Server.

Avid Training Services

Avid makes lifelong learning, career advancement, and personal development easy and convenient. Avid understands that the knowledge you need to differentiate yourself is always changing, and Avid continually updates course content and offers new training delivery methods that accommodate your pressured and competitive work environment.

For information on courses/schedules, training centers, certifications, courseware, and books, please visit www.avid.com/support and follow the Training links, or call Avid Sales at 800-949-AVID (800-949-2843).

1 Automatic Server Failover Introduction

This chapter covers the following topics:

- [Server Failover Overview](#)
- [How Server Failover Works](#)
- [Installing the Failover Hardware Components](#)
- [Clustering Terminology](#)

Server Failover Overview

The automatic server failover mechanism in Avid Interplay allows client access to the Interplay Engine in the event of failures or during maintenance, with minimal impact on the availability. A failover server is activated in the event of application, operating system, or hardware failures. The server can be configured to notify the administrator about such failures using email.

The Interplay implementation of server failover uses Microsoft® clustering technology. For background information on clustering technology and links to Microsoft clustering information, see [“Clustering Terminology” on page 39](#).

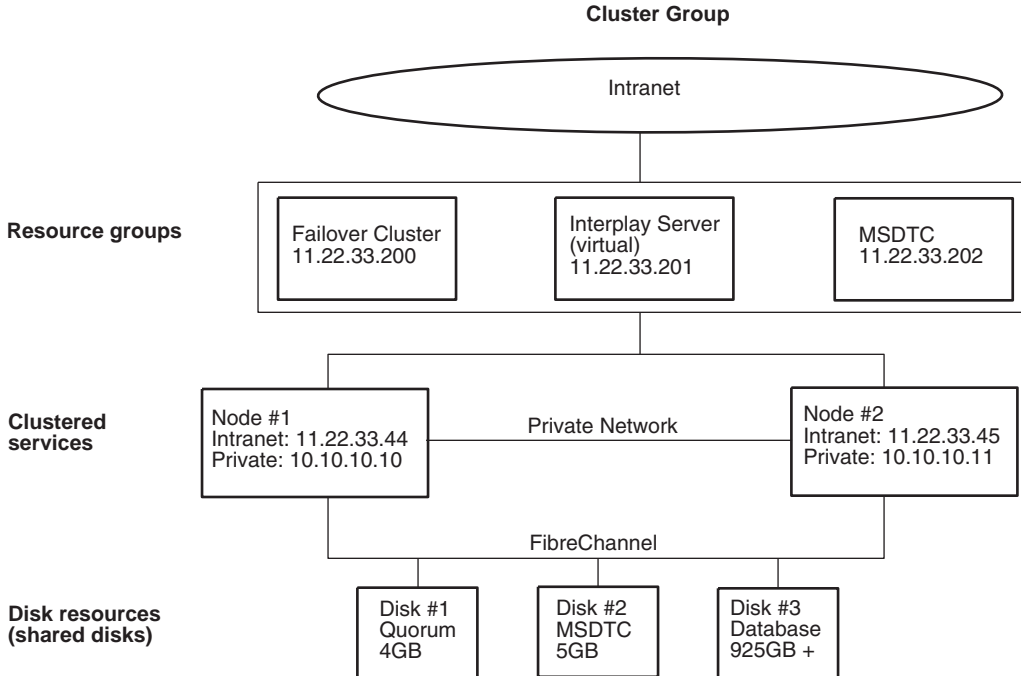


Additional monitoring of the hardware and software components of a high-availability solution is always required. Avid delivers Interplay preconfigured, but additional attention on the customer side is required to prevent outage (for example, when a private network fails, RAID disk fails, or a power supply loses power). In a mission critical environment, monitoring tools and tasks are needed to be sure there are no silent outages. If another (unmonitored) component fails, only an event is generated, and while this does not interrupt availability, it might go unnoticed and lead to problems. Additional software reporting such issues to the IT administration lowers downtime risk.

The failover cluster is a system made up of two server nodes and a shared-storage device connected over Fibre Channel. These are to be deployed in the same location given the shared access to the storage device. The cluster uses the concept of virtual servers to specify groups of resources that failover together.

1 Automatic Server Failover Introduction

The following diagram illustrates the components of a cluster group, including sample IP addresses. For a list of required IP addresses and node names, see “[List of IP Addresses and Network Names](#)” on page 43.



If you are already using clusters, the Avid Interplay Engine will not interfere with your current setup.

How Server Failover Works

Server failover works on two different levels:

- Failover in case of hardware failure
- Failover in case of network failure

Hardware Failover Process

When the Microsoft cluster service is running on both systems and the server is deployed in cluster mode, the Interplay Engine and its accompanying services are exposed to users as a virtual server. To clients, connecting to the clustered virtual Interplay Engine appears to be the same process as connecting to a single, physical machine. The user or client application does not know which node is actually hosting the virtual server.

When the server is online, the resource monitor regularly checks its availability and automatically restarts the server or initiates a failover to the other node if a failure is detected. The exact behavior can be configured using the Windows Cluster Administrator console. Because clients connect to the virtual network name and IP address, which are also taken over by the failover node, the impact on the availability of the server is minimal.

Network Failover Process

The cluster resource monitors one primary network that connects the virtual server to the intranet. If the primary network fails, the virtual server (and thus both cluster nodes) will go offline. Avid supports a configuration that uses connections to two public networks (VLAN 10 and VLAN 20) on a single switch. However, in this configuration Windows clustering technology binds multiple virtual IP addresses to VLAN 10 as the primary network, and if VLAN 10 fails the virtual server will go offline.

For a high degree of protection against network outages, Avid supports a configuration that uses two network switches, each connected to a shared primary network (VLAN 30) and protected by a failover protocol. If one network switch fails, the virtual server remains online through the other VLAN 30 network and switch.

These configurations are described in the next section.

Server Failover Configurations

There are three supported configurations for integrating a failover cluster into an existing network:

- A cluster in an Avid Unity ISIS environment that is integrated into the intranet through two layer-3 switches (VLAN 30 in Zone 3). This “redundant-switch” configuration protects against both hardware and network outages and thus provides a higher level of protection than the dual-connected configuration.
- A cluster in an Avid Unity ISIS environment that is integrated into the intranet through two public networks (VLAN 10 and VLAN 20 in Zone 1). This “dual connected” configuration protects against hardware outages and network outage on VLAN 20 but, because only VLAN 10 is monitored, does not protect against a outage on the other network.
- A cluster in an Avid Unity MediaNetwork environment that is integrated into the intranet through a single public network. This configuration protects against hardware outages. It relies on a single public network, and so does not protect against a network outage.

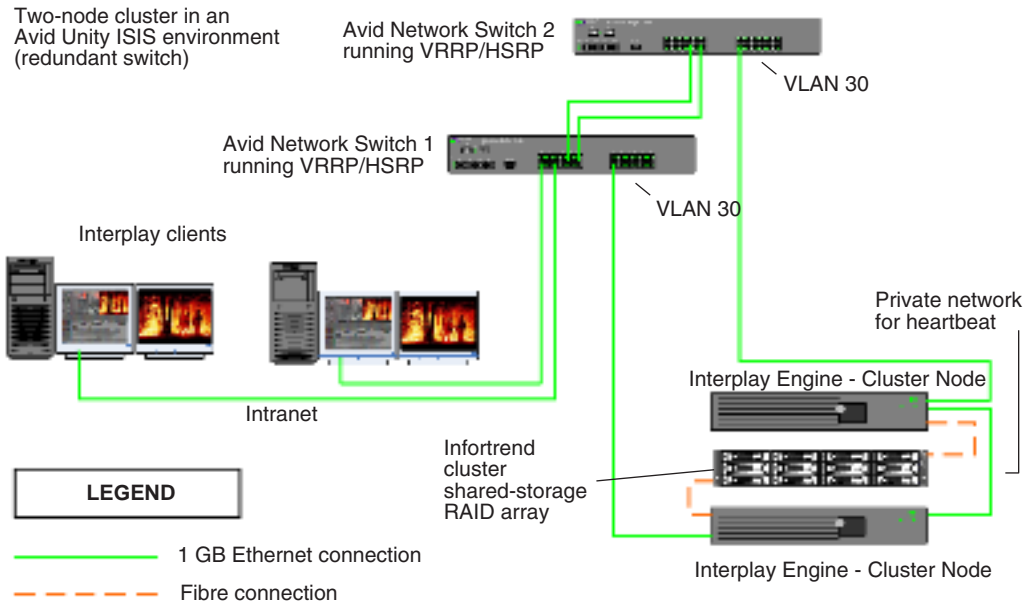
1 Automatic Server Failover Introduction

Redundant-Switch Configuration

The following diagram illustrates the failover cluster architecture for an Avid Unity ISIS environment that uses two layer-3 switches. These switches are configured for failover protection through either HSRP (Hot Standby Router Protocol) or VRRP (Virtual Router Redundancy Protocol). The cluster nodes are connected to two subnets (VLAN 30), each on a different switch. If one of the VLAN 30 networks fails, the virtual server remains online through the other VLAN 30 network and switch.



This guide does not describe how to configure redundant switches for an Avid Unity ISIS media network. Configuration information is included in the Avid Unity ISIS Switch Reference Guide, which is available for download from the Avid Customer Support Knowledge Base at www.avid.com/onlinesupport.

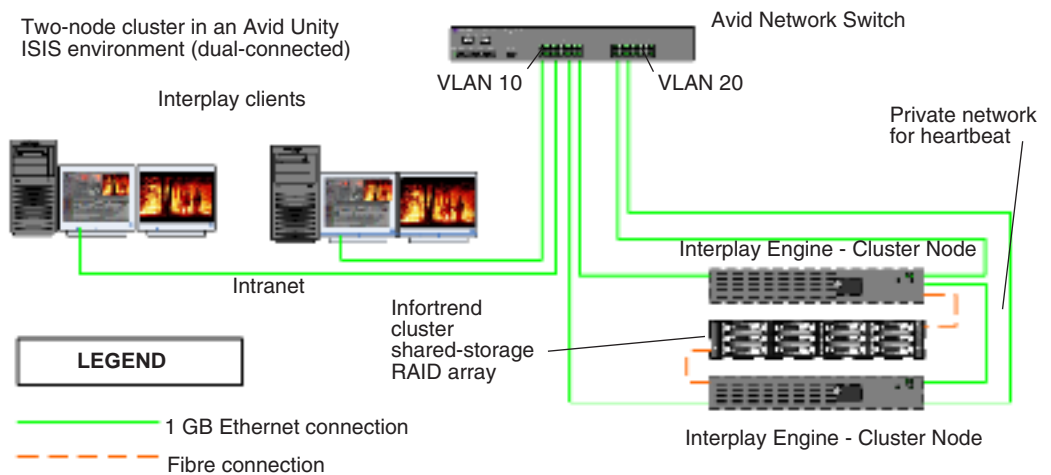


The following table describes what happens in the redundant-switch configuration as a result of an outage:

Type of Outage	Result
Hardware (CPU, network adapter, memory, cable, power supply) fails	The cluster detects the outage and triggers failover to the remaining node. The Interplay Engine is still accessible.
Network switch 1 (VLAN 30) fails	External switches running VRRP/HSRP detect the outage and make the gateway available as needed. The Interplay Engine is still accessible.
Network switch 2 (VLAN 30) fails	External switches running VRRP/HSRP detect the outage and make the gateway available as needed. The Interplay Engine is still accessible.

Dual-Connected Configuration

The following diagram illustrates the failover cluster architecture for an Avid Unity ISIS environment. In this environment, each cluster node is “dual-connected” to the network switch: one network interface is connected to the VLAN 10 subnet and the other is connected to the VLAN 20 subnet. In this configuration Windows clustering technology binds multiple virtual IP addresses to VLAN 10 as the primary network, and if VLAN 10 fails the virtual server goes offline.



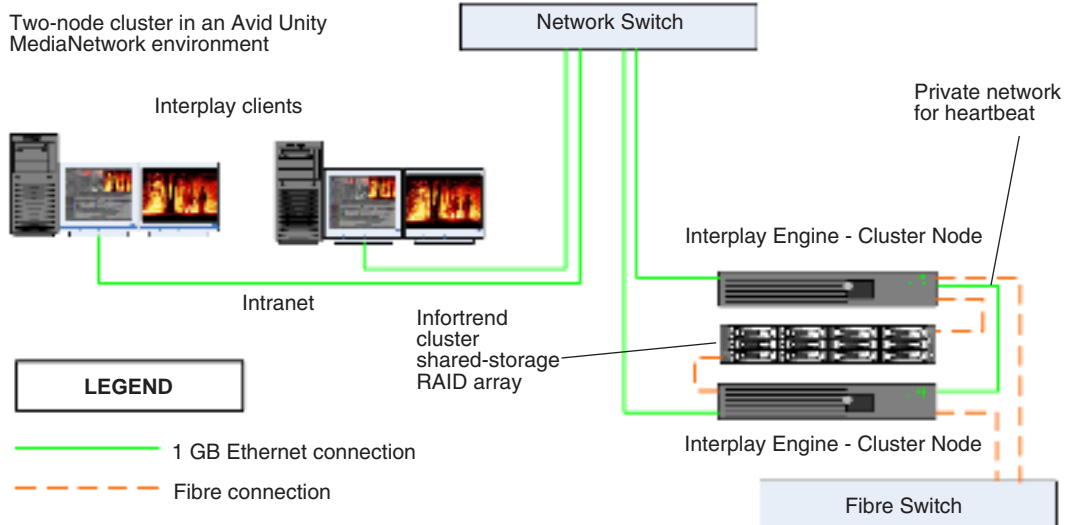
1 Automatic Server Failover Introduction

The following table describes what happens in the dual-connected configuration as a result of an outage:

Type of Outage	Result
Hardware (CPU, network adapter, memory, cable, power supply) fails	The cluster detects the outage and triggers failover to the remaining node. The Interplay Engine is still accessible.
Left ISIS VLAN (VLAN10, primary network) fails	The cluster detects the outage and triggers failover but detects that the second node is also disconnected from the left network, and both clusters fail. The Interplay Engine is not accessible.
Right ISIS VLAN (VLAN 20) fails	The Interplay Engine is still accessible through the left network.

Avid Unity MediaNetwork Configuration

The following diagram illustrates the failover cluster architecture for an Avid Unity MediaNetwork environment. In this environment, each cluster node is connected to a network switch through a single public network.



The following table describes what happens in the MediaNetwork configuration as a result of an outage:

Type of Outage	Result
Hardware (CPU, network adapter, memory, cable, power supply) fails	The cluster detects the outage and triggers failover to the remaining node. The Interplay Engine is still accessible.
Network switch fails	The cluster detects the outage and triggers failover but detects that the second node is also disconnected from the network, and both clusters fail. The Interplay Engine is not accessible.

Server Failover Requirements

You should make sure the server failover system meets the following requirements.

Hardware

A dual-server failover cluster-capable system with an Infortrend® cluster shared-storage RAID disk set is needed. The automatic server failover system was developed on and tested with the following:

- Intel Server Chassis SR2500 Packaged Cluster, which is the recommended hardware:
<http://www.intel.com/design/servers/chassis/sr2500/>
- Intel Server Chassis SR2400 Packaged Cluster:
<http://www.intel.com/design/servers/chassis/sr2400/>

The servers in a cluster are connected using one or more cluster shared-storage buses and one or more physically independent networks acting as a heartbeat.

Server Software

Two licenses of Windows Server 2003 Enterprise Edition or Windows Server 2003 Datacenter Edition are needed.

Space Requirements

The default disk configuration for the cluster shared RAID array is as follows:

- Quorum disk - 4GB
- MSDTC disk - 5GB
- Database disk - 925GB or larger

1 Automatic Server Failover Introduction

Antivirus Software

You can run antivirus software on a cluster, if the antivirus software is cluster-aware. For information about cluster-aware versions of your antivirus software, contact the antivirus vendor. If you are running antivirus software on a cluster, make sure you exclude these locations from the virus scanning: Q:\ (Quorum disk), C:\Windows\Cluster, and S:\Workgroup_Databases (database).

Functions You Need To Know

Before you set up a cluster in an Avid Interplay environment, you should be familiar with the following functions:

- Microsoft Windows Active Directory domains and domain users
- Microsoft Windows clustering (current version, as there are changes from prior version)
- Disk configuration (format, partition, naming)
- Network configuration

Installing the Failover Hardware Components

A failover cluster system includes the following components:

- Two Interplay Engine nodes or two Interplay Archive nodes (two SR2400 servers or two SR2500 servers)
- One Infortrend cluster shared-storage RAID array (one Infortrend A16F-R221 or one Infortrend A16F-R2431)

The following topics provide information about installing the failover hardware components for the supported configurations:


- [“SR2400 Slot Locations” on page 21](#)
- [“SR2500 Slot Locations \(for Infortrend A16F-R221\)” on page 22](#)
- [“SR2500 Slot Locations \(for Infortrend A16F-R2431\)” on page 23](#)
- [“Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-R221” on page 24](#)
- [“Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-R2431” on page 27](#)
- [“Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R221” on page 29](#)
- [“Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R2431” on page 32](#)

- “Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R221” on page 34
- “Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R2431” on page 37

SR2400 Slot Locations

The SR2400 is supported as a server for the Interplay applications. This section describes the slot locations that are specific to the Interplay components in a cluster configuration.

Use the following figure and table as guides to configuring an SR2400 system.



On the SR2400, all boards must be installed starting in the top slot, and the second board must be in the middle slot. The second board cannot be in the bottom slot with the middle slot left open.

SR2400 Back Panel Configuration for Avid Unity Environment


Slot	Avid Unity ISIS	Avid Unity MediaNetwork
3	Intel Pro 1000MT	ATTO
2	QLogic® Card	QLogic Card
1	Empty	Intel Pro 1000MT ^a


- a. Unity MediaNetwork environment: the Pro 1000MT card is shipped in slot 3 (top). You must move the card to slot 1 (bottom) and install the ATTO card in slot 3 (top). The Pro 1000MT is not used in an Unity MediaNetwork environment.

SR2500 Slot Locations (for Infortrend A16F-R221)

The SR2500 is supported as a server for the Interplay applications. This section describes the slot locations that are specific to the Interplay components in a cluster configuration that uses the Infortrend model A16F-R221 shared-storage RAID array.

Use the following figure and table as guides to configuring an SR2500 system.



 *It is important to match the slot locations in the following tables because they match the order that the drivers are loaded on the SR2500 Recovery DVDs.*

SR2500 Back Panel Configuration for Avid Unity Environment


Slot Type	Slot	Avid Unity ISIS	Avid Unity MediaNetwork
PCI-X	3	Empty	ATTO
	2	Empty	Empty
	1	QLogic Card ^a to Infortrend A16F-R221	QLogic Card ^a to Infortrend A16F-R221
PCIe	NA	NA	NA
	2	Intel Pro 1000PT	Intel Pro 1000PT
	1	Empty	Empty

- a. The SR2500 server might ship with the QLogic card in PCI-X slot 2 (middle). You must move the QLogic card to PCI-X slot 1 (bottom), because this configuration matches the order that the drivers are loaded on the SR2500 Recovery DVDs.

SR2500 Slot Locations (for Infortrend A16F-R2431)

The SR2500 is supported as a server for the Interplay applications. This section describes the slot locations that are specific to the Interplay components in a cluster configuration that uses the Infortrend A16F-R2431 shared-storage RAID array.

Use the following figure and table as guides to configuring an SR2500 system.



It is important to match the slot locations in the following tables because they match the order that the drivers are loaded on the SR2500 Recovery DVDs.

SR2500 Back Panel Configuration for Avid Unity Environment

Slot Type	Slot	Avid Unity ISIS	Avid Unity MediaNetwork
PCI-X	3	Empty	Empty
	2	Empty	Empty
	1	ATTO FC-41XS to Infortrend A16F-R2431	ATTO FC-41XS to Infortrend A16F-R2431
PCIe	NA	NA	NA
	2	Intel Pro 1000PT	Intel Pro 1000PT
	1	Empty	ATTO FC-41EL to MediaNetwork


Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-221

Make the following cable connections to add a failover cluster to an Avid Unity ISIS environment, using the redundant-switch configuration with an Infortrend A16F-R221 RAID array:

- First cluster node:
 - Left on-board network interface connector to layer-3 switch 1 (VLAN 30)
 - QLogic card connector to RAID array, Fibre Channel 1 left connector
- Second cluster node:
 - Left on-board network interface connector to layer-3 switch 2 (VLAN 30)
 - QLogic card connector to RAID array, Fibre Channel 0 left connector
- Right connector on PCI adapter network interface in the first cluster node to right connector on PCI adapter network interface in second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left).

You can implement this configuration using either SR2400 servers or SR2500 servers. The following illustrations show the connections for each type of server.

Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, SR2400, Infortrend A16F-R221




LEGEND

- 1GB Ethernet connection
- - - Fibre connection

1 Automatic Server Failover Introduction

Failover Cluster Connections: Avid Unity ISIS Environment, Redundant-Switch Configuration, SR2500, Infortrend A16F-R221



Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, Infortrend A16F-R2431


Make the following cable connections to add a failover cluster to an Avid Unity ISIS environment, using the redundant-switch configuration with an Infortrend A16F-R2431 RAID array:

- First cluster node:
 - Left on-board network interface connector to layer-3 switch 1 (VLAN 30)
 - ATTO 41XS card connector to RAID array, Fibre Channel 0 top-left connector
- Second cluster node:
 - Left on-board network interface connector to layer-3 switch 2 (VLAN 30)
 - ATTO 41XS card connector to RAID array, Fibre Channel 1 bottom-right connector
- Right connector on PCI adapter network interface in the first cluster node to right connector on PCI adapter network interface in second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left)

You can implement this configuration using SR2500 servers. The following illustration shows the connections for these servers.

1 Automatic Server Failover Introduction

Failover Cluster Connections: Avid Unity ISIS, Redundant-Switch Configuration, SR2500, Infortrend A16F-R2431



Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R221


Make the following cable connections to add a failover cluster to an Avid Unity ISIS environment, using the dual-connected configuration with an Infortrend A16F-R221 RAID array:

- First cluster node:
 - Left on-board network interface connector to ISIS left subnet (VLAN 10)
 - Right on-board network interface connector to ISIS right subnet (VLAN 20)
 - QLogic card connector to RAID array, Fibre Channel 1 left connector
- Second cluster node:
 - Left on-board network interface connector to ISIS left subnet (VLAN 10)
 - Right on-board network interface connector to ISIS right subnet (VLAN 20)
 - QLogic card connector to RAID array, Fibre Channel 0 left connector
- Right connector on PCI adapter network interface in the first cluster node to right connector on PCI adapter network interface in second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left)


You can implement this configuration using either SR2400 servers or SR2500 servers. The following illustrations show the connections for each type of server.

1 Automatic Server Failover Introduction

Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, SR2400, Infortrend A16F-R221



Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, SR2500, Infortrend A16F-R221




Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, Infortrend A16F-R2431

Make the following cable connections to add a failover cluster to an Avid Unity ISIS environment, using the dual-connected configuration with an Infortrend A16F-R2431 RAID array:

- First cluster node:
 - Left on-board network interface connector to ISIS left subnet (VLAN 10)
 - Right on-board network interface connector to ISIS right subnet (VLAN 20)
 - ATTO 41XS card connector to RAID array, Fibre Channel 0 top-left connector
- Second cluster node:
 - Left on-board network interface connector to ISIS left subnet (VLAN 10)
 - Right on-board network interface connector to ISIS right subnet (VLAN 20)
 - ATTO 41XS card connector to RAID array, Fibre Channel 1 bottom-right connector
- Right connector on PCI adapter network interface in the first cluster node to right connector on PCI adapter network interface in second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left)

You can implement this configuration using SR2500 servers. The following illustration shows the connections for these servers.

Failover Cluster Connections: Avid Unity ISIS, Dual-Connected Configuration, SR2500, Infortrend A16F-R2431



Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R221

Make the following cable connections to add a failover cluster to an Unity MediaNetwork environment, using the Infortrend A16F-R221RAID array:


- First cluster node:
 - Left on-board network interface connector to Ethernet[®] public network on the Avid network switch
 - QLogic card connector to RAID array, Fibre Channel 1 left connector
 - ATTO card connector to Unity MediaNetwork FC switch
- Second cluster node:
 - Left on-board network interface connector to Ethernet public network on the Avid network switch
 - QLogic card connector to RAID array, Fibre Channel 0 left connector
 - ATTO card connector to Unity MediaNetwork FC switch
- Right on-board network interface connector on the first cluster node to right on-board network interface connector on the second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left)



SR2400 servers ship with an Intel Pro 1000 MT card in slot 3 (top). You need to move this card to slot 1 (bottom). Then add an ATTO host bus adapter in slot 3 (top). See “SR2400 Slot Locations” on page 21.

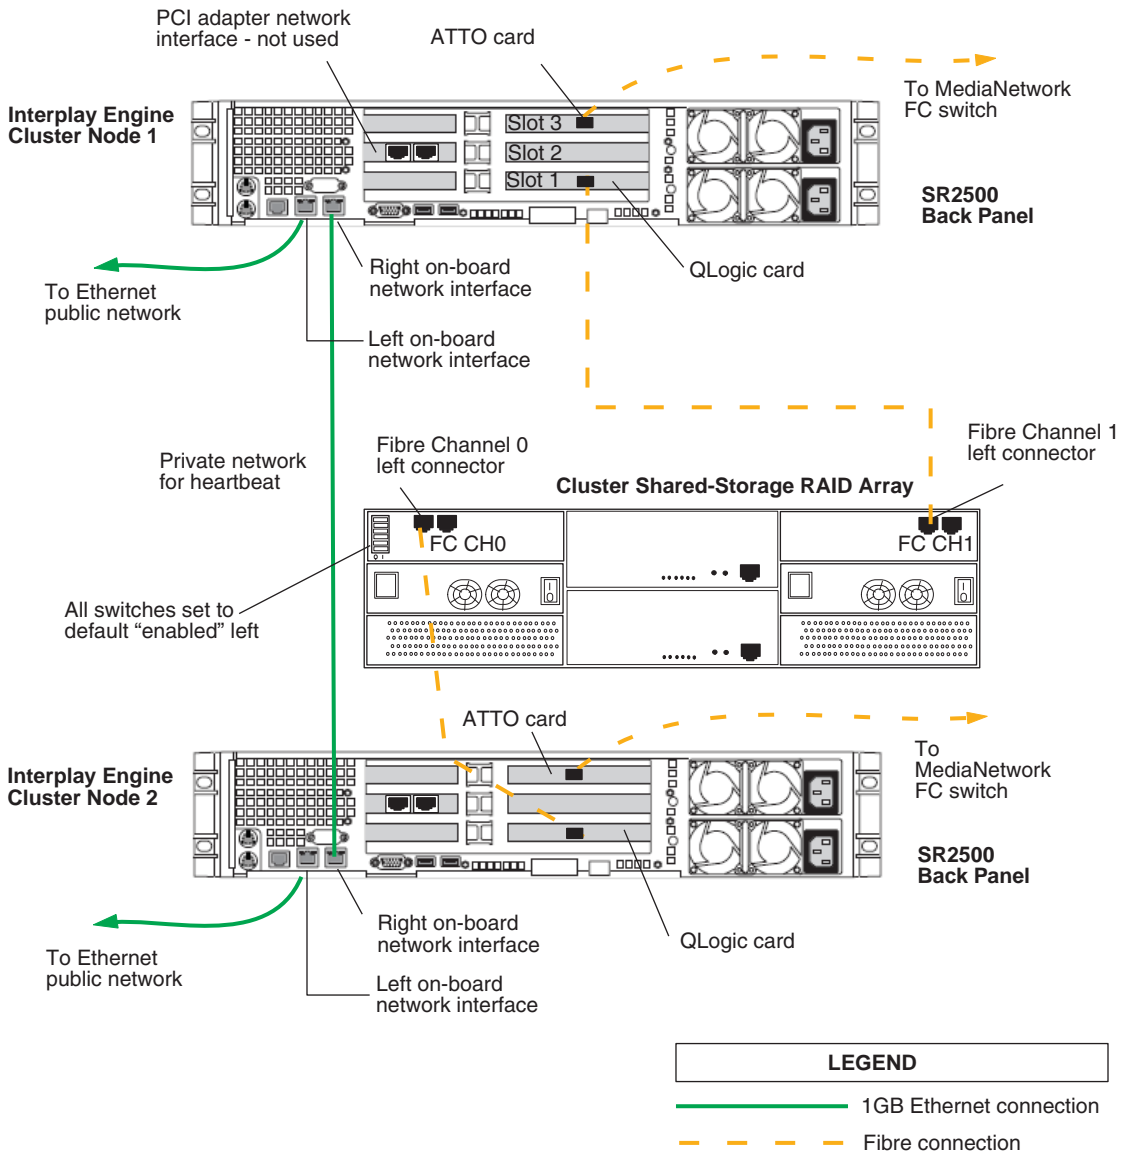
You can implement this configuration using either SR2400 servers or SR2500 servers. The following illustrations show the connections for each type of server.

Failover Cluster Connections: Avid Unity MediaNetwork, SR2400, Infortrend A16F-R221



1 Automatic Server Failover Introduction

Failover Cluster Connections: Avid Unity MediaNetwork, SR2500, Infortrend A16F-R221



Failover Cluster Connections: Avid Unity MediaNetwork, Infortrend A16F-R2431

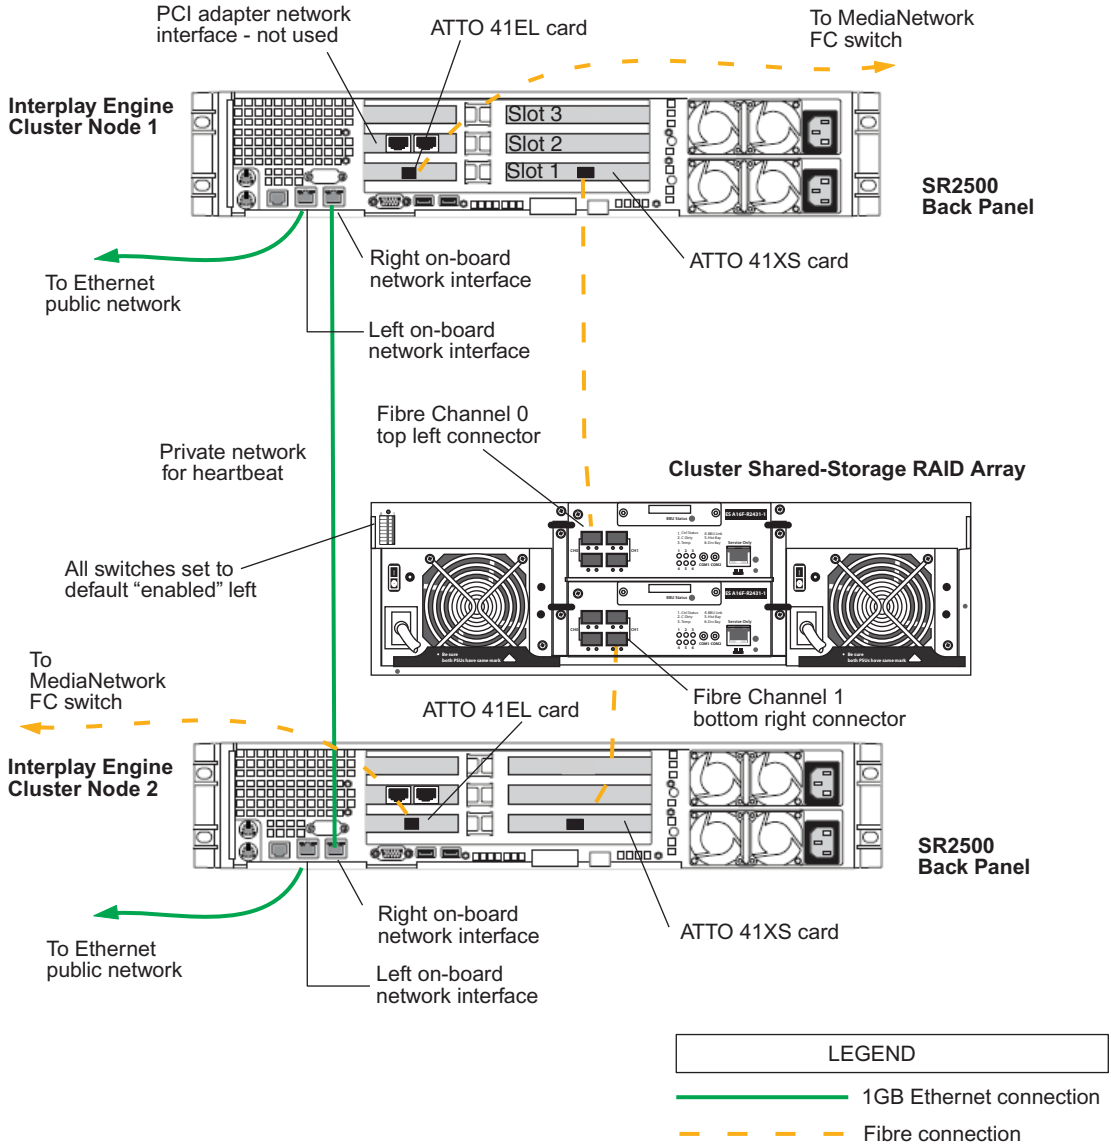
Make the following cable connections to add a failover cluster to an Unity MediaNetwork environment.

- First cluster node:
 - Left on-board network interface connector to Ethernet public network on the Avid network switch
 - ATTO 41XS card connector to RAID array, Fibre Channel 0 top-left connector
 - ATTO 41EL card connector to Unity MediaNetwork FC switch
- Second cluster node:
 - Left on-board network interface connector to Ethernet public network on the Avid network switch
 - ATTO 41XS card connector to RAID array, Fibre Channel 1 bottom-right connector
 - ATTO 41EL card connector to Unity MediaNetwork FC switch
- Right on-board network interface connector on the first cluster node to right on-board network interface connector on the second cluster node (private network for heartbeat)
- All switches on the cluster shared-storage RAID array are in the default “enable” position (left)

You can implement this configuration using SR2500 servers. The following illustration shows the connections for these servers.

1 Automatic Server Failover Introduction

Failover Cluster Connections: Avid Unity MediaNetwork, SR2500, Infortrend A16F-R2431



Clustering Terminology

Clustering is not always straightforward, so it is important that you get familiar with the terminology of server clusters before you start. A good source of information is the Microsoft Technology Center for Clustering Services under:

<http://www.microsoft.com/windowsserver2003/technologies/clustering/default.mspx>

Detailed architecture documentation can be found here:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/servercluster.mspx>

Here is a brief summary of the major concepts and terms:

- *Nodes*: Individual computers in a cluster configuration.
- *Cluster service*: The group of components on each node that perform a cluster-specific activity.
- *Resource*: Cluster components (hardware and software) that are managed by the cluster service. Resources are physical hardware devices such as disk drives, and logical items such as IP addresses and applications.
- *Online resource*: A resource that is available and is providing its service.
- *Quorum resource*: A special common cluster resource. This resource plays a critical role in cluster operations.
- *Resource group*: A collection of resources that are managed by the cluster service as a single, logical unit.

1 Automatic Server Failover Introduction

2

Creating a Microsoft Failover Cluster

This chapter describes the processes for creating a Microsoft failover cluster for automatic server failover. It is crucial that you follow the instructions given in this chapter completely, otherwise the automatic server failover will not work.

This chapter covers the following topics:

- [Server Failover Installation Overview](#)
- [Before You Begin the Server Failover Installation](#)
- [Preparing the Server for the Cluster Service](#)
- [Configuring the Cluster Service](#)
- [Configuring Rules for the Cluster Networks](#)
- [After Setting Up the Cluster](#)
- [Installing the Distributed Transaction Coordinator](#)

Instructions for installing the Interplay Engine are provided in [“Installing the Interplay Engine for a Failover Cluster”](#) on page 83.

Server Failover Installation Overview

Installation and configuration of the automatic server failover consists of the following major tasks:

- Make sure that the network is correctly set up and that you have reserved IP host names and static IP addresses (see [“Before You Begin the Server Failover Installation”](#) on page 42).
- Prepare the servers for the cluster service (see [“Preparing the Server for the Cluster Service”](#) on page 47). This includes configuring the nodes for the network and formatting the drives.
- Configure the cluster service (see [“Configuring the Cluster Service”](#) on page 61, [“Configuring Rules for the Cluster Networks”](#) on page 70, and [“After Setting Up the Cluster”](#) on page 72).
- Install the Distributed Transaction Coordinator (MSDTC group) (see [“Installing the Distributed Transaction Coordinator”](#) on page 76).

2 Creating a Microsoft Failover Cluster

- Install the Interplay Engine on both nodes (see “[Installing the Interplay Engine for a Failover Cluster](#)” on page 83).
- Test the complete installation (see “[Testing the Complete Installation](#)” on page 104).



Do not install any other software on the cluster machines except the Interplay engine. For example, Media Indexer software needs to be installed on a different server. For complete installation instructions, see the Avid Interplay Software Installation and Configuration Guide.

For more details about server clusters, see the Microsoft document “[Guide to Creating and Configuring a Server Cluster under Windows Server 2003](#),” available at:

<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/confclus.mspx>

Before You Begin the Server Failover Installation

Before you begin the installation process, you need to do the following:

- Make sure all cluster hardware connections are correct. See “[Installing the Failover Hardware Components](#)” on page 20.
- Make sure that the facility has a network that is qualified to run Active Directory and DNS services.
- Determine the subnet mask, the gateway, DNS, and WINS server addresses on the network.
- Install and set up an Avid Unity client on both servers. See the *Avid Unity MediaNetwork File Manager Setup Guide* or the *Avid Unity ISIS System Setup Guide*.
- Create or select two domain user accounts:
 - Cluster Service Account (Server Execution User): Create or select an account (sometimes called the cluster user account) that is used to start the cluster service and is also used by the Interplay Engine service. This account must be a domain user and it must be a unique name that will not be used for any other purpose. The procedures in this document use *sqauser* as an example of a Cluster Service Account. This account is automatically added to the Local Administrators group on each node by the Interplay Engine software during the installation process.

The Server Execution User is critical to the operation of the Interplay Engine. If necessary, you can change the name of the Server Execution User after the installation. For more information, see “[Troubleshooting the Server Execution User Account](#)” and “[Re-creating the Server Execution User](#)” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide* and the Interplay ReadMe.

For information on creating a cluster user account, see the Microsoft document “Guide to Creating and Configuring a Server Cluster under Windows Server 2003.”

<http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/clustering/confclus.msp>.

- Cluster Installation and Administration Account: Create or select a user account to use during the installation process. This user account must be a domain user account with privileges to add servers to the domain. Also use this account to log in to and administer the system.



Do not use the same username and password for the Cluster Service Account and the Cluster Installation Account. These accounts have different functions and require different privileges.

- Create an Avid Unity user account with read and write privileges. This account is not needed for the installation of Interplay Engine, but is required for the operation of Interplay Engine. The user name and password must match the user name and password of the Cluster Service Account.
- Make sure the network includes an Active Directory domain before you install or configure the cluster.
- Reserve static IP addresses for all network interfaces and host names. See “[List of IP Addresses and Network Names](#)” on page 43.

List of IP Addresses and Network Names

You need to reserve IP host names and static IP addresses on the in-network DNS server before you begin the installation process. The number of IP addresses you need depends on your configuration:

- An Avid Unity ISIS environment with a redundant-switch configuration requires 5 IP addresses
- An Avid Unity ISIS environment with a dual-connected configuration requires 8 IP addresses
- An Avid Unity MediaNetwork environment requires 5 IP addresses.

The following table provides a list of example names that you can use when configuring the cluster. The procedures in this chapter use these example names.



Make sure that these IP addresses are outside of the range that is available to DHCP so they cannot automatically be assigned to other machines.

2 Creating a Microsoft Failover Cluster



If your Active Directory domain or DNS includes more than one cluster, to avoid conflicts, you need to make sure the cluster names, MSDTC names, and IP addresses are different for each cluster.



All names must be valid and unique network host names.

IP Addresses and Node Names: ISIS Redundant-Switch Configuration

Node or Service	Item Required	Example Name	Where Used
First Cluster Node	<ul style="list-style-type: none">• 1 Host Name• 1 ISIS IP address - public• 1 IP address - private (Heartbeat)	SECLUSTER1	See “Configuring the Cluster Service on the First Node” on page 62 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Second Cluster Node	<ul style="list-style-type: none">• 1 Host Name• 1 ISIS IP address - public• 1 IP address - private (Heartbeat)	SECLUSTER2	See “Configuring the Cluster Service on the Second Node” on page 67 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Cluster service	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 ISIS IP address (virtual IP address)	SECLUSTER	See “Configuring the Cluster Service on the First Node” on page 62.
MSDTC service — Distributed Transaction Coordinator	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 ISIS IP address (virtual IP address)	CLUSTERMSDTC	See “Assigning a Network Name to the MSDTC Group” on page 79.
Interplay Engine service	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 ISIS IP address - public (virtual IP address)	SEENGINE	See “Specifying the Interplay Engine Details” on page 87 and “Specifying the Interplay Engine Service Name” on page 89.

IP Addresses and Node Names: ISIS Dual-Connected Configuration

Node or Service	Item Required	Example Name	Where Used
First Cluster Node	<ul style="list-style-type: none"> • 1 Host Name • 2 ISIS IP addresses - public (one for left and one for right) • 1 IP address - private (Heartbeat) 	SECLUSTER1	See “Configuring the Cluster Service on the First Node” on page 62 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Second Cluster Node	<ul style="list-style-type: none"> • 1 Host Name • 2 ISIS IP addresses - public (one for left and one for right) • 1 IP address - private (Heartbeat) 	SECLUSTER2	See “Configuring the Cluster Service on the Second Node” on page 67 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Cluster service	<ul style="list-style-type: none"> • 1 Network Name (virtual host name) • 1 ISIS IP address (virtual IP address) 	SECLUSTER	See “Configuring the Cluster Service on the First Node” on page 62.
MSDTC service — Distributed Transaction Coordinator	<ul style="list-style-type: none"> • 1 Network Name (virtual host name) • 1 ISIS IP address (virtual IP address) 	CLUSTERMSDTC	See “Assigning a Network Name to the MSDTC Group” on page 79.
Interplay Engine service	<ul style="list-style-type: none"> • 1 Network Name (virtual host name) • 2 ISIS IP addresses - public (one for left and one for right) (virtual IP address) 	SEENGINE	See “Specifying the Interplay Engine Details” on page 87 and “Specifying the Interplay Engine Service Name” on page 89.

2 Creating a Microsoft Failover Cluster

IP Addresses and Node Names: MediaNetwork Configuration

Node or Service	Item Required	Example Name	Where Used
First Cluster Node	<ul style="list-style-type: none">• 1 Host Name• 1 MediaNetwork IP address - public• 1 IP address - private (Heartbeat)	SECLUSTER1	See “Configuring the Cluster Service on the First Node” on page 62 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Second Cluster Node	<ul style="list-style-type: none">• 1 Host Name• 1 MediaNetwork IP address - public• 1 IP address - private (Heartbeat)	SECLUSTER2	See “Configuring the Cluster Service on the Second Node” on page 67 and “Creating a Resource Group for the Distributed Transaction Coordinator” on page 77.
Cluster service	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 MediaNetwork IP address (virtual IP address)	SECLUSTER	See “Configuring the Cluster Service on the First Node” on page 62.
MSDTC service — Distributed Transaction Coordinator	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 MediaNetwork IP address (virtual IP address)	CLUSTERMSDTC	See “Assigning a Network Name to the MSDTC Group” on page 79.
Interplay Engine service	<ul style="list-style-type: none">• 1 Network Name (virtual host name)• 1 MediaNetwork IP address - public (virtual IP address)	SEENGINE	See “Specifying the Interplay Engine Details” on page 87 and “Specifying the Interplay Engine Service Name” on page 89.

Preparing the Server for the Cluster Service

Before you configure the cluster service, you need to complete the tasks in the following procedures:

- “Setting the QLogic HBA Link Speed” on page 47
- “Increasing the Boot Delay” on page 49
- “Setting the ATTO Link Speed” on page 50
- “Renaming the Local Area Network Interface on Each Node” on page 52
- “Removing Unnecessary Windows Components” on page 51
- “Renaming the Local Area Network Interface on Each Node” on page 52
- “Configuring the Private Network Adapter on Each Node” on page 55
- “Configuring the Binding Order Networks on Each Node” on page 58
- “Configuring the Public Network Adapter on Each Node” on page 60
- “Joining Both Servers to the Active Directory Domain” on page 60
- “Configuring the Cluster Shared-Storage RAID Disks on Each Node” on page 60

Setting the QLogic HBA Link Speed

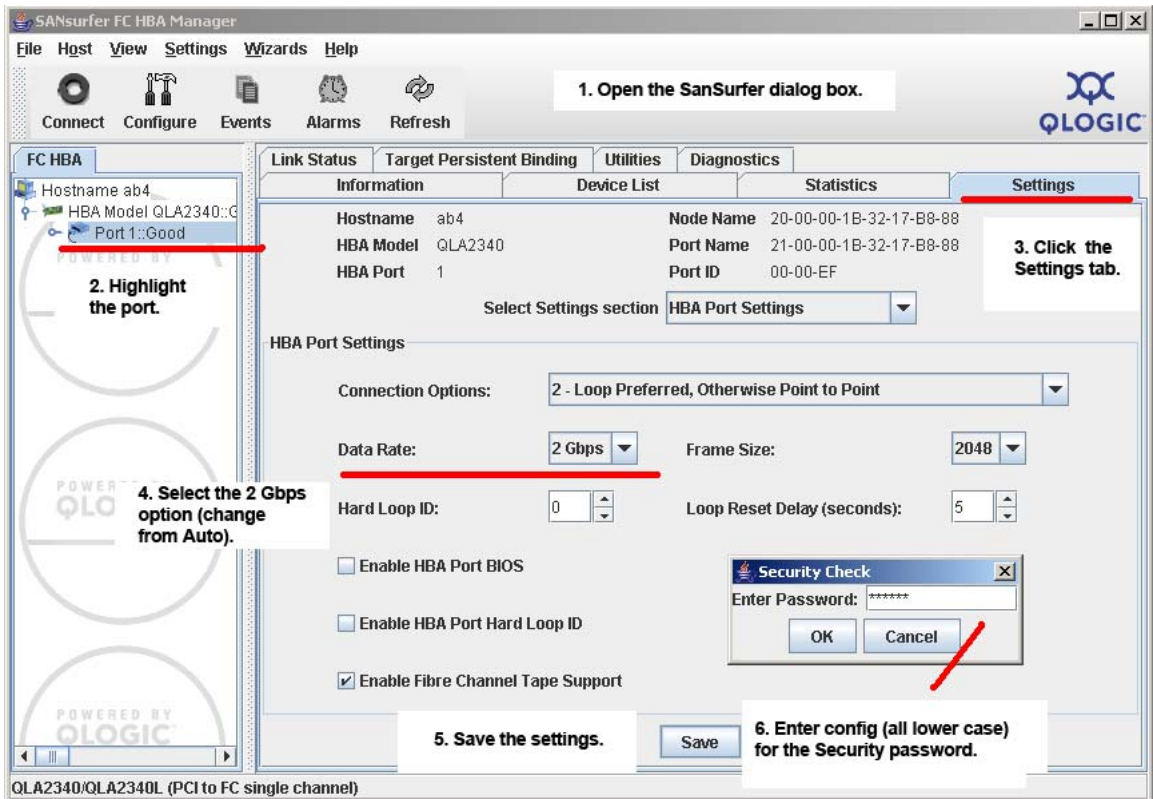
To avoid possible problems with the Infortrend RAID array (Model A16F-R221), Avid recommends that you change the QLogic HBA link speed (data rate) from the default setting to 2 Gbps. You need to specify this setting on both the SR2400 server and the SR2500 server. Change the setting by using the SAN Surfer utility on both nodes.

To set the QLogic HBA link speed:

1. On the first node, click Start, and select Programs > QLogic Management Suite > San Surfer.

The San Surfer FC HBA Manager dialog box opens.

2 Creating a Microsoft Failover Cluster



2. In the left pane, select Port 1.
3. Click the Settings tab.
4. In the HBA Port Settings section, click the arrow pointer for the Data Rate list and change the default setting from Auto to 2 Gbps.
5. Click Save.
6. When prompted for a password, enter `config` and click OK.
7. On the other node, repeat steps 1 through 6.
8. Verify that the SAN Surfer data rate is set to 2 Gbps on both nodes.

Increasing the Boot Delay

Increasing the the timeout value in the boot.ini file increases the time it takes for the server to boot. This boot delay might avoid a problem if the Infortrend requires a longer than usual time for a self-test and initializing. If a server boots before the Infortrend is ready, the cluster might fail.

Avid recommends that you offset the startup time of each node's operating system (see [“Setting the Startup Times on Each Node” on page 73](#)).

To increase the timeout value in the boot.ini file:

1. Start the first node and log in to Windows.
2. At the command prompt, type the following command and press Enter:

```
bootcfg /timeout 60
```

This command changes the boot timeout delay to 60 seconds, after which the default operating system is loaded.

3. Repeat these steps on the second node, but set the timeout delay to 120 seconds.

2 Creating a Microsoft Failover Cluster

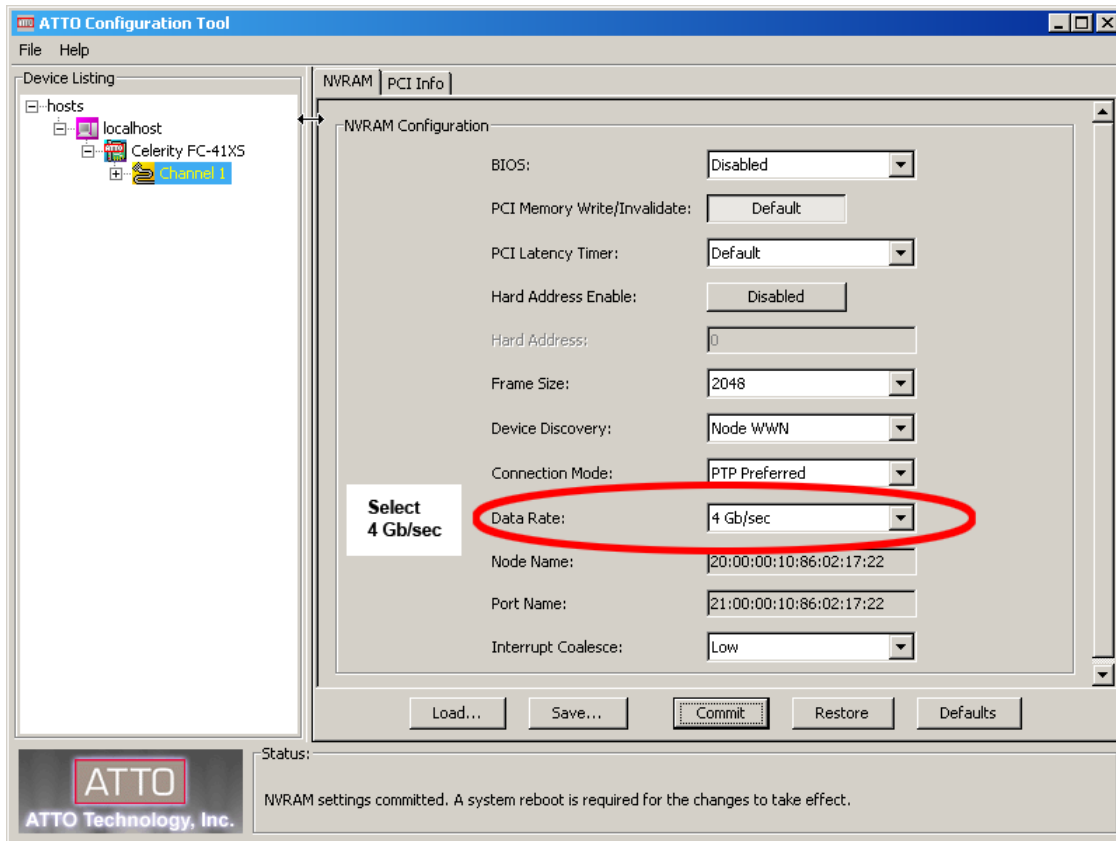
Setting the ATTO Link Speed

To avoid possible problems with the Infortrend RAID array (Model A16F-R2431), Avid recommends that you change the ATTO link speed (data rate) from the default setting to 4 Gbps. You need to specify this setting on both the SR2400 server and the SR2500 server. Change the setting by using the ATTO Configuration Tool on both nodes.

To set the ATTO link speed:

1. On the first node, click Start, and select Programs > ATTO Configuration Tool > Configuration Tool.

The ATTO Configuration Tool dialog box opens.



2. In the left pane, navigate to the appropriate channel on your host adapter.

The NVRAM tab opens.

3. Click the arrow pointer for the Data Rate list and change the default setting from Auto to 4 Gb/sec.

4. Click Commit.
5. Reboot the system.
6. Open the Configuration tool again and verify that the Data Rate is set to 4 Gb/sec.
7. On the other node, repeat steps 1 through 6.

Removing Unnecessary Windows Components

You need to remove some unnecessary Windows components before you configure the cluster service. Which components you remove depends on the type of cluster you are configuring.

- For a cluster that will be used as an Interplay Engine, remove the following from each server:
 - Internet Information Services (IIS)
 - IIS-Admin
 - Internet Explorer Enhanced Security Configuration
 - Microsoft SQL Native Client
- For a cluster that will be used as an Interplay Archive Engine, remove the following from each server:
 - Internet Information Services (IIS)
 - IIS-Admin
 - Internet Explorer Enhanced Security Configuration



Microsoft SQL Native Client is required for SGL archive solutions and should be correctly installed and configured before creating a cluster that will be used as an Archive Engine.

To remove unnecessary Windows components:

1. On one of the servers, click Start and select Control Panel > Add or Remove Programs.
2. **(Interplay Engine only)** If Microsoft SQL Native Client is listed, select it and click Change/Remove.

2 Creating a Microsoft Failover Cluster

3. Click Add/Remove Windows Components.
 - a. If a box is checked for Internet Explorer Enhanced Security Configuration, click the check box to remove the check.
 - b. Select Application Server and click Details.
 - c. If a box is checked for Internet Information Services (IIS), click the check to remove the check.
 - d. Click OK, then click Next. At the end of the process, click Finish.
4. Repeat the procedure on the other server.

Renaming the Local Area Network Interface on Each Node

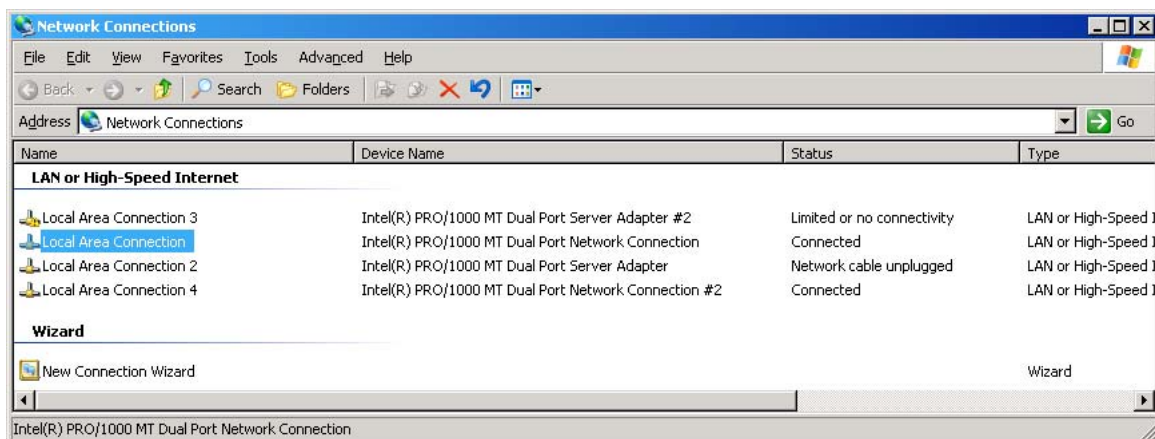
You need to rename the LAN interface on each node to appropriately identify each network. Although you can use any name for the network connections, Avid suggests that you use the naming conventions provided in the table in the following procedure.

Make sure you use the same name on both nodes. The names and network connections on both nodes must match.

To rename the local area network connections:

1. Open the Network Connections window.
 - a. Click Start and select Control Panel.
 - b. Right-click Network Connections, and select Open.

The Network Connections window opens.



2. Right-click one of the listed network connections and select Rename.

You need to match the numbered connection with the appropriate device. For example, you can start by determining which connection refers to the left on-board network interface and select that connection.




Both nodes must use identical network interface names. Although you can use any name for the network connections, Avid suggests that you use the naming conventions provided in the following table.

3. Depending on your Avid Unity network and the device you selected, type a new name for the network connection and press Enter.

Use the following illustration and table for reference. The illustration uses an SR2400 in an Avid Unity ISIS environment as an example.

2 Creating a Microsoft Failover Cluster




Naming Network Connections

Network Interface	Avid Unity ISIS New Names (Redundant-switch configuration)	Avid Unity ISIS New Names (Dual-connected configuration)	Avid Unity MediaNetwork New Names
Left on-board network interface	<p>Public</p> <p>This is a public network connected to network switch</p>	<p>Left-subnet number</p> <p>This is a public network connected to network switch.</p> <p>Use the subnet number of the interface. The examples in this document use Left-74</p>	<p>Public</p> <p>This is a public network connected to network switch</p>
Right on-board network interface	<p>Not used</p>	<p>Right-subnet number</p> <p>This is a public network connected to network switch.</p> <p>Use the subnet number of the interface. The examples in this document use Right-75.</p>	<p>Private</p> <p>This is a private network used for the heartbeat between the two servers in the cluster</p>
Left PCI adapter network interface	<p>Not used (Disabled)</p>	<p>Not used (Disabled)</p>	<p>Not used (Disabled)</p>
Right PCI adapter network interface.	<p>Private</p> <p>This is a private network used for the heartbeat between the two servers in the cluster.</p>	<p>Private</p> <p>This is a private network used for the heartbeat between the two servers in the cluster.</p>	<p>Not used (Disabled)</p>

- Repeat steps 2 and 3 for each network connection.

The following Network Connections window shows the new names used in an Avid Unity ISIS environment.



- Close the Network Connections window.

Configuring the Private Network Adapter on Each Node

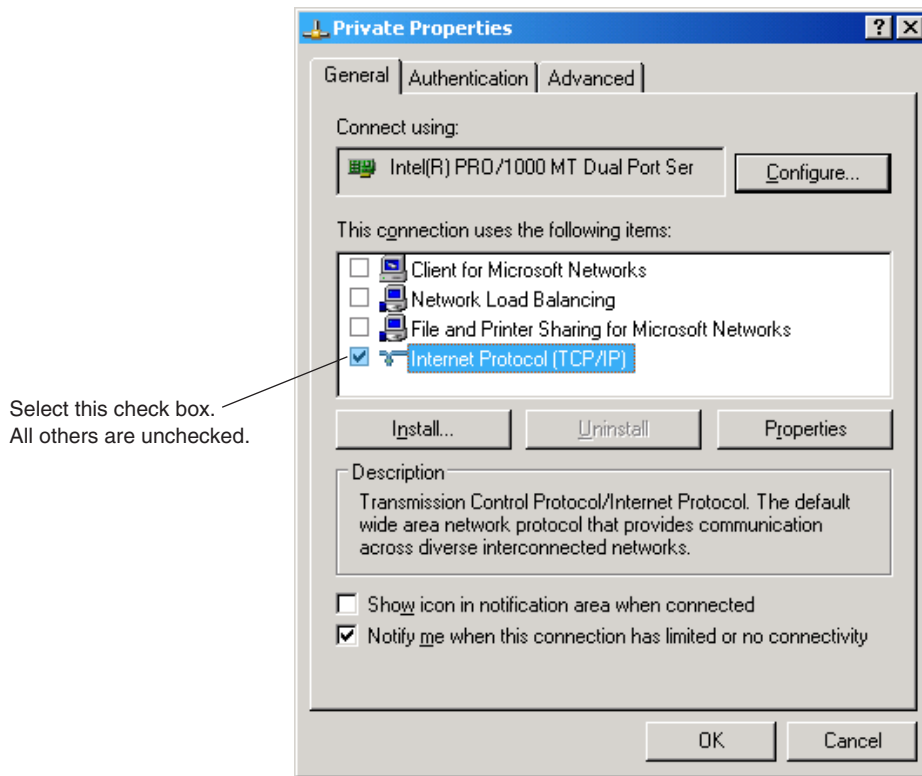
To configure the private network adapter for the heartbeat connection:

- Open the Network Connections window.
- Right-click the Private network connection and select Properties.

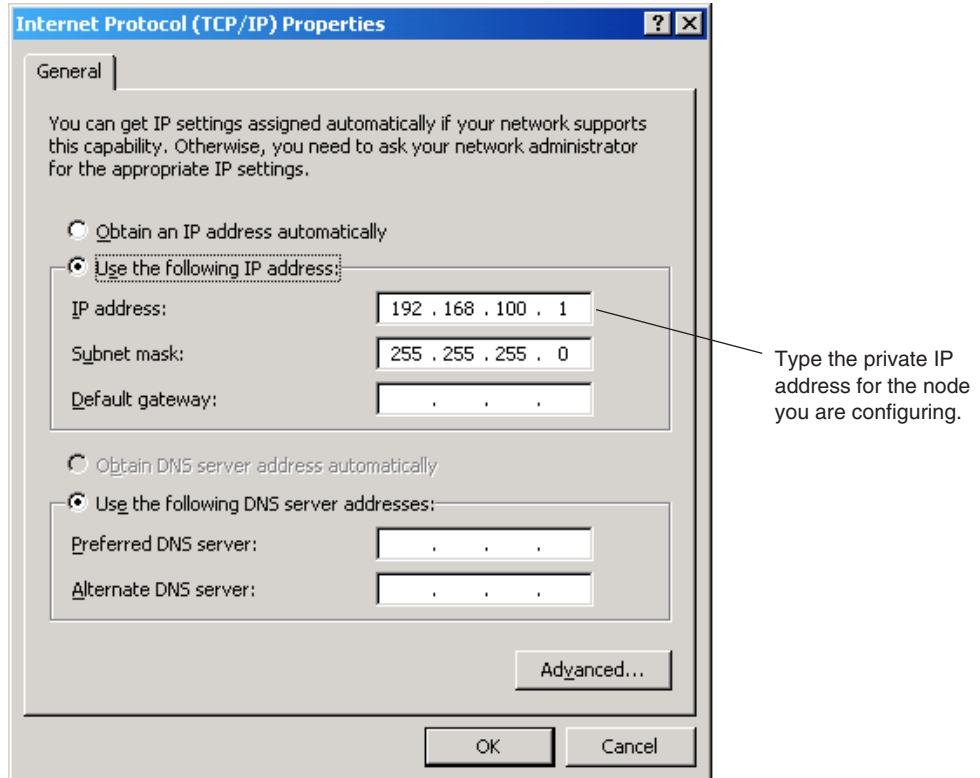
The Private Properties dialog box opens.

2 Creating a Microsoft Failover Cluster


3. On the General tab, click the Internet Protocol (TCP/IP) check box. Make sure all other components are unchecked.




4. Select Internet Protocol (TCP/IP) and click Properties.
The Internet Protocol (TCP/IP) Properties dialog box opens.



5. On the General tab of the Internet Protocol (TCP/IP) Properties dialog box:
 - a. Select “Use the following IP address.”
 - b. IP address: type the IP address for the Private network connection for the node you are configuring. See [“List of IP Addresses and Network Names”](#) on page 43.

 *When performing this procedure on the second node in the cluster make sure you use the static private IP addresses for that node. In this example, use 192. 168. 100. 2.*

- c. Subnet mask: type the subnet mask address


 *Make sure you use a completely different IP address scheme from the one used for the public network.*

- d. Make sure the “Default gateway” and “Use the Following DNS server addresses” text boxes are empty.

6. Click Advanced.

The Advanced TCP/IP Settings dialog box opens.

2 Creating a Microsoft Failover Cluster



7. On the DNS tab, make sure no values are defined and that the “Register this connection’s addresses in DNS” and “Use this connection’s DNS suffix in DNS registration” are not selected.
8. On the WINS tab, do the following:
 - ▶ Make sure no values are defined in the WINS addresses area.
 - ▶ Uncheck “Enable LMHOSTS Lookup”.
 - ▶ Select “Disable NetBIOS over TCP/IP.”
9. Click OK.

A message might be displayed stating “This connection has an empty primary WINS address. Do you want to continue?” Click Yes.

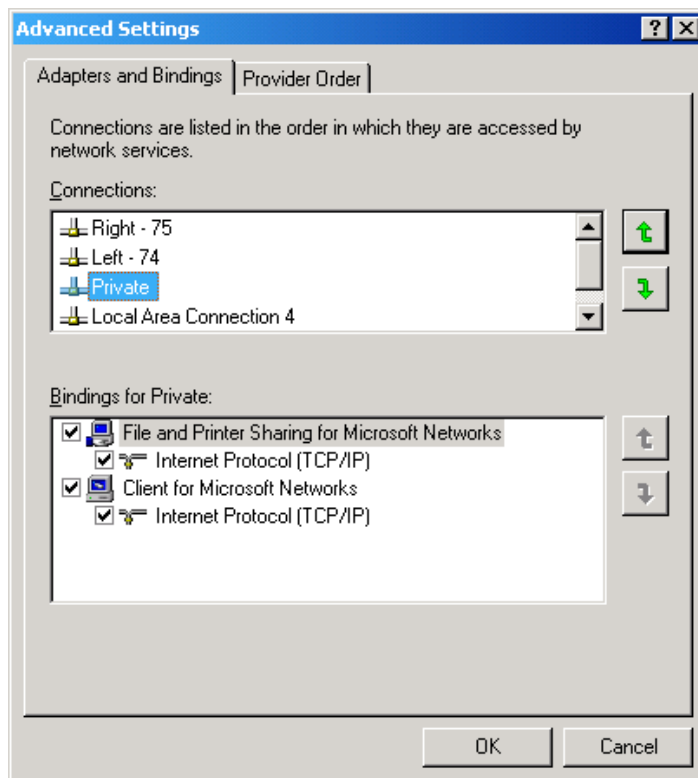
10. Repeat this procedure on the other node in the cluster, using the static private IP addresses for that node.

Configuring the Binding Order Networks on Each Node

Repeat this procedure on each node and make sure the configuration matches on both nodes.

To configure the binding order networks:

1. On one node, open the Network Connections window.
2. Select Advanced > Advanced Settings.
3. In the Connections area, use the arrow controls to position the network connections in the following order:
 - For a redundant-switch configuration in an Avid Unity ISIS environment, use the following order
 - Public
 - Private
 - For a dual-connected configuration in an Avid Unity ISIS environment, use the following order, as shown in the illustration:
 - Right
 - Left
 - Private
 - Local Area Connection 4



2 Creating a Microsoft Failover Cluster

- For an Avid Unity MediaNetwork environment use the following order:
 - Public
 - Private
- 4. Click OK.
- 5. Repeat this procedure on the other node and make sure the configuration matches on both nodes.

Configuring the Public Network Adapter on Each Node

Make sure you configure the IP address network interfaces for the Public Network Adapter as you normally would. For examples of public network settings, see [“List of IP Addresses and Network Names”](#) on page 43.

Joining Both Servers to the Active Directory Domain

After configuring the network information, join the two servers to the Active Directory domain. You can then use your domain credentials for the Cluster Installation Account (see [“Before You Begin the Server Failover Installation”](#) on page 42).

Configuring the Cluster Shared-Storage RAID Disks on Each Node

Both nodes must have the same configuration for the cluster shared-storage RAID disk. When you configure the disks on the second node, make sure the disks match the disk configuration you set up on the first node.

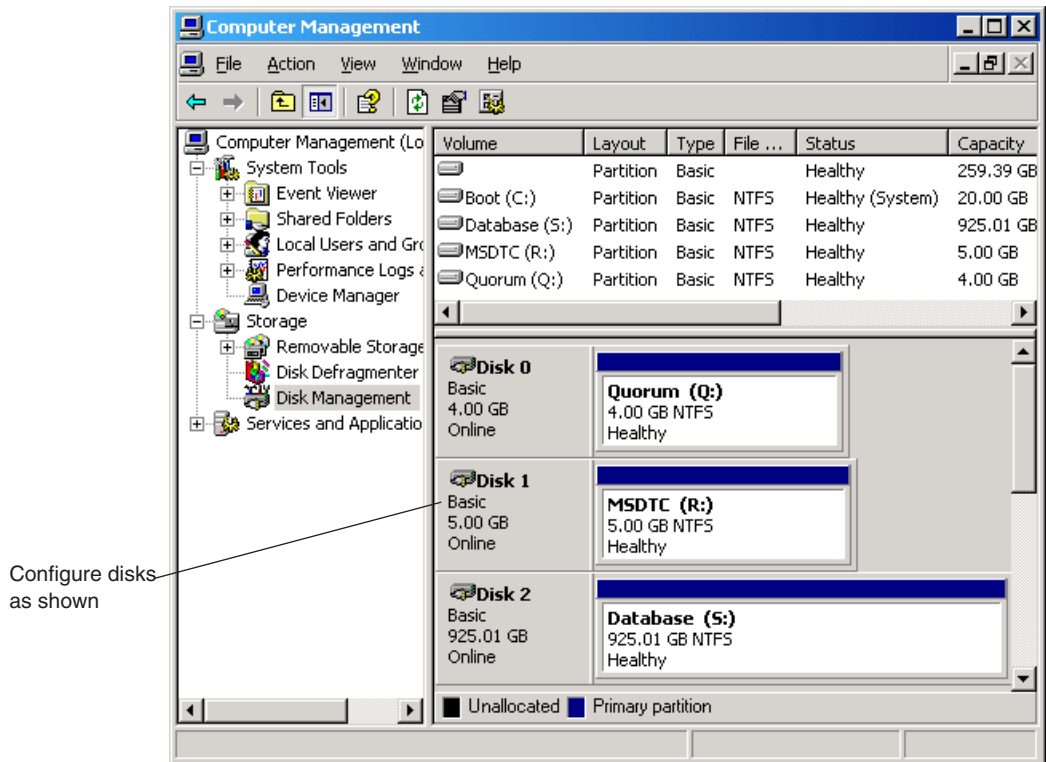


Before you create the partitions on the cluster nodes, make sure the cluster shared-storage RAID disks were pre-configured (mirror, stripe, etc.) by the vendor. Make sure the disks are Basic and not Dynamic.

To configure the disks on each node:

1. Shut down the server node you are not configuring at this time.
2. Open the Disk Management tool.
3. Initialize the disks, if not already initialized, by right-clicking the disk and selecting Initialize Disk.
4. Use Quick Format to configure the disks as partitions, using the following names and drive letters:
 - Quorum (Q:) 4GB
 - MSDTC (R:) 5GB
 - Database (S:) 925GB

The following illustration shows the required names and drive letters.



5. Verify you can access the disk and that it is working by creating a file and deleting it.
6. Shut down the first node and start the second node.
7. On the second node, assign drive letters and names. You do not need to format the disks.
 - a. Open the Disk Management tool. Right-click the partition, select Change Drive Letter, and enter the appropriate letter. Repeat these actions for the other partitions.
 - b. Open My Computer. Select a drive, right-click, select Rename, and enter the appropriate name. Repeat these actions for the other drives.

Configuring the Cluster Service

Take the following steps to configure the cluster service:

1. Turn off the second node.
2. Configure the first node using the New Server Cluster Wizard. See [“Configuring the Cluster Service on the First Node”](#) on page 62

2 Creating a Microsoft Failover Cluster

3. Validate the cluster service installation on the first node. See [“Validating the Cluster Service on the First Node”](#) on page 67.
4. Turn on the second node. Leave first node turned on.
5. Configure the second node using Add Cluster Computers Wizard. See [“Configuring the Cluster Service on the Second Node”](#) on page 67.

Configuring the Cluster Service on the First Node

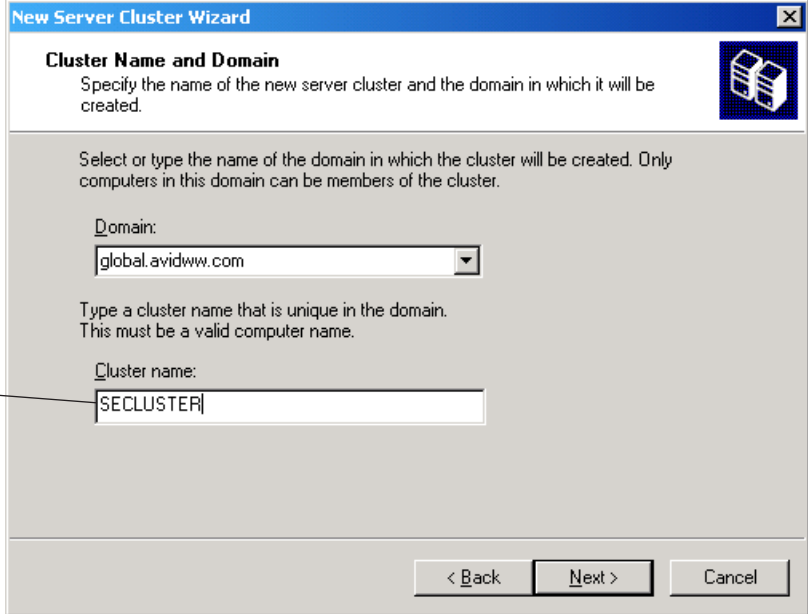
To configure the cluster service on the first node:

1. Turn off the server for the node you are not configuring at this time.
2. Make sure all storage devices are turned on.
3. Click Start and select All Programs > Administrative Tools > Cluster Administrator.

The Open Connection to Cluster dialog box opens.



4. Select “Create new cluster” from the Action menu.
5. Make sure you have the prerequisites to configure the cluster, as shown in the New Server Cluster Wizard Welcome window.
6. Click Next.
7. In the Cluster Name and Domain dialog box, do the following:
 - Domain: select the name of your Active Directory domain
 - Cluster name: type the Cluster service name, for example SECLUSTER — see [“List of IP Addresses and Network Names”](#) on page 43.



The screenshot shows a Windows dialog box titled "New Server Cluster Wizard" with a sub-header "Cluster Name and Domain". The main text reads: "Specify the name of the new server cluster and the domain in which it will be created." Below this, there is a dropdown menu for "Domain:" with "global.avidww.com" selected. The next section says: "Select or type the name of the domain in which the cluster will be created. Only computers in this domain can be members of the cluster." Below that is a text box for "Cluster name:" containing "SECLUSTER". An arrow points from the text "Type the Cluster service name." to this text box. At the bottom are buttons for "< Back", "Next >", and "Cancel".

Cluster Name and Domain

Specify the name of the new server cluster and the domain in which it will be created.

Select or type the name of the domain in which the cluster will be created. Only computers in this domain can be members of the cluster.

Domain:
global.avidww.com

Type a cluster name that is unique in the domain.
This must be a valid computer name.

Cluster name:
SECLUSTER

Type the Cluster service name.

< Back Next > Cancel

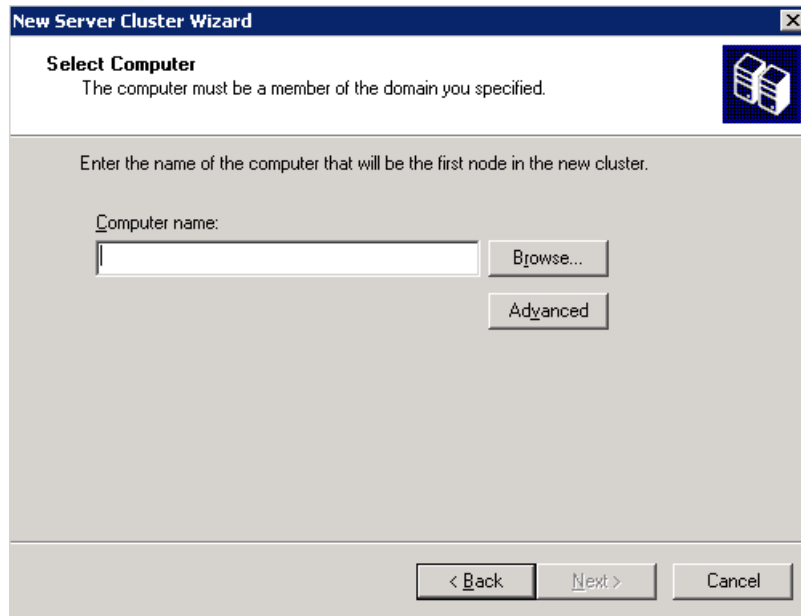
8. Click Next.

The Select Computer dialog box opens.



You might be prompted for an account. If so, use a domain user account, such as the Cluster Installation Account referred to in "Before You Begin the Server Failover Installation" on page 42. Do not use the Cluster Service Account (Service Execution User).

2 Creating a Microsoft Failover Cluster

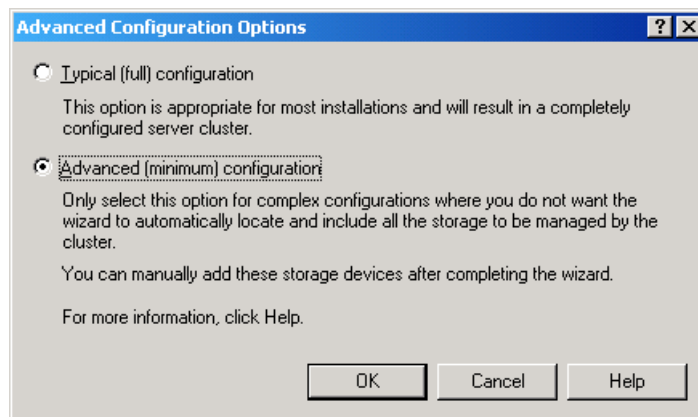


9. In the Select Computer dialog box, in the Computer name text box, type the Cluster node host name of the first node.

For example, use SECLUSTER1. See [“List of IP Addresses and Network Names” on page 43](#).

10. Click Advanced.

The Advanced Configuration Options dialog box opens.



11. Select Advanced (minimum) configuration, and click OK.

12. Click Next.

The setup process analyzes the node for hardware or software problems that might cause problems during installation. A warning icon displays next to “Checking Cluster feasibility.” In this case, the warnings do not indicate a problem.

13. Click Next after the analyze is complete and the Task Complete bar is green.
14. In the IP Address dialog box, type the Cluster Service IP address (virtual IP address) in the IP Address text box. See [“List of IP Addresses and Network Names” on page 43](#). (Do not type the MSDTC service IP address or the Interplay Engine service IP address.)
15. Click Next.
16. In the Cluster Service Account dialog box, type the cluster user name and password, and select the domain.

This is the Cluster Service Account (Server Execution User) used to start the cluster service. It is also used by the Interplay Engine. It must be a unique name that will not be used for any other purpose. See [“Before You Begin the Server Failover Installation” on page 42](#). Check that the account is part of the domain, and that the name and password are correct, by logging into the domain.

Type the Cluster Service Account user name.

New Server Cluster Wizard

Cluster Service Account
Enter login information for the domain account under which the cluster service will be run.

User name: sqausef

Password: ●●●●●●●●

Domain: global.avidww.com

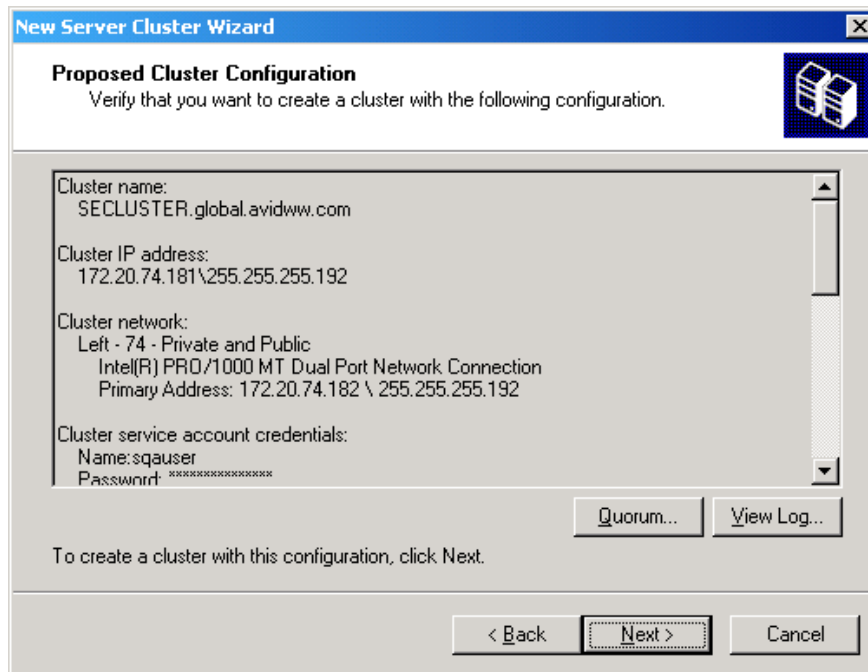
This account will be given local administrative rights on all nodes of this cluster to allow for proper operation.

< Back Next > Cancel

2 Creating a Microsoft Failover Cluster


17. Click Next.

The Proposed Cluster Configuration dialog box opens.



18. Click Quorum.

The Cluster Configuration Quorum dialog box opens.



19. Select Disk Q: from the menu, and click OK.
20. Review the summary on the Proposed Cluster Configuration dialog box to verify all the information for creating the cluster is correct.
21. Click Next. The Creating the Cluster dialog box opens.
22. Review any errors during the cluster creation.



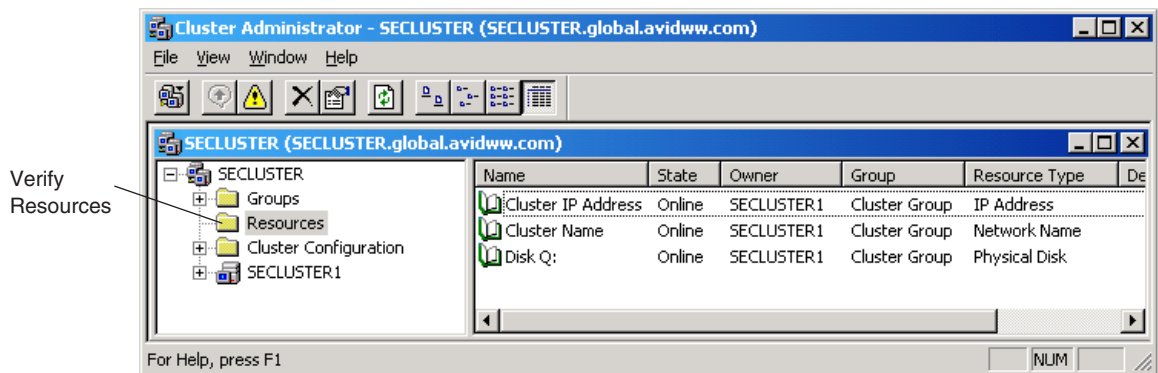
If red errors display, check the Cluster Service ISIS IP address you entered in step 14.

23. Click Next.
24. Click Finish.

Validating the Cluster Service on the First Node

To validate the first node cluster installation:

1. Click Start and select Programs > Administrative Tools > Cluster Administrator.
2. In the left pane, click Resources to make sure all resources are online.



Configuring the Cluster Service on the Second Node

To configure the cluster service on the second node:

1. Make sure the first node is on and all storage devices are turned on.
2. Turn on the server for the second node.
3. In the first node, click Start and select Programs > Administrative Tools > Cluster Administrator.
4. Select File > New > Node.
The Add Node Wizard opens.
5. Click Next.

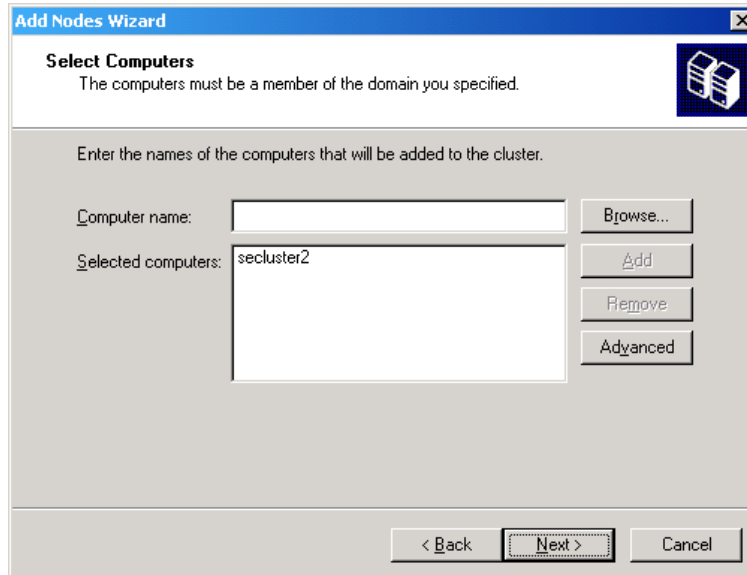


You might be prompted for an account. If so, use a domain user account, such as the Cluster Installation Account referred to in “Before You Begin the Server Failover Installation” on page 42. Do not use the Cluster Service Account (Service Execution User).

2 Creating a Microsoft Failover Cluster

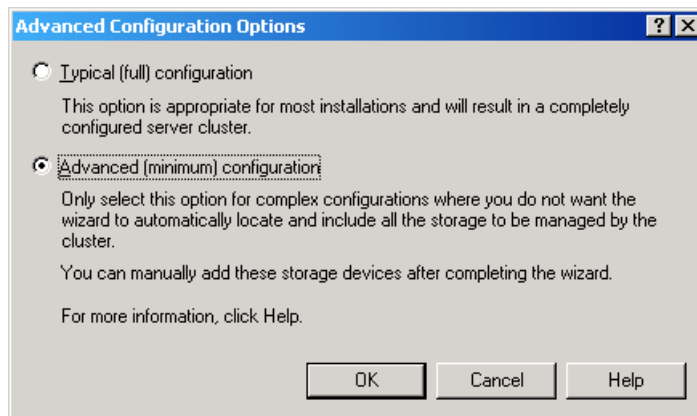
6. In the Select Computers dialog box, in the Computer name text box, type the Cluster node host name of the second node and click Add.

For example, use SECLUSTER2. See “List of IP Addresses and Network Names” on page 43.



7. Click Advanced.

The Advanced Configuration Options dialog box opens.



8. Select Advanced (minimum) configuration, and click OK.

9. Click Next.

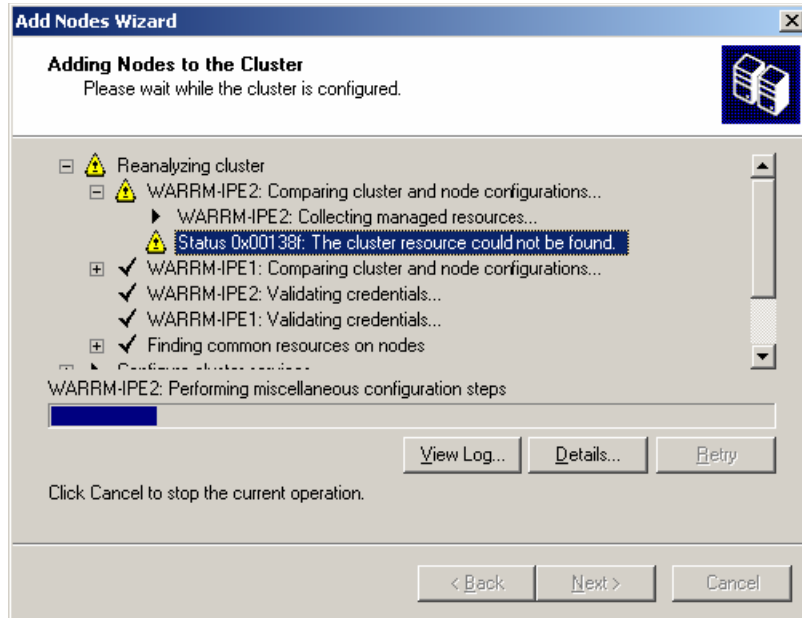
The setup process analyzes the node for hardware or software problems that might cause problems during installation. A warning icon displays next to “Checking Cluster feasibility.” In this case, the warnings do not indicate a problem.

10. Click Next after the analyze is complete and the Task Complete bar is green.
11. Type the password for the cluster service account. This account is used to start the cluster service.
12. Click Next.
13. In the Proposed Cluster Configuration dialog box, review the summary to verify all the information for creating the cluster is correct.
14. Click Next.

The Adding Nodes to the Cluster dialog box opens.

15. Review any errors during the cluster creation.

A warning icon displays next to “Reanalyzing cluster.” In this case, the warnings do not indicate a problem.



16. Click Next.
17. Click Finish.

Configuring Rules for the Cluster Networks

After the networks are configured on each node and the cluster service is configured, you need to configure the network roles to determine the function within the cluster.

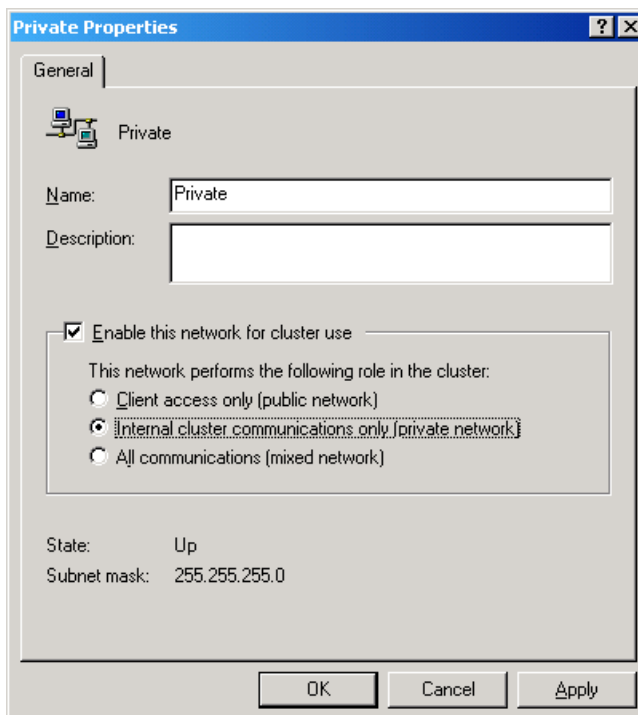


This procedure uses Left-74 and Right-75 as examples of the public networks. If you are installing a dual-connection configuration, replace the numbers with your subnet numbers.

To configure the rules for the cluster networks:

1. Click Start and select Programs > Administrative Tools > Cluster Administrator.
2. In the left pane, click Cluster Configuration > Networks, and right-click Private and select Properties.
3. Select “Internal cluster communications only (private network).”

The Private network (virtual cluster) is used for the Heartbeat.



4. Click OK.
5. In the left pane, click Cluster Configuration > Networks, and right-click either Public or Left-74 and select Properties.

6. In the Public or Left-74 Properties dialog box, verify these options:
 - Name: Left-74
 - Enable this network for cluster use
 - All communications (mixed network)
7. Click OK.
8. If you are installing a dual-connection configuration, in the left pane, click Cluster Configuration > Networks, and right-click Right-75 and select Properties.
9. In the Right-75 Properties dialog box, verify these options:
 - Name: Right-75
 - Enable this network for cluster use
 - All communications (mixed network)
10. Click OK.

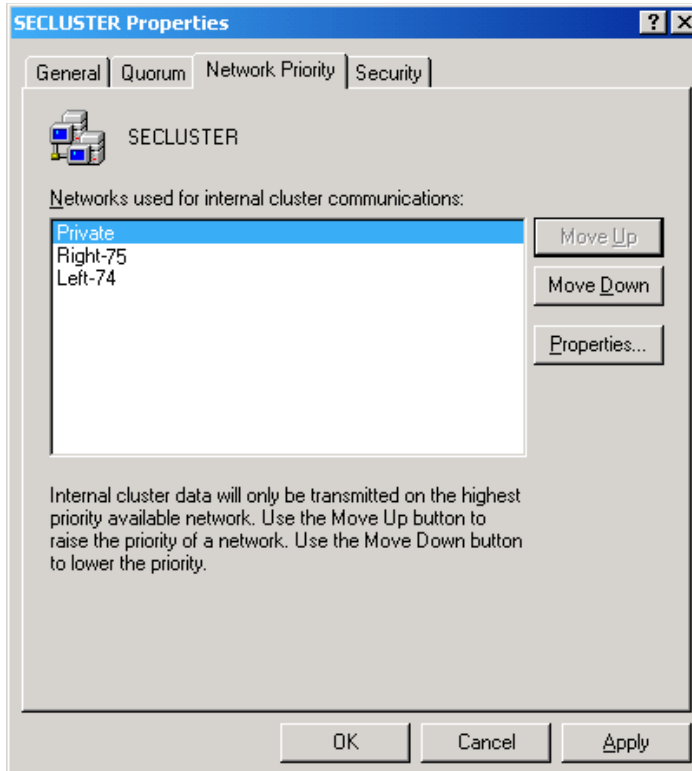
Prioritizing the Heartbeat Adapter

After you configure network roles for how the cluster service uses the network adapter, you need to prioritize the order in which they are used for intra-cluster communications. The cluster service will use the next network adapter in the list when it cannot communicate by using the first network adapter.

To prioritize the heartbeat adapter:

1. Click Start and select Programs > Administrative Tools > Cluster Administrator.
2. In the left pane, right-click the cluster name at the top of the list and select Properties.
3. Click the Network Priority tab.

2 Creating a Microsoft Failover Cluster



4. Verify the Private network is at the top of the list. You can use the Move Up and Move Down buttons to change the priority order.
5. Click OK.

After Setting Up the Cluster

After you finish setting up the cluster you need to verify that the quorum disk is using disk Q, set the startup times for each node, and test the cluster installation.

The following sections provide procedures for these tasks:

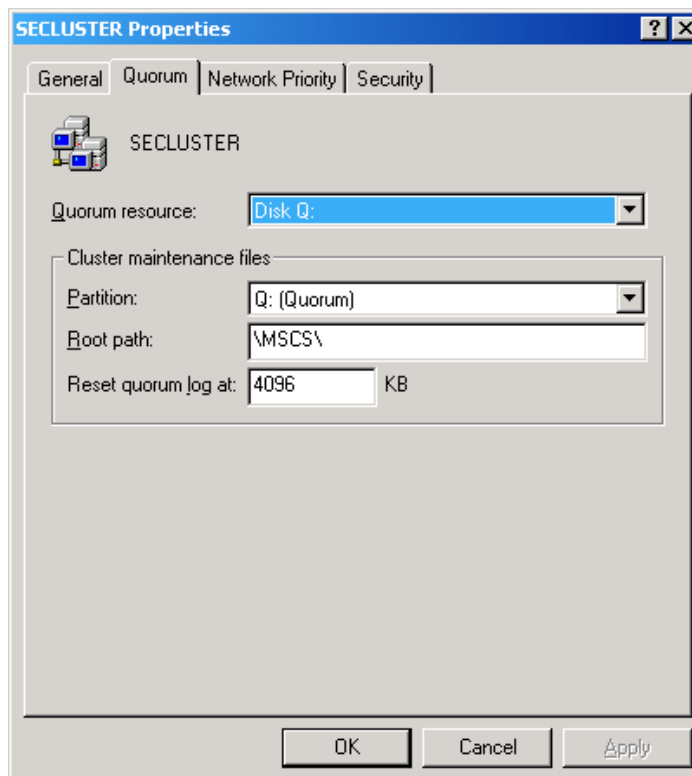
- [“Verifying the Quorum Disk” on page 73](#)
- [“Setting the Startup Times on Each Node” on page 73](#)
- [“Testing the Cluster Installation” on page 75](#)

Verifying the Quorum Disk

The Cluster Configuration Wizard automatically selects the disk used as the quorum device. Check to make sure the quorum device is using disk Q.

To verify the quorum disk:

1. On either node, click Start and select Programs > Administrative Tools > Cluster Administrator.
2. In the left pane, right-click the cluster name at the top of the list and select Properties.
3. Click the Quorum tab and make sure Quorum resource displays Disk Q.



4. Click OK.

Setting the Startup Times on Each Node

Avid recommends that you offset the startup time of each node's operating system used during the power up of the cluster.

2 Creating a Microsoft Failover Cluster

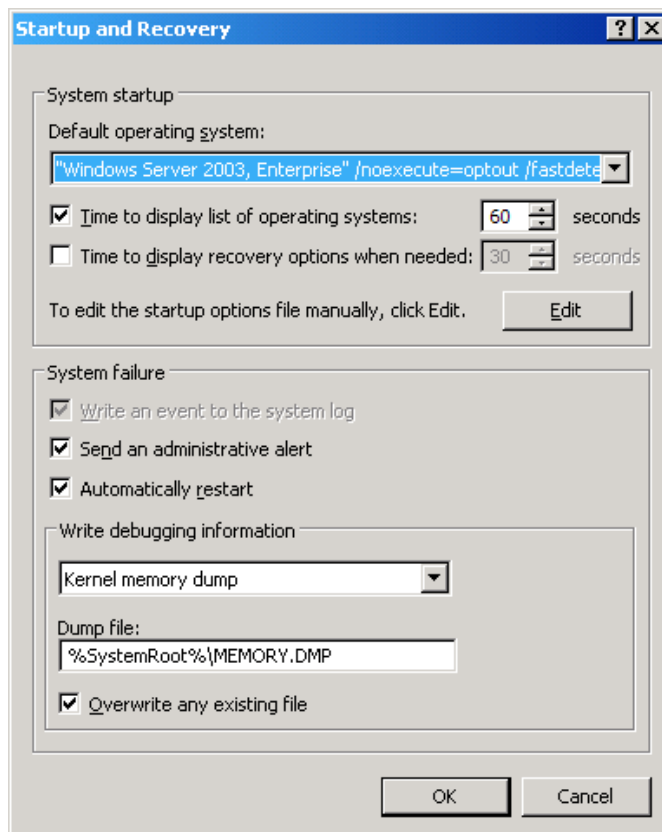


This setting should display the value you set in the boot.ini file (see “Increasing the Boot Delay” on page 49). Change this setting on the second node.

To set the time for displaying the list of operating systems:

1. On the first node, click Start and right-click My Computer and select Properties.
2. Click the Advanced tab.
3. In the Startup And Recovery area, click Settings.

The Setup and Recovery dialog box opens.



4. Select “Time to display list of operating systems.”
5. Make sure the time is set to 60 seconds.
6. Click OK.
7. Repeat this procedure on the second node, but set the time to 120 seconds.

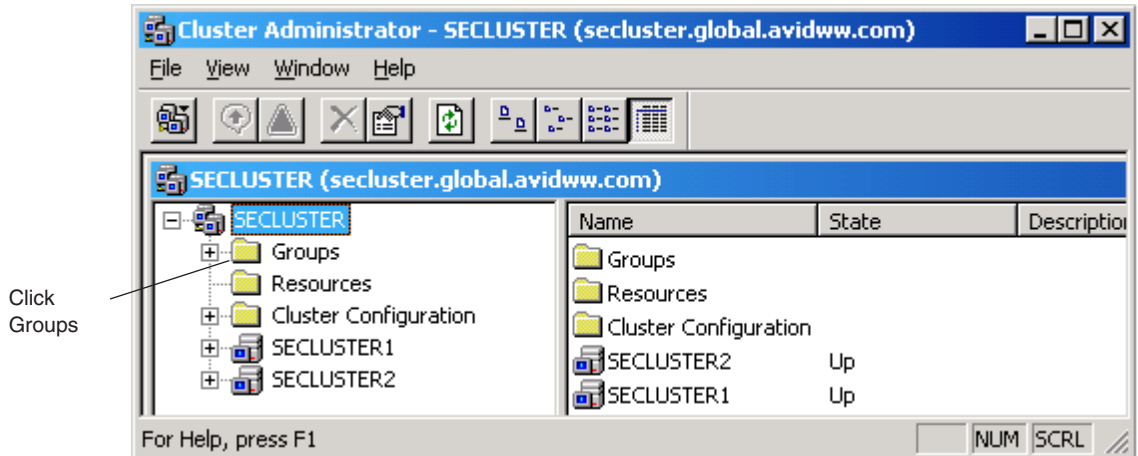
Testing the Cluster Installation

You must test the cluster installation to make sure the failover process is working.

To verify that resources will failover:

1. Click Start, and select Programs > Administrative Tools > Cluster Administrator.

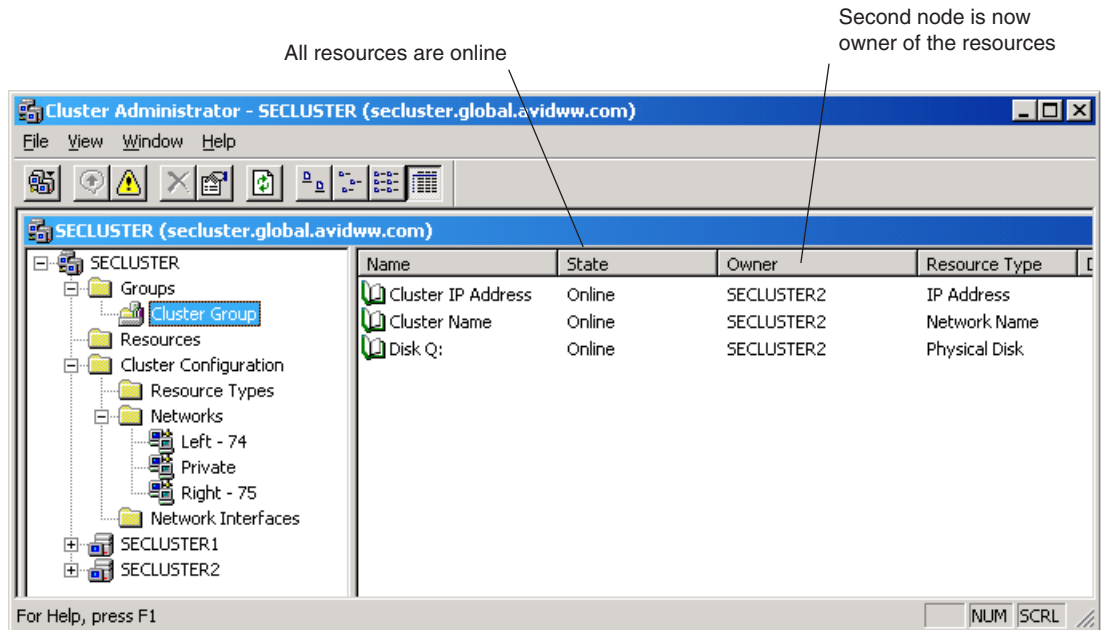
The Cluster Administrator window opens.



2. In the left pane, open the Groups folder, right-click Cluster Group, and select Move Group.

The group and all its resources are moved to the other node. Disk Q is brought online on the second node. Make sure the window displays that the second node is now the owner of the Resources and that all resources are online.

2 Creating a Microsoft Failover Cluster



3. Move the group back to node 1 after you finish testing the cluster installation.
4. Close the Cluster Administrator.

Configuration of the cluster service on all nodes is complete and the cluster is fully operational. You can now install cluster resources, such as file shares, cluster aware services such as Distributed Transaction Coordinator.

Installing the Distributed Transaction Coordinator

Interplay Engine requires DCOM services in the cluster. To allow DCOM services in the cluster, create a resource group for the Distributed Transaction Coordinator. This resource group needs its own physical 5GB disk, an IP address and a network name (MSDTC). Finish the group by adding a resource of the Distributed Transaction Coordinator type.

The following sections provide procedures for creating a resource group for the Distributed Transaction Coordinator by using the Cluster Administrator tool.

- [“Creating a Resource Group for the Distributed Transaction Coordinator” on page 77](#)
- [“Assigning an IP Address to the MSDTC Group” on page 78](#)
- [“Assigning a Network Name to the MSDTC Group” on page 79](#)

- “Creating a Physical Resource for the MSDTC Group” on page 80
- “Assigning Distributed Transaction Coordinator Resource to the MSDTC Group” on page 81

When performing these procedures Avid suggests you use the same entries shown in the procedure. These entries are from the list in section “List of IP Addresses and Network Names” on page 43.

For more information about Distributed Transaction Coordinator, see the Microsoft Knowledge Base article addressing this topic (301600):

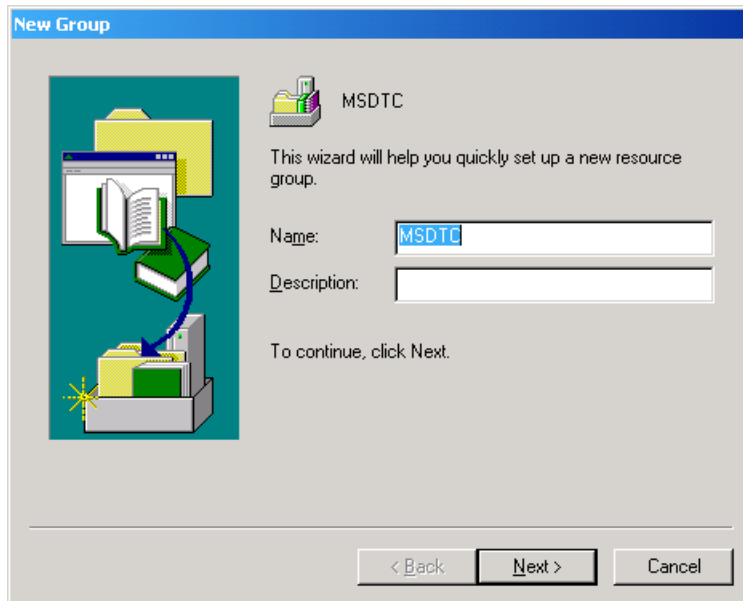
<http://support.microsoft.com/default.aspx?scid=kb;en-us;301600>

Creating a Resource Group for the Distributed Transaction Coordinator

To create a resource group named MSDTC for the Distributed Transaction Coordinator:

1. Click Start and select Programs > Administrative Tools > Cluster Administrator.
2. Select File > New > Group.

The New Group Wizard opens.



2 Creating a Microsoft Failover Cluster

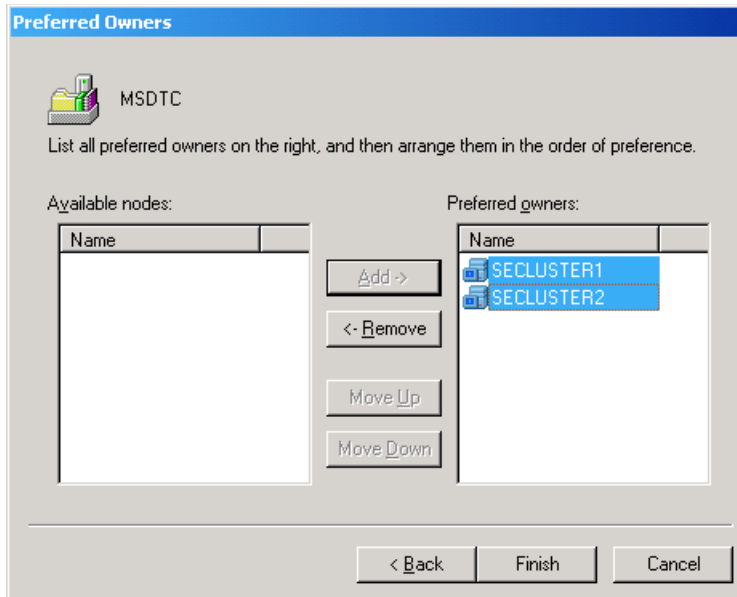
3. In the Name text box, type MSDTC.

You can use any name for the group name, however Avid suggests you use MSDTC.

4. Click Next.

The Preferred Owners dialog box opens.

5. Select both owners in the Available nodes list and add them to the Preferred owners list.



6. Click Finish.

The group is created.

Assigning an IP Address to the MSDTC Group

To assign an IP address to MSDTC group:

1. In the Cluster Administrator, right-click MSDTC group and select New > Resource.
2. Complete the New Resource dialog box as follows:
 - Name: MSDTC IP
 - Resource Type: IP Address
 - Group: MSDTC

3. Complete the Possible Owners dialog box as follows:
 - Add the cluster server host names to the Possible owners lists. For example, SECLUSTER1 and SECLUSTER2. See [“List of IP Addresses and Network Names” on page 43](#).
4. Complete the Dependencies dialog box as follows:
 - Leave the Resource dependencies list empty.
5. Complete TCP/IP Address Parameters dialog box as follows:
 - Address: type the IP address of the MSDTC service. See [“List of IP Addresses and Network Names” on page 43](#).
 - Subnet mask: displays subnet for the network
 - Network: select the default network connection: *Right-subnet* or *Left-subnet* (for ISIS), or *Public* (for MediaNetwork).
 - Select Enable NetBIOS for this address
6. Click Finish.

Assigning a Network Name to the MSDTC Group

To assign a network name to MSDTC group:

1. In the Cluster Administrator, right-click MSDTC group and select New > Resource.
2. Complete the New Resource dialog box as follows:
 - Name: MSDTC NAME
 - Resource Type: Network Name
 - Group: MSDTC
3. Click Next.
4. Complete the Possible Owners dialog box as follows:
 - Add the cluster server host names to the Possible owners lists. For example, SECLUSTER1 and SECLUSTER2. See [“List of IP Addresses and Network Names” on page 43](#).
5. Click Next.
6. Complete the Dependencies dialog box as follows:
 - Add MSDTC IP to the Resource Dependencies list.
7. Click Next.

2 Creating a Microsoft Failover Cluster

8. Complete Network Name Parameters dialog box as follows:
 - Name: Type the virtual host name for the MSDTC group, for example, CLUSTERMSDTC. See [“List of IP Addresses and Network Names” on page 43](#). Make sure to use a unique name for each Interplay Engine cluster on the network.
 - Uncheck “DNS Registration Must Succeed.”
 - Check “Enable Kerberos Authentication” unless you are using Windows NT 4 domain functionality.

If you are using Kerberos authentication, make sure the Kerberos time is in sync with the Active Directory controller (plus or minus five minutes) or the authentication will fail.
9. Click Finish.

Creating a Physical Resource for the MSDTC Group

To create a physical disk resource for MSDTC group:

1. In the Cluster Administrator, right-click MSDTC group and select New > Resource.
2. Complete the New Resource dialog box as follows:
 - Name: MSDTC DISK R
 - Resource Type: Physical disk
 - Group: MSDTC
3. Click Next.
4. Complete the Possible Owners dialog box as follows:
 - Add the cluster server host names to the Possible owners lists. For example, SECLUSTER1 and SECLUSTER2. See [“List of IP Addresses and Network Names” on page 43](#).
5. Click Next.
6. Complete the Dependencies dialog box as follows:
 - Leave the Resource dependencies list empty.
7. Click Next.
8. Complete Disk Parameters dialog box as follows:
 - Select: R: (MSDTC)
9. Click Finish.

Assigning Distributed Transaction Coordinator Resource to the MSDTC Group

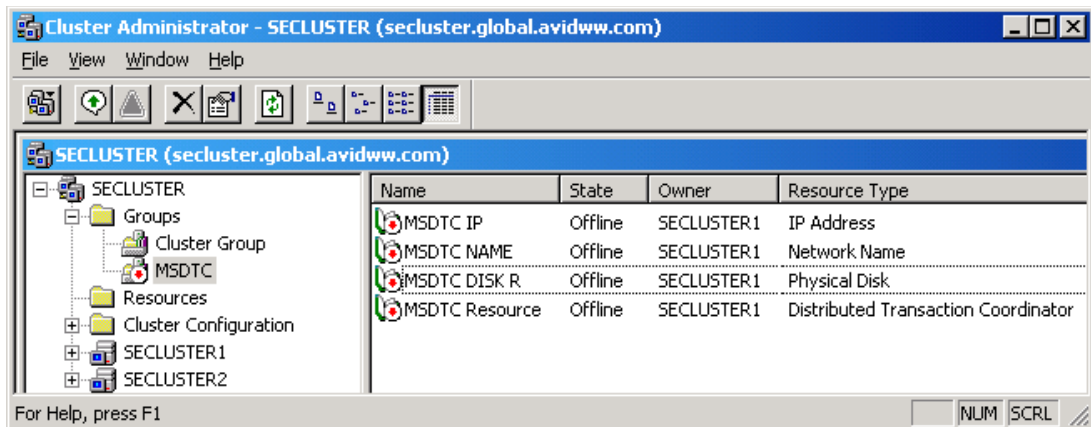
To assign Distributed Transaction Coordinator Resource to MSDTC group:

1. In the Cluster Administrator, right-click MSDTC group and select New > Resource.
2. Complete the New Resource dialog box as follows:
 - Name: MSDTC Resource
 - Resource Type: Distributed Transaction Coordinator
 - Group: MSDTC
3. Click Next.
4. Complete the Possible Owners dialog box as follows:
 - Add the cluster server host names to the Possible owners lists. For example, SECLUSTER1 and SECLUSTER2. See [“List of IP Addresses and Network Names”](#) on page 43.
5. Click Next.
6. Complete the Dependencies dialog box as follows:
 - Add MSDTC DISK R and MSDTC NAME to the Resource dependencies list.
7. Click Finish.

2 Creating a Microsoft Failover Cluster

Bringing the MSDTC Online

The following illustration shows the Cluster Administrator after you complete the setup of the MSDTC group.



To bring the MSDTC online:

1. Initialize the MSDTC Log file by doing the following:
 - a. Bring MSDTC DISK R online: right-click MSDTC DISK R and select Bring Online.
 - b. In the Command Window, run the following command on the node that is the owner to reset the log: `msdtc -resetlog`
2. Bring MSDTC Group online by right-clicking MSDTC, and selecting Bring Online.



If you are running Active Directory on the cluster nodes, the MSDTC Resource might fail to run on the backup domain controller. If this occurs, see the following Microsoft article:

<http://support.microsoft.com/kb/900216/en-us>.

3 Installing the Interplay Engine for a Failover Cluster

After you set up and configure the cluster, you need to install the Interplay Engine software on both nodes. The following topics describe installing the Interplay Engine and other final tasks:

- [Disabling Any Web Servers](#)
- [Installing the Interplay Engine on the First Node](#)
- [Installing the Interplay Engine on the Second Node](#)
- [Bringing the Interplay Engine Online](#)
- [Testing the Complete Installation](#)
- [Updating a Clustered Installation \(Rolling Upgrade\)](#)
- [Uninstalling the Interplay Engine on a Clustered System](#)

Disabling Any Web Servers

The Interplay Engine uses an Apache web server that can only be registered as a service if no other web server (for example, IIS) is serving the port 80 (or 443). Stop and disable or uninstall any other http services before you start the installation of the server. You must perform this procedure on both nodes.



If you followed the procedures in this document no action is required, since the only web server installed at this point is the IIS and it is disabled.

Installing the Interplay Engine on the First Node

The following sections provide procedures for installing the Interplay Engine on the first node. For a list of example entries, see [“List of IP Addresses and Network Names” on page 43](#).

- [“Preparation for Installing on the First Node” on page 84](#)
- [“Starting the Installation and Accepting the License Agreement” on page 85](#)
- [“Installing the Interplay Engine Using Custom Mode” on page 85](#)
- [“Bringing the Disk Resource Online” on page 99](#)



Shut down the second node while installing Interplay Engine for the first time.

Preparation for Installing on the First Node

You are ready to start installing the Interplay Engine on the first node. During setup you must enter the following cluster-related information:

- **Virtual IP Address:** the Interplay Engine service IP address of the resource group. For a list of example names, see [“List of IP Addresses and Network Names” on page 43](#).
- **Subnet Mask:** the subnet mask on the local network.
- **Public Network:** the name of the public network connection.
 - For a redundant-switch ISIS configuration, type **Public**.
 - For a dual-connection ISIS configuration, type **Left-subnet**. For a dual-connection configuration, you set the other public network connection after the installation. See [“Bringing the Disk Resource Online” on page 99](#).
 - For a MediaNetwork configuration, type **Public**.

To check the public network connection on the first node, open the Network Connections panel in the Windows Control Panel and look up the name there.

- **Shared Drive:** the letter for the shared drive that holds the database. Use S: for the shared drive letter.
- **Cluster Service Account User and Password (Server Execution User):** the domain account that is used to run the cluster. See [“Before You Begin the Server Failover Installation” on page 42](#).



Shut down the second node while installing Interplay Engine for the first time.



When installing the Interplay Engine for the first time on a machine with cluster services, you are asked to choose between clustered and regular installation. The installation on the second node (or later updates) reuses the configuration from the first installation without allowing you to change the cluster-specific settings. In other words, it is not possible to change the configuration settings without uninstalling the Interplay Engine.

Starting the Installation and Accepting the License Agreement

To start the installation:

1. Insert the Avid Interplay installation DVD.
A start screen opens.
2. Double-click Install Avid Interplay Engine to begin the Avid Interplay Engine Installation Wizard, which guides you through the installation.
The Welcome dialog box opens.
3. Close all Windows programs before proceeding with the installation.
4. Information about the installation of Apache is provided in the Welcome dialog box. Read the text and then click Next.
The License Agreement dialog box opens.
5. Read the license agreement information and then accept the license agreement by selecting “I accept the agreement.” Click Next.
The Specify Installation Type dialog box opens.
6. Continue the installation as described in the next topic.

Installing the Interplay Engine Using Custom Mode

The first time you install the Interplay Engine on a cluster system, you should use the Custom installation mode. This lets you specify all the available options for the installation. This is the recommended option to use.

The following procedures are used to perform a Custom installation of the Interplay Engine:

- [“Specifying Cluster Mode During a Custom Installation” on page 86](#)
- [“Specifying the Interplay Engine Details” on page 87](#)
- [“Specifying the Interplay Engine Service Name” on page 89](#)
- [“Specifying the Destination Location” on page 90](#)
- [“Specifying the Default Database Folder” on page 90](#)
- [“Specifying the Share Name” on page 91](#)
- [“Specifying the Configuration Server” on page 92](#)

3 Installing the Interplay Engine for a Failover Cluster

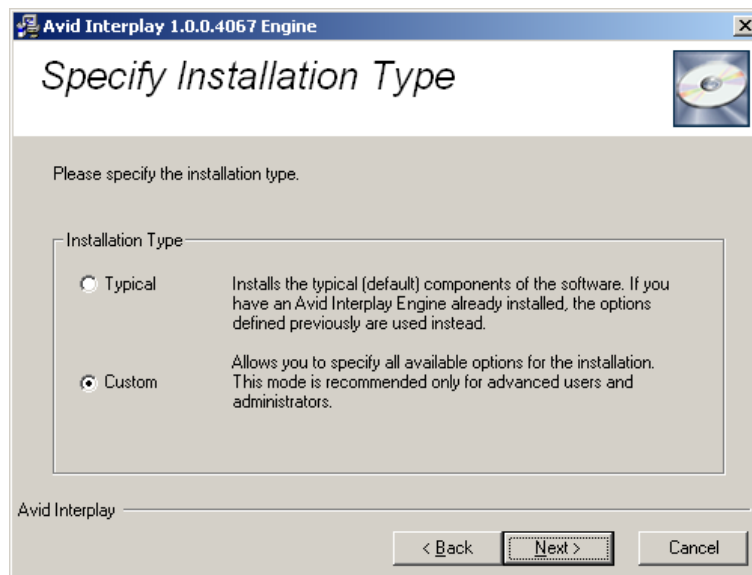
- “Specifying the Server User” on page 94
- “Specifying the Server Cache” on page 95
- “Enabling Email Notifications” on page 96
- “Installing the Interplay Engine for a Custom Installation on the First Node” on page 98

For information about updating the installation, see “Updating a Clustered Installation (Rolling Upgrade)” on page 106.

Specifying Cluster Mode During a Custom Installation

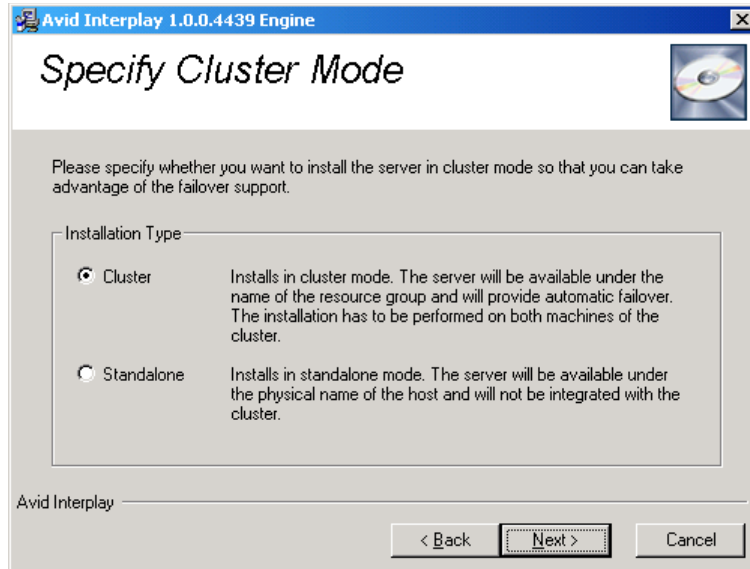
To specify cluster mode:

1. In the Specify Installation Type dialog box, select Custom.



2. Click Next.

The Specify Cluster Mode dialog box opens.



3. Select Cluster and click Next to continue the installation in cluster mode.
The Specify Interplay Engine Details dialog box opens.

Specifying the Interplay Engine Details

In this dialog box, provide details about the Interplay Engine.

3 Installing the Interplay Engine for a Failover Cluster

Avid Interplay Engine

Specify Interplay Engine Details

Setup will create the resource group with the parameters specified here. The virtual IP address is going to be used for the resource group of the server. The drive must be a shared resource that is mounted using the same drive letter on the other machine.

Interplay Engine Information

Virtual IP Address:
not the Cluster IP address: 172.20.74.180

Subnet Mask: 255.255.255.192

Public Network: Left-74

Shared Drive: S:

Avid Interplay

< Back Next > Cancel

To specify the Interplay Engine details:

1. Type the following values:

- Virtual IP address: This is the Interplay Engine service IP Address, *not* the Cluster service IP address. For a list of example names, see [“List of IP Addresses and Network Names”](#) on page 43.
- Subnet Mask: The subnet mask on the local network.
- Public Network: For a redundant-switch ISIS configuration or MediaNetwork configuration, type Public. For a dual-connected ISIS configuration, type *Left-subnet*. For a dual-connected configuration, you set the other public network connection after the installation. See [“Bringing the Disk Resource Online”](#) on page 99. For MediaNetwork, type Public.

To check the public network connection on the first node, open the Network Connections panel in the Windows Control Panel and look up the name there.

- Shared Drive: The letter of the shared drive that is used to store the database. Use S: for the shared drive letter.



Make sure you type the correct information here, as this data cannot be changed afterwards. Should you require any changes to the above values later, you will need to uninstall the server on both nodes.

2. Click Next.

The Specify Interplay Engine Name dialog box opens.

Specifying the Interplay Engine Service Name

In this dialog box, type the name of the Interplay Engine service.

Avid Interplay Engine

Specify Interplay Engine Name

Please specify the names that are going to be associated with the Avid Interplay Engine.

The specified Network Name will be used to register the clustered installation in the DNS. This should be a new, unused name and not the name of any machine.

The Server Name field can be used to enter a static IP address if you use Avid Interplay Clients on MAC or Unix systems and these can not resolve the Network Name.

Interplay Engine Name

Network Name:
not the MS Cluster name

Server Name:

seengine

seengine

Avid Interplay

< Back Next > Cancel

To specify the Interplay Engine name:

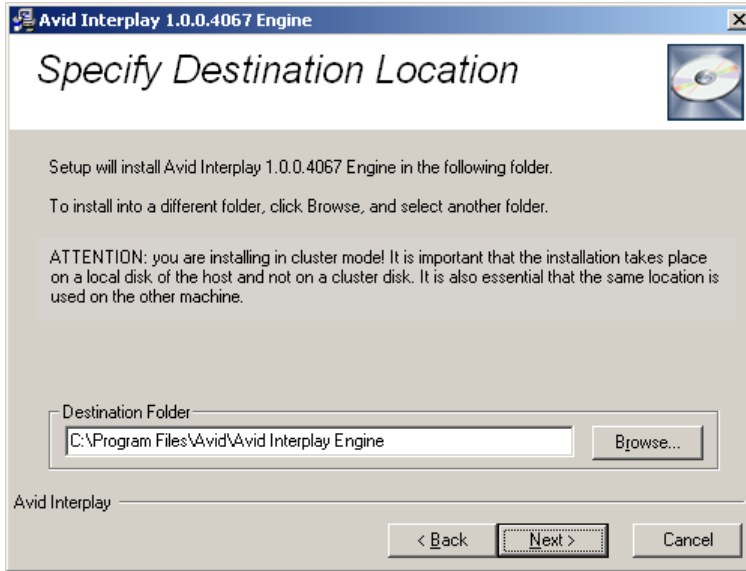
- Specify the public names for the Avid Interplay Engine service by typing the following values:
 - The Network Name will be associated with the virtual IP Address that you entered in the previous Interplay Engine Details dialog box. This is the Interplay Engine service name (see [“List of IP Addresses and Network Names” on page 43](#)). It must be a new, unused name, and must be registered in the DNS so that clients can find the server without having to specify its address.
 - The Server Name is used by clients to identify the server. If you only use Avid Interplay Clients on Windows computers, you can use the Network Name as the server name. If you use several platforms as client systems, such as Macintosh[®] and Linux[®] you need to specify the static IP address that you entered for the resource group in the previous dialog box. Macintosh systems are not always able to map server names to IP addresses. If you type a static IP address, make sure this IP address is not provided by a DHCP server.
- Click Next.

3 Installing the Interplay Engine for a Failover Cluster

The Specify Destination Location dialog box opens.

Specifying the Destination Location

In this dialog box specify the folder in which you want to install the Interplay Engine program files.



To specify the destination location:

1. Avid recommends that you keep the default path C:\Program Files\Avid\Avid Interplay Engine.



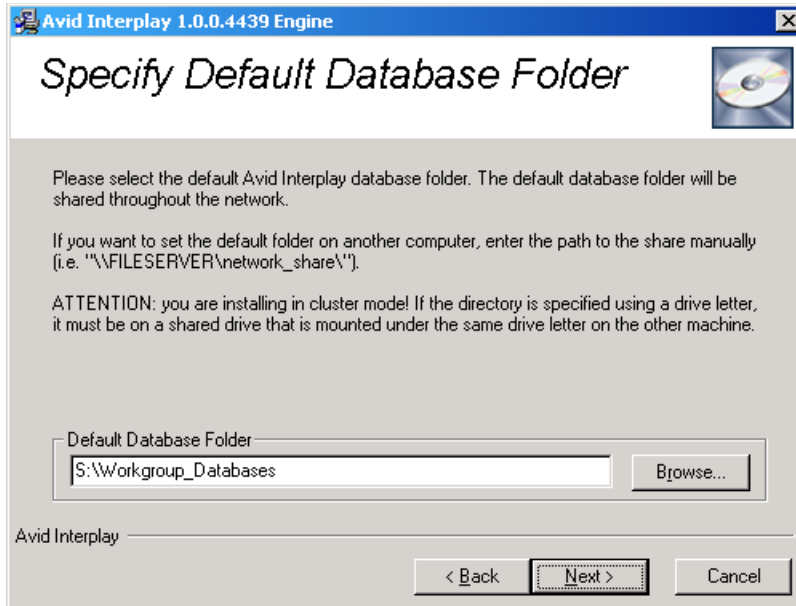
Under no circumstances attempt to install to a shared disk; independent installations are required on both nodes. This is because local changes are also necessary on both machines. Also, with independent installations you can use a rolling upgrade approach later, upgrading each node individually without affecting the operation of the cluster.

2. Click Next.

The Specify Default Database Folder dialog box opens.

Specifying the Default Database Folder

In this dialog box specify the folder where the database data is stored.



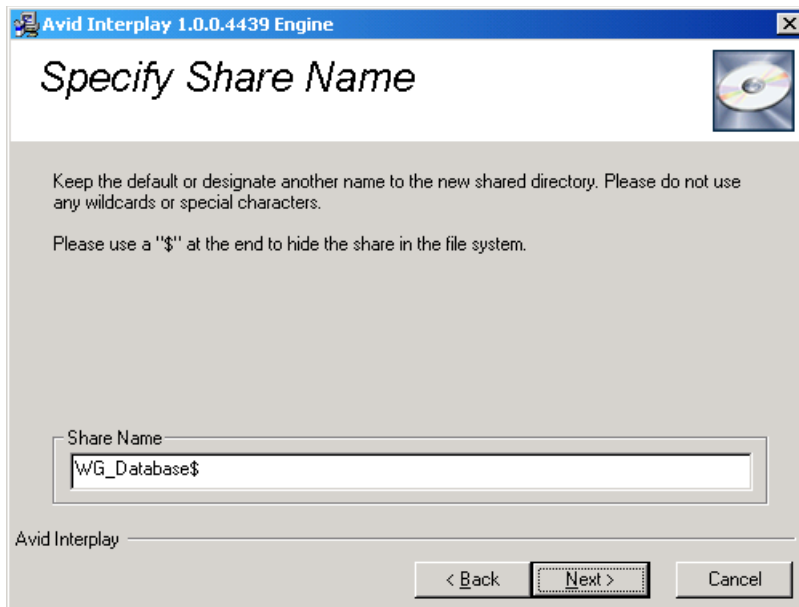
To specify the default database folder:

1. Type S:\Workgroup_Databases. Make sure the path specifies the shared drive (S:).
This folder should reside on the shared drive that is owned by the resource group of the server. Avid strongly recommends using the shared drive resource so that it can be monitored and managed by the cluster service. The drive must be assigned to the physical drive resource that is mounted under the same drive letter on the other machine.
2. Click Next.
The Specify Share Name dialog box opens.

Specifying the Share Name

In this dialog box specify a share name to be used for the database folder.

3 Installing the Interplay Engine for a Failover Cluster



To specify the share name:

1. Accept the default share name.

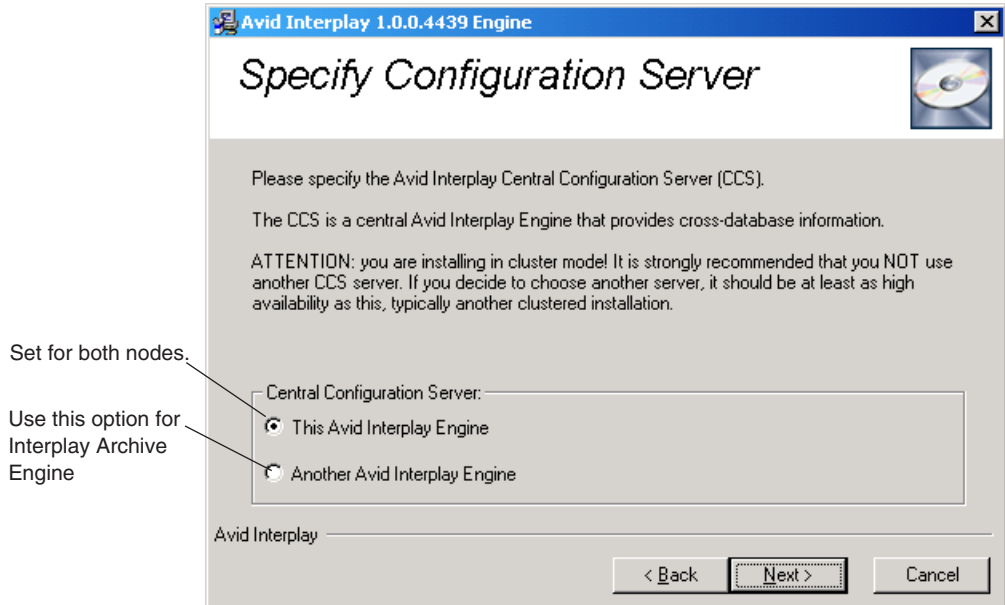
Avid recommends you use the default share name WG_Database\$. This name is visible on all client platforms, such as Windows 98, Windows ME, Windows NT Windows 2000 and Windows XP. The “\$” at the end makes the share invisible if you browse through the network with the Windows Explorer. For security reasons, Avid recommends using a “\$” at the end of the share name. If you use the default settings, the directory S:\Workgroup_Databases is accessible as \\InterplayEngine\WG_Database\$.

2. Click Next.

This step takes a few minutes. When finished the Specify Configuration Server dialog box opens.

Specifying the Configuration Server

In this dialog box, indicate whether this server is to act as a Central Configuration Server.



A Central Configuration Server (CCS) is an Avid Interplay Engine with a special module that is used to store server and database-spanning information. For more information, see the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.

To specify the server to act as the CCS server:

1. Select either the server you are installing or a previously installed server to act as the Central Configuration Server.

Typically you are working with only one server, so the appropriate choice is “This Avid Interplay Engine,” which is the default.

If you need to specify a different server as the CCS (for example, if an Interplay Archive Engine is being used as the CCS), select “Another Avid Interplay Engine.” You need to type the name of the other server to be used as the CCS in the next dialog box.



Only use a CCS that is at least as high availability as this cluster installation, typically another clustered installation.

If you specify the wrong CCS, you can change the setting later on the server machine in the Windows Registry. See [“Automatic Server Failover Tips and Rules”](#) on page 109.

2. Click Next.

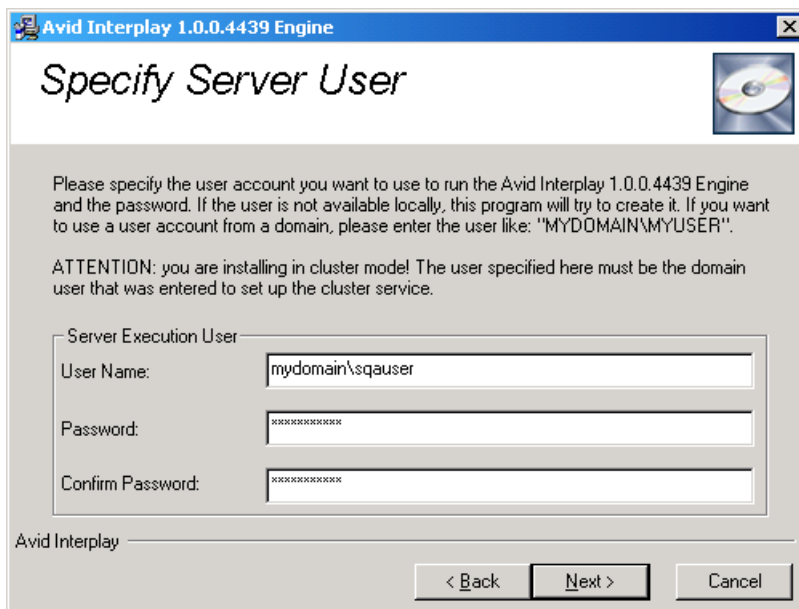
The Specify Server User dialog box opens.

3 Installing the Interplay Engine for a Failover Cluster

Specifying the Server User

In this dialog box, define the Cluster Service account (Server Execution User) used to run the Avid Interplay Engine.

The Server Execution User is the Windows domain user that runs the Interplay Engine and the cluster service. This account is automatically added to the Local Administrators group on the server. This account must be the one that was used to set up the cluster service. See “Before You Begin the Server Failover Installation” on page 42.



Avid Interplay 1.0.0.4439 Engine

Specify Server User

Please specify the user account you want to use to run the Avid Interplay 1.0.0.4439 Engine and the password. If the user is not available locally, this program will try to create it. If you want to use a user account from a domain, please enter the user like: "MYDOMAIN\MYUSER".

ATTENTION: you are installing in cluster mode! The user specified here must be the domain user that was entered to set up the cluster service.

Server Execution User

User Name:

Password:

Confirm Password:

Avid Interplay

< Back Next > Cancel

To specify the Server Execution User:

1. Type the Cluster Service Account user login information.



The installer cannot check the username or password you type in this dialog. Make sure that the password is set correctly, or else you will need to uninstall the server and repeat the entire installation procedure. Avid does not recommend changing the Server Execution User in cluster mode afterwards, so choose carefully.



When typing the domain name do not use the full DNS name such as mydomain.company.com, because the DCOM part of the server will be unable to start. You should use the NetBIOS name, for example, mydomain.

2. Click Next.

The Specify Preview Server Cache dialog box opens.

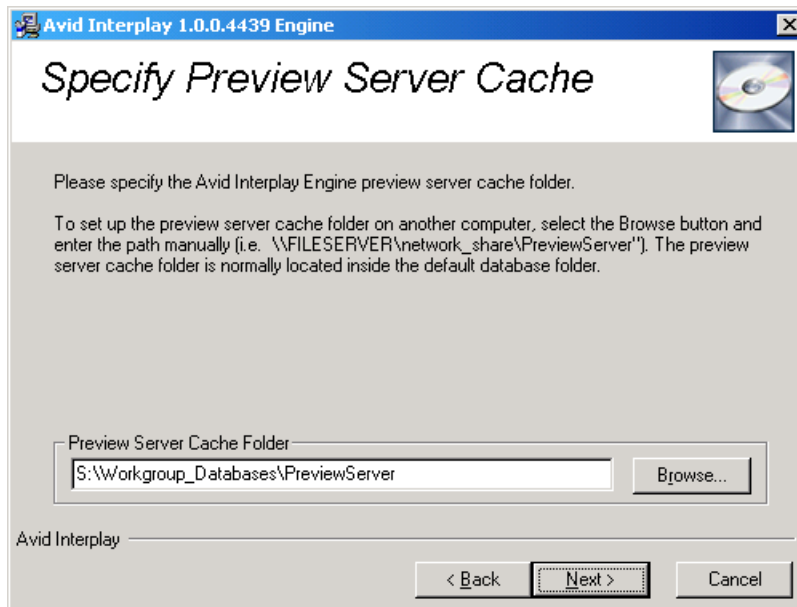
If necessary, you can change the name of the Server Execution User after the installation. For more information, see “Troubleshooting the Server Execution User Account” and “Re-creating the Server Execution User” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide* and the Interplay ReadMe.

Specifying the Server Cache

In this dialog box, specify the path for the cache folder.



For more information on the Preview Server cache and Preview Server configuration, see “Avid Workgroup Preview Server Service” in the Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide.



To specify the server cache folder:

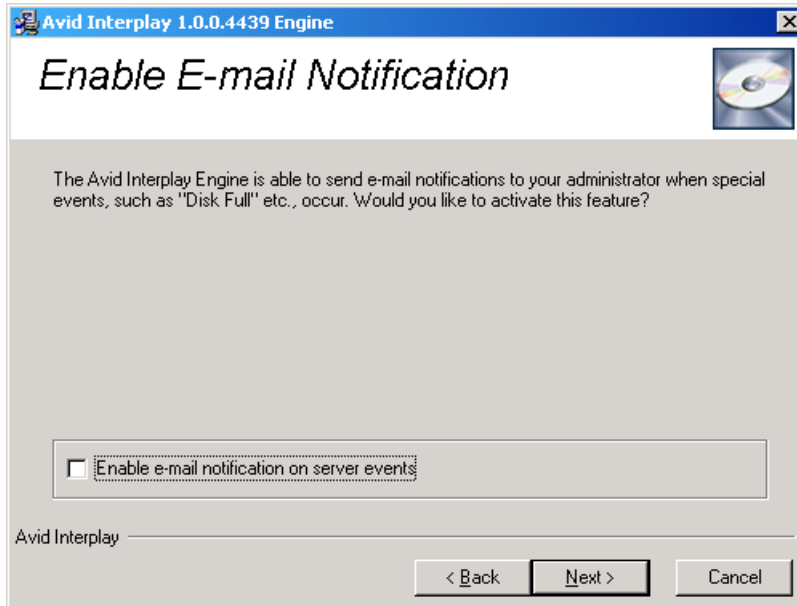
1. Type or browse to the path of the server cache folder. Typically, the default path is used.
2. Click Next.

The Enable Email Notification dialog box opens if you are installing the Avid Interplay Engine for the first time.

3 Installing the Interplay Engine for a Failover Cluster

Enabling Email Notifications

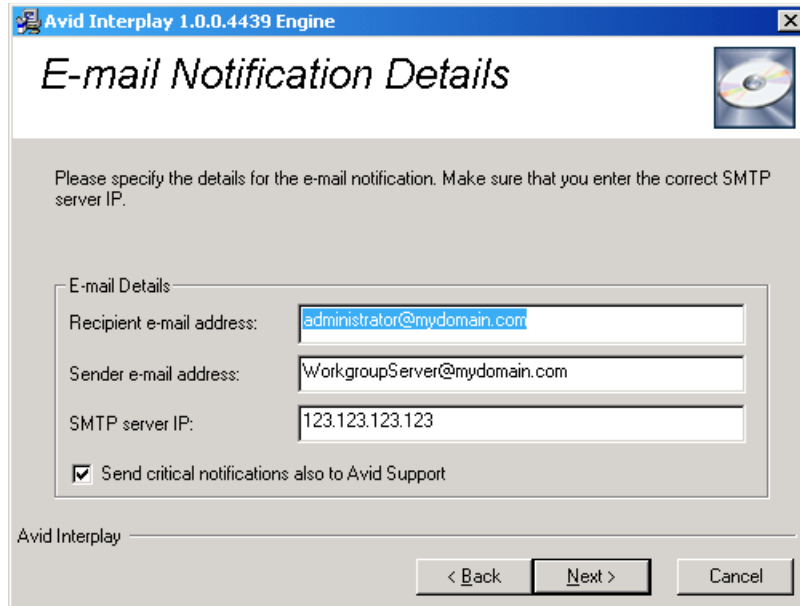
The first time you install the Avid Interplay Engine, the Enable Email Notification dialog box opens. The email notification feature sends emails to your administrator when special events, such as “Cluster Failure,” “Disk Full,” and “Out Of Memory” occur. Activate email notification if you want to receive emails on special events, server or cluster failures.



To enable email notification:

1. (Option) Select Enable email notification on server events.

The Email Notification Details dialog box opens.



2. Type the administrator's email address and the email address of the server, which is the sender.

If an event, such as “Resource Failure” or “Disk Full” occurs on the server machine, the administrator receives an email from the sender's email account explaining the problem, so that the administrator can react to the problem. You also need to type the static IP address of your SMTP server. The notification feature needs the SMTP server in order to send emails. If you do not know this IP, ask your administrator.

3. If you also want to inform Avid Support automatically using email if problems arise, select “Send critical notifications also to Avid Support.”
4. Click Next.


The installer modifies the file Config.xml in the Workgroup_Data\Server\Config\Config directory with your settings.

The Ready to Install dialog box opens.

3 Installing the Interplay Engine for a Failover Cluster

Installing the Interplay Engine for a Custom Installation on the First Node

In this dialog box, begin the installation of the engine software.



To install the Interplay Engine software:

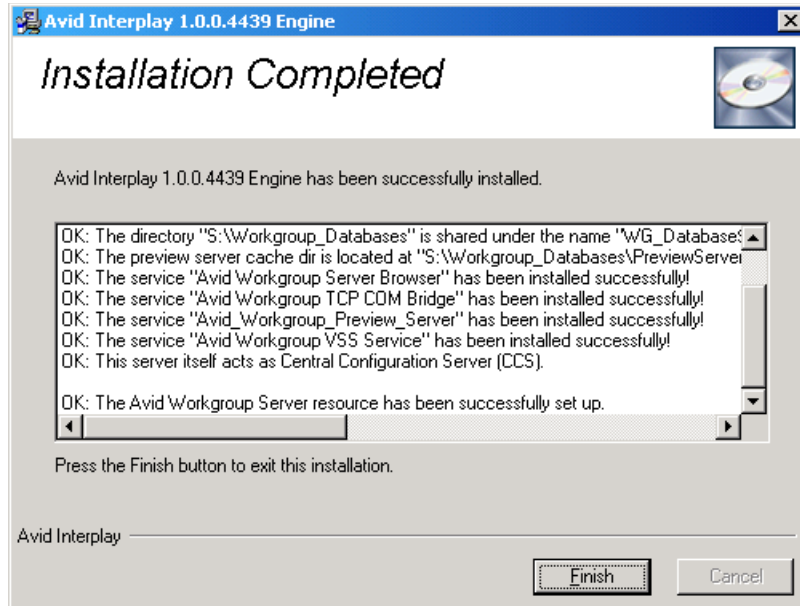
1. Click Next.

Use the Back button to review or change the data you have entered. You can also terminate the installer using the Cancel button, because no changes have been done to the system yet.

The first time you install the software, a dialog box opens and asks if you want to install the Sentinel driver. This driver is used by the licensing system.

2. Click Continue.

The Installation Completed dialog box opens after the installation is completed.



3. Do one of the following:
 - ▶ Click Finish.
 - ▶ Analyze and resolve any issues or failures reported.
4. Click OK if prompted for a restart the system.

The installation procedure requires the machine to restart (up to twice). For this reason it is very important that the other node is shut down, otherwise the current node loses ownership of the Avid Workgroup resource group. This applies to the installation on the first node only.



Subsequent installations should be run as described in “Updating a Clustered Installation (Rolling Upgrade)” on page 106 or in the Avid Interplay ReadMe.

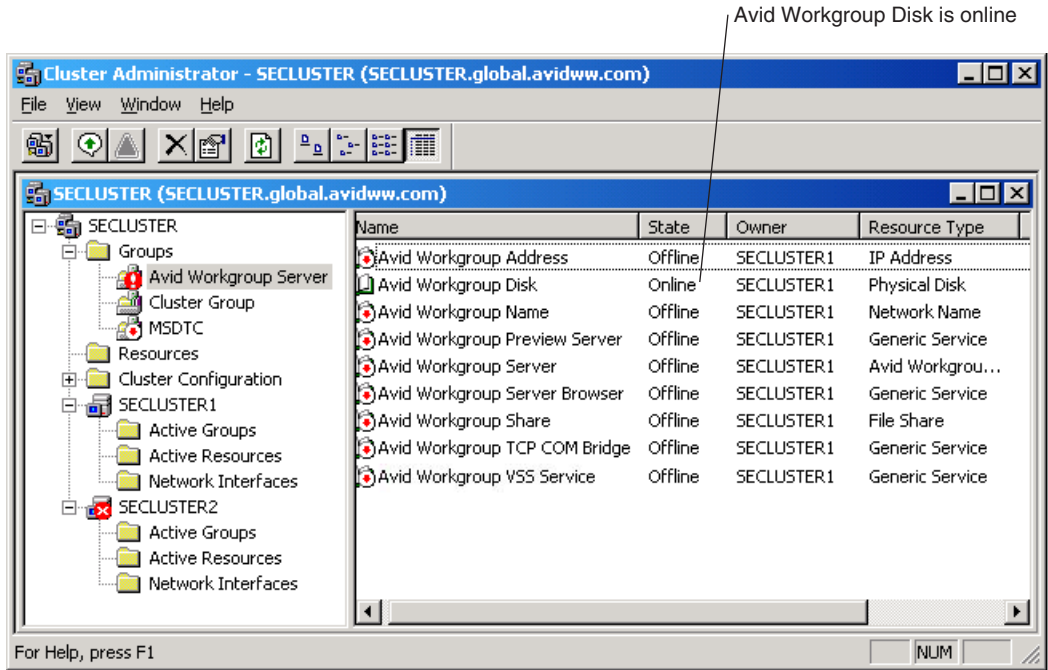
Bringing the Disk Resource Online

To bring the Disk Resource online:

1. After the installation is complete, start the Cluster Administrator tool by clicking Start and selecting Programs > Administrative Tools > Cluster Administrator.
2. Open the Avid Workgroup Server resource group.

The list of resources should look similar to those in the following illustration.

3 Installing the Interplay Engine for a Failover Cluster



The Avid Workgroup Disk resource should be online and all other resources offline.

3. Bring the disk resource online manually before continuing if necessary.

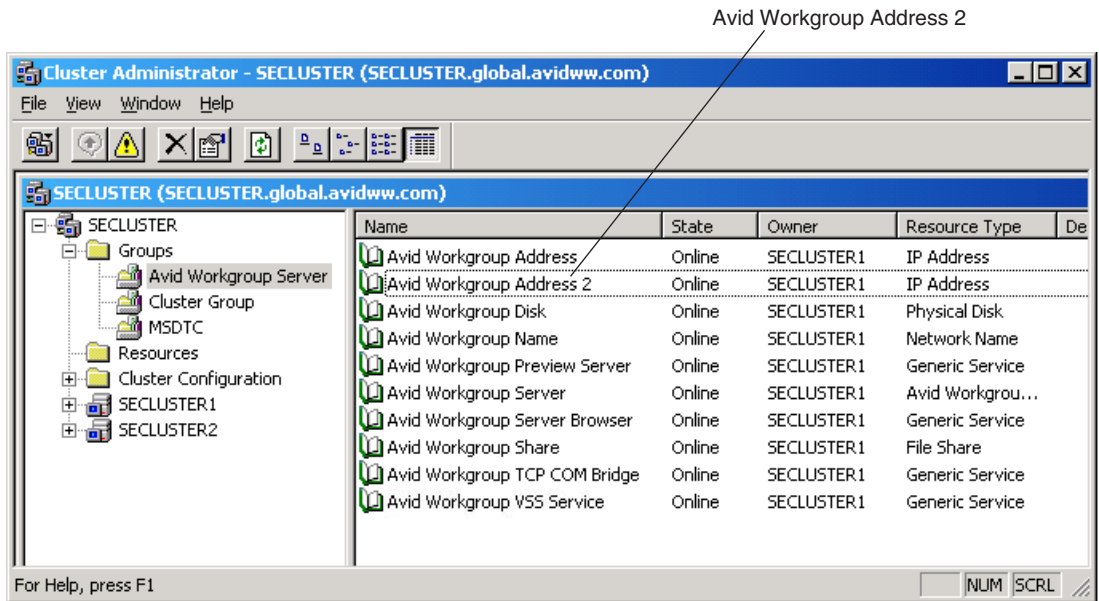


Avid does not recommend starting the server at this stage yet, since it is not installed on the other node and a failover would be impossible.

4. **(Avid Unity ISIS dual-connected configuration only)** Add the IP address on the second subnet for the Interplay Engine.
 - a. In the Cluster Administrator, right-click Avid Workgroup Server and select New > Resource.
 - b. Complete the New Resource dialog box as follows:
 - Name: Avid Workgroup Address 2
 - Resource Type: IP Address
 - Group: Avid Workgroup Server
 - c. Complete the Possible Owners dialog box as follows:
 - Add the cluster server host names to the Possible owners lists. For example, SECLUSTER1 and SECLUSTER2. See [“List of IP Addresses and Network Names”](#) on page 43.

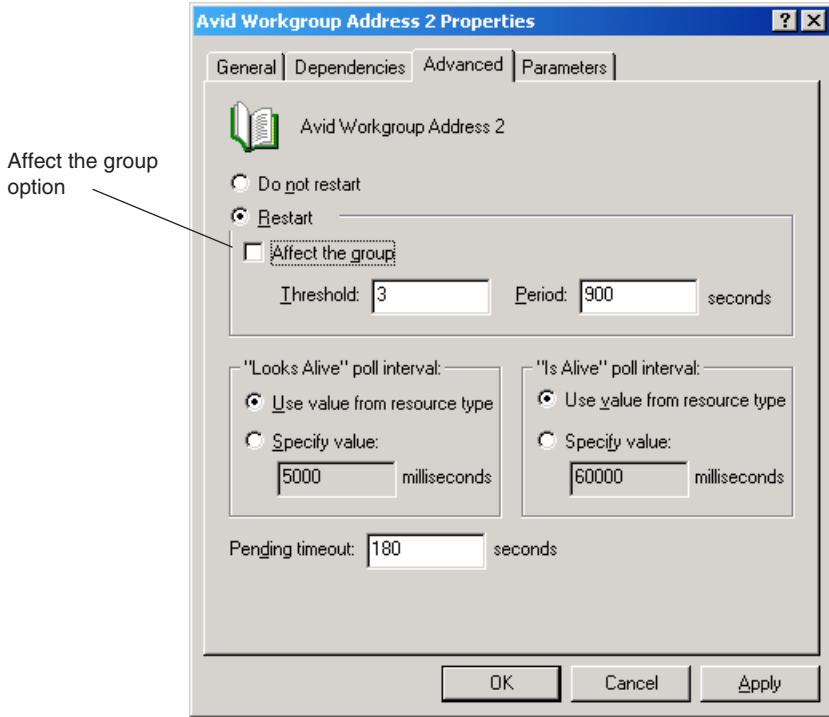
- d. Complete the Dependencies dialog box as follows:
 - Leave the Resource dependencies list empty.
- e. Complete TCP/IP Address Parameters dialog box as follows:
 - Address: type the second Interplay Engine service Avid Unity ISIS IP address. See [“List of IP Addresses and Network Names” on page 43](#)
 - Subnet mask: displays subnet for the second subnet network
 - Network: select a network connection: *Right-subnet*
 - Select Enable NetBIOS for this address
- f. Click Finish.

The following illustration shows the new entry.



- g. Right-click Avid Workgroup Address 2 and select Properties.
- h. Click the Advanced tab.
- i. Deselect Affect the Group as shown in the following illustration.

3 Installing the Interplay Engine for a Failover Cluster



- j. Click OK.
5. When the installation is complete, leave this node running so that it maintains ownership of the resource group and proceed to [“Installing the Interplay Engine on the Second Node” on page 102.](#)

Installing the Interplay Engine on the Second Node

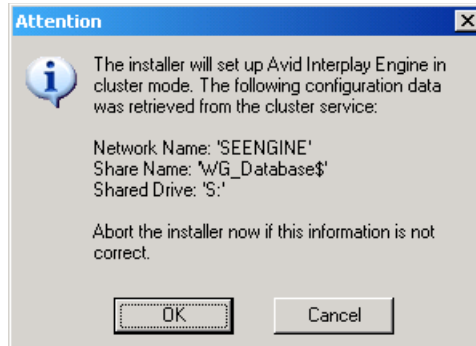
To install the Interplay Engine on the second node:

1. Leave the first machine running so that it maintains ownership of the resource group and start the second node.



Do not attempt to move the resource group over to the second node, or similarly, do not shut down the first node while the second is up, before the installation is completed on the second node.

2. Perform the installation procedure for the second node as described in [“Installing the Interplay Engine on the First Node” on page 84.](#) In contrast to the installation on the first node, the installer automatically detects all settings previously entered on the first node. The Attention dialog box opens.



3. Click OK.
4. The same installation dialog boxes will open that you saw before, except for the cluster related settings that only need to be entered once. Enter the requested information and allow the installation to proceed.



Make sure you use the installation mode that you used for the first node and enter the same information throughout the installer. Using different values results in a corrupted installation.

5. The installation procedure requires the machine to restart (up to twice). Allow the restart as requested.

Bringing the Interplay Engine Online

To bring the Interplay Engine online:

1. Click Start and select Programs > Administrative Tools > Cluster Administrator.
2. Click Groups and right-click Avid Workgroup Server and select Bring on line.

All resources are now online.

Installing a Permanent License

During Interplay Engine installation a temporary license for one user is activated automatically so that you can administer and install the system. There is no time limit for this license. A permanent license is provided by Avid in the form of a file (*.nxn) on a CD-ROM.

For more information on managing licenses, see the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.

3 Installing the Interplay Engine for a Failover Cluster

To install a permanent license on a clustered system:

1. Start the first node.
2. Insert the Avid Interplay Licenses CD-ROM into a CD drive.
3. Start and log in to the Interplay Administrator.
4. In the Server section of the Interplay Administrator window, click the Licenses icon.
5. Click the Import license button.
6. Browse for the .nxx file on the CD-ROM.
7. Select the file and click Open.
You should see information about the permanent license in the License Types area.
8. Close the Interplay Administrator.
9. Force a failover of the cluster, so that the second node is now active.
10. Repeat steps 3 through 8 on the second node.

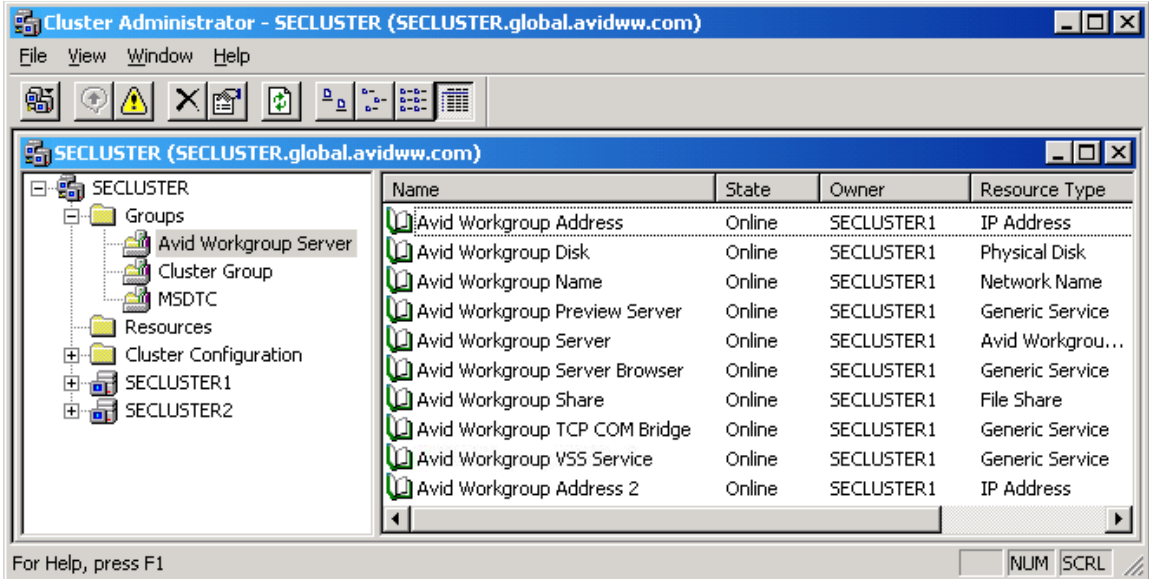
Testing the Complete Installation

After you complete all the previously described steps, you are now ready to test the installation. Make yourself familiar with the Cluster Administrator and review the different failover-related settings.

To test the complete installation:

1. To start the server, bring the resource group online; this starts the Interplay Engine and its affiliated services.

After starting the Avid Interplay Engine on the first node, the Cluster Administrator should look similar to the following figure.



2. Start an Interplay Administrator, install the licenses if needed, create a test database and add some files to it. If the other node is also running, you are ready to test the failover functionality.
3. Initiate a failover by moving the resource group; do this through the context menu of the resource group. Failures can also be simulated, again through the context menu of the appropriate resource.



Failures do not necessarily initiate failover.

4. You might also want to experiment by terminating the Interplay Engine manually using the Windows Task Manager (NxNServer.exe). This is also a good way to get familiar with the failover settings which can be found in the Properties Panel of the Avid Workgroup resource, under the Advanced tab.
5. Look at the related settings of the resource group. If you need to change any configuration files, make sure that the Avid Workgroup Disk resource is online; the configuration files can be found on the resource drive in the Workgroup_Data folder.

Updating a Clustered Installation (Rolling Upgrade)

A major benefit of a clustered installation is that you can perform “rolling upgrades.” You can keep a node in production while updating the installation on the other, then move the resource over and update the second node as well.



For information about updating specific versions of the Interplay Engine and a cluster, see the Avid Interplay ReadMe. The ReadMe describes an alternative method of updating a cluster, in which you lock and deactivate the database before you begin the update.

When updating a clustered installation, the settings that were entered to set up the cluster resources cannot be changed. Additionally, all other values must be reused, so Avid strongly recommends choosing the Typical installation mode. Changes to the fundamental attributes can only be achieved by uninstalling both nodes first and installing again with the new settings.

Make sure you follow the procedure in this order, otherwise you might end up with a corrupted installation.

To update a cluster:

1. Determine which node is active.
 - a. Select Control Panel > Administrative Tools > Cluster Administrator.
 - b. Open the Groups folder and check the owner column for the Avid Workgroup Server.

Consider this the first node.

2. Make sure this node is also the owner of the Cluster and the MSDTC groups. If these groups are not on the active node, right-click each group and select Move Group.
3. Run the Interplay Engine installer to update the installation on the non-active node (second node). Select Typical mode to reuse values set during the previous installation on that node. Restart as requested and continue with Part 2 of the installation. The installer will ask you to restart again after Part 2.



Do not move the Avid Workgroup Server resource group to the second node yet.

4. Make sure that first node is active. Run the Interplay Engine installer to update the installation on the first node. Select Typical mode so that all values are reused.

5. During the installation, the installer displays a dialog box that asks you to move the Avid Workgroup Server group to the second node. Move the group, then click OK in the installation dialog box to continue. Restart as requested and continue with Part 2 of the installation. The installer will ask you to restart again after Part 2.
6. You might want to test the final result of the update by moving the server back to the first node. The Interplay Administrator can be used to display the version of the server.

After completing the above steps, your entire clustered installation is updated to the new version. Should you encounter any complications or face a specialized situation, contact Avid Support as instructed in [“If You Need Help” on page 10](#).

Uninstalling the Interplay Engine on a Clustered System

To uninstall the Avid Interplay Engine, use the Avid Interplay Engine uninstaller, first on the inactive node, then on the active node.

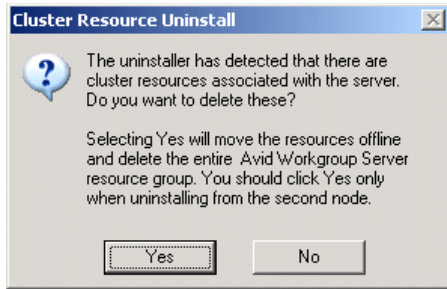


The uninstall mechanism of the cluster resources only functions properly if the names of the resources or the resource groups are not changed. Never change these names.

To uninstall the Interplay Engine:


1. If you plan to reinstall the Interplay Engine and reuse the existing database, create a complete backup of the AvidWG database and the _InternalData database in S:\Workgroup_Databases. For information about creating a backup, see “Creating and Restoring Database Backups” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.
2. (Dual-connected configuration only) Remove the second network address within the Avid Workgroup Server group.
 - a. In the Cluster Administrator, right-click Avid Workgroup Server.
 - b. Right-click Avid Workgroup Address 2 and select Remove.
3. Make sure that both nodes are running before you start the uninstaller.
4. On the inactive node (the node that does not own the Avid Workgroup Server resource group), start the uninstaller by selecting Programs > Avid > Avid Interplay Engine > Uninstall Avid Interplay Engine.
5. When you are asked if you want to delete the cluster resources, click No.

3 Installing the Interplay Engine for a Failover Cluster



6. When you are asked if you want to restart the system, click Yes.
7. At the end of the uninstallation process, if you are asked to restart the system, click Yes.
8. After the uninstallation on the inactive node is complete, wait until the last restart is done. Then open the Cluster Administrator on the active node and make sure the inactive node is shown as online. (The nodes are shown in the lower part of the tree on the left side of the Cluster Administrator.)
9. Start the uninstallation on the active node (the node that owns the Avid Workgroup Resource Group).
10. When you are asked if you want to delete the cluster resources, click Yes.

A confirmation dialog box opens.



11. Click Yes.
12. When you are asked if you want to restart the system, click Yes.
13. At the end of the uninstallation process, if you are asked to restart the system, click Yes.
14. After the uninstallation is complete, but before you reinstall the Interplay Engine, rename the folder S:\Workgroup_Data (for example, S:\Workgroup_Data_Old) so that it will be preserved during the reinstallation process. In case of a problem with the new installation, you can check the old configuration information in that folder.



If you do not rename the Workgroup_Data, the reinstallation might fail because of old configuration files within the folder. Make sure to rename the folder before you reinstall the Interplay Engine.

4 Automatic Server Failover Tips and Rules

This chapter provides some important tips and rules to use when configuring the automatic server failover.

Don't Access the Machines Directly

Don't access the machines (nodes) directly. Use the virtual network name or IP address that has been assigned to the Interplay Engine resource group (see [“List of IP Addresses and Network Names” on page 43](#)). Never use the actual physical names or IP addresses of the machines that are part of the cluster.

Make Sure to Connect to the Interplay Engine Resource Group

The network names and the virtual IP addresses resolve to the physical machine they are being hosted on. For example, it is possible to mistakenly connect to the Interplay Engine using the network name or IP address of the cluster group (see [“List of IP Addresses and Network Names” on page 43](#)). The server is found using the alternative address also, but only while it is online on the same node. Therefore, under no circumstances connect the clients to a network name other than what was used to set up the Interplay Engine resource group.

Do Not Rename Resources

Do not rename resources. The resource plugin, the installer, and the uninstaller all depend on the names of the cluster resources. These are assigned by the installer and even though it is possible to modify them using the cluster administrator, doing so corrupts the installation and is most likely to result in the server not functioning properly.

Do Not Install the Interplay Engine Server on a Shared Disk

The Interplay Engine must be installed on the local disk of the cluster nodes and not on a shared resource. This is because local changes are also necessary on both machines. Also, with independent installations you can later use a rolling upgrade approach, upgrading each node individually without affecting the operation of the cluster. The Microsoft documentation is also strongly against installing on shared disks.

Do Not Change the Interplay Engine Server Execution User

The domain account that was entered when setting up the cluster (the Cluster Service Account —see “[Before You Begin the Server Failover Installation](#)” on page 42) also has to be the Server Execution User of the Interplay Engine. Given that you cannot easily change the cluster user, the Interplay Engine execution user has to stay fixed as well. For more information, see “[Troubleshooting the Server Execution User Account](#)” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.

Do Not Edit the Registry While the Server is Offline

If you edit the registry while the server is offline, you will lose your changes. This is something that most likely will happen to you since it is very easy to forget the implications of the registry replication. Remember that the registry is restored by the resource monitor before the process is put online, thereby wiping out any changes that you made while the resource (the server) was offline. Only changes that take place while the resource is online are accepted.

Do Not Remove the Dependencies of the Affiliated Services

The TCP-COM Bridge, the Preview Server, and the Server Browser services must be in the same resource group and assigned to depend on the server. Removing these dependencies might speed up some operations but prohibit automatic failure recovery in some scenarios.

Consider Disabling Failover When Experimenting

If you are performing changes that could make the Avid Interplay Engine fail, consider disabling failover. The default behavior is to restart the server twice (threshold = 3) and then initiate the failover, with the entire procedure repeating several times before final failure. This can take quite a while.

Changing the CCS

If you specify the wrong Central Configuration Server (CCS), you can change the setting later on the server machine in the Windows Registry under:

HKEY_LOCAL_MACHINE/Software/Avid Technology/Workgroup/DatabaseServer

The string value CMS specifies the server. Make sure to set the CMS to a valid entry while the Interplay Engine is online, otherwise your changes to the registry won't be effective. After the registry is updated, stop and restart the server using the Cluster Administrator (in the Administration Tools folder in Windows).

Specifying an incorrect CCS can prevent login. See “Troubleshooting Login Problems” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.

For more information, see “Understanding the Central Configuration Server” in the *Avid Interplay Engine and Avid Interplay Archive Engine Administration Guide*.

4 Automatic Server Failover Tips and Rules

Index

A

- Active Directory domain
 - adding cluster servers [60](#)
- Antivirus software
 - running on a failover cluster [19](#)
- Apache web server
 - on failover cluster [83](#)
- ATTO card
 - setting boot delay [49](#)
 - setting link speed [50](#)
- Avid
 - online support [10](#)
 - training services [12](#)
- Avid Unity environment
 - ISIS failover cluster connections SR2400 (illustration) [27](#)
 - ISIS failover cluster connections SR2500 (illustration) [26](#), [28](#), [32](#)
 - MediaNetwork failover cluster connections SR2500 (illustration) [36](#), [38](#)
 - SR2400 server slot locations (failover cluster) [21](#)
 - SR2500 server slot locations (failover cluster) [22](#), [23](#)
- Avid Unity ISIS
 - connections for failover cluster [24](#), [29](#)
 - failover cluster configurations [15](#)
 - failover cluster connections SR2400 (illustration) [27](#)
 - failover cluster connections SR2500 (illustration) [26](#), [28](#), [32](#)
- Avid Unity MediaNetwork
 - connections for failover cluster [34](#)
 - failover cluster configuration [15](#)
 - failover cluster connections SR2500 (illustration) [36](#), [38](#)

B

- Binding order networks
 - configuring [58](#)

C

- Central Configuration Server (CCS)
 - changing for failover cluster [109](#)
 - specifying for failover cluster [92](#)
- Cluster
 - overview [13](#)
 - See also* Failover cluster
- Cluster group
 - partition [60](#)
- Cluster installation
 - updating [106](#)
- Cluster Installation account
 - defined [42](#)
- Cluster network
 - configuring rules [70](#)
- Cluster service
 - configuring [62](#)
 - configuring on second node [67](#)
 - defined [39](#)
 - specifying name [62](#)
 - validating on first node [67](#)
- Cluster Service account
 - defined [42](#)
 - Interplay Engine installation [94](#)
 - specify name [62](#)
- Cluster system
 - monitoring [13](#)

D

- Database folder
 - default location (failover cluster) [90](#)
- Distributed Transaction Coordinator
 - assigning resource to MSDTC group [81](#)
 - creating MSDTC resource group [77](#)
 - installing [76](#)
 - resource group for [76](#)
- Dual-connected cluster configuration [15](#)

E

- Email notification
 - setting for failover cluster [96](#)

F

- Failover cluster
 - configurations [15](#)
 - connections in Avid Unity ISIS [24](#), [29](#)
 - connections in Avid Unity MediaNetwork [34](#)
 - hardware and software requirements [19](#)
 - system components [14](#)
 - system overview [13](#)

H

- Hardware
 - requirements for failover cluster system [19](#)
- Heartbeat adapter
 - prioritizing [71](#)
- Heartbeat connection
 - configuring [55](#)

I

- Importing
 - license [103](#)
- Installation (failover cluster)
 - testing [75](#)
- Installing
 - Distributed Transaction Coordinator [76](#)
 - Interplay Engine (failover cluster) [85](#)
- Interplay Engine
 - Central Configuration Server, specifying for failover cluster [92](#)

- cluster details [89](#)
- cluster information for installation [87](#)
- default database location for failover cluster [90](#)
- installing on first node [84](#)
- Server Execution User, specifying for failover cluster [94](#)
- share name for failover cluster [91](#)
- Interplay Portal
 - viewing [11](#)
- IP addresses (failover cluster)
 - assigning to MSDTC group [78](#)
 - private network adapter [55](#)
 - public network adapter [60](#)
 - required [43](#)

L

- License
 - agreement (failover server) [85](#)
- License requirements
 - failover cluster system [19](#)
- Licenses
 - importing [103](#)
 - permanent [103](#)

M

- MSDTC resource group
 - creating [77](#)
 - creating physical disk [80](#)

N

- Network connections
 - naming for failover cluster [52](#)
- Network interface
 - renaming LAN for failover cluster [52](#)
- Network name
 - assigning MSDTC group [79](#)
- Network names
 - examples for failover cluster [43](#)
- Node
 - defined [39](#)
 - name examples [43](#)
 - setting startup time [73](#)

O

Online resource
defined [39](#)
Online support [10](#)

P

Partition
for the cluster group [41](#)
Permanent license [103](#)
Port
for Apache web server [83](#)
Private network adapter
configuring [55](#)
Public Network
for failover cluster [87](#)
Public network adapter
configuring [60](#)

Q

QLogic card
setting link speed [47](#)
Quorum disk [41](#)
configuring [62](#)
verifying [73](#)
Quorum resource
defined [39](#)

R

RAID array
configuring for failover cluster [60](#)
Redundant-switch cluster configuration [15](#)
Registry
editing while offline [109](#)
Resource group
connecting to [109](#)
defined [39](#)
services [109](#)
Resources
defined [39](#)
renaming [109](#)
Rolling upgrade (failover cluster) [106](#)

S

Server
setting startup time on each node [73](#)
Server cache
Interplay Engine cluster installation [95](#)
Server Execution User
changing [109](#)
specifying for failover cluster [94](#)
Server Failover
overview [13](#)
See also Failover cluster
Service name
examples for failover cluster [43](#)
Services
dependencies [109](#)
Shared drive
configuring for failover cluster [60](#)
specifying for Interplay Engine [87](#)
Slot locations
SR2400 server (failover cluster) [21](#)
SR2500 server (failover cluster) [22, 23](#)
Software
requirements for failover cluster system [19](#)
SR2400 server
slot locations (failover cluster) [21](#)
SR2500 server
slot locations (failover cluster) [22, 23](#)
Subnet Mask [87](#)

T

Training services [12](#)
Troubleshooting [10](#)
server failover [109](#)

U

Uninstalling
Interplay Engine (failover cluster) [107](#)
Updating
cluster installation [106](#)

V

Virtual IP address
 for Interplay Engine (failover cluster) [87](#)
Virtual Server Address [84](#)

W

Web servers
 disabling [83](#)
Windows Cluster Administrator
 console [14](#)



Avid
75 Network Drive
Burlington, MA 01803-2756 USA

Technical Support (USA)
Visit the Online Support Center at
www.avid.com/support

Product Information
For company and product information,
visit us on the web at www.avid.com