



Telestream



Installation and Configuration Guide

Release: 8.3.x

Revision: 1.3

Copyrights and Trademark Notices

Specifications subject to change without notice. Copyright © 2023 Telestream, LLC and its Affiliates. Telestream, CaptionMaker, Cerify, DIVA, Episode, Flip4Mac, FlipFactory, Flip Player, Gameshow, GraphicsFactory, Kumulate, Lightspeed, MetaFlip, Post Producer, Prism, ScreenFlow, Split-and-Stitch, Switch, Tempo, TrafficManager, Vantage, VOD Producer, and Wirecast are registered trademarks and Aurora, ContentAgent, Cricket, e-Captioning, Inspector, iQ, iVMS, iVMS ASM, MacCaption, Pipeline, Sentry, Surveyor, Vantage Cloud Port, CaptureVU, Cerify, FlexVU, PRISM, Sentry, Stay Genlock, Aurora, and Vidchecker are trademarks of Telestream, LLC and its Affiliates. All other trademarks are the property of their respective owners.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

Contents

Telestream Contact Information 17

Preface 18

- Audience 19
- Documentation Accessibility 19
- Access to Telestream Support 19
- Related Documents 19
- Document Updates 20

Overview 21

- Release Overview 22
- New Terminology 22
- Port Utilization 23
- Enhanced Features and Functionality 24
- Media Storage Formats 25
 - AXF (Archive eXchange Format) 25
 - Native File and Folder Support 25
 - Tape Groups 25
 - Disk Arrays 26
- Complex Objects 27
 - Complex Objects and FTP 28
- DIVA Core Software Components 30
 - Manager 31
 - Actor 31
 - Client APIs 31
 - REST API 32
 - Database 32
 - Metadata Database 33
 - Notification Service (RabbitMQ) 33
 - System Management App 33
- Additional Software Components 34
 - Robot Manager 34

VACP (Video Archive Communications Protocol) Service	34
SPM (Storage Policy Manager)	35
SNMP (Simple Network Management Protocol) Agent	35
DIVA Connect	35
Watch Folder Monitor	35
Transcoder Support	36
Avid Connectivity	36
Avid DHM (Data Handler Module) Interface	36
Avid DET (Dynamically Extensible Transfer) Interface	37
Archive Manager Interface	37
Analytics App	37
OTU (Object Transfer Utility)	37
DIVA View	37
DIVAmigrate	38
DIVA Core Utilities	39
System Management App	39
Robot Manager Utilities	39
Backup Service	39
Scandrive Utility	40
Tape Reading Utility	40
DIVAscript	40
RDTU (Recover Damaged Tape Utility)	40

DIVA Core Security 41

General Security Principles	42
Keep Software up to Date	42
Restrict Network Access to Critical Services	42
Run as DIVA User and use Principle of Least Privilege Where Possible	42
Monitor System Activity	42
Keep up to Date on Latest Security Information	43
Secure Installation	44
Understand Your Environment	44
Which Resources Need to be Protected?	44
From whom are the resources being protected?	45
What will happen if the protections on strategic resources fails?	45
Recommended Deployment Topologies	45
Separate Metadata Network	45
FC Zoning	45
Safeguard SAN Disks Configuration Access	45
Install the DIVA Core Package	45
DIVA Core Tape Security	46
Backups	46
Post-Installation Configuration	46
Security Features	47
The Security Model	47
Authentication	47
Access Control	47

Tape Group Encryption	48
S3 Server-side Encryption	49
SSL (Secure Sockets Layer) and Authentication	52
External Certificate Authorities	52
Security Tools	52
DIVA Core API Changes	52
Secure Communication with Core Database	54

Database Installation and Configuration 56

DIVA Core Databases and Backup Service (BKS) Overview	57
Complex Objects	58
DIVA Core Backup Service (BKS)	59
DBAgent	60
Backup Initiator	61
Workflows	62
Database Service Failover	64
Core Database	65
Core Metadata Database	65
DIVA Core Backup Service Recommended Practices	65
Installing, Upgrading, and Configuring the Core SQL Database and Backup Service	67
General DIVA Core and Database Upgrade Processes	67
Starting ORACLE 11g and DIVA Core 7.6	67
Upgrading to ORACLE 11g and DIVA Core 8.x	67
Upgrade from ORACLE 11g to ORACLE 19c and DIVA Core 8.x	67
Exporting the Database Dump Files	68
Export the Database Dump Files Using DIVADBinStaller	68
Example	68
Export the Database Dump Files Using sqlplus	69
Importing the Database Dump Files	70
Import the Database Dump File Using DIVADBinStaller	70
Example - Same Source Server User Name	70
Example - Different Source Server User Name	70
Import the Database Dump File Using sqlplus	70
Uninstalling the Core Database Server (if required)	72
Uninstalling the Core Database Server in Windows	72
Uninstalling the Core Database Server in Linux	72
Installing the Core Database Server in Windows	73
Installing the Core Database Server in Linux	73
Prerequisites for Installing the Core Database: Configure Shared Memory	73
Prerequisite for Installing the Core Database: Creating Drive Partitions	74
Installing the Core Database Server	75
Installing the DIVA Database User and Schema	76
Using DIVADBinStaller for DIVA Core 8.3	77
Manually Create the Database User and Schema for 7.6.1 and earlier	78
Secure Communications with Core Database	79
Migrating Core Database Server from 11.2 to 12.1	80
Preparing the Source Server Computer (Core Manager with Oracle Database	

11.2)	80
Updating the Destination Server (Core Manager with Oracle Database 12.1)	81
Installing and Configuring the DIVA Core Backup Service	82
DIVA Core Backup Service Overview	82
Installing the DIVA Core Backup Service Software	82
Installing BKS and DBAgent	83
Configuring the DIVA Core Backup Service	86
Backup Interval Overrun	91
Backup Service Running Normally	92
Backup Service Not Currently Running	92
Backup Service Failed to Start	92
Uninstalling BKS and DBAgent	92
Monitoring the DIVA Core Backup Service	93
Monitoring Minimum Disk Space	95
Email Notifications	95
Metadata (non-SQL) Database Configuration	97
Configuring the Metadata Database	97
Sizing the Metadata Database	98
Migrating an MDDB (Flat File Metadata Database) to MDS (Metadata Service)	98
Upgrading and Migrating in Linux	100
Installing DIVA Core Notification Service (RabbitMQ)	103
New Windows Installation	103
Upgrading Windows Installations	106
New Linux Installation	107
Upgrading Linux Installations	108
Troubleshooting	109
Core Database Failure Scenarios and Recovery Procedures	109
Non-failover Scenarios	109
Failover Scenarios	110
Failover Procedures	110
Metadata Database Failure Scenarios	113
Identifying Failure Scenarios, Causes, and Resolutions	113
Scenario 1: Metadata Database Storage Disk Failure	113
Scenario 2: Metadata Database File Corruption	114
Scenario 3: Lost or Manually Deleted Metadata Database File	115
Scenario 4: Failure to Backup Metadata Database to All Backup Systems	116
Scenario 5: Failure of the Metadata Database Backup to One Backup System	116
Core Manager Will Not Start	117
DIVA Core Backup Service Will Not Start	117

Cluster Manager Installation 118

Overview	119
Related Documentation	119
Prerequisites	120
Oracle Fail Safe Integration with Windows	120
Real Application Clusters Integration with Windows	120

DIVA Core Cluster Solution	121
Installation Requirements	122
Hardware Requirements	122
Software Requirements	123
Network Requirements	123
Example IP Addresses and Host Names	124
Domain Account Requirements	124
Granting Domain User Permissions to Create the Cluster	125
Granting Microsoft Cluster Object Permissions to Create the Cluster Role	126
Microsoft Cluster Configuration	127
Configuring the Operating System	127
Configuring the Microsoft Cluster Server Cluster	127
Installing the Windows Failover Server Clustering Feature	127
Enabling the Remote Registry Service	128
Registering the Required Host Names to the DNS Manager	128
Creating the Windows Failover Cluster Resources	129
Validating the Nodes Configuration for MSCS Clustering	131
Testing the Configuration	131
Performing a Manual Cluster Failover Test from the Failover Cluster Manager	132
DIVA Core and Oracle Fail Safe Configuration	133
Installing DIVA Core Prerequisites	133
Creating the Diva Role	133
Installing Oracle Database on Node1	134
Installing Oracle Database on Node2	135
Configuring Oracle Fail Safe	136
Installing Oracle Fail Safe	136
Oracle Fail Safe 4.2.1 References:	136
Verifying the Oracle Fail Safe Installation	137
Preparing the Database Installation for the Cluster	137
Configuring Oracle Fail Safe	138
Additional Required Cluster Configurations	139
Installing DIVA and Adding Services to the Role	140
Creating the DIVA Database User	140
Add DIVA Services to the Cluster	140
Additional Notes	141
Cluster Configuration Examples	142
Testing the Configuration	146
Performing a Manual Cluster Failover of the Core Resources	146
Performing a Cluster Failover Test by Restarting the Active Cluster Node	146
Moving a Configured Role to Another Cluster Node	146
Maintenance	148
Manually Placing a Service Offline	148
Adding a Network for Client Access	148
Rebuilding the Cluster after a Node Hardware Failure	149
Evicting a Failed Node	149
Preparing New Hardware	150
Joining a New Node Server to a Cluster	150

Installing DIVA Core	150
Installing and Configuring Oracle Fail Safe	150
Replacing an HBA (Host Bus Adapter)	151
Configuring Windows Firewall with Advanced Security	151
Cluster-Aware Updating	154

DIVA Core Installation 155

Software Component Relationships	156
Software Component Distribution	157
Installing the DIVA Core System	159
Installation Overview	159
Downloading the Software	160
Installing DIVA Core for Windows	161
Manually Creating the Database User and Schema for 7.6.1 and earlier	162
Installing DIVA Core for Linux	163
Prerequisites and Initial Set-up	163
Installing FTP Services	164
Installing DIVA Core 8.3 for Linux	164
Installing the DIVA Core Services	165
Creating System Management App Shortcut	166
Starting, Stopping, and Accessing DIVA Core in Linux	166
System Management App Installation and Configuration	168
Navigation Menu	168
Backend Support	168
Importing the DIVA Core License	170
Using the MDDB (Metadata Database) Services	170
Flashnet Migration Tool	172
Avid AM (Archive Manager) Updater Tool	173
System Requirements	174
Installation and Configuration	174

DIVA Core Configuration 176

Configuration Overview	177
Module Configuration Files	177
DIVA Core Databases	178
Metadata Database	178
Metadata Database Sizing	179
Environment Variables	179
SSL Authentication and Security	180
External Certificate Authorities	181
Prerequisites	182
Database Installation and Configuration	182
DIVA Core Installer and Database Schemas	182
REST API Installation and Configuration	182
Core Manager Configuration	183
DIVA Core System Management App and Configuration Utility	184

	185
Configuration Utility Overview	185
Configuration Utility Frame Buttons	185
DIVA Core System Management App Profiles and Passwords	185
	186
Setting Profile Passwords in the System Management App	186
Changing the Database Logging Level in the Configuration Utility	186
System Management App and Configuration Utility Tabs Overview	188
System Tab	188
System Tab Frames	188
Actor Configuration in the Database	189
Robots Tab	191
Robots Tab Frames	191
Disks Tab	192
Disks Tab Frames	192
Drives Tab	192
Drives Tab Frames	193
Tapes Tab	193
Tapes Tab Frames	193
Sets, Tape Groups & Media Mapping Tab	194
Sets, Tape Groups & Media Mapping Tab Frames	195
Configuring Clone Tape Groups	197
Analytics App Tab	199
Configuration Frame	199
Event Definitions Frame	200
Metrics Definitions Frame	200
Default Events and Metrics Configuration	201
Sample Metrics Definition	202
Media Tab	203
Storage Plans Tab	203
Storage Plans Tab Frames	204
Slots Tab	205
Slots Tab Frame	205
Manager Setting Tab	206
Media Configuration	206
Metadata Database Configuration	206
Keystore Configuration	206
License Tab	207
Cloud Storage Configuration	208
Common Parameters	208
Configuring Azure Blob Storage Accounts	212
Configuring Google Cloud Storage Accounts	214
Configuring Oracle Cloud Infrastructure Accounts	215
Cloud Replicated Bucket Scanning	217
Automatic Scan Restart Configuration	217
New Access Type for Actor-Disk Connections	217
System Management App Cloud Bucket Scanning Support	218
Object Auto-Indexing	220

Currently Supported Disks and Functionality	220
Duplicate Objects	220
Trial Mode	221
Configuration	221
Cloud Array Parameters used by Object Auto-Indexing	223
Preventing Auto-Indexing and AXF Archiving on the same Array	225
System Management App Object Auto-Indexing Support	226
Disk Storage Configuration	229
Defining Actor to Disk Connections	232
Actor to Disk Interfaces and Mount Points	233
Local Interface	233
Remote Interface	234
BML Interface	234
FTP Interface	234
MetaSAN Interface	235
Simulation Interface	235
Object Storage Servers	235
Media Storage Configuration	237
Adding a Robot Manager	237
Database Configuration Synchronization	238
Robot Manager-ACS Association	238
Defining Tape Capacity and Block Sizes	238
DIVA Core General Settings	242
Checksums	242
Media	242
Objects	243
SMTP Notifications	243
Security	246
Licensing Configuration	247
Synchronizing Media and Drive Compatibility with the Database	248
Synchronizing Media Types with the Database	248
Synchronizing Drive Types with the Database	249
Synchronizing the Library Drive List with the Database	249
Manually Identifying Drive Serial Numbers	251
Synchronizing the Library Tapes with the Database	253
Clearing Unused Tapes	254
Creating Tape Groups	254
Tape Group Encryption	255
Creating Tape Sets	256
Remapping Media	256

Component Configuration 257

Configuration Overview	258
Module Configuration Files	259
Manager Configuration	260
Configuring the Local Manager	260
Basic Settings	261

Database Settings	262
Advanced Settings	263
Logging Settings	274
Configuring Request Priorities	274
Rerouting Destinations (restore_translations.conf)	275
Controlling the Manager	276
Installing and Removing the Manager Service in Windows	276
Installing and Removing the Manager Service in Linux	277
Managing the Manager Service	277
Logging Manager Activity	278
Confirming System Connectivity	279
Confirming Remote Client to Manager Connectivity	279
Confirming Manager to Actors Connectivity	279
Confirming Manager to Robot Manager Connectivity	279
Manager Failover Procedures	280
Cluster Failovers	281
Actor Configuration	282
Configuration Overview	282
Configuring the Local Actor (actor.conf)	283
Configuring DIVA Core Partial File Restore	283
Defining and Declaring Actors	294
Advanced Actor Settings	296
Configuring Actor to Drive Connections	298
Defining Core Proxy Actors	298
Resource Selection and Manager-Actor Communication	299
Cloning Actors and Tapes	299
Logging Actor Activity	300
Installing and Uninstalling Actor Services in Windows	300
Installing and Uninstalling Actor Services in Linux	301
Actor Service Management Functions	301
Launching the Actors	302
Robot Manager Configuration	303
Configuration Overview	303
SCSI Connected Managed Storage	304
Fiber Channel HBA (Host Bus Adapter) and SCSI Persistent Binding	304
Determining the SCSI Library Connection	305
Sony ODA Drives	306
Configuration File Adjustments	308
System Management App Settings and Information	309
System Management App Settings and Information	309
Additional Information	309
Configuring Direct Attached SCSI Managed Storage	310
Common Settings for SCSI-based Managed Storage	310
Configuring ACSLS Attached Managed Storage	312
Configuring LibAttach	312
Testing the LibAttach Connectivity to ACSLS	312
Firewall Support	312
robotmanager.conf Common Options	313

Configuring Sony PetaServe Managed Storage	315
robotmanager.conf Common Options	315
Configuring ADIC Managed Storage with SDLC	318
robotmanager.conf Common Options	318
Configuring Simulated Managed Storage (for DIVA Core Simulators)	319
robotmanager.conf Common Options	319
Robot Manager Command Options	321
Installing and Uninstalling the Robot Manager Services in Windows	321
Installing and Uninstalling the Robot Manager Services in Linux	323
Robot Manager Service Management Functions	323
Testing the Robot Manager Library Interface	323
Starting, Stopping, and Restarting the Robot Manager	324
Testing the Robot Manager Library Control	324
Configuring the Robot Manager at the System Level	325
Logging Robot Manager Activity	326
Configuring Media and Drive Types	326
SCSI_drive_types and ACSLS_drive_types	328
SCSI_tape_types and ACSLS_tape_types	328
ADIC_media_types	328
DIVAmigrate Installation and Configuration	329
DIVAmigrate Embedded Utility Overview	329
Installing DIVAmigrate	329
Windows Files and Folders	329
Linux Files and Directories	330
Configuring the DIVAmigrate Service	330
Configuring the Logging Settings	333

Additional Functionality 335

Checksum Support Configuration	336
Overview	336
Global Checksum Parameters	336
Configuring Checksum Support for Servers	337
Configuring Checksum Support for Arrays and Disks	338
Configuring Checksum Support for Tape Groups	338
Configuring Checksum Support for Actors	339
AXF and TEXT Genuine Checksum Modes	339
Configuring AXF Genuine Checksum Mode	339
DIVA Core System Management App Settings	340
Configuring TEXT Genuine Checksum Mode	340
System Management App Settings	341
Selecting the Root File Path	341
Transcoder Installation and Configuration	342
Transcoder Overview	342
Upgrading from Telestream Vantage 5.0 and Earlier	342
Installing Telestream Vantage	342
Installing the Telestream License	342
Configuring DIVA Core and Transcoders	343

Preparing a Fixed Mount Point for Linux-based Actors (optional)	345
Configuring the Transcoder and Actor on a Single Computer	346
Configuring the Transcoder and Actor on Separate Computers	346
Configuring Telestream Vantage	347
Creating the Output Path	347
Creating a Minimum Vantage Workflow	348
Creating a Complex Vantage Workflow	349
Configuring Transcoders	350
Configuring Source and Destination Servers	350
Disk Auto-Discovery	351
Auto-Discovery Configuration	351
Auto-Discovery Security	353
Auto-Discovery Install, Start, Stop and Uninstall	356
Auto-Discovery Logging	356

Frequently Asked Questions 357

General DIVA Core Questions	358
DIVA Core Database and Backup Service Questions	361

Appendix A: Core Options and Licensing 364

Appendix B: Secure Deployment Checklist 365

Appendix C: Server Guide 366

General Parameters	368
Files Path Root Parameter	368
Root Path Parameter	368
UNIX Style Paths	370
Windows Style Paths	371
Metasource Parameter	371
Connect Options Parameter	371
Quality of Service (qos=)	371
Server FTP User Log In (-login)	372
Server Swift (-oracle_storage_class)	373
Server CIFS User Log In (-user)	373
Server Password (-pass)	373
Server Connection Port (-port)	373
Deleting Source Server Content after Archiving (-allow_delete_on_source)	373
Archiving and Restoring Filename and Path Renaming Rules (-arch_renaming, -	

rest_renaming, -arch_path_renaming, -rest_path_renaming)	374
Using a Temporary Filename when Restoring to a Destination	377
Skipping Files During Restore (-rest_ignoring)	377
Ignoring File Relative Paths (-ignore_relative_path)	378
Archiving Files in a Specific Order (-file_order)	378
Specifying the Transcode Format (-tr_archive_format, -tr_restore_format)	380
Specifying a Transcoder Name (-tr_names)	380
Restoring Metadata (-rest_metadata, -rm)	381
Restricting the Number of Actors to Retry (-num_actors_to_retry)	381
MSS Server in MXF Mode (-mxf)	381
FTP Socket Window Size (-socket_window_size)	382
FTP Socket Block Size (-socket_block_size)	382
FTP Passive Mode Transfers (-pasv)	382
Restoring in AXF Mode (-axf)	383
Specifying Connection Timeouts (-list_timeout, -transfer_timeout, -control_timeout)	383
Alto Disk Archive Integration	384
Configuration	384
System Management App Support	386
Avid MSS (Program Stream) Servers	387
MSS with Independent Storage	388
MSS with Shared Storage	388
MSS with Shared Storage in MXF Mode	389
Using MSS with DIVA_archiveVirtualObject	389
Avid Airspace Servers	390
Avid Transfer Manager DHM Interface	391
Avid Transfer Manager DET Interface	393
SeaChange BMS and BMC Servers	394
SeaChange BML Servers	396
SeaChange BMLe and BMLex Servers	398
Leitch vR Series Servers	401
Leitch Nexio Servers	402
Grass Valley Profile Servers	403
Grass Valley UIM Gateway	405
Grass Valley K2 Servers	407
Grass Valley M-Series iVDR Servers	409
Sony MAV70 Servers	410
Omneon Spectrum MediaDirector Servers (QuickTime)	411
Omneon MediaGrid Content Storage System	413
Quantel Power Portal Gateway	415
Sony Hyper Agent Servers	417
Standard FTP and SFTP Servers	419
Local Source Servers	421
Disk and CIFS Source Servers	422
Metasources	423
Expedat Servers	429

Appendix D: Dynamic Configuration Changes 432

Updates in the Manager Configuration	433
Updates in the System Management App System Page	435
Networks Area	435
Sites Area	435
Servers Area	435
Actors Area	436
Transcoders Area	437
Updates in the System Management App Robots Page	438
Robot Managers Area	438
Media Compatibility Area	438
Robot Managers-ACS Area	438
Updates in the System Management App Disks Page	439
Arrays Area	439
Disks Area	439
Actor-Disk Connections Area	439
Object Storage Accounts Area	440
Updates in the System Management App Drives Page	441
Drives Area	441
Managed Storage Area	441
Drive Properties Area	441
Actor-Drives Area	441
Updates in the System Management App Tapes Page	442
Updates in the System Management App Sets, Tape Groups & Media Mapping Page	442
Updates in the System Management App Analytics App Page	442
Updates in the System Management App Storage Plans Page	442
Updates in the System Management App Slots Page	443
Event Fields	443
Metrics Definitions	451
Configuration Parameter Defaults and Values	486

Appendix E: ADIC SDLC Installation and Configuration 487

SDLC Server	488
Prerequisites	488
Configuration	488
SDLC Client	490
Installation	490
Configuration	490
Using dasadmin Commands	491
Troubleshooting	493

Appendix F:
Backup Service and DBAgent Configuration 494

Sample BKS Configuration File **495**

Sample DBAgent Configuration File **499**

Glossary 501

Telestream Contact Information

To obtain product information, technical support, or provide comments on this guide, contact us using our web site, email, or phone number as listed below.

Resource	Contact Information
DIVA Core Technical Support	<p>Web Site: https://www.telestream.net/telestream-support/diva/support.htm</p> <p>Depending on the problem severity, we will respond to your request within 24 business hours. For P1, we will respond within 1 hour. Please see the Maintenance & Support Guide for these definitions.</p> <ul style="list-style-type: none"> • Support hours for customers are Monday - Friday, 7am - 6pm local time. • P1 issues for customers are 24/7.
Telestream, LLC	<p>Web Site: www.telestream.net</p> <p>Sales and Marketing Email: info@telestream.net</p> <p>Telestream, LLC 848 Gold Flat Road, Suite 1 Nevada City, CA USA 95959</p>
International Distributor Support	<p>Web Site: www.telestream.net</p> <p>See the Telestream Web site for your regional authorized Telestream distributor.</p>
Telestream Technical Writers	<p>Email: techwriter@telestream.net</p> <p>Share comments about this or other Telestream documents.</p>

Preface

This book describes initial and general installation and configuration of the DIVA Core Suite system. The manual assumes a working knowledge of the Windows and Linux operating systems, and additional concepts such as networking, RAID, tape drives, and fiber channel technologies.

Note: This book has been updated for release 8.3.1

Topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Access to Telestream Support](#)
- [Related Documents](#)
- [Document Updates](#)

Audience

This document is intended for the Installation Team, System Administrators, and system users.

Documentation Accessibility

For information about our commitment to accessibility, visit the Support Portal located at:

<https://www.telestream.net/telestream-support/diva/support.htm>

Access to Telestream Support

Customers that have purchased support have access to electronic support through the Support Portal located at:

<https://www.telestream.net/telestream-support/diva/support.htm>

Related Documents

For more information, see the DIVA Core documentation set for this release located at:

<https://www.telestream.net/telestream-support/diva/support.htm>

For information on Cloud Storage visit the following links:

Metered and non-metered Oracle Cloud Storage:

<http://docs.oracle.com/en/cloud/get-started/subscriptions-cloud/csgsg/>

Up to date Oracle Cloud information:

<http://docs.oracle.com/cloud/latest/>

EMC ECS (Elastic Cloud Storage)

<https://www.delltechnologies.com/ru-by/learn/data-storage/ecs.htm>

Amazon S3 Cloud Storage

<https://aws.amazon.com/s3/>

Scality Zenko Integration

<https://www.zenko.io/what-is-zenko/>

Cloudian

<https://cloudian.com/>

NetApp StorageGrid

<https://www.netapp.com/cloud-services/>

Alibaba OSS

<https://www.alibabacloud.com/product/oss>

For further assistance:

<https://www.telestream.net/telestream-support/diva/support.htm>

Document Updates

The following table identifies updates made to this document.

Date	Update
December 2022	Updated book for release 8.3.1: <ul style="list-style-type: none"> • Added Object Auto-Indexing information. • Added Library Storage Configuration information. • Updated various content to clarify where what settings are to be configured; Configuration Utility or System Management App. • Added or updated various screenshots.
January 2023	<ul style="list-style-type: none"> • Updated copyright dates. • Updated Appendix F: <ul style="list-style-type: none"> – Updates to BKS and DBAgent configuration files.

Overview

This chapter gives an overview of the DIVA Core 8.3.x release.

Topics:

- [Release Overview](#)
- [New Terminology](#)
- [Port Utilization](#)
- [Enhanced Features and Functionality](#)
- [Media Storage Formats](#)
- [Complex Objects](#)
- [DIVA Core Software Components](#)
- [Additional Software Components](#)
- [DIVA Core Utilities](#)

Release Overview

The DIVA Core architecture enables integration of many different types of servers and technologies, for example Broadcast Video Servers, Storage Area Networks, and Enterprise Tape Managed Storage. The DIVA Core installation varies from site to site, therefore the exact configuration of your specific DIVA Core platform is not described in this book.

Notes: The File System Interface is not released with DIVA Core 8.x and is only supported by special request.
DIVA Command has been deprecated.

See [Appendix A: Core Options and Licensing](#) for detailed information.

New Terminology

The following terminology has been updated to reflect standardization efforts across all DIVA and Kumulate applications. There will be some variations in the documentation compared to the interface until everything is switched over to the new terminology; the documentation uses the new terms wherever possible.

Note: DIVA command has been deprecated.

- Running Requests are now called Jobs
- Request History is now called Job History
- Libraries are now called Managed Storage
- Datahub is now called Actor
- Proxyhub is now called Proxy Actor
- DIVA Core and DIVA Manager are now called DIVA Core / Core / Core Manager
- Category is now called Collection
- Source/Destination is now called Unmanaged Storage Repository
- Storage Repository is now called Managed Storage Repository
- Group is now called Tape Group
- Link is now called Storage Link
- Storage Plan Manager is now called Storage Policy Manager
- Drop Folder Monitor (DFM) is now called Watch Folder Monitor (WFM)
- DIVA Configuration Utility and Control Panel are now called System Management App
- DIVA Analytics and DIVAProtect are now called Analytics App

Port Utilization

The following table lists the standard ports used by the DIVA Core system. Contact Technical Support for assistance in necessary.

DIVA Core Service	Default Port	Needed by External Applications	Description and Notes
FTP	21 / TCP	Yes	Port depends on configuration
SSH	22 / TCP	Yes	Linux hosts only
HTTP	80 / TCP	Yes	DIVA View
SQLNet	1521 / TCP	Yes	Manager database access
RDP (Microsoft Terminal Services)	3389 / TCP	Yes	Remote desktop access
Flip Factory	9000 / TCP	Yes	Typically installed on Manager Server
Actor	9900 / TCP 8800 / UDP	No	Typically installed on Actor Server
Auto Discovery Data	7443 / TCP	No	Typically installed on Manager Server
Auto Discovery Publisher	8443 / TCP 11443 / TCP	No	Typically installed on Manager Server
AMC	6101 / TCP	Yes	For Avid AMC, typically installed on Manager Server
DB Agent	1878 / TCP	No	Typically installed on Manager Server
Database Backup service (BKS)	1877 / TCP 1876 / TCP	No	Typically installed on Manager Server
REST API Discovery (DIVA Core 8.0 and later)	8761 / TCP	Yes	Typically installed on Manager Server
REST API Data Service (DIVA Core 8.0 and later)	13443 / TCP	Yes	Typically installed on Manager Server
REST API Gateway (DIVA Core 8.0 and later)	8765 / TCP	Yes	Typically installed on Manager Server
REST API DIVA Connect Adapter (DIVA Core 8.0 and later)	17443 / TCP	No	Typically installed on Manager Server

DIVA Core Service	Default Port	Needed by External Applications	Description and Notes
Manager	9000 / TCP 8000 / TCP 12443 / TCP	Yes	Typically installed on Manager Server. 9000 / TCP is an unsecure port 8000 / TCP is a secure port
Metadata Service	1776 / TCP 1777 / TCP	No	Typically installed on Manager Server
MongoDB	27017 / TCP	No	Typically installed on Manager Server
Oracle Database	1521 / TCP 1522 / TCP	Yes	Typically installed on Manager Server
DIVAmigrate	9191 / TCP	Yes	Typically installed on Manager Server
Robot Manager	8500 / TCP	No	Typically installed on Manager Server
VACP	5010 / TCP	Yes	Typically installed on Manager Server
DIVA Connect DB Sync	9802 / TCP	No	Typically installed on DIVA Connect Server
DIVA Connect Client Adapter	9801 / TCP 7101 / TCP	No Yes	Typically installed on DIVA Connect Server
DIVA Connect Manager Adapter	9800 / TCP	No	Typically installed on Manager Server
Web Services Admin Server	9443 / TCP	Yes	Typically installed on Manager Server
Web Services Application Server (WS API 2.1)	9763 / TCP	Yes	Typically installed on Manager Server
Enterprise Connect Admin Server (WS API)	7001 / TCP	Yes	Typically installed on Manager Server
Enterprise Connect Application Server (WS API 2.2)	9443 / TCP	Yes	Typically installed on Manager Server

Enhanced Features and Functionality

Refer to the DIVA Core Release Notes in the DIVA Core documentation library at:

<https://www.telestream.net/telestream-support/diva/support.htm>

Media Storage Formats

This section describe the media storage formats available in this DIVA Core release.

AXF (Archive eXchange Format)

Archive eXchange Format is an open format that supports interoperability among disparate content storage systems and ensures the content's long-term availability no matter how storage or file system technology evolves.

An AXF object is an IT-centric file container that can encapsulate any number, and any type, of files in a fully self-contained and self-describing package. The encapsulated package contains its own internal file system, which shields data from the underlying operating system and storage technology. It's like a file system within a file that can store any type of data on any type of storage media.

Tape groups or disk arrays used by Complex Object requests must be in an AXF format, because Complex objects cannot be stored in Legacy format. Because all Complex objects are written in the AXF format, any instance of a Complex object will also be in the AXF format.

Native File and Folder Support

Users can see their files and folders in native format on archive devices rather than as an AXF container files. Files and folders on storage devices like Object storage can also be accessed. This access opens the archive to the use of third party software to perform operations on the archive (for example, metadata collection, face recognition, transcoding, and so on).

Tape Groups

In DIVA Core, a Tape Group or Disk Array has a media format parameter that indicates which storage media format to use when creating Archived Objects. Set the media format to either DIVA Core Legacy Format or the AXF Format. This setting can be changed at any time and does not influence content already stored. It is possible to have more than one storage media format within tape groups and disk arrays.

A DIVA Core Object instance is only written in one media format. Therefore, if an Object spans tapes, each tape used as part of an Object instance will be written in the same media format. An Object can contain multiple instances, each of which can be stored in either Legacy or AXF format.

Although a tape group can contain more than one storage format, an individual tape has at most one storage media format. The format of a tape instance is the format of the tape on which the instance resides. All instances on a tape must have the same format.

The media format for an empty tape is assigned when the first Object on that tape is written. The tape is assigned the format of the tape group that appears in the request.

After the media format for a tape is assigned, it cannot be changed unless all Objects on the tape are deleted. After deletion of all Objects from a tape, the tape's format becomes unassigned until content is again written to the tape. If the tape was in use, the tape format cannot change unless it is empty and cleared.

Both Legacy and AXF formatted tapes can exist in the same group. Nevertheless, Objects in AXF format will only be written to AXF formatted tapes, and Objects in Legacy format will only be written to Legacy formatted tapes, even though they are in the same tape group.

Note: A Repack request will always write the destination tape in the same media format as the source tape.

Similarly, tape spanning operations will always use the same format across all tapes storing spanned Objects. If an instance spans across multiple tapes, then all tapes used to span the content will have the same format.

Disk Arrays

Unlike tapes, disks do not have a format. DIVA Core allows storing Objects in different media formats on the same disk. If a disk contains Objects in Legacy format, and that disk is then assigned to an AXF formatted array, it will still contain Objects in Legacy format. However, new Objects written to the disk will be in AXF format.

If a disk instance is non-complex and permanent (not a cache instance), it is stored in the format of the Destination Server array. If a cache instance is non-complex, it is stored in the format of the group specified in the request.

Use the Copy To Group, or Copy As New requests to migrate Objects from Legacy media format to AXF media format (or back). However some AXF Objects cannot be copied to the Legacy format; copying Objects from Legacy format to AXF format does not present any issues. In DIVA Core the only limitation on copying an Object instance from AXF format to Legacy format is the Complex Object feature.

Complex Objects

Complex Objects have significantly expanded the object component boundaries, allowing up to one million files and ten thousand folders per object.

Note: The minimum server operating system for using Complex Objects is Windows Server 2016.

Complex Objects maintain information about files and folders in the archive. They store subtotals for each folder, including the total number of files and subfolders within the folder, and the total size of all files within the folder and within any subfolders.

DIVA Core uses the configurable Complex Object Threshold parameter during archival to determine whether a new object should be complex based on the number of components. This value is set in the manager.conf configuration file. If the number of components is greater than the Complex Object Threshold, the object becomes a Complex Object. After an object is identified as a Complex Object it will always be complex; even if it is copied using the Copy As command, or imported using the Export/Import Utility.

It is recommended that the threshold remain at the default value (1,000 components) unless there is a specific reason to adjust the value. Contact Technical Support for assistance as required.

A Complex Object differs from a non-complex object in several key ways. For example, the file and folder metadata information of a Complex Object is stored in a file, not in the DIVA Core Database. The file contains the file names, folder names, checksums, and files sizes. The files are located in the Metadata Database root directory. Complex Objects must be stored in AXF format whether on tape or on disk.

Complex Objects can contain hundreds of thousands of files. However, some DIVA Core API commands (for example, GetObjectInfo) will not return the entire set of files. Instead, these commands return a single placeholder file which prevents downstream applications from being overwhelmed by file and folder information. Also, the entire set of files on a tape are not displayed in the System Management App Object Properties and Tapes screens, only a single placeholder file is shown. The DIVA Core API includes a command to return all of the files and folders within a Complex Object. See the appropriate DIVA Core API documentation in the DIVA Core documentation libraries for details.

DIVA Connect does not currently support replication of Complex Objects.

The following features do not support Complex Objects:

- Delete on Source Server option
- Verify on Restore (VFR) checksum feature
- Verify on Archive (VFA) checksum feature
- deleteFile API call
- getObjectListbyFileName API call

- GetByFilename API call (for Avid connectivity)
- DeleteByFilename API call (for Avid connectivity)

Complex Objects and FTP

When archiving Complex Objects using the FTP protocol, and an FTP Client with default settings (FileZilla is recommended), the transfer will typically fail when archiving any object with more than approximately 3,900 files.

Occasionally, during the directory scan, the Actor connection times out before the size of the object can be computed. More often, a request terminates in the middle of the transfer because the FTP server is consuming all of the available sockets.

Add the following parameters in the Server Command Options or in the Options of the command itself to resolve timeout issues:

```
-transfer_timeout 1200  
-list_timeout 600
```

See [Appendix C: Server Guide](#) for detailed parameter information.

Use the following procedure to include the parameters in the Server page in System Management App:

1. Open the DIVA Core System Management App.
2. Navigate to the System page.
3. Double-click the desired Server in the Servers page to open the edit dialog box.
4. Add the two parameters (*-transfer_timeout 1200* and *-list_timeout 600*) in the Connect Options field.
5. Click OK to save the changes.
6. Notify the Manager of the changes using the Control+N key combination.

We recommend setting the following corresponding parameters in the FileZilla server under General Settings:

Connections Timeout = 600

No Transfer Timeout = 1200

1. Open the FileZilla server interface.
2. Click the Server Options icon on the tool bar.
3. Adjust the settings in the General Settings area.

If requests terminate unexpectedly during transfers, adjust the Windows Registry parameters as follows:

1. Open regedit.
2. Modify (or create) the following values under
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters:

```
TcpTimedWaitDelay = 10
```

MaxUserPort = 90000

- a.** If the desired registry parameter does not exist, right-click and create a new double word (DWORD) value.
 - b.** If the parameter does exist, double-click it and enter the values.
- 1.** Restart the computer to enable the new registry values.

DIVA Core Software Components

DIVA Core includes the DIVA Core software components discussed in the following subsections. All DIVA Core components support Oracle Linux 7 x86_64 and later. See the DIVA Core Supported Environments Guide for information about certain limitations when running in the Linux environment.

Long path names are supported on both Windows and Linux. Absolute path names are supported on both Windows and Linux to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems (only).

All Windows batch files (.bat) have corresponding shell scripts (.sh) in Linux. Windows paths must be substituted with Linux paths when operating on Linux. For example, the Windows path C:\DIVA\Program equates to /home/diva/DIVA/Program in the Linux environment.

Note: Linux commands, paths, and file names are case-sensitive.

Archive and restore operations of symbolic Storage Links are now supported in Linux. Shortcuts created using the Windows operating system are not represented as symbolic Storage Links because they are treated as files. Only symbolic Storage Links created on the UNIX platform are archived and represented as symbolic Storage Links in DIVA Core.

The Java and C++ APIs file list returned from a `getFilesAndFolders` call includes symbolic Storage Links, and the export and import operations type attribute contains the letter S to represent a symbolic Storage Link.

The following features require Windows-based Actors:

- DIVA Core Avid Connectivity
- Transcoder integration
- Tape Reading Utility

On Linux Actors, standard commands like DD and MT are alternatives of Tape Reading Utility. Linux Actors support QuickTime, GXF, and MXF and MPEG2 Transport stream wrappers for DIVA Core Partial File Restore (video).

See [Appendix A: Core Options and Licensing](#) for detailed information.

Due to degraded performance, Windows IIS and FileZilla FTP Source Servers and Destination Servers cannot be used for Complex Objects. The software only supports Linux-based FTP servers when operating in a Linux environment. The Windows IIS and FTP servers cannot accommodate large numbers of files.

Manager

The Core Manager is the main component in a DIVA Core system. All archive operations are controlled and handled by the Core Manager. Operation requests are sent by initiator applications through the DIVA Core Client API. DIVA Core supports Main and Backup Core Managers. See [Appendix A: Core Options and Licensing](#) for detailed information.

Actor

The Core Actor is the data mover between devices in the network. Actor supports interfacing and data transfer between many different types of devices.

All Actor operations are initiated and coordinated by the Core Manager through a TLS 1.2 secure connection. Key benefits of the distributed design of the Core Actors are:

- Expand the archive subsystem to increase the overall bandwidth by adding more Actors to the system. See [Appendix A: Core Options and Licensing](#) for detailed information.
- Share SAN based disk and tape drive resources among multiple Actors.
- In combination with the Core Manager, multiple Actors provide scalability, load balancing, redundancy, and failover. Take individual Actors offline for maintenance without shutting down the DIVA Core system.

Note: UNC paths are supported for SMB Servers and managed disks if the UNC path is mounted directly on Windows Actors.

DIVA Core 7.5 and later supports archive and restore of empty files and folders. Empty files and folders are only supported by AXF. When Legacy format is in use, DIVA Core reports an error if an empty file or folder is discovered during the transfer.

Client APIs

The DIVA Core Client APIs are a set of functions enabling external applications, acting as clients, to use the services offered by the DIVA Core system.

A library of client functions is provided with the selected API and must be linked to each DIVA Core client application. These functions encapsulate client commands into DIVA Core request messages sent over a TCP/IP connection to the Core Manager.

Currently available APIs include REST API, C++, Java, and Web Services (DIVA Enterprise Connect). Refer to the appropriate DIVA Core API documentation and the DIVA Enterprise Connect documentation for more information.

REST API

DIVA Core exposes its functionality through a REST interface. It is self-contained in DIVA Core 8.0 and all future DIVA Core releases. In the 8.0 release, the API is used exclusively by the DIVA Core Web Application.

Note: Telestream recommends using the REST API rather than the previous existing APIs (that is, DIVA Core Enterprise Connect, DIVAS, Java and C++). Although all previous APIs will remain available, the REST API offers new and enhance features.

See the REST API Programmer's Guide for detailed information.

Database

The DIVA Core software is bundled with an Oracle database installation. The database stores all information relating to the DIVA Core system including its configuration. SQL queries used by the Manager are optimized to support configurations with up to 58 million components.

In DIVA Core 7.5 and later, the JDBC Thin Driver enables replacing the Oracle SID setting with the Oracle Service Name.

When installing DIVA Core in a 64-bit environment, the latest 64-bit DIVA Core Oracle release must be installed to use 64-bit support.

DIVA Core 8.3 supports Oracle Database 11.2.0.4 or greater.

Oracle 19c can be used with DIVA Core 8.0 and later, and supports the following Windows and Linux Oracle packages:

- OracleDivaDB_3-2-0_19_3_0_0_0_SE2_Windows_64-bit
- OracleDivaDB_3-3-0_19_3_0_0_0_SE2_Linux_x86_64

Note: The Oracle database is not intended to be modified directly by customers, but rather by using Oracle utilities. Direct modification of this database by customers through Oracle utilities is not supported.

Metadata Database

DIVA Core stores object metadata separately from the Core Database in the DIVA Core Metadata Database (MongoDB). The metadata database contains files stored in a file system local to the Core Manager. The files are located in the Metadata Database root folder. This storage method enables DIVA Core to effectively operate with large volumes of files, folders and other metadata.

The metadata database is very high performance, and has almost unlimited scalability. Treat the metadata database with the same caution as the Core Database, and it must be backed up at regular intervals through the DIVA Core Backup Service.

Note: MongoDB, in its default configuration, can use up to half the available RAM minus 1GB on the server on which it is installed. Plan the location of MDS MongoDB installation accordingly.

Notification Service (RabbitMQ)

RabbitMQ has been integrated into the DIVA Core windows installer starting with release 8.2.0.91 (and later). The DIVA Core installer identifies it as Notification Service instead of RabbitMQ because RabbitMQ is just an implementation.

The Notification Service is required for the System Management App to function properly.

System Management App

The DIVA Core System Management App is installed as part of the DIVA Core 8.3 installer (or later). It is hosted by the Manager Service. Installing the Manager and the REST API Data services will automatically set it up; see the [REST API Installation and Configuration](#) section for instructions.

The DIVA Core System Management App connects to both the Core Manager and the Core Database. Use it to monitor, control, and supervise operations in DIVA Core. Multiple System Management App instances can be operated simultaneously from any computer that has TCP/IP connectivity to both the Core Manager and the Core Database.

The System Management App is not intended for the intensive archive operations of a DIVA Core system. Archive operations are typically initiated to DIVA Core from a Broadcast Automation or MAM (Media Asset Management) system.

See the DIVA Core Operations Guide in the DIVA Core documentation library for more information on using the interface.

The refresh rate for the System Management App is set in the Manager Setting page of the System Management App in the GUI: Dashboard Refresh Delay field.

Additional Software Components

Additional modules are available to expand the DIVA Core system capabilities. Most of these options are currently covered in separate documents, but are briefly described here for completeness.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Robot Manager

DIVA Core can be used to only manage disk storage, but storage capacity can be further expanded by adding one or more tape Managed Storage. In these cases, the Core Robot Manager module provides an intermediate software layer for the Core Manager to interact with many different types of tape Managed Storage. It is connected to the Core Manager through TCP/IP.

See [Appendix A: Core Options and Licensing](#) for detailed information.

This distributed architecture provides substantial flexibility including:

- Managed Storage controlled using a SCSI interface are limited by the cable length. Because the connection to the Core Robot Manager from the Core Manager is over TCP/IP, the library does not need to be co-located near the Core Manager host computer.
- Enabling installation of multiple, or dissimilar, Managed Storage by configuring additional Core Robot Manager modules.
- Enabling rapid development to support new types or models of Managed Storage.
- Restart the robotics interface without needing to restart the Core Manager.

The Core Robot Manager interfaces with the library using either a direct interface to the library itself (through native SCSI, or SCSI over Fiber Channel), or through an intermediate Ethernet connection to the manufacturer's own library control software.

VACP (Video Archive Communications Protocol) Service

The Video Archive Communications Protocol is developed by Harris Automation Solutions and used by some automation systems for interfacing to an archive system. DIVA Core has its own API for communicating with the Core Manager, which is not compatible with VACP.

To provide interoperability without the need to redevelop the archive interface at the automation level, this module is provided to act as an interface to convert VACP commands from the attached automation system to DIVA Core API commands on computers that have TCP/IP connectivity to DIVA Core.

SPM (Storage Policy Manager)

The DIVA Core Storage Policy Manager provides automatic migration and life cycling of material within the archive, based on the rules and policies defined in the SPM configuration. The DIVA Core DSM (Disk Space Monitor) works with SPM to delete material from SPM managed arrays (based on disk space watermarks).

SNMP (Simple Network Management Protocol) Agent

The DIVA Core Simple Network Management Protocol interface supports status and activity monitoring of different DIVA Core components. DIVA Core MIB (Management Information Base) is provided to third party SNMP monitoring applications. The SNMP Agent uses the Windows SNMP Service and has not been ported to the Linux environment.

DIVA Connect

DIVA Connect provides DIVA Core client authentication and authorization. It can act as an intermediate gateway between DIVA Core components (for example the VACP converter) or third party applications and the Core Manager, and can restrict that component or application from access to the DIVA Core system.

DIVA Connect is a powerful feature that allows multiple DIVA Core platforms to exchange archive resources and content, whether the archive systems are local to each other or remote.

The DIVA Connect is used in DIVA Connect installations and is the portal for multiple DIVA Core systems to communicate with each other. See the DIVA Connect Installation, Configuration, and Operations Guide for more information.

Watch Folder Monitor

The DIVA Core Watch Folder Monitor provides automatic monitoring of newly created files in multiple local directories or FTP folders (or combinations thereof). One file, or multiple files, per DIVA Core object are supported. When a new file is identified, WFM issues an archive request automatically to DIVA Core to archive the new file. After the files are successfully archived, they are then automatically deleted from the Source Server. Refer to the DIVA Core Watch Folder Monitor User's Guide for more information.

When WFM is used in a Linux environment to monitor an FTP folder, it must be configured as in the following example:

User

diva

User Home Directory

`/ifs`

Folder to be monitored

`/ifs/folder1`

A correct WFM configuration with these parameters is:

`ftp://diva:password@host_ip/folder1`

An incorrect WFM configuration with these parameters is:

`ftp://diva:password@host_ip/ifs/folder1`

Transcoder Support

The Core Actor can integrate with a transcoder engine to provide real time transcoding of material as it is archived or restored, or to create objects from already existing content within the archive. Currently, integration to BitScream products, and Telestream Flip Factory, and Telestream Vantage are supported. However, multiple transcoders are only supported for Vantage.

Note: DIVA Core 7.5 ended technical support for Telestream Flip Factory. We will provide best efforts to assist customers to transition to other transcoding solutions.

Linux-based Actors only support Telestream Vantage for transcoding operations.

DIVA Core assumes a local transcoder address of 127.0.0.1 if a transcoder address is not specified in the transcoder's working directory.

The Promedia Carbon (formerly Rhozet) transcoder is supported in DIVA Core. Select the transcoder type "tre" from the System Management App to use this transcoder. Both the Name and GUID are supported as options for Presets and Profiles format types.

Avid Connectivity

The following sections describe general Avid connectivity with DIVA Core.

See the DIVA Core Avid Connectivity User's Guide in the DIVA Core documentation library for more information. Also see [Appendix A: Core Options and Licensing](#) for detailed information.

Avid DHM (Data Handler Module) Interface

The Avid DHM (Data Handler Module) interface support in DIVA Core enables finished content to be shared between post-production Avid environments and On Air Video servers. This eliminates the need for tape based content exchange. Timecode based

Partial File Restores of content to On Air environments, and finished Avid Sequence submissions to On Air servers are key to the DHM functionality offered within DIVA Core. DHM support is implemented in DIVA Core TMC (Transfer Manager Communicator).

Avid DET (Dynamically Extensible Transfer) Interface

The Avid DET (Dynamically Extensible Transfer) interface support in DIVA Core allows storage expansion of Avid Unity infrastructures and enables editors to move native Avid content in and out of the DIVA Core storage system. Partially edited content stored within DIVA Core through the Avid DET interface can be later restored to Unity, and an editor can then resume editing at the point where they stopped. DIVA Core stores these files in native Avid format. DET support is implemented in DIVA Core TMC.

Archive Manager Interface

An interaction between the Avid Archive Manager solution and DIVA Core is implemented in a separate service called AMC (Archive Manager Communicator). AMC handles Archive, Restore, Partial File Restore, and Delete commands from the Avid Archive Manager using DIVA Core to store Avid content in its native MXF OP1 Atom format.

Analytics App

The Analytics App option is a utility that collects operational statistics from the DIVA Core system to monitor and maintain the archive's subcomponents (servers, media, drives, tapes, and so on). Analysis of these statistics allows both proactive and reactive maintenance of the DIVA Core system. See the DIVA Core Analytics App User's Guide for more information.

OTU (Object Transfer Utility)

The Object Transfer Utility is an optional feature of the System Management App providing a drag and drop interface to archive and restore content between DIVA Core and a (supported) Source Server or Destination Server. See [Appendix A: Core Options and Licensing](#) for detailed information.

DIVA View

DIVA View is a Graphical User Interface for DIVA Core 7.4 and later. The web interface can locate and work with videos remotely, from any computer on the network. Administrators monitor and maintain the system and user accounts with minimal effort from anywhere they have access to the network. See the DIVA View documentation for details.

DIVAmigrate

DIVAmigrate is installed as part of the DIVA Core Suite's standard installation. It is located in the %DIVA_HOME%\Program\ folder, and runs as a Windows Service.

DIVA Core Utilities

The following sections describe utilities available in the DIVA Core system

System Management App

DIVA Command has been deprecated and is replaced by the System Management App starting with this release to configure a DIVA Core system. It can be run on any computer that has TCP/IP connectivity to the host running the Core Database.

Caution: The System Management App is intended only for experienced users. Incorrect or incomplete changes in the System Management App can adversely affect DIVA Core operations, possibly delete data from the archive, or prevent the Manager from running. If unsure about making changes, contact Telestream Support for assistance before attempting to make alterations to the system configuration.

The System Management App primarily connects to the Core Database, and for some tasks, directly to the Robot Managers. After launching the utility you must first connect to the database to edit the DIVA Core system configuration. Although used primarily for configuration of DIVA Core, some operational functions are also performed from the System Management App.

Robot Manager Utilities

During configuration and troubleshooting of the library and its tape drives, DIVA Core provides both a command-line interface and GUI utility to send commands directly to the tape library through the Robot Manager. These utilities are not (and must not be) used while the Core Manager is running because this can adversely affect archive operations. See [Appendix A: Core Options and Licensing](#) for detailed information.

Backup Service

The DIVA Core Backup Service ensures reliability and monitoring of both the DIVA Core Database and Metadata Database backups.

The DIVA Core Backup Service component is installed as an integral part of the standard DIVA Core system installation. The component is typically installed on the same server as the Core Manager and Core Database. The DIVA Core Backup Service enables configuration of scheduled backups through its configuration file. The DIVA Core Backup Service manages and monitors the entire backup process.

See [DIVA Core Backup Service \(BKS\)](#) for more information.

Scandrive Utility

This utility is provided on both Windows and Linux platforms. It assists in obtaining detailed device information such as serial numbers, firmware releases, and SCSI information from tape Managed Storage or tape drives for use in the DIVA Core configuration.

Tape Reading Utility

Caution: This utility must not be used while the Core Manager is running.

This utility is provided on both Windows and Linux platforms and is primarily used with the Robot Manager Client utilities to send manual Eject commands to a tape drive connected to an Actor. This utility also provides advanced tape based operations, such as tape formatting, but should only be used under guidance from Technical Support.

The Tape Reading Utility is only supported by Windows-based Actors. Standard commands must be used, for example, DD and MT when operating in a Linux environment.

DIVAscript

This utility allows DIVA Core C++ API commands to be executed using UNIX or DOS based scripts. It is designed to run automated tasks for testing rather than for any intensive uses. There is no Linux-based DIVAscript release.

RDTU (Recover Damaged Tape Utility)

The DIVA Core Recover Damaged Tape Utility is designed to recover object instances contained on a damaged tape. The utility can recover instances that have valid copies on other available media (that is, internal tape, or a connected disk or array) within a local or remote DIVA Core system.

DIVA Core Security

These general principles of DIVA Core application security should be adhered to for a secure system. Also see the [Port Utilization](#) section for additional detailed port requirements.

Topics:

- [General Security Principles](#)
- [Secure Installation](#)
- [Security Features](#)

General Security Principles

Keep Software up to Date

Stay current with the version of DIVA Core that being run. Current versions of the software are available for download at the Software Delivery Cloud located at:

<https://www.telestream.net/telestream-support/diva/support.htm>

Restrict Network Access to Critical Services

DIVA Core uses the following TCP/IP ports:

- Core Robot Manager uses 8500 / tcp
- Core Manager uses 8000 / tcp for secure connections (this is the default), and 9000 / tcp to accommodate legacy versions of the DIVA Core API to connect to the Core 8.3 Manager.
- DIVA Core Backup Service uses 9300 / tcp
- DIVA Connect uses 9500 / tcp
- Core Actor uses 9900 / tcp
- DIVA Core Migrate Service uses 9191 / tcp
- Core Proxy Actor uses 8800 / udp
- REST API uses https://localhost:8765 or https://127.0.0.1:8765

Run as DIVA User and use Principle of Least Privilege Where Possible

Do not run DIVA Core services using an Administrator (or root) operating system user account. Always run all DIVA Core services using a dedicated operating system user (or Tape Group) named DIVA.

The DIVA Core System Management App provides three fixed user profiles (Administrator, Operator, and User). The Administrator and Operator accounts require a password to obtain access. An Administrator and Operator password must be assigned in the System Management App before using these profiles.

Passwords are created during installation and configuration for both the Administrator and Operator accounts. The passwords must be changed every 180 days (minimum) thereafter. Passwords must be made available for Technical Support if needed.

Monitor System Activity

Monitor system activity to determine how well DIVA Core is operating and whether it is logging any unusual activity. Check the log files located in the installation directory under /Program/log/.

Keep up to Date on Latest Security Information

For security information and alerts for a large variety of software products, see <http://www.us-cert.gov>.

The primary way to keep up to date on security matters is to run the most current release of the DIVA Core software.

Secure Installation

This section outlines the planning process for a secure installation and describes several recommended deployment topologies for the systems.

Understand Your Environment

To better understand security needs, the following questions must be asked:

Which Resources Need to be Protected?

Many of the resources in the production environment can be protected. Consider the type of resources that to protect when determining the level of security to provide.

When using DIVA Core, protect the following resources:

Primary Data Disk

There are Data Disk and Cache Disk resources used to build DIVA Core systems. They are typically local or remote disks connected to the DIVA Core systems. Independent access to these disks (other than by DIVA Core) presents a security risk. This type of external access might be from a rogue system that reads or writes to these disks, or from an internal system that accidentally provides access to these disk devices.

Database Disk, Metadata Disk, and Backup Disks

There are Database Disk, Metadata Disk and Backup Disk resources used to build DIVA Core systems with Complex Objects. They are typically local or remote disks connected to the DIVA Core systems. Independent access to these disks (other than by DIVA Core) presents a security risk. This type of external access might be from a rogue system that reads or writes to these disks, or from an internal system that accidentally provides access to these disk devices.

DIVA Core Tapes

It is a security risk to allow independent access to tapes, typically in a tape library controlled by DIVA Core systems, where data is written.

Export Tape Metadata

Tape Metadata dumps that are created from export operations contain data and metadata. This data and metadata permissions must be restricted to only the Administrator (or root) operating system account, or the DIVA operating system user (or Tape Group) during a routine export or import activity.

Configuration Files and Settings

DIVA Core system configuration settings must be protected from operating system level non-administrator users. Making the configuration files writable to non-administrative operating system users presents a security risk, therefore, these file permissions must be restricted to only the Administrator (or root) operating system account, or the DIVA operating system user (or Tape Group).

From whom are the resources being protected?

In general, the resources described in the previous sections must be protected from all non-administrator access on a configured system, or from a rogue external system that can access these resources through the WAN or FC Fabric.

What will happen if the protections on strategic resources fails?

Protection failures against strategic resources can range from inappropriate access (that is, access to data outside of normal DIVA Core operations) to data corruption (writing to disk or tape outside of normal permissions).

Recommended Deployment Topologies

This section describes how to install and configure an infrastructure component securely. Consider the following points when installing and configuring DIVA Core:

Separate Metadata Network

For connections between DIVA Core services components, connection to Metadata Database, and the connection from its clients, provide a separate TCP/IP network and switch hardware that is not connected to any WAN. Because the metadata traffic is implemented using TCP/IP, an external attack on this traffic is theoretically possible. Configuring a separate metadata network mitigates this risk and also provides enhanced performance. If a separate network is infeasible, at least deny traffic to the DIVA Core ports from the external WAN and any untrusted hosts on the network. See [Restrict Network Access to Critical Services](#).

FC Zoning

Use FC Zoning to deny access to the DIVA Core disks connected through the Fiber Channel from any server that does not require access to the disks. Preferably, use a separate FC switch to physically connect only to the servers that require access.

Safeguard SAN Disks Configuration Access

SAN RAID disks can usually be accessed for administrative purposes through TCP/IP or more typically HTTP. Protect the disks from external access by limiting the administrative access to SAN RAID disks to systems only within a trusted domain. Also, change the default password on the disk arrays.

Install the DIVA Core Package

First, install only required DIVA Core services for your environment. For example, if you do not plan to run the GUI or System Management App from a system, then deselect them in the list of components to be installed during installation. The default DIVA Core

installation directory permissions and owners must be restricted to only the Administrator (or root) account, or the DIVA operating system user (or Tape Group).

DIVA Core Tape Security

Prevent external access to DIVA Core tapes inside a tape library controlled by the DIVA Core system. Unauthorized access to DIVA Core tapes can compromise or destroy user data.

Backups

Set up and perform database backups using the DIVA Core Backup service. Permissions for the backup dump must be restricted to only the Administrator (or root) operating system account, or the DIVA operating system user (or Tape Group).

Post-Installation Configuration

After installing any of the DIVA Core, go through the security checklist in [Appendix B: Secure Deployment Checklist](#).

Security Features

To avoid potential security threats, customers operating DIVA Core must be concerned about authentication and authorization of the system. These security threats can be minimized by proper configuration and by following the post-installation checklist in [Appendix B: Secure Deployment Checklist](#).

The Security Model

The critical security features that provide protections against security threats are:

Authentication

Ensures that only authorized individuals are granted access to the system and data.

Authorization

Access control to system privileges and data. This feature builds on authentication to ensure that individuals get only appropriate access.

Tape Group Encryption

Tape drive encryption securely supports bulk tape migration between DIVA Core systems.

SSL Authentication and Secure Communications

DIVA Core 8.3 includes SSL Authentication for services, and to secure DIVA Core internal and API communications. Certificate authentication provides unique identification and secure communication for each DIVA Core Service in a network.

Authentication

The DIVA Core System Management App provides three fixed user profiles (Administrator, Operator and User). The Administrator and Operator accounts require a password to obtain access. An Administrator and/or Operator password must be assigned in the System Management App before using these profiles.

Both the Administrator and Operator account passwords must be changed every 180 days (or before). Passwords must be made available for Technical Support if needed.

Access Control

Access control in DIVA Core is divided into three profiles. The Administrator and Operator accounts require a password to obtain access. An Administrator and/or Operator account password must be assigned in the System Management App before using these profiles.

User

After the connection to the Core Manager is established, the System Management App will only allow the user to monitor DIVA Core operations, and retrieve data from the database. This is known as the User Profile. Not all functions that issue commands to DIVA Core are accessible while in the User profile mode, enabling situations where monitoring is required but no commands are permitted to be sent to DIVA Core.

Administrator

To issue requests to DIVA Core, such as archive or restore requests, or to eject a tape from a library, the Administrator Profile must be used. The Administrator Profile is password protected. The password for this profile must be assigned in the System Management App before using the profile. For more information, refer to the DIVA Core 8.3 Customer documentation located at <https://www.telestream.net/telestream-support/diva/support.htm>.

Operator and Advanced Operator

In addition to User Profile permissions, the Operator Profile provides access to the Object Transfer Utility and requires a password configured in the System Management App before using the profile. Both Operator and Advanced Operator profiles in the System Management App can now optionally enable privileges for canceling and changing the priority of requests. The options are defined in the Manager Configuration panel of the System Management App. By default, this option is disabled.

Tape Group Encryption

The DIVA Core 8.3 release includes tape drive encryption that securely supports bulk tape migration between DIVA Core systems.

After enabling encryption on a tape group, all additional tapes added to the group will also be encrypted. However, any existing tapes in the group remain unencrypted if encryption was previously disabled.

Enabling encryption on a tape group generates an encryption key, which is also encrypted. The encryption key can be changed at any time. New tapes added to the group after the change will use the new encryption key. The existing tapes that were already encrypted will continue to use the original key. Therefore, tapes in the same tape group can have different encryption keys. The Manager must be notified of the change when updating the encryption key.

Disabling encryption (after it is already enabled) only affects additional tapes added to the group, and the existing tapes remain encrypted.

S3 Server-side Encryption

In DIVA Core 8.3, three new options can be added to the Storage options field in order to configure S3 server-side encryption.

Option Syntax	Description
-server_side_encryption=<KMS or AES256>	This option must be set in order to enable S3 server-side encryption. DIVA Core supports SSE-S3, SSE-OSS and KSM. Set this option to AES256 if you want to use SSE-S3 or SSE-OSS. Set this option to KMS if you need to use KMS.
-kms_key_id=<kms key ID>	If -sse is set to KMS, this option can be set if you want DIVA Core to use a specific encryption key identified by its key id. This option is supported by aws:s3 and alibaba:oss.
-bucket_key_enabled (aws:s3 specific)	If -sse is set to KMS, -bucket-key-enabled can be set if the kms key id to be used is associated with the bucket. This option is supported by aws:s3 but not alibaba:oss. With Alibaba, there is no need to configure server side encryption on the DIVA Core side if the bucket is configured with a default server-side encryption.

The -kms-key-id and -bucket-key-enabled parameters are mutual exclusive. If both are specified by mistake, DIVA Core will use -kms-key-id and ignore -bucket-key-enabled.

Core Manager-Actor Communications

The options previously described are converted by the Core Manager into new elements of the transfer request message.

Example:

```
<TransferRequest ...>
...
<DiskPath><https://s3.amazonaws.com</DiskPath>>
<Proxy></Proxy>
<Login></Login>
<Password></Password>
<ServiceName>S3</ServiceName>
<IdentityDomain>us-east-1</IdentityDomain>
<ThreadsPerTransfer>5</ThreadsPerTransfer>
```

```
<PartSize>5</PartSize>
<DisableETagVerification>>false</DisableETagVerification>
<ContainerName>diva-b6e27aae40abf4f8f04d755d18fdb82c256de0d6-
00000000</ContainerName>
<StorageClass>STANDARD</StorageClass>
<VirtualHostedStyle>>false</VirtualHostedStyle>
<ServerSideEncryption>kms</ServerSideEncryption>
<KmsKeyId>1234abcd-12ab-34cd-56ef-1234567890ab</KmsKeyId>
<BucketKeyEnabled>>true</BucketKeyEnabled>
<AccountName>AWS_STD_172.16.10.192_12E76933EF49</AccountName>
<StorageOptions>-storage_class STANDARD -storage_location S3 -
restore_tier NONE -virtual_hosted_style false -
disable_etag_verification false -server_side_encryption AES256 -
aws_kms_key_id true -bucket_key_enabled true</StorageOptions>
</DeviceDisk>
...
</TransferRequest>
```

Enforce SSE-S3 or SSE-KMS at Bucket Level

When using SSE, the SSE option must be set to AES256 or aws:kms. It is recommended to add a bucket policy to prevent uploads without valid encryption headers.

When there is a bucket policy preventing uploads without encryption, or bad encryption header, the S3 server returns an HTTP error 403. Here is the error message returned by the Core Actor when it happens:

```
Error during disk instance closure divastorage_GREG_STD_000002
[StorageCloudAPI error [16ba7b81-9dad-12d1-80b4-10c040243241.axf
upload failed, http response code: 403 [AccessDenied: Access
Denied]]]
```

Checking if the Files are Encrypted

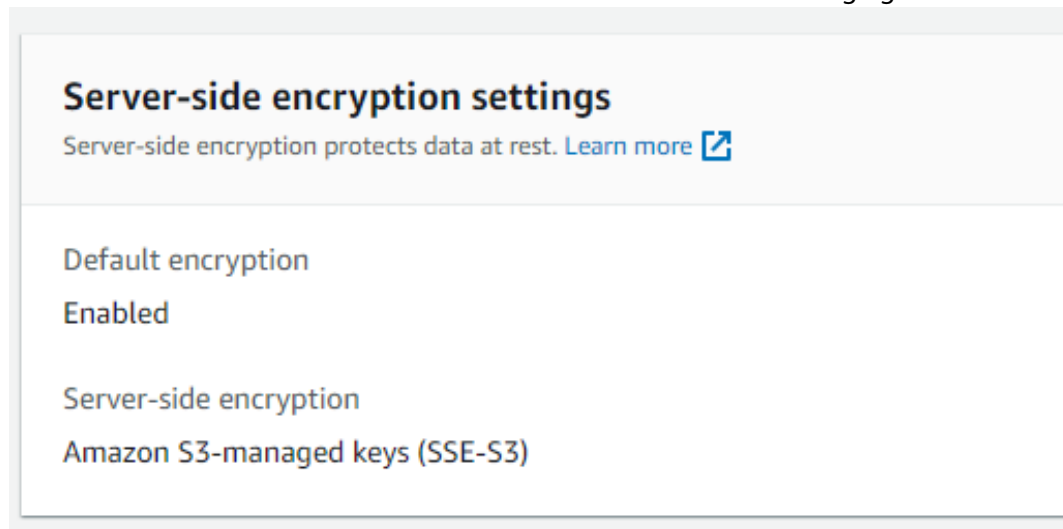
With Alibaba OSS, this is visible in the detail view of a particular file as shown in the following figure.

The screenshot displays the following details for the file:

- File Name:** 13ba7b81-9dad-12d1-80b4-10c040243241/200MB
- ETag:** A0D437FFE898FBD18F506644C1643BD0-20
- Validity Period (Seconds):** 300
- HTTPS:** Enabled (toggle switch is on)
- URL:** <https://diva-gregory-std-000000.oss-us-east-1.aliyuncs.com/13ba7b81-9dad-12d1-80b4-10c040243241/200MB?Expires=1658426272&OSSAccessKeyId=TMP.3KhTgPhUhBj6KUzt1htWJYLnWXB39PENKYIzpfNyTj2UmrCFuhNMpyB7EYkWuXAz3pQ2D7BjXPzxeuevdihoQ5bN4fuCGb&Signature=UeTuGc9jqplFnVZd1MjCAiYIb4I%3D>
- Storage Class:** application/octet-stream
- File ACL:** Inherited from Bucket
- Storage Class:** Archive
- Server-side Encryption:** KMS


A yellow warning banner at the bottom of the details section reads: "The object cannot be downloaded because it is an Archive or Cold Archive object or is being restored. Restore the object completely first."

Similar information is available with AWS S3 as shown in the following figure.



System Management App Support

SSE for S3 buckets can be enabled in the System Management App Configuration page. 'Enable Server Side Encryption' will only appear for S3 Cloud Buckets. If toggled on, 'Server Side Encryption', 'KMS Key Id', and 'Bucket Key Enabled' will display.



SSL (Secure Sockets Layer) and Authentication

DIVA Core 8.3 includes [SSL Certificate Authentication](#) for authentication of services, and securing the internal and API communications in DIVA Core. Certificate authentication provides unique identification and secure communications for each DIVA Core service in a network.

DIVA Core 8.3 includes a Default Root [CA \(Certificate Authority\)](#) called DIVA_CA. The DIVA_CA Certificate Authority is a self-signed authority that signs all SSL certificates for the DIVA Core services. Every DIVA Core service now has its own password protected private key and a [SSL \(Secure Sockets Layer\)](#) certificate signed by the DIVA_CA authority.

Certificate authentication functions similar to identification cards like passports and drivers licenses. For example, passports and drivers licenses are issued by recognized government authorities. SSL certificates are signed by a recognized CA. An SSL certificate verifies the identity of its owner. When the SSL certificate is presented to others, it helps verify the identity of its owner based on the quality of the contents of the certificate.

An external third party CA (for example, VeriSign, Comodo, and so on) can be used to generate and sign your certificates.

External Certificate Authorities

External third party CAs (for example, VeriSign, Comodo, and so on) are usable with DIVA Core. The external CA must create a [CSR \(Certificate Signing Request\)](#) for DIVA_CA, signed by the third party CA, and the third party certificate must be added to the [Trust Store](#) to satisfy the [SSL Certificate Chain](#).

Security Tools

The DIVA Core 8.3 release includes a security tool as follows:

- Windows: DivaSecurityTool.bat
- Linux: DivaSecurityTool.sh

The tool is located in the %DIVA_HOME%/security/bin directory.

DIVA Core API Changes

The DIVA Core APIs include changes to establish secure communication with the Core Manager. The Core Manager is backward compatible with earlier Java, C++ and Web Services APIs to establish connections over regular sockets. The DIVA Core 8.3 (and later) Java and C++ API releases can establish Manager communications using secure, or unsecure, sockets.

The Java API includes new parameters added to the SessionParameters class to facilitate secure connections to the Manager Service.

Exporting and importing encrypted tapes is also available using the Java API.

See the Java API Readme for the location of the Java API documentation.

The C++ API `DIVA_SSL_initialize` call is added to set the environment for secure communication with the Manager service. See the *DIVA Core C++ API Programmer's Guide* for detailed information.

The Java and C++ APIs initiators both use the default keys and certificates under the `%DIVA_API_HOME%/lib/security` subfolder when connecting to the Manager.

DIVA Enterprise Connect connects to the Manager Service through the unsecure 9000 / tcp port. See the *DIVA Enterprise Connect Installation, Configuration, and Operations Guide* for detailed information.

The Manager Service is backward compatible with earlier releases of DIVA Connect, Java API, C++ API, and Web Services API, and establishes the connection over regular sockets.

Dual Ports

All internal DIVA Core services can only connect to secure ports. The System Management App will report an SSL Handshake Timeout if you attempt to connect to the non-secure port.

SSL (Secure Sockets Layer) and Authentication

DIVA Core consist of services in Java and C++. The format in how certificates and keys are represented are different in each. DIVA Core has the keys and certificates for JAVA services in a Java Keystore file, and in PEM (Privacy Enhanced Mail) format files for the C++ services.

The Manager can simultaneously support two communications ports - one secure, and one unsecure. The default secure port number is 8000 and the unsecure default port number is 9000.

All internal DIVA Core 8.x services (System Management App, Migration Utility, Actor, SPM, WFM, SNMP, Robot Manager, RDTU, and Migration Services) can only connect to secure ports. The System Management App will report an SSL Handshake Timeout if you attempt to connect to the non-secure port. Clients using the Java or C++ API are allowed to connect to either port.

The following is a relative snippet from the Manager configuration file:

```
# Port number on which the DIVA Manager is waiting for incoming
connections.
# Note: If you are using a Sony library and plan to execute the
DIVA Manager
# on the same machine as the PetaSite Controller (PSC) software, be
aware
# that the PSC server uses the 9000 port and that this cannot be
modified.
# In that situation, you have to use a different port for the DIVA
Manager.
# This same warning applies to FlipFactory which uses ports 9000
and 9001.
# The default value is 9000.
DIVAMANAGER_PORT=9000
```

```
# Secure port number on which the DIVA Manager is waiting for
incoming connections.
# The default value is 8000.
DIVAMANAGER_SECURE_PORT=8000
```

A new folder called %DIVA_API_HOME%/security is added to the DIVA Core API installation structure as follows:

```
%DIVA_API_HOME%
  security
    conf
```

The conf folder contains the SSLSettings.conf file that is used to configure the SSL handshake timeout.

See the DIVA Core Java API documentation included with the API, and the C++ API Programmer's Guide for detailed information.

Secure Communication with Core Database

With DIVA Core 7.6.1, a new DIVAOracle package version 3-1-0 was created:

- Windows: OracleDivaDB_3-1-0_12_2_0_1_0_SE2_Windows_64-bit
- Linux: OracleDivaDB_3-1-0_12_2_0_1_0_SE2_OEL7_x86_64

This new package includes the following

- Secure DIVA Core Database listener listening on port 1522, additional on top of the regular unsecured listener listening on port 1521.
- Core Database wallet for storing the Trust Certificate and DIVADatabaseServer Certificates. During installation DIVADatabaseServer.jks holding the default DIVA_CA trust certificate and Default DIVADatabaseServer certificate is import into the Core Database wallet for enabling the secure communication.
- This new package also creates a secure TNSNames LIB5SSL which enables any DIVA services to connect to the Core Database securely over SSL connecting to the new secure Core Database listener listening on port 1522 using the TNSNames.

New Entry in TNSNames.ora:

```
LIB5SSL =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCPS) (HOST = HOSTNAME) (PORT = 1522))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = LIB5.WORLD)
    )
  )
```

Valid parameter values are:

- TRUE - When set to TRUE, the DIVAMANAGER_DBPORT in the Manager, Migrate, and configuration file must point to the secure port of the Core Database.
- FALSE (default)

the System Management App also supports connecting securely to the database. SPMSERVICE can connect securely only using TNSNames.

Database Installation and Configuration

This chapter describes installation and configuration of the Core Databases and the DIVA Core Backup Service.

Topics:

- [DIVA Core Databases and Backup Service \(BKS\) Overview](#)
- [Installing, Upgrading, and Configuring the Core SQL Database and Backup Service](#)
- [Metadata \(non-SQL\) Database Configuration](#)
- [Installing DIVA Core Notification Service \(RabbitMQ\)](#)
- [Troubleshooting](#)

DIVA Core Databases and Backup Service (BKS) Overview

The Backup Service as a whole is comprised of two types of services, DIVA Core Backup Service (BKS) and one or more DBAgents. Both services have REST APIs such that they can be integrated with a UI component. The main backup service controls command execution, DIVA Core archives, synchronization and configuration. Each database implementation is in managed code and a minimal amount of scripting is utilized to future proof the solution. Backup configurations are also agnostic of the data contained within them such that the solution can be applied to any type of application database you would like to backup assuming the routines to do so are implemented.

At the system level, settings that relate to the overall operation of each DIVA Core component and their interaction are configured and retained by an DIVA Core Database. This is commonly known (and will be referred to in this document) as the Core Database (or just simply as the database).

User modification of this database is performed through the System Management App. It is only intended for experienced users and caution should be exercised when altering settings. An incorrect setting can impede DIVA Core operations or prevent the Core Manager from starting successfully. Contact Technical Support for assistance if unsure about making a particular change.

When launched, the Core Manager obtains the DIVA Core system configuration from the database. However, it does not poll the database for changes made through the System Management App. Therefore, the database must be notified of any changes made. This is performed using the Notify Manager in the System Management App.

Most changes to the configuration can be completed while the Core Manager is running. There are a small number of configuration changes that require a restart of the Core Manager to become effective. A full list of changes that can be made to the system configuration dynamically while the Manager is running is listed in [Appendix D: Dynamic Configuration Changes](#).

The System Management App also does not dynamically poll the database for changes that are made when the Manager is running. In such cases, click the Update button in the utility to refresh the information displayed from the database.

The System Management App can be installed on any computer that has TCP/IP connectivity to the database and a supported Java Runtime Environment installed. DIVA Core release 8.3 requires the Java Runtime Environment 64-bit (build 1.8.1_45-b14), to be installed to launch the System Management App successfully.

In some cases, a network firewall between the two can prevent a connection. For complete operation and functionality of the System Management App, the Oracle Listener Port (typically 1521) and the Core Robot Manager Ports (typically 8500 and higher) must be opened in the firewall. Full functionality of the System Management App also requires that the Core Manager Port (typically 9000) is open.

DIVA Core uses a Metadata Database to support Complex Object workflows. The DIVA Core Backup Service ensures reliability and monitoring of both the DIVA Core Database backups and Metadata Database backups.

The Core Metadata Database has very high performance and almost unlimited scalability. The Metadata Database should be treated with the same caution as the Core Database. It should be backed up at regular intervals through the DIVA Core Backup Service.

Technical Support highly recommends that the Metadata Database files are stored on a RAID disk array. The Metadata Database should not be on a standard disk due to decreased performance and the real-time backup functionality that a RAID array offers the system.

Metadata Database files stored on a standard disk are vulnerable to data loss if a single disk failure until the information is replicated through the DIVA Core Backup Service. Storing the Metadata Database files on a RAID array isolates the data from this type of failure.

The information stored in the Core Database is already stored on a RAID-1 array and is not subject to data loss if a single disk fails.

Notes: DIVA Core 8.3 supports minimum Oracle Database Server 11.2.0.4.

Starting with DIVA Core 8.3, BKS also optionally supports the PostgreSQL.

Caution: See the DIVA Core Supported Environments Guide to confirm disk partitioning and recommended block sizes before proceeding.

Complex Objects

By default, objects archived with more than 1,000 files are considered Complex Objects. Complex Objects have Metadata stored in both the Core Database and Metadata Database. Configure the threshold on the number of files before an object is considered complex in the Manager service configuration file. Complex Objects can only be stored in AXF format within the DIVA Core system. The DIVA Core Backup Service must be used to back up the Core Database and Metadata Database when Complex Object workflows are used.

DIVA Core Backup Service (BKS)

Caution: The DIVA Core Backup Service is strictly required to be used when using Complex Objects. The DIVA Core Backup Service is the only component backing up the Metadata Database and removing outdated Metadata files. When a Delete request for a Complex Object is sent and processed, the data is removed from the Core Database, but the Metadata Database file is not deleted. It is removed by the Backup Service after the configured clean up period (defined by the Recovery Period parameter) has been reached. Users should have an elevated awareness of error messages from the Backup Service.

The DIVA Core Backup Service is now referred to as BKS and ensures reliability and monitoring of both the Core Database and Metadata Database backups. The DIVA Core Backup Service enables configuration of scheduled backups through its configuration file. The DIVA Core Backup Service manages and monitors the entire backup process.

Many replication locations may be configured through the BKS. These locations may be a local path or a UNC path, however the primary backup location must be local as it is used as the source of replication to all other locations. Each location may be configured with a URL to the DBAgent endpoint for that location. This is only necessary if that location is managing a remote database, in which case the database should be listed under the Managed Databases list. Any database in a Managed Database list will be part of the automated backup system and are eligible for restores or failovers.

A source name must be provided for any location that manages a database with DBAgent. This allows the BKS to make calls to DIVA to restore archived backups directly to the related database server for a restoration or failover to process.

Notes: The primary location must have the same source name that is provided in the Backup Settings section of the configuration file.

Configuration within DIVA Core must point to the base directory of the corresponding location.

One of the primary responsibilities of the BKS is to maintain a ledger of backups for each database it manages. These ledgers are located in the BKS log directory in the same folder structure the backups themselves. The default location is

`C:\DIVA\Program\log\backup_service\Ledgers\\\Ledger.json`

These ledgers can be queried through the API and are the primary reporting structure for the active backup or restoration state of a given database. If a ledger is lost or deleted, it will be automatically created on the restart of the BKS based off of the primary backup locations contents.

Each backup is check-summed through MD5 and logged in the database ledger for each database. After a backup occurs it is replicated across all of the backup locations.

After replication, if configured to do so, an archive is made using a call to the DIVA Core API to persist the backups to tape storage. The source in DIVA Core is configured in the location itself under the Source Name parameter. The name of the object will be DatabaseBackups_<Unix timestamp of the archive> and the Collection will be Backups. This only occurs after every full backup, after which a cleanup task will delete any archives that exceed the retention period.

If a database or system failure (or both) occurs, where restoring from a system backup is necessary, restoration of a stored backup is done manually and should only be performed by Technical Support personnel.

Core Database backups and Metadata Database backups are incrementally replicated to one or more remote back up systems by the DIVA Core Backup Service. depending on your configuration.

The Backup Service files are located in the \$DIVA_HOME\Program\conf\backup_service\appsettings.json folder.

One of the primary responsibilities of the BKS is to maintain a ledger of backups for each database it manages. These ledgers are located in the BKS log directory in the same folder structure the backups themselves. The default location is

```
C:\DIVA\Program\log\backup_service\Ledgers\<Database type>\<Database profile>\Ledger.json
```

These ledgers can be queried through the API and are the primary reporting structure for the active backup or restoration state of a given database. If a ledger is lost or deleted, it will be automatically created on the restart of the BKS based off of the primary backup locations contents.

Each backup is check-summed through MD5 and logged in the database ledger for each database. After a backup occurs it is replicated across all of the backup locations. After replication, if configured to do so, an archive is made using a call to the DIVA Core API to persist the backups to tape storage. The source in DIVA Core is configured in the location itself under the Source Name parameter. The name of the object will be DatabaseBackups_<Unix timestamp of the archive> and the Collection will be Backups. This only occurs after every full backup, after which a cleanup task will delete any archives that exceed the retention period.

See [Appendix F: Backup Service and DBAgent Configuration](#) for a sample BKS configuration file.

Caution: Do not change the Metadata Location parameter when the system is running.

DBAgent

The DBAgent Service performs database specific tasks (that is, backups and restores), monitors their progress, and reports disk usage. Database maintenance functionality can easily be added if necessary, but only the specific backup tasks are currently implemented. Any number of DBAgents may be installed and configured, but only one

per server/container. This is to support multi-server installations and automate access control. The DBAgent also exposes a REST API that the Backup Service will call to check the status of a backup, initiate backups/restores/failovers, and monitor disk space for configured mount points.

Configuration for the DBAgent is purposefully minimal with the only the space monitoring and backup location being required. The majority of the configuration resides in the BKS. By default, mount point configuration will monitor the backup location, the C, E, and F drives as expected by the default DIVA installation. These can be expanded to monitor other locations if necessary and can trigger alerts to DIVA when those locations are reaching their space thresholds.

A state file is created in the log directory of the DBAgent for a given database request. Backup request state files are stored in the BackupHistory directory, while respectively, Restore request state files will be in the RestoreHistory directory. These files are actively updated as the backup or restore progresses to completion and are used to gather statuses about a given action. These state files include a full log of the action itself and any files that have been created as a result of the backup process.

See [Appendix F: Backup Service and DBAgent Configuration](#) for a sample BKS configuration file.

Backup Initiator

A command line initiator is included in the bin installation folder. This program is a simple wrapper around the BKS API to perform backups, restores, and failovers. However it does not wait for their completion. It will offer four options when executed:

- Backup
- Restore
- Failover
- Quit

The user will select the related function they would like to perform from the additional options as follows:

Backup

1. <Database 1 -x>
2. Back
3. Quit

Restore

1. <Database 1-x>
 - a. <List of restore points 1-x>
 - b. Back
 - c. Quit
2. Back

3. Quit

Failover

1. <Eligible failover databases 1-x>
 - a. X -> Y
 - <List of restore points 1-x>
 - Back
 - Quit
 - b. Y -> X
 - c. Back
 - d. Quit
2. Back
3. Quit

Workflows

The following subsections describe the BKS workflows.

Archive Workflow

BKS will begin to archive backups after configuration is complete.

An archive consists of a full backup and all of the related incremental backups. Each of these files contains a Unix timestamp within their filename for the BKS to identify the correct files required to perform a restoration. The object created within DIVA Core will have a name that follows this format along with a fixed category/collection:

Object Name: <Name of the DB Profile>__<Unix timestamp>_to_<Unix timestamp>

Category/Collection: DB_BackupArchive

This allows the BKS to identify a related archives range of times.

Note: Because an archive will need the entirety of its incremental backups to be present, an archive will not process until the next full backup is performed. This will create a lag time of one day using the default configuration; this could be longer depending on the configured full backup interval.

Gold Archives

A gold archive is a permanent backup that is kept once per the PermanentRetentionPeriod in days.

These archives are saved per database and therefore could be at different intervals depending on when the database was configured and when backups commenced. It is recommended that if you are configuring multiple databases for a given application

(for example, DIVA) that you make all configuration changes at the same time so that these Gold Archives for each database have a related timeframe.

Archive Ledger

In order for the BKS to keep track of what archives are available for restoration it keeps a ledger of every archived backup that it creates. This ledger is copied to all backup locations and its contents are emailed if email notifications are setup in DIVA. The ledger is located here:

`<Location Path>\Backups\ArchiveLedger.json`

This ledger is automatically generated from DIVA Core if it does not exist, or is deleted, and will contain records for both regular archives and gold archives.

Restore Workflow

In general, restoration is handled automatically when either a Restore or Failover request is made from the API or the Initiator.exe application. During this request the BKS performs the following steps:

1. Checks the managed Backup files on a given backup location to determine if they can satisfy the request.
2. Next, BKS checks archive restoration directory to determine if there are files there that will satisfy the request.

`<Location Path>\Restore\FromArchive\...`

3. If not, BKS checks the archive ledger to determine if any archive on the list can be restored to the above location for the request to proceed.

After the Restore or Failover request succeeds, the related files within the FromArchive directory are deleted. If the request fails for any reason, the files within this directory are preserved to attempt the action again.

Manual Restoration

Manually placing the backup files within the FromArchive directory allows the previous restore process to be achieved manually without the request to DIVA. The files must be copied with the same relative paths that the archived object would restore them in. This is the same relative path that is contained within the Backups folder. You can copy the Backups folder contents to this directory from another system to achieve the same result.

Note: The FromArchive directory is not monitored by any process and will only be cleaned up upon the successful completion of a Restore or Failover request; this way it can hold old backups that would normally be removed by the retention window.

Database Service Failover

Caution: These procedures are critical and sensitive. They should only be performed under the control of a Telestream Support Technician.

If a database or system failure occurs, where restoring from a system backup is necessary, restoration of a stored backup is accomplished using the following outlined procedures.

A failover command is very similar to the restore command though it does not guarantee the database will be up if it fails to process. During failover it is assumed that the existing data at the locations database is invalid and will be deleted prior to the failover script execution. You can failover to the same database or a different database with the same configuration.

It is recommended that failover only be used on an in-place database if the database is corrupted and in an unrecoverable state. In the case of failover to another server, the backup files from the source database are used and the existing backups for the target are essentially invalid (although you can use them to failover to itself if necessary). The verification of a compatible database is done at the BKS service before the command is issued to the DBAgent.

Use the following procedure to configure a standby server for failover:

1. Add the configuration in a new database profile and install a DBAgent on that standby server.
2. Add a location in the configuration that points to the main backup point for that server and add the DBAgent URL to this location configuration.

Note: Do not add the database profile to the list of managed databases unless active backups are to be taken.

The new location will automatically be synchronized from the primary location such that all the backups are ready to be used if you need to failover.

Use the following procedures to perform the failover:

1. Add the failover target to the managed list of databases for the target location.
2. Send the failover command with the source and target database profiles, along with a timestamp of the recovery.
3. Remove the source server from managed databases so it does not make active backups.

Also, recovery from the loss of a Backup Service in case the server that it was running on is down by installing the backup service on another location that it was replicating to. You will need to reconfigure any existing database profiles as well as the locations associated with any prior backup locations you were replicating. This needs to be done in a stepwise fashion such that the new primary backup location can catalog all the

backups into new ledgers before attempting any replication to remote locations. After this is complete, the failover procedure is the same.

Core Database

By default, the DIVA Core Backup Service generates a full database backup every 24 hours, and an incremental backup every 15 minutes. The backup files are compressed with 7zip tool with the .gz extension. See [Prerequisites for Installing the Core Database: Configure Shared Memory](#) for a list of prerequisites.

Core Metadata Database

The Metadata Database is a binary file in the file system. To support the Recovery Window for the Metadata Database, the DIVA Core Backup Service uses the following techniques:

- Whenever a new Complex Object is archived, the Manager creates Complex Object Metadata files in the Metadata Database Path configured in the System Management App.
- By default, the DIVA Core Backup Service backs up Metadata files inside the Metadata Database every 15 minutes. The Metadata file is transferred to all backup systems shortly after creation so that file alterations do not influence the backup copies.

Note: If there is a failure backing up to one of the configured Backup Systems, the Backup Service will continue to retry the failed backup until all backups to all configured Backup Systems are successful. Metadata Files are not marked as being successfully backed up until the backup to all configured Backup Systems is successful.

- During every Metadata Database backup, the Backup Service searches for any Complex Object Metadata files that are not backed up, and replicates them to all of the FBM_BACKUP_REMOTE_DESTINATIONS configured in the configuration file.

Technical Support recommends having the same Metadata Database Location on all main and remote backup Destination Servers. For example, if the Metadata Database Location is set to H:\metaback\, on the main system, the Backup Service must copy the Metadata Database backups to the same location on all remote backup Destination Servers. If the paths are different, the Metadata Database Location must be updated in the Core Database after a Core Database restore during failover. See [Core Database Failure Scenarios and Recovery Procedures](#) for more details.

DIVA Core Backup Service Recommended Practices

The following are recommended practices for the DIVA Core Backup Service:

- The Backup Service must be installed on the same server as the Core Manager and Core Database.

- At least two Backup Systems are always required to store backups. Core Actor computers can serve dual purposes and be used as both backup computers and Actor computers.
- Oracle Incremental backups should be performed every 15 minutes.
- Metadata Database backups should be performed every 15 minutes.
- The Backup Recovery Window should be set to value greater than, or equal to, 10 days.
- The Backup Clean-up function should be performed every 24 hours.
- Oracle Full Backups should be performed every 24 hours.
- If required, restoration of a system backup must only be performed by Technical Support.
- Core Database data files, Core Database backups, and the Metadata Database must be stored on RAID disk array.
- Equal backup disk space must be allocated on the main and all remote backup systems.

Installing, Upgrading, and Configuring the Core SQL Database and Backup Service

General DIVA Core and Database Upgrade Processes

The following subsections describe general upgrading instructions and notes.

Starting ORACLE 11g and DIVA Core 7.6

Use the following process to start Oracle 11g in DIVA Core 7.6:

1. Copy the ORACLE Database package .iso onto the server.
2. Mount iso image.
3. Run *InstallEngine.cmd*.
4. Run *InstallDatabase-large.cmd*.
5. Install DIVA Core.

Upgrading to ORACLE 11g and DIVA Core 8.x

Upgrade to DIVA Core 8.x on top of Oracle 11g. This will update DB schemas to DIVA Core 8.x schemas. See the appropriate DIVA Core 8.x Installation and Configuration Guide for details on upgrading to DIVA Core release 8.x.

Upgrade from ORACLE 11g to ORACLE 19c and DIVA Core 8.x

Use the following process to upgrade from Oracle 11g to Oracle 19c and upgrading DIVA Core 8.3 to DIVA Core 8.x:

1. Now that the DIVA Core 8.x binaries are installed and the database schema updated, we can proceed and upgrade the database engine.
2. Follow the instructions from DIVA Core 8.3 Installation and Configuration Guide / Uninstalling the Core Database Server in Windows.
3. Stop all DIVAx services.
4. Export the Database Dump File using data pump:

```
C:\DIVA\Program\Database\DBInstaller\bin>DIVADBInstaller.bat --  
dbuser=DIVA --dbpass=lib5 --syspass=lib5 --dbhost=localhost --  
dbport=1521 --requesttype=backuprequest --  
dbdumpdirectory=H:\<dump_directory>\
```

For Core Database package releases 2.3.4 Oracle 11g and earlier, use the following commands from 203211-010-Database_2-3-4_Oracle_11-2-0-4-7_Windows_64-bit.iso image in G:\Tools\uninstall:

```
uninstall_database.cmd  
uninstall_engine.cmd
```

5. Check that C:\app directory is empty.

6. Proceed with ORACLE 19c bundle installation:

OracleDivaDB_3-2-0_19_3_0_0_0_SE2_Windows_64-bit\Install.bat

7. Import the Database Dump File using data pump (this step creates the user, schema, and data).

```
C:\DIVA\Program\Database\DBInstaller\bin>DIVADBInstaller --
dbuser=DIVA --dbpass=lib5 --syspass=lib5w0rld --
dbhost=localhost --dbport=1521 --requesttype=restorererequest --
dbdumpdirectory=H:\<dump_directory> --
dbdumpfilename=<dump_filename>
```

Where <dump_directory> is the directory you exported the dump file to, and <dump_filename> is similar to DIVA_8_2_in_11g_via_DIVADBInstaller_DIVA_03_09_2022_15-56-11.DMP.

8. Start all DIVAx services (or reboot server).**9.** Check that DIVA Core is operational by performing some Archive and Restore Object tests.

Exporting the Database Dump Files

There are two methods for exporting the dump files:

- [Export the Database Dump Files Using DIVADBInstaller](#)
- [Export the Database Dump Files Using sqlplus](#)

Export the Database Dump Files Using DIVADBInstaller

For 8.3, using DIVADBInstaller you can backup the database using data-pump export.

The backup dump file has the naming convention

USERNAME_Month_Date_Year_Hour-Minute-Second.DMP (for example:

DIVAUSER_07_11_2018_12-32-11.DMP).

To create a backup execute DIVADBInstaller with the following:

- Set --requesttype as *backup*
- Set the --dbdumpdirectory location for the backup. If --dbdumpdirectory is omitted, it will default to H:/ for Windows and /u04 for Linux.

Note: In the following example:

divapass = DIVA user password

syspass = System user password

Example

```
DIVADBInstaller --dbuser=DIVA --dbpass=divapass --syspass=syspass
--dbhost=localhost --dbport=1521 --requesttype=backuprequest --
dbdumpdirectory=H:\Backup
```

Export the Database Dump Files Using sqlplus

For 7.6.1 or earlier releases, you must perform the following procedure on the Source Server:

1. Open sqlplus and log in as the sys user.
2. Execute the following commands to create the directory object:

```
CREATE OR REPLACE DIRECTORY {directory_virtualobject_name} AS
{'TargetPath'};

GRANT READ,WRITE ON DIRECTORY {directory_virtualobject_name} TO
{source_server_username};
```

Windows Example:

```
CREATE OR REPLACE DIRECTORY diva_dpump_dir AS
'H:\Support\DUMPS';
GRANT READ,WRITE ON DIRECTORY diva_dpump_dir TO DIVA;
exit;
```

Linux Example:

```
CREATE OR REPLACE DIRECTORY diva_dpump_dir AS '/u05/support/
DUMPS';
GRANT READ,WRITE ON DIRECTORY diva_dpump_dir TO DIVA;
exit;
```

3. Open a command prompt and execute the following command to export to the dump file:

```
expdp {source_server_username}/{source_server_user_password}
schemas={source_server_username} flashback_time= systimestamp
DIRECTORY={directory_virtualobject_name}
dumpfile={dump_file_name} logfile={log_file_name}
```

Windows and Linux Example:

Note: In the following example:

divapass = DIVA user password
syspass = System user password

```
expdp DIVA/password schemas=DIVA flashback_time=systimestamp
directory=diva_dpump_dir dumpfile=diva_db.dmp
logfile=diva_exp.log
```

Importing the Database Dump Files

There are two methods for importing the database dump files:

- [Import the Database Dump File Using DIVADBInstaller](#)
- [Import the Database Dump File Using sqlplus](#)

Import the Database Dump File Using DIVADBInstaller

For 8.3, using DIVADBInstaller you can restore or import the DIVA database from a previous state if required. using the `--requesttype=restorerequest`, and the dump file name using `--dbdumpfilename`.

To restore or import the database, execute DIVADBInstaller with the following:

- Set the `--requesttype` to `restorerequest`
- Specify the location of the dump file using `--dbdumpdirectory`
- Specify the name of the dump file using `--dbdumpfilename`

If the source server username of the dump file is different from the `--dbuser`, you must also specify the source server username using `--dbimportfromuser`.

Note: In the following example:

```
divapass = DIVA user password
syspass = System user password
```

Example - Same Source Server User Name

```
DIVADBInstaller --dbuser=DIVA --dbpass=divapass --syspass=syspass
--dbhost=localhost --dbport=1521 --requesttype=restorerequest --
dbdumpdirectory=H:\Dump --dbdumpfilename=DIVA_07_11_2018_12-32-
11.DMP
```

Example - Different Source Server User Name

```
DIVADBInstaller --dbuser=DIVA --dbpass=divapass --syspass=syspass
--dbhost=localhost --dbport=1521 --requesttype=restorerequest --
dbdumpdirectory=H:\Dump --dbdumpfilename=DIVA_07_11_2018_12-32-
11.DMP --dbimportfromuser=DIVA_75
```

Import the Database Dump File Using sqlplus

For 7.6.1 or earlier releases, perform the following procedures on the Destination Server:

1. Open sqlplus and log in as the sys user.
2. Execute the following commands to create the directory object:

```
CREATE OR REPLACE DIRECTORY {directory_object_name} AS
{'TargetPath'};
```

```
GRANT READ,WRITE ON DIRECTORY {directory_virtualobject_name} TO
{destination_server_username};
```

Windows Example:

```
CREATE OR REPLACE DIRECTORY diva_dpump_dir AS
'H:\Support\DUMPS';
GRANT READ,WRITE ON DIRECTORY diva_dpump_dir TO DIVA;
exit;
```

Linux Example:

```
CREATE OR REPLACE DIRECTORY diva_dpump_dir AS '/u05/support/
DUMPS';
GRANT READ,WRITE ON DIRECTORY diva_dpump_dir TO DIVA;
exit;
```

3. Open a command window and copy the exported dump file to the {'TargetPath'}.
For example: H:\Support\DUMPS (Windows) or /u05/support/DUMPS (Linux)
4. Navigate to the %DIVA_HOME%\program\database\core\install folder in your DIVA Core installation.
5. Create a Core Database user with the following command:

Note: In the following example:

```
divapass = DIVA user password
syspass = System user password
```

Windows: *create_diva_user.bat syspass DIVA2 divapass -useronly*

Linux: *create_diva_user.sh syspass DIVA2 divapass -useronly*

6. Execute the import command as follows:

```
impdp {destination_server_username}/{user_password}
transform=OID:n:type DIRECTORY={directory_virtualobject_name}
dumpfile={dump_file_name} table_exists_action=replace
REMAP_SCHEMA={source_server_username}:{destination_server_user
name} logfile={log_file_name}
```

Example:

```
impdp DIVA2/pass transform=OID:n:type DIRECTORY= diva_dpump_dir
dumpfile= diva_db.dmp table_exists_action=replace
REMAP_SCHEMA=DIVA:DIVA2 logfile=diva_imp.log
```

Uninstalling the Core Database Server (if required)

Before installing the new Core Oracle Database, you may be required to uninstall the existing database and database engine. If Core Database is already installed on the computer, then you must remove the existing database and database engine.

Uninstalling the Core Database Server in Windows

Use the following procedure to uninstall the existing database in Windows environments:

Caution: Use the same Core Database package to uninstall the database that was used to install it.

1. Stop all running DIVA Core services.
2. Export the existing database contents using the procedures previously described.

Caution: Confirm the export completed successfully before continuing.

3. Extract the original database .zip file used to perform the installation.
4. For Core Database package releases 2.3.4 and earlier, use the following commands in Oracle Bundle ISO mount point \Tools\uninstall subdirectory in the exact sequence shown:

```
uninstall_database.cmd  
uninstall_engine.cmd
```

5. For Core Database packages release 3.0.0 and later, execute `C:\app\Oracle\product\12.1.0\db_home1\deinstall\deinstall.bat` and follow the displayed instructions.

Uninstalling the Core Database Server in Linux

Use the following procedure to uninstall the existing database (package release 3.0.0 and later) in a Linux environment:

1. Log in as the Oracle operating system user.
2. Open a terminal window.
3. Export the existing Core Database.
4. Execute `$ORACLE_HOME/deinstall/deinstall` and follow the displayed instructions.

Installing the Core Database Server in Windows

You must log in to the computer as an Administrator. After you have backed up and uninstalled the existing database (see the previous sections), use the following procedure to install the new database:

1. Locate the latest release of the DIVAOracle database package for Windows and unzip it.
2. Execute *install.bat* to start the installation.
3. Follow the prompts through the wizard to complete the installation.
4. Import the previously exported data into the new database using the procedure previously described.

Assuming no errors occurred, you have successfully installed the database and imported the existing data from the original database.

Installing the Core Database Server in Linux

Before running the installer verify the following is complete:

- Yum is configured to connect to the latest release of Oracle Linux.
- The recommended partitions for the Core Database exist. Technical Support recommends partitions that dedicate the space to the Core Database.
 - /u01 partition for the Oracle Binaries
 - /u02 partition for the Oracle Database files (8 KB cluster size recommended)
 - /u03 partition for the Oracle Archive Logs (4 KB cluster size recommended)
 - /u04 partition for the Oracle database backups (64 KB cluster size recommended)

To begin installation, locate the latest release of the DIVAOracle database package for Linux, execute it as root, and follow the displayed instructions.

Prerequisites for Installing the Core Database: Configure Shared Memory

If the shared memory on the server where the Core Database is installed is less than 16 GB, you must set it to at least 70 percent of your RAM.

1. Use the following command to confirm the computer's RAM size:

```
# free -m
```

The output will look similar to the following:

```

total      used      free      shared  buff/cache   available
Mem: 15791    186      15456         8        148        15516
Swap:16380     0       16380

```

2. Use the following command to check your shared memory setting in MB:

```
# df -m /dev/shm
```

The output will look similar to the following:

```
Filesystem 1M-blocks Used Available Use% Mounted on
tmpfs      7896      0      7896   0% /dev/shm
```

3. To change the size of shared memory you must add the following line into `/etc/fstab`. The setting must not exceed the size of your installed memory. You must restart the computer after making this change for it to take affect.

For example, the following command will increase the size of `/dev/shm` to 11GB:

```
tmpfs /dev/shm tmpfs defaults,size=11g 0 0
```

Prerequisite for Installing the Core Database: Creating Drive Partitions

First you must configure the drive partitions for the Core Database as follows:

1. Navigate to Applications > Utilities.
2. Click Disks from the menu.
3. Locate your disk in the Disks dialog box. Selecting the disk will display the Device Name.
4. In Linux you must add the disk (that you want to add partitions to) to the partition table using the `fdisk` utility. For example, `fdisk /dev/xvdb1`. You can use the `g` and `w` options to add it to the partitions table.
5. Click the Plus button on the right side of the Disks dialog box to add a partition.
6. When the Create Partition dialog box appears create the following four partitions. For each partition leave the Erase option and Type option at their default settings, and then click Create. Repeat this step for each partition.

/u01

This partition must be 10 GB in Linux. Use the operating system default block size.

/u02

This partition must be 30 GB in Linux. Technical Support recommends using an 8 KB cluster size.

/u03

This partition must be 5 GB in Linux. Technical Support recommends using a 4 KB cluster size.

/u04

This partition must be either 100 GB or all of the remaining disk space. Technical Support recommends using a 64 KB cluster size.

7. When you are done creating the partitions and returned to the Disks dialog box, click the Gears icon on the right side of the screen.
8. Click Edit Mount Options.
9. Change Automatic Mount Options to OFF.

10. Select the Mount at startup check box.
11. Enter the appropriate mount point in the Mount Point field for that specific partition (/u01, /u02, /u03, /u04).
12. Click OK.
13. When this is completed successfully, all four partitions are identified and displaying their appropriate mount points in the Disks dialog box.

Use the following procedure for the Managed Disk partition (this must be 54 GB):

1. Locate the Managed Disk in the Disks dialog box.
2. Click the Gears icon on the right side of the screen.
3. Click Format.
4. Leave all settings at their defaults, but enter */managed* in the Mount Point field.
5. Click Format.
6. When asked, click Format to confirm that you want to format the disk.
7. Click the Gears icon.
8. Click Edit Mount Options.
9. Change Automatic Mount Options to *OFF*.
10. Select the Mount at startup check box.
11. Enter */managed* in the Mount Point field.
12. Confirm that the Filesystem Type is set to *ext4*.
13. Click OK.

Installing the Core Database Server

Verify you have completed the following:

- [Prerequisites for Installing the Core Database: Configure Shared Memory](#)
- [Prerequisite for Installing the Core Database: Creating Drive Partitions](#)

After completing the prerequisites, use the following procedure to install the Core Database Server:

1. Open a terminal console.
2. If you run in a Virtual Machine confirm that your host name is in the `/etc/hosts` file using the following command:

```
gedit /etc/hosts
```

If the hosts file looks similar to this:

```
127.0.0.1    localhost localhost.localdomain localhost4
localhost4.localhostdomain4
::1        localhost localhost.localdomain localhost6
localhost6.localhostdomain6
```

You must replace `localhost` with your host name. For example, if the host name is `clefvm015L`, it will look like this:

```
127.0.0.1   celfvm015L localhost.localdomain localhost4
localhost4.localdomain4
::1       celfvm015L localhost.localdomain localhost6
localhost6.localdomain6
```

3. If you made changes to the host file save the changes and exit gedit.
4. Change to the directory of the shell script for the Core Database Package.
5. Change the permissions on the shell script using the following command to make it an executable file:

```
chmod +x OracleDivaDB_3-0-0_12_1_0_2_0_SE2_OEL7_x86_64.sh
```

6. Execute the script as follows:

```
./OracleDivaDB_3-0-0_12_1_0_2_0_SE2_OEL7_x86_64.sh
```

If an operating system account has already been created, you may be asked whether you want to change the password. Follow the prompts if you require a password change for this account.

7. When prompted for a SYS account password, ensure you use a secure password.

If at some point during the installation you receive the following error:

```
[FATAL] [INS-35172] Target database memory (5181MB) exceeds
available shared memory (3866MB) on the system
```

You must run the commands below to extend your tmpfs partition (if it is still not large enough):

8. Check the current size of the tmpfs partition:

```
df -h /dev/shm
```

9. Extend the amount of the target database memory size as follows:

- a. Execute `gedit/etc/fstab`.

- b. Add the following line to the bottom of the file:

```
tmpfs /dev/shm tmpfs defaults,size=6G 0 0
```

- c. Save the file and exit gedit.

10. Execute the following commands:

```
umount tmpfs
mount -a
```

11. If the commands in Step 3 do not work, restart the computer and run the `df-h /dev/shm` command again to check that the size of tmpfs has actually increased.
12. Run the Oracle Database shell script again.

Installing the DIVA Database User and Schema

DIVA Core 8.3 has DIVADBinstaller which can install a new DIVA database or upgrade an existing DIVA database on the Core Database Server. For 7.6.1 or earlier releases, you must manually create the user.

- [Using DIVADBinstaller for DIVA Core 8.3](#)
- [Manually Create the Database User and Schema for 7.6.1 and earlier](#)

Using DIVADBIInstaller for DIVA Core 8.3

Verify Core Database Version

Verify the existing Core Database Server release before upgrading a system to DIVA Core 8.3. The Core Database Server must be at a minimum of 11.2.0.4. You can verify the release level by navigating to C:\app\oracle and opening the VERSION.TXT file. The release number is displayed in the file.

Installer Location

The database installer DIVADBIInstaller.bat (Windows) or DIVADBIInstaller.sh (Linux) can be found under <DIVA_HOME>/Database/DBInstaller/bin.

DIVADBIInstaller Parameters

Parameter	Description
--dbuser=<username>	DIVA Database username (required)
--dbpass=<password>	DIVA Database username password (required)
--syspass=<syspassword>	SYS Database username password (required)
--requesttype=<requesttype>	Request type to executed can be one of the following: <ul style="list-style-type: none"> installrequest: performs a fresh install upgraderequest: upgrades the DIVA database backuprequest: performs a datadump export of the DIVA database restorerequest: performs a datadump import to the DIVA database user form the file mentioned in --dbdump-filename If Requesttype is omitted, its defaults to installrequest if the user does not exist or upgraderequest if the user already exists.
--requestname=<requestname>	Given a custom name for the Request execution. Optional and defaults to the system timestamp.
--dbhost=<databaseHost>	Database hostname or IP address. Optional and defaults to localhost.

Parameter	Description
--dbport=<databasePort>	Database port. Optional and defaults to 1521.
--db servicename=<dbServiceName>	Database service name. Optional and defaults to lib5.world.
--dbsecureconnect=<"TRUE FALSE">	Enables secure connection to Database. Optional and defaults to FALSE.
--dbdumpdirectory=<dbdumpdirectory>	Database dump directory. Optional and defaults to H:/ for Windows and /u04 for Linux
--dbdumpfilename=<dbdumpfilename>	Database dump filename.
--dbimportfromuser=<dbimportfromuser>	Dump filename source_server username if different than --dbuser, Mandatory only for --requesttype=importrequest. Defaults to NULL.

Note: In the following examples:
 divapass = DIVA user password
 syspass = System user password

Example Fresh Installation of DIVA Database

```
DIVADBInstaller --dbuser=DIVA --dbpass=divapass --syspass=syspass
--dbhost=localhost --dbport=1521 --db servicename=lib5.world --
requesttype=installrequest
```

Example Upgrade Installation of DIVA Database

```
DIVADBInstaller --dbuser=DIVA --dbpass=divapass --syspass=syspass
--dbhost=localhost --dbport=1522 --db servicename=lib5.world --
dbsecureconnect=TRUE --requesttype=upgarderequest
```

The database installer always backs up the existing user using data-pump export before upgrading. The backup dump file is under the --dbdumpdirectory location on the Database server. If you omit --dbdumpdirectory, it will default to H:/ in Windows and /u04 in Linux.

Manually Create the Database User and Schema for 7.6.1 and earlier

Note: If upgrading release 7.2.2 and lower using 8.3 installer, you must manually update the Actor configuration and Actor Partial Restore configuration in the database using the config utility.

The database user must be created using the DIVA operating system user account. Use the following procedure to create the database user:

1. Open a terminal console.
2. Change to the DIVA_HOME/Program/Database/Core/Install directory.
3. Execute *create_diva_user.bat* (Windows) or *create_diva_user.sh* (Linux), which creates the given DIVA database user and its associated tables

Usage:

```
create_diva_user syspasswd username userpasswd
oracle_connection [-useronly|-tablesonly] [-custom_tablespaces
tables_tablespace indexes_tablespace temp_tablespace]
```

```
create_diva_user {DIVA|SYS} current_password new_password [-
orapwd]
```

Parameter Definitions:

- *syspasswd*: Password of the Oracle sys account
- *username*: Username to create
- *userpasswd*: Associated user password
- *oracle_connection*: Oracle TNS service name or Oracle connection string (such as IP_ADDRESS:PORT/ORACLE_SERVICE_NAME)
- *DIVA|SYS*: Mention either DIVA or SYS to reset the respective password in the password file
- *new_password*: New password
- *current_password*: Current password. If there is no current database password, then enter the new password for the *is* parameter.
- *-useronly*: Only creates the database user and no database objects
- *-tablesonly*: Only creates the database objects for the given user.
- *-custom_tablespaces*: Use of custom tablespaces
 - *tables_tablespace*: tablespace for tables
 - *indexes_tablespace*: tablespaces for indexes
 - *temp_tablespace*: database temp tablespace
- *-orapwd*: Option to reset/generate password file.

Secure Communications with Core Database

With DIVA 7.6.1, a new DIVAOracle package version 3-1-0 was created:

- Windows: OracleDivaDB_3-1-0_12_2_0_1_0_SE2_Windows_64-bit
- Linux: OracleDivaDB_3-1-0_12_2_0_1_0_SE2_OEL7_x86_64

This new package included the following:

1. Secure Core Database listener listening on port 1522, additional on top of the regular unsecured listener listening on port 1521.
2. Core Database wallet for storing the Trust Certificate and DIVADatabaseServer Certificates. During installation DIVADatabaseServer.jks holding the default DIVA_CA trust certificate and Default DIVADatabaseServer certificate is import into the Core Database wallet for enabling the secure communication.
3. This new package also creates a secure TNSNames LIB5SSL which enables any DIVA services to connect to the Core Database securely over SSL connecting to the new secure Core Database listener listening on port 1522 using the TNSNames.

New Entry in TNSNames.ora:

```
LIB5SSL =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCPS) (HOST = HOSTNAME) (PORT = 1522))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = LIB5.WORLD)
    )
  )
```

DIVAMANAGER_DB_SECURE_CONNECT was added to the Manager, and Migrate configuration files to enable secure communication to database using Hostname/IP_Address and port. This parameter has no effect if using DIVAMANAGER_TNSNAME parameter in the configuration file.

Valid parameter values are:

- TRUE: When set to TRUE, the DIVAMANAGER_DBPORT in the Manager, and Migrate configuration files must point to the secure port of the Core Database.
- FALSE (default)

The System Management App also supports connecting securely to the database. SPMSERVICE can connect securely only using TNS names.

Migrating Core Database Server from 11.2 to 12.1

Use the following procedures to migrate DIVA Core releases with Oracle 11g installed. Typically this procedure is performed to upgrade installations with legacy DIVA Core installations to a current release.

Preparing the Source Server Computer (Core Manager with Oracle Database 11.2)

Use the following procedure to export the Core Manager and file system data from the Source Server:

1. Stop all running DIVA Core services, and then export the database to a dump file. See [Exporting the Database Dump Files](#).
2. Copy the dump file from the Source Server to the target computer.

Updating the Destination Server (Core Manager with Oracle Database 12.1)

Use the following procedure to import the Core Manager and file system data to the Destination Server:

1. Stop all running DIVA Core services.
2. Install Oracle 12.1 on the Destination Server. See [Installing the Core Database Server in Windows](#), or [Installing the Core Database Server in Linux](#) for instructions depending on your operating system environment.
3. Import the database dump file on the Destination Server. See [Importing the Database Dump Files](#).

Installing and Configuring the DIVA Core Backup Service

Use the following procedures to install and configure the DIVA Core Backup Service.

DIVA Core Backup Service Overview

The DIVA Core Backup Service enables configuration of scheduled backups through its configuration file, and manages and monitors the entire backup process.

Note: It is **strictly required** to use the DIVA Core backup service when using Complex Objects.

The service uses existing DIVA Core Backup scripts (these scripts use the Oracle RMAN tool) to generate full database backups, and incremental database backups of the Core Database. Generated Core Database backup files and Metadata Database files created by the Manager (when Complex Objects are created) are incrementally replicated by the Backup Service to remote backup servers.

Installing the DIVA Core Backup Service Software

The DIVA Core Backup Service component is installed as an integral part of the standard DIVA Core system installation. You must install the component on the same server as the Core Manager and Core Database. Also, the Backup Service does not support installation with the Manager and Core Database installed on separate computers.

You must configure the DIVA Core Backup Service to replicate files across multiple backup servers for redundancy. Therefore, you must identify the following systems before installation for successful use of the DIVA Core Backup Service:

- Which computer is called Backup System 1 (required)
- Which computer is called Backup System 2 (required)
- Which additional computers are called Backup System additional_number. The additional_number identifies additional backup server numbering, for example Backup System 3, or Backup System 4. This is optional and only required to have more than two backup systems.

You must ensure the Database check box is selected on the Choose Components screen during DIVA Core installation to install the DIVA Core Backup Service.

Installing BKS and DBAgent

Use the following command-line interfaces to install BKS and DBAgent:

Windows

- *backup_service.bat [command] [options]*

Where command is one of the following:

install (or -i)

Installs the module as a system service.

Options:

Option	Description
-log	Path to log directory. Default: ..\..\log\backup_service
-conf	Path to configuration directory. Default: ..\..\conf\backup_service
-httpport	Port to listen for http connections. Default: 1876
-httpsport	Port to listen for https connections. Default: 1877
-certpath	Path to certificate located on disk.
-user	Username to install the service under. Blank entries will be installed as LocalSystem.
-path	Password for the provided user.

-uninstall (or -u)

To remove the executable as a system service.

-start

Starts the module.

-stop

Stops the module if it is currently running.

-restart

Stops and subsequently starts the module.

-status

Determines whether the module is running.

-version (or -v)

Displays the module version information and exits.

-help (or either -h or -?)

Displays help information and exits.

- *db_agent.bat [command] [options]*

Where command is one of the following:

install (or -i)

Installs the module as a system service.

Options:

Option	Description
-log	Path to log directory. Default: ..\..\log\dbagent
-conf	Path to configuration directory. Default: ..\..\conf\dbagent
-httpport	Port to listen for http connections. Default: 1876
-httpsport	Port to listen for https connections. Default: 1877
-certpath	Path to certificate located on disk.
-user	Username to install the service under. Blank entries will be installed as LocalSystem.
-path	Password for the provided user.

-uninstall (or -u)

To remove the executable as a system service.

-start

Starts the module.

-stop

Stops the module if it is currently running.

-restart

Stops and subsequently starts the module.

-status

Determines whether the module is running.

-version (or -v)

Displays the module version information and exits.

-help (or either -h or -?)

Displays help information and exits.

Linux

- `backup_service.sh [command] [options]`

Where command is one of the following:

start

Starts the MetadataService.

stop

Stops the MetadataService.

status

Displays the status of the MetadataService.

stopdb

Stops MongoDB.

startdb

Starts MongoDB.

installdb

Installs MongoDB

Options:

Option	Description
-datadir	Path to the data directory to store the MongoDB database. Default: /u02/MongoData
-port	Port for MongoDB to listen on. Default: 27017

upgradedb

Upgrades existing MongoDB installation.

uninstalldb

Uninstalls MongoDB if installed locally.

version (or -v)

Displays the version string.

help (or -h)

Displays the command line syntax.

- `db_agent.sh [command] [options]`

Where command is one of the following:

start

Starts the DBAgent.

stop

Stops the DBAgent.

status

Displays the status of the DBAgent.

version (or -v)

Displays the version string.

help (or -h)

Displays the command line syntax.

Configuring the DIVA Core Backup Service

The Backup Service configuration file is monitored to allow for live updating of the configuration through the API or by direct manipulation without requiring restarting the service. By default the configuration is located here:

```
$DIVA_HOME\Program\conf\backup_service\appsettings.json
```

This path can be modified during service installation. The Backup Service contains all of the required information to connect to a database and passes that information on to the DBAgent when an action is required. The DBAgent itself also has a configuration, but it contains relatively few values.

The following are the relevant sections of the configuration file located as follows:

Note: All of the related settings can also be modified through the REST API.

Backup Settings

The majority of the archive configuration is done within the Backup Settings configuration section. You can configure the number of days to keep a daily archive, the number of days between the creation of a gold backup (an archive that is stored in perpetuity), the name of the storage media, and the source in DIVA Core of the primary backup location.

```
"DatabaseBackup": {
  "Enabled": false,
  "FullBackupInterval": {
    "ExecutionPeriod": "Daily",
    "TimeOfDay": "00:00:00",
    "InstancesInPeriod": [ 0 ]
  },
  "IncrementalPeriod": 15,
  "FullBackupFileRetention": 10,
  "FullBackupArchiveRetention": 30, <=== IN DAYS
```

```

    "ArchiveMediaGroup": "<some media, disk, or storage plan>",
<=== UPDATE
    "PermanentRetentionPeriod": 180, <=== IN DAYS
    "ArchiveSourceName": "<Source name for primary backup
location>", <=== UPDATE
    "BackupExecutionTimeout": 120,
    "RestoreExecutionTimeout": 120,
    "StatusPollingPeriod": 3,
    "StatusReportingInterval": 1440
  }

```

DIVA Core API Settings

A valid API configuration must be provided for automatic archive, restoration, and events to be sent to DIVA Core. This can be configured in the DIVACore API Settings section.

Typically, only the password must be added; although the URL may require updating if the Core Manager location is on a different system than the BKS.

```

"DIVACoreAPISettings": {
  "Url": " https://127.0.0.1:8765/",
  "User": "sysadmin",
  "Password": "changeit", <=== PASSWORD IS ENCRYPTED UPON BKS
STARTUP
  "TimeoutInMs": 20000
}

```

See [Appendix F: Backup Service and DBAgent Configuration](#) for a sample BKS configuration file.

You must set the following parameters in the System Management App's Manager Setting page. You must set the Metadata Database file location to an existing, valid location. The Manager uses this value to save the Metadata Database files. For example, F:\META_DATABASE_ROOT\.

Complex Objects Metadata Database Location

This is the path to the Metadata Database. There is no default path specified. The path must exist, and is validated by the Core Manager and the Backup Service. You must use a drive with ample storage. See [Sizing the Metadata Database](#) for information on calculating space requirements.

This parameter is not reloadable and is only checked one time when the Manager and the Backup Service services start. If you make any changes to this parameter you must restart the Manager and Backup Service.

Database Backup Notification

You select the desired notification level from the list as follows. The default setting is ERRORS AND WARNINGS. You must restart connected System Management Apps if any changes are made to this parameter.

ERRORS AND WARNINGS

Errors and warnings are recorded in the event log. This is the default setting.

ERRORS

Errors and warnings are recorded in the event log.

DISABLED

All of the errors and warnings are recorded in the event log.

Enable Metadata Database Feature

The Core Manager can archive Complex Objects and Backup Service can backup up the Metadata Database only when you enable this parameter (the check box is selected). When disabled (the check box is deselected) Core Manager cannot archive Complex Objects and the Backup Service cannot backup the Metadata Database. This parameter must be left at the default enabled setting.

This parameter is not reloadable and is only checked one time when the Manager and the Backup Service services start. You must restart the Manager and Backup Service services if any changes are made to this parameter.

If the `BACKUP_SERVICE_MANAGE_METADATA_BACKUPS` is set to Y (indicating yes, or enabled) in the Backup Service configuration file, the values of Enable Metadata Database Feature and Complex Objects Metadata Database Location in the System Management App is validated when the Backup Service starts. If the Enable Metadata Database Feature parameter is set to N (indicating no, or disabled), or the Complex Objects Metadata Database Location is invalid, the Backup Service will fail to start.

You must set the following values on the Manager Setting page of the System Management App before starting the Manager and Backup Service services:

DIVAMANAGER_HOST

This parameter identifies the name of the computer where the Manager is installed. The default value is localhost.

DIVAMANAGER_PORT

This parameter identifies the port number the Manager is listening on for connections. The default value is 9000.

SERVICE_NAME

This parameter identifies the name of the Windows service. The default value is Core Backup.

SERVICE_PORT

This parameter identifies the port number where the service is running. The default value is 9300. You must change this value if it conflicts with other services.

DIVAMANAGER_DBHOST

This parameter identifies the IP address of the database to connect to from the Manager.

DIVAMANAGER_DBPORT

This parameter identifies the port number of the database to connect to from the Manager. The Core Database installation uses the Oracle default 1521 port number.

DIVAMANAGER_DBUSER

This parameter identifies the database username; typically diva.

DIVAMANAGER_DBSID

This parameter identifies the Oracle Database SID (typically lib5) to connect to from the Manager.

BACKUP_SERVICE_MANAGE_DATABASE_BACKUPS

This parameter enables or disables backup of the Core Oracle Database. The default value is Y (indicating yes, or enabled).

BACKUP_SERVICE_MANAGE_METADATA_BACKUPS

This parameter enables or disables backup of the Core Metadata Database. The default value is N (indicating no, or disabled).

CYGWIN_BIN_DIRECTORY

This parameter identifies the location of the CYGWIN installation. The default is C:\cygwin\bin.

DB_BACKUP_LOCATION

This parameter identifies the location of the Oracle Database backup files. The default location is H:/oraback/lib5.

DB_BACKUP_REMOTE_DESTINATIONS

This parameter identifies the location of the Core Database remote backup destinations. All remote destinations must be a service module name, following by a folder name. The backups must not be copied to the root of the module. Multiple destinations are allowed and must be delimited by commas.

FULL_BACKUP_START_HOUR_24

This parameter identifies the hour of day to perform a full database backup when the service is initially started. If the service is started later than the configured value, the full backup will occur at this hour on the following day. The default value is midnight; 0 hours.

FULL_BACKUP_START_MINUTE

This parameter identifies the number of minutes after the FULL_BACKUP_START_HOUR_24 hour to start the full backup. The default value is 0 minutes.

FULL_BACKUP_FREQUENCY_HOURS

This parameter identifies the frequency to execute a full backup of the database. The default value is every 24 hours.

INCREMENTAL_FREQUENCY_MINUTES

This parameter identifies the frequency to execute an incremental backup of the database. The default value is every 15 minutes.

The Backup Service will automatically determine if a full backup is required.

If the FBM_FREQUENCY_MINUTES parameter is not set, then this value is also used to notify the Manager how often to expect a message from the DIVA Core Backup Service. If a message is not received by the Manager within the incremental minutes, all connected System Management Apps are notified that the DIVA Core Backup Service may not be running. This event is then recorded in the event log. If the FBM_FREQUENCY_MINUTES is set, the Backup Service uses the lowest parameter value to notify the Manager how often to expect a message from the DIVA Core Backup Service.

By default, the Manager expects a message from the Backup Service within 15 minutes after the start of the Manager service. After the Backup Service is started and connected to the Manager, the Manager expects a message within every INCREMENTAL_FREQUENCY_MINUTES, or FBM_FREQUENCY_MINUTES value identified in the Backup Service configuration file.

FBM_FREQUENCY_MINUTES

This parameter identifies the frequency to execute a Metadata Database backup to all remote metadata backup destinations. The default value is every 15 minutes.

If the INCREMENTAL_FREQUENCY_MINUTES parameter is set, the Backup Service uses the lowest parameter value to notify the Manager how often to expect a message from the Backup Service.

A Metadata Database backup is executed when the services start.

DB_FBM_RECOVERY_WINDOW_DAYS

This parameter identifies the recovery window period for the Core Database and Metadata Database. This value indicates how many days of backups must be retained. Obsolete backup copies are then deleted. The default is 10 days.

The DIVA Core Backup Service sets this value using the RMANRecoveryWindow.bat file included in the DIVA Core Backup Service bin folder. If this batch file is missing the DIVA Core Backup Service will not start.

CLEANUP_START_HOUR_24

This parameter identifies the hour of the day for initial start of the Backup Service clean up process to delete the obsolete backup copies. The default value is 2 (representing 2:00 AM).

CLEANUP_START_MINUTE

This parameter identifies the number of minutes after CLEANUP_START_HOUR_24 to start the clean up process. The default value is 0 (representing the top of the hour).

CLEANUP_FREQUENCY_HOURS

This parameter identifies the frequency to run the clean up process. The default value is every 24 hours.

See [Monitoring the DIVA Core Backup Service](#) for additional monitoring and notification options and configuration.

Backup Interval Overrun

A Backup Interval Overrun occurs when a specific backup is taking a longer time to complete beyond the next scheduled iteration.

The following example is called a Backup Interval Overrun because the Backup Service must run the next incremental backup by 12:15 PM, but it cannot because the backup process started at 12:00 PM is still running.

1. The Oracle Incremental Backup is schedule to run every 15 minutes:
`INCREMENTAL_FREQUENCY_MINUTES = 15`
2. The incremental backup starts at 12:00 PM and runs at the value set for the INCREMENTAL_FREQUENCY_MINUTES parameter; in this case every 15 minutes.
3. At 12:15 PM the incremental backup is incomplete and still running, causing a Backup Interval Overrun.

The DIVA Core Backup Service sends a Backup Timeout Warning to the Manager when a Backup Interval Overrun occurs. The Manager records the warning in the event log. If a Backup Timeout occurs three consecutive times, the timeout warning messages are elevated to an error message.

IMPORTANT: You must take immediate and necessary action to modify the backup's frequency by updating the configuration file to avoid future Backup Interval Overrun occurrences

Note: Updating the configuration file requires a Backup Service restart.

Backup Service Running Normally

When the Backup Service is running, the following information is displayed when the *status* command is executed:

- Running release of the service
- IP address and port the service is running on
- System statistics
- Operating system information
- Memory information
- Disk array information
- Database backup statistics including:
 - Last executed backup command and the current status
 - Number of Metadata Database files backed up
 - A list of the last 25 Metadata files backed up including the object name and creation date

The information output to the console is also saved in the logs directory. This file, and the log files, must be included when submitting issues to Technical Support.

Backup Service Not Currently Running

When the Backup Service is not currently running, the following information is displayed when the *status* command is executed:

- Running release of the service
- IP address and port the service runs on
- An extract from the DIVA Core Backup Service log files from the last error, or irrecoverable error, reported

Backup Service Failed to Start

If the Backup Service fails to start identify the cause of the failure, correct the issue, and then try to start the service again. If you require assistance contact Technical Support.

Uninstalling BKS and DBAgent

The DIVA Installer does not support uninstalling BKS or DBAgent, so uninstalling these has always been done manually by using scripts provided in each DIVA component.

Use the following commands to uninstall BKS and DBAgent respectively:

```
backup_service.bat uninstall
db_agent.bat uninstall
```

Monitoring the DIVA Core Backup Service

The DIVA Core Backup Service notifies the Manager about all backup errors and warnings.

You use the list menu to the right of the Suppress Alerts label to snooze alerts. The list menu enables you to snooze the error or warning as follows: Never (never allow this message type to be snoozed), One Hour, Three Hours, and Eight Hours. The system snoozes the specific message type displayed in the dialog box and suppresses future messages for the same error or warning. Snoozing a message dialog box has no effect on the currently displayed error or warning; it only affects future messages about the same error or warning that has been snoozed.

All messages generated by the Backup Service are also written to the Database Event Log and marked as Backup Service Messages.

Events in the Logged Events panel may be filtered using the filter check boxes and fields to reduce the number of entries being viewed simultaneously. The following figure shows that the screen has been filtered to show only Warnings and Errors because their associated check boxes are selected in the filter area. It is readily apparent there are three warning events that have been logged about the Database Backup Manager timing out during an incremental backup attempt. If the timeout occurs again, the warning is elevated to an error (after three warnings) and displayed in red (rather than yellow).

The screenshot shows the 'Logged Events' interface with the following details:

- Filter Area:**
 - Event ID: *
 - Severity: Information, Warnings, Errors
 - Description: *
 - Dates: Begin: 2012-03-01 12:50, End: 2012-04-01 13:20
 - Request ID: *
 - enable
- Event List:**

Event ID	Severity	Description	Date
941262	Warning	Warning message received from the Database Backup Manager: Incremental Database Backup timeout, process is taking longer to complet....	01/03/2012 12:51:48
943536	Warning	Warning message received from the Database Backup Manager: Incremental Database Backup timeout, process is taking longer to complet....	01/03/2012 13:20:36
944733	Warning	Warning message received from the Database Backup Manager: Incremental Database Backup timeout, process is taking longer to complet....	01/03/2012 13:35:36

Error messages are prefixed with the process that generated the error or warning, and where applicable, post fixed with the start of the process and elapsed time. The elapsed time is the time the process ran before generating the error.

The following table describes the different warning and error notifications.


Message Type	Code	User Message	Posted to Manager
SUCCESS	0	Completed successfully	Yes, informational
RUN	1	Running	No, internal only
ERROR	2	Failure: Refer to the Backup Service logs for more details.	Yes, error
TIMEOUT	3	Timeout: The process is taking longer to complete than the configured intervals. The Backup Service continues to display timeout messages as a warning. If the timeout occurs three consecutive times, the message will be elevated to an error message and displayed.	Yes, warning
STARTUP_FAILURE	4	DIVA Core Backup Service failed to start. Refer to the Backup Service logs for more details.	Yes, error
INITIALIZE	5	Scheduling Backups	No, internal only
TIMEOUTERROR	6	Timeout: The process is taking longer to complete than the configured interval.	Yes, error
CONFIGERROR	1000	Invalid Configuration Error. Refer to the Backup Service logs for more details.	Yes, error
METADATALOCATIONERROR	6000	The Metadata Database Location does not exist. Refer to the Backup Service logs for more details.	Yes, error
CLEANUPFBMFILEERROR	7000	The Metadata Database file deletion failed. Refer to the Backup Service logs for more details.	Yes, error
CLEANUPFBMFILEWARNING	7001	Failed deleting the Metadata Database.	Yes, error
DBCONNECTERROR	9000	Database connection error. Refer to the Backup Service logs for more details.	Yes, error
SQLERROR	9001	Database SQL error. Refer to the Backup Service logs for more details.	Yes, error
DBROLLBACKERROR	9002	Database Rollback error. Refer to the Backup Service logs for more details.	Yes, error

Monitoring Minimum Disk Space

The `DISK_MIN_SPACE_THRESHOLD_PERCENT` is a notification threshold percentage of the available space for each drive accessible by the Manager. The default value is 5 percent. For example, `DISK_MIN_SPACE_THRESHOLD_PERCENT=25` sets the notification threshold to 25 percent. This function does not monitor removable media and drives.

When the configured threshold of available space on the media is reached, warning notifications are sent out. After the available space reaches 80 percent of the designated percentage an error message is sent out.

When the configured percentage is reached, a dialog box will be displayed as shown in the following figure.



The Suppress Alerts list at the bottom of the dialog box functions identically to the other warning and error dialog boxes. In the previous figure a warning was issued to notify the operator that the `DISK_MIN_SPACE_THRESHOLD_PERCENT` was reached.

Snoozing this alert causes no additional disk space warnings or errors to be displayed. Clicking OK without setting a suppression level enables future alerts for this particular warning to be displayed.

In the previous figure, when 80 percent of the threshold percentage is reached (2.4 GB on C drive and 24.8 GB on D drive), this dialog turns into an error rather than a warning.

Email Notifications

The DIVA Core Backup Service incorporates the ability to send out emails for issues arising from the process of backing up the Core Database and Metadata Database files. In order to take advantage of this feature, DIVA Core must be configured to connect to an SMTP mail provider. The email notifications are configured through the System Management App under the Manager Setting page.

Use the following procedure to enable email notifications:

1. Open the System Management App and connect to the database.
2. Click the Manager Setting page.
3. Set the values for the following email notification parameters as required:

Caution: If the following parameters are misconfigured entries into the Manager Event Log will be made. However, email notification will not be sent.

Enable E-Mail Notification

If you select the check box (enabled), the Manager attempts to send out email using the configured values.

(SMTP) Outgoing Mail Host

Enter the URL of the email provider for outgoing mail in the (SMTP) Outgoing Mail Host field. This is provided by your Email Administrator.

(SMTP) Outgoing Mail Port

The port value is port 25 by default. However, many email providers are using a different port for security reasons. The correct port number is provided by your Email Administrator. Enter the correct port number in the (SMTP) Outgoing Mail Port field.

E-Mail Subject

Enter the value to be used in the E-Mail Subject field if an email subject is not provided when an error is generated.

(SMTP) Outgoing Mail Required Authentication

Many email providers require you to log in to the email server to allow sending emails. You must select the (SMTP) Outgoing Mail Required Authentication check box, and provide a valid account name and password (using the following two fields) if required to log in to the email server.

Account Name

Enter the full senders email address in the Account Name field if the (SMTP) Outgoing Mail Required Authentication check box is selected.

Account Password

You must enter the password associated with the senders email address in the Account Password field if you have entered an email address in the Account Name field.

DIVA Core System Administrator's E-mail Address

Enter the full email address for the DIVA Core System Administrator in the DIVA Core System Administrator's E-mail Address field so they receive a copy of any email notifications.

Notification E-Mail Recipients

You must enter the full email addresses for anyone who should receive the email notifications in the Notification E-Mail Recipients field. This should be a comma-delimited list with no spaces.

After you have configured the values, if the Manager is already running you must notify the Manager of any changes. When the Manager starts, or when it receives notifications from the System Management App, it reads the configured values and attempts to send out a test email. If the test is successful, all recipients on the Notification E-Mail Recipients list will receive a Test Successful email notification. Otherwise, they will receive an email notifying them of any error that occurred.

Events are logged in the Logged Events panel of all connected System Management Apps.

Metadata (non-SQL) Database Configuration

This section describes configuration of the Core Metadata Database and includes the following information:

- [Configuring the Metadata Database](#)
- [Sizing the Metadata Database](#)
- [Backup Interval Overrun](#)
- [Backup Service Running Normally](#)
- [Migrating an MDDDB \(Flat File Metadata Database\) to MDS \(Metadata Service\)](#)
- [Core Database Failure Scenarios and Recovery Procedures](#)

Configuring the Metadata Database

You must set the following two parameters on the Manager Setting page of the System Management App to enable Complex Object workflows and Metadata Database backups:

Enable Metadata Database

Select this check box to enable use of the Metadata Database.

Metadata Database Location

Enter an empty directory path that exists in the file system in the Metadata Database Location field.

Note: Changes made to these parameters require you to restart the Manager and Backup Service. When it is necessary to change the Metadata location, you must confirm that you have copied all of the Metadata files from the old location to the new location.

Technical Support **highly recommends** that you store the Metadata Database files on a RAID disk array. The Metadata Database should not be on a standard disk due to decreased performance and the real-time backup functionality that a RAID array affords the system.

Metadata Database files stored on a standard disk are vulnerable to data loss if a single disk failure occurs until the information is replicated with the DIVA Core Backup Service. Storing the Metadata Database files on a RAID array isolates the data from these types of failures.

Sizing the Metadata Database

Note: MongoDB, in its default configuration, can use up to half the available RAM minus 1GB on the server on which it is installed. You have to plan the location of MDS MongoDB installation accordingly.

You can use the following formula as a rough guide to determine the minimum disk space required to support the Metadata Database:

$$(100 + \text{avg_path_file_name_size}) * 1.15 * \text{avg_number_component_files} * \text{number_virtual_objects}$$

When planning, enough Metadata Database disk space should be allocated to ensure expected, or unexpected, growth of your environment. You must allocate the same disk space for the Metadata Database on all of the remote backup systems.

Example:

avg_path_file_name_size = 60

this/nested/subdir01/As_The_World_Turns_24fps_scenes1-10.avi

avg_number_component_files = 200,000

This is the average number of files and folders within the Complex Object.

num_objs = 50,000

This is the number of Complex Objects to be archived.

In this example, the recommended minimum disk space allotment would be for a Metadata Database size of approximately 1.67 TB.

Migrating an Mddb (Flat File Metadata Database) to MDS (Metadata Service)

If Mddb Migration failed during a DIVA Upgrade using the DIVA Installer and retrying did not work due to MDS and Rest API services incorrectly installed, migration can be

performed manually using DIVA/Program/Manager/bin/DBMigrate.BAT/ (or DBMigrate.sh for Linux).

You must remember the folder that contains the MDDDB database files (that is, the Complex Objects Metadata Database Location setting in DIVA Database) before upgrade. This setting is removed automatically during a Database Upgrade to 8.3 and later.

If you enter DBMigrate.BAT without arguments, or with the -h parameter, you will see the following:

```
Select Administrator: Windows PowerShell
DIVA Metadata DB Migration
Usage:
DBMigration.bat help (or) -userId ***** -userPassword ***** -metadataRoot C:\MDDDBData
where:
  help          (or -h) displays this information and exits
  -userId (required) authorized user having access to diva manager API
  -userPassword (required) password for authorized user
  -metadataRoot (required) need to provided original metadata root directory.
  -restapiUrl (optional) The root URL of the rest API (default = https://127.0.0.1:8765).

  service exit code of 0 indicates successful start of migration and for successful completion
  service exit code of 1 indicates migration in progress
  service exit code of 2 indicates a problem with the migration, a description of the problem will be outputted to the console.
PS C:\DIVA\810\site1\Program\Manager\bin>
```

For example:

```
DBMigration -userId [REST API user name] -userPassword [password]
-metadataRoot "C:\DIVA\metadata"
```

The migration will begin the first time this is executed. Subsequent calls will provide the current status until complete as shown in the following figure:

```
Administrator: Windows PowerShell
PS C:\DIVA\810\site1\Program\Manager\bin> .\DBMigration -userId string -userPassword string -metadataRoot "C:\DIVA\810\site1\metadata"
Migration started at 2021-10-14 14:23:07.895 has been created, 4 complex objects will be migrated with 2053 total components - 2000
PS C:\DIVA\810\site1\Program\Manager\bin> .\DBMigration -userId string -userPassword string -metadataRoot "C:\DIVA\810\site1\metadata"
Migration started at 2021-10-14 14:23:07.895, 2 objects out of 4 migrated, 621 components moved. - 2003
PS C:\DIVA\810\site1\Program\Manager\bin> .\DBMigration -userId string -userPassword string -metadataRoot "C:\DIVA\810\site1\metadata"
Migration started at 2021-10-14 14:23:07.895 completed, 4 objects out of 4, moved 2053 components out of 2053 - 2004
PS C:\DIVA\810\site1\Program\Manager\bin>
```

The first call created the migration and returned 2000. The next call shows two of four objects migrated and returned 2003. This indicates the migration is still in progress. A typical migration will show this response many times. The final call shows the migration is complete with a return of 2004.

Calling the script in this way does not show the exit codes of 0, 1 or 2. To see these exit codes create a batch file as follows to make the call:

```
call DBMigration -userId string -userPassword string -metadataRoot
"C:\DIVA\810\site1\metadata"
echo ERRORLEVEL %ERRORLEVEL%
```

Upgrading and Migrating in Linux

When upgrading DIVA on Linux, you must run the DIVA installer as root:

```
[root@ip-172-16-10-241 ~]# /cifs/Releases/01_DIVA/V8-1/8.1.0.228_NOT_TESTED/DIV
Core-8.1.0.228.sh

Please specify install directory [/home/diva/DIVA]:
The install directory already exists are you upgrading [y/n]: y
Upgrade Install Confirmed
Unpacking to /home/diva/DIVA/__stage__...
Extracting to /home/diva/DIVA/__stage__
Changing permission
```

After specifying DIVA home path, answer *y* to start the upgrade process; the DIVA installer will prompt you to install Metadata Database (which is MongoDB starting with the DIVA 8.1 release) on localhost.

```
Do you want to upgrade (or install if NOT installed) metadata database (MDDDB)? [
y/n]:
```

Answer *y* to install MongoDB on localhost. Answer *n* to skip it, but in that case MongoDB must be manually installed on another server before starting the MDDDB migration.

```
Please specify MDDDB Data folder [/u04/MDDDBData]:
Please specify MDDDB service port [27017]:
```

If *y* is answered, DIVA installer will ask for two parameters to install and setup MDDDB (that is, MongoDB). The first is where data files are stored and the second is port. Default values are offered by the DIVA installer if you hit enter without answering anything.

Proceed the upgrade process like normal. DIVA installer will detect if MDDDB migration is required and if yes, then after the Database schema upgrade, DIVA installer will prompt you to start MDDDB Migration.

```
DIVA Software Upgrade SUCCEEDED
Do you want to upgrade DIVA database [y/n]: y
Upgrading DIVA Database Schema...
```

You must answer *y* to upgrade the DIVA database or the installer will not prompt you to migrate the flat file MDDDB.

```
DIVA services must be started in order to migrate flat file database. Starting t
hem now, this may take a while...
```

If the DIVA installer detects that MDDDB migration is required, it will start all DIVA services first because the migration functionality is implemented as a Rest API endpoint in the Manager service. This step is automatic and no user interactions are required.

```
MDS must be installed in order to migrate flat file database. You must install it right now on localhost or insatll it manually on another server before proceeding to the next step.

Install MDS on localhost <y/n> (y) :

RestAPI services must be installed in order to migrate flat file database. You must install them right now on localhost or insatll them manually on another server before proceeding to the next step.

Install RestAPI services on localhost <y/n> (y) :
```

If MDS and RestAPI services are not installed on localhost, the DIVA Installer will offer you the option to install MDS (Metadata Service) and the Rest API services on localhost. If you choose to not install them on localhost, they must be installed on other servers before proceeding to next step because DIVA installer will prompt you to enter the MDS and RestAPI Service URL string.

```
Please enter RestAPI user id (sysadmin) :

Please enter RestAPI user password (lib5) :

Please enter DIVA RestAPI URL (https://127.0.0.1:8765) :

Please enter MDS URL (https://127.0.0.1:1777) :

Validatiing user entered information, this may take a while...
```

Before starting the MDDDB migration, DIVA installer will prompt you to answer the questions shown in the previous figure and answered in the following list:

RestAPI user by default

sysadmin

RestAPI user password by default

Core's Oracle Database sys user password. This default password was already entered during the Database Upgrade.

RestAPI URL

This should be the RestAPI Gateway's service URL. It defaults to localhost. You must provide the correct value if you choose to install this on another machine.

MDS URL

This is the metadata service URL. It defaults to localhost. You must provide the correct value if you choose to install on another machine.

If wrong values are provided validation will fail and DIVA installer will ask all these prompts again until the migration Request is successfully submitted.

```
Please enter MDS URL (https://127.0.0.1:1777):

Validating user entered information, this may take a while...
Migration started at 2021-07-05 19:38:07.748, 0 objects out of 10 migrated, 0 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 0 objects out of 10 migrated, 0 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 0 objects out of 10 migrated, 0 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 0 objects out of 10 migrated, 0 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 1 objects out of 10 migrated, 5002 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 1 objects out of 10 migrated, 5002 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 1 objects out of 10 migrated, 5002 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 1 objects out of 10 migrated, 5002 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 2 objects out of 10 migrated, 10004 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 2 objects out of 10 migrated, 10004 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 2 objects out of 10 migrated, 10004 components moved. - 2003

Migration started at 2021-07-05 19:38:07.748, 2 objects out of 10 migrated, 10004 components moved. - 2003
```

After the migration Request is successfully submitted, the DIVA installer will query the status every three seconds and display the progress of how many Complex Objects have been migrated.

```
Migration started at 2021-07-05 19:38:07.748 completed, some objects could not be migrated. Migrated 10 objects out of 20, moved 50020 component - 2005
```

The DIVA installer will complete and exit after the migration has completed. If migration failed (due to various reasons), the installer will prompt you to retry and the migration will resume from where it left off.

```
Migration Failed: Migration started at 2021-07-12 15:31:17.836 did not fully complete cause of system errors, 1 objects out of 20 migrated, 5002 components moved, system failed migration for 19 archives. - 5006

Do you want to retry <y/n> (n): y

Please enter MDS URL (https://127.0.0.1:1777):
```

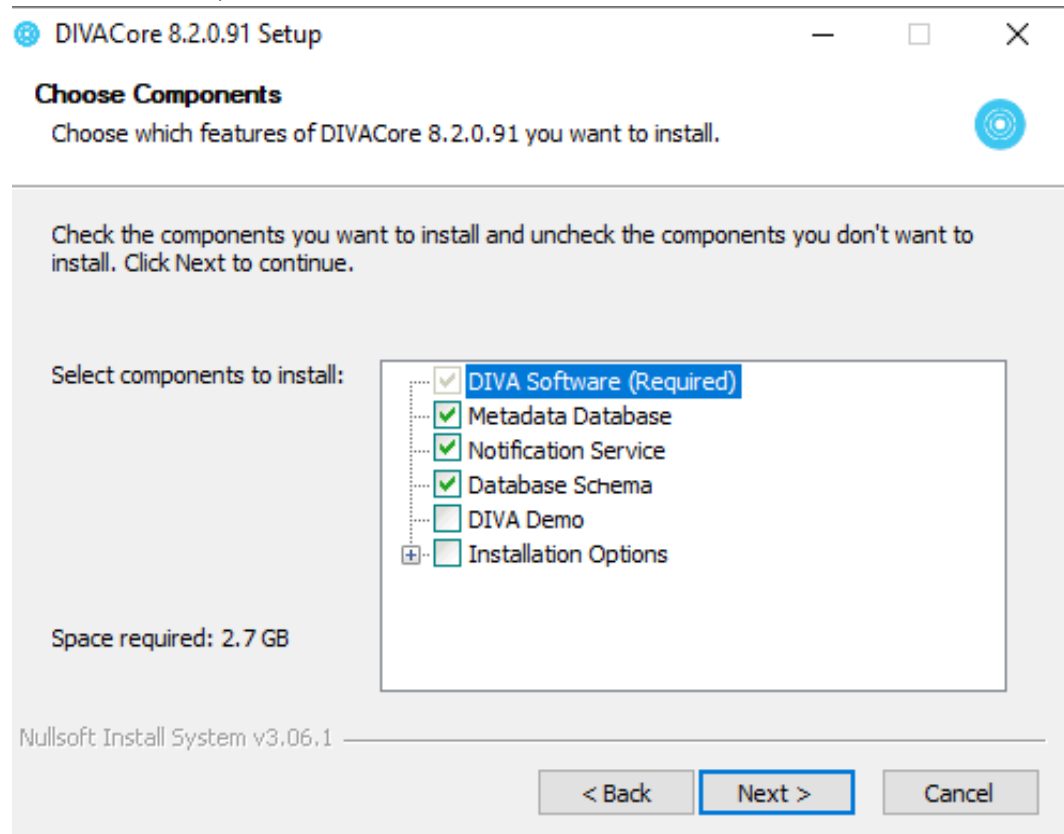
Installing DIVA Core Notification Service (RabbitMQ)

RabbitMQ has been integrated into DIVA windows installer starting with release 8.2.0.91 and later. It is the database used for the DIVA Core Notification Service.


New Windows Installation

The DIVA Core installer mentions RabbitMQ as Notification Service instead of RabbitMQ because RabbitMQ is just an implementation. Use the following procedure to install the Notification Service (RabbitMQ):

1. When choosing components to install, select the Notification Service in the DIVA Core installer, then click Next.



- The Notification Service has its own database and by default will be stored inside the DIVA\Program\rabbitmq folder.



The path to this folder is stored inside a text file under DIVA\rabbitmq_server\etc\notificationDataDir.txt. This file is used if DIVA Core is upgraded in the future so the installer knows where this path is. However, the actual setting for RabbitMQ service is in the Windows registry. A user cannot change the value in this file and expect RabbitMQ to start using a new data directory. The user must modify registry if a different directory is desired).





















- Click Install.

The RabbitMQ application (the binary) folder is the DIVA\rabbitmq_server folder.

RabbitMQ will be installed as a service after the DIVA Core installation completes.

Program Compatibility Assistant Service	This service provides support for the Program Compatibility Assistant (P...	Running	Manual	Local System
Quality Windows Audio Video Experience	Quality Windows Audio Video Experience (qWave) is a networking platfo...		Manual	Local Service
RabbitMQ	Multi-protocol open source messaging broker	Running	Automatic	Local System
Radio Management Service	Radio Management and Airplane Mode Service	Running	Manual	Local Service
Recommended Troubleshooting Service	Enables automatic mitigation for known problems by applying recomm...		Manual	Local System
Remote Access Auto Connection Manager	Creates a connection to a remote network whenever a program referenc...		Manual	Local System

The `install_rabbitmq.bat` and `uninstall_rabbitmq.bat` are not meant to be used by regular users. Support and admin users may edit these files to understand how RabbitMQ was installed or uninstalled.

Name	Date modified	Type	Size
 CollectSysInfo.bat	2021-03-01 9:28 AM	Windows Batch File	23 KB
 CollectSysInfo.ksh	2019-08-03 9:11 AM	KSH File	8 KB
 DivaApi.dll	2015-12-18 2:27 AM	Application exten...	873 KB
 DIVAConfigurationPrinter.bat	2021-09-16 10:19 AM	Windows Batch File	5 KB
 DivaScript CLI reference.docx	2016-10-15 11:56 AM	Microsoft Word D...	62 KB
 DivaScript CLI reference.pdf	2016-10-15 11:56 AM	Adobe Acrobat D...	264 KB
 DivaScript.exe	2016-10-17 11:14 AM	Application	181 KB
 DivaService.exe	2017-10-27 7:18 AM	Application	38 KB
 FindMetadataFile.bat	2021-09-16 10:19 AM	Windows Batch File	4 KB
 FlashnetMigration.bat	2021-09-16 10:19 AM	Windows Batch File	3 KB
 Gather_Activity_Statistics.bat	2021-09-16 10:19 AM	Windows Batch File	51 KB
 GetVersion.exe	2019-08-03 5:21 AM	Application	49 KB
 <code>install_rabbitmq.bat</code>	2021-09-16 10:19 AM	Windows Batch File	2 KB
 lynxLocalDelete.bat	2021-09-16 10:19 AM	Windows Batch File	6 KB
 rdtu.bat	2021-09-16 10:19 AM	Windows Batch File	1 KB
 servicetag.vbs	2019-08-03 9:11 AM	VBScript Script File	1 KB
 <code>uninstall_rabbitmq.bat</code>	2021-09-16 10:19 AM	Windows Batch File	1 KB
 wrapper.dll	2021-09-16 10:19 AM	Application exten...	366 KB
 wrapper.exe	2021-09-16 10:19 AM	Application	677 KB
 wrapper.jar	2021-09-16 10:19 AM	Executable Jar File	122 KB

The `install_rabbitmq.bat` batch file requires erlang installer to be placed inside the C:\DIVA folder; this script does not work as delivered. The DIVA Core installer also does not use it directly; it is just there to document the install procedure.

However, the `uninstall_rabbitmq.bat` batch file can be used as delivered to uninstall RabbitMQ service. To execute the file (also not intended to be used by regular DIVA Core users), run `uninstall_rabbitmq.bat C:\DIVA`.

Note: The `uninstall_rabbitmq.bat` will also uninstall Erlang runtime. After it is uninstalled, you cannot reinstall it by running `install_rabbitmq.bat` unless you have a copy of Erlang installer in the C:\DIVA folder.

DIVA Core Installer also creates the RabbitMQ `advanced.config` file, which is stored in the `DIVA\Program\rabbitmq\advanced.config` folder. This file enables a secure websocket connection that is required for the System Management App Running Requests page to display new requests in real-time.

Upgrading Windows Installations

Upgrade will install the newer version over the old version by overwriting files that already exist.

Choose Components

Choose which features of DIVACore 8.2.0.888 you want to install.

Check the components you want to install and uncheck the components you don't want to install. Click Next to continue.

Select components to install:

- Upgrade DIVA Software (Required)
- Upgrade Metadata Database
- Upgrade Notification Service
- Upgrade Database Schema
- Installation Options

Space required: 2.7 GB

Jullsoft Install System v3.06.1

< Back Next > Cancel

During upgrade, the user can change the location of data folder but by default, installer will take the value from the DIVA\rabbitmq_server\etc\notificationDataDir.txt file. If the user does not want to change it, it will be use the same folder as the previous RabbitMQ installation.

Other than the data folder setting, no other settings are preserved after performing the upgrade; every setting will be reset back to defaults. There are no settings that the user should change in RabbitMQ.

After a successful installation use the following URL to access the RabbitMQ admin console and DIVA Core installer will always create a default admin user (username wsuser and password changeit): <http://127.0.0.1:15672>.



Username: *

Password: *

Login

New Linux Installation

The DIVA Core Linux installer will prompt the user whether Notification Service should be installed. However, everything is installed at the OS level in Linux (instead of under a DIVA folder in Windows) and the installer uses all default settings. Use the following procedure to perform a new installation in Linux:

1. Execute `sudo rabbitmqctl status` to display the location of the following:
 - RabbitMQ data directory
 - RabbitMQ log directory
 - RabbitMQ configuration file (that is, the location of `advanced.config` file, which is in the `/etc/rabbitmq/advanced.config` directory)

```

's RPM package ...
Adding rabbitmq-server tcp port 15672 to firewalld default zone
Adding rabbitmq-server_SSL SSL port 15673 to firewalld default zone
Enabled rabbitmq-server as a service...
Started rabbitmq-server as a service...

Succeeded to install DIVA Notification service. Setting up wsuser account ...
Diva Upgrade Tool is running and its logs (if any) can be found under: '/home
e/diva/DIVA/Program/log/diva_upgrade' folder.
Response(IsSuccess=true, FailureReason=null)
Succeeded to setup wsuser account.
    
```

If you select n to not let DIVA Core installer install RabbitMQ it can be installed later Linux using the `divaservice` script using the following commands:

- `ds configure rabbitmq-server`
 - This command will install Erlang, RabbitMQ then install it as a systemctl managed service then start it.
- `ds install rabbitmq-server`
 - This command is exactly same as `ds configure rabbitmq-server` except it does not start the service when complete.
- `ds uninstall rabbitmq-server`
 - This command will uninstall the systemctl service but will not uninstall Erl and RabbitMQ runtimes.
- `ds list`
 - This command shows all the DIVA Core services and displays whether they are running; see `rabbitmq-server` service in the following screenshot:

```

[diva@ip-172-16-10-241 ~]$ ds list

DIVATestLog      running
DIVAWebAPI       running
rabbitmq-server  running
    
```

- `ds status/start/stop/restart rabbitmq-server`
 - These commands will show status, start, stop or restart the service similar to other DIVA Core services.

After a successful installation use the following URL to access the RabbitMQ admin console and DIVA Core installer will always create a default admin user (username `wsuser` and password `changeit`): `http://<ip_of_linux_server>:15672`.

The Erlang runtime and RabbitMQ software are rpm packages self-contained inside DIVA Core installer, they can be found under the `DIVA/erl` and `DIVA/rabbitmq_server` directories after the fresh installation is complete.



Upgrading Linux Installations

During upgrade, DIVA Core Linux installer runs the following commands:

```
ds uninstall rabbitmq-server
ds configure rabbitmq-server
```

The script uninstalls and then reinstalled RabbitMQ. The RabbitMQ data files are not deleted during the upgrade; any user configuration and/or queue configuration changes are not affected by the upgrade.

Any failures to remove the `rabbitmq-server` service is normal because the `ds uninstall` command has already removed the service, but the RabbitMQ uninstall script will attempt to uninstall it again.

```
Failed to stop rabbitmq-server.service: Unit rabbitmq-server.service not loaded.
warning: file /usr/lib/systemd/system/rabbitmq-server.service: remove failed: No such file or directory
warning: /home/diva/DIVA/erl/erlang-24.3.2-1.el8.x86_64.rpm: Header V4 RSA/SHA256 Signature, key ID cc4bbe5b: NOKEY
warning: /home/diva/DIVA/rabbitmq_server/rabbitmq-server-3.9.13-1.el8.noarch.rpm: Header V4 RSA/SHA512 Signature, key ID 6026dfca: NOKEY
[/usr/lib/tmpfiles.d/rabbitmq-server.conf:1] Line references path below legacy directory /var/run/, updating /var/run/rabbitmq - /run/rabbitmq; please update the tmpfiles.d/ drop-in file accordingly.
[/usr/lib/tmpfiles.d/subscription-manager.conf:1] Line references path below legacy directory /var/run/, updating /var/run/rhsm - /run/rhsm; please update the tmpfiles.d/ drop-in file accordingly.
```

The upgrade process should look exactly same as a fresh install except it uninstalls the existing installation before updating.

Troubleshooting

This section describes basic troubleshooting methods and includes the following information:

- [Metadata Database Failure Scenarios](#)
- [Core Manager Will Not Start](#)
- [DIVA Core Backup Service Will Not Start](#)

Core Database Failure Scenarios and Recovery Procedures

There are two types of failure scenarios; non-failover, and failover.

Non-failover Scenarios

If the Main Core Manager computer is still fully operational, and there has been no RAID Disk failure, you can restore and recover the DIVA Core system and its database from failure without moving the Core Manager or database to a Backup Core Manager computer.

The following are non-failover scenarios and recovery actions (in sequence) to correct them. Contact Technical Support if you require assistance or need to restore from a backup.

Manager Failure

1. Restart the Manager
2. Apply a cumulative patch (if available) and restart the Manager
3. Upgrade your DIVA Core installation

Core Database Instance Failure

1. Restart the Oracle instance
2. Reinstall Oracle and restore the database from a backup

Core Database Data File Corruption

Restore the data file from an Oracle Secure Backup.

Core Database Parameter File or Control File Corruption

Restore the parameter file, or control file, from an Oracle Secure Backup.

DIVA Core Online Redo Logs Corruption

Restore the database using an Oracle Secure Backup.

DIVA Core Archive Redo Logs Corruption

Shut down the database and perform a full backup.

Failover Scenarios

If the main Core Manager computer fails, is not operational, or a RAID disk fails, you must restore and recover the Core Manager and database on the Backup Core Manager computer to restore DIVA Core back to an operational state.

The following are failover scenarios. The recovery actions are the same for all of the listed scenarios.

See [Manager Failover Procedures](#) to get the system back online and contact Technical Support if you require assistance or need to restore from a backup.

The following are possible failures that require failover recovery actions:

- Main Core Manager Computer Failure
- RAID Disk Failure where Oracle Data Files are Stored
- RAID Disk Failure where Oracle RMAN Backups are Stored
- RAID Disk Failure where Metadata Database Files are Stored

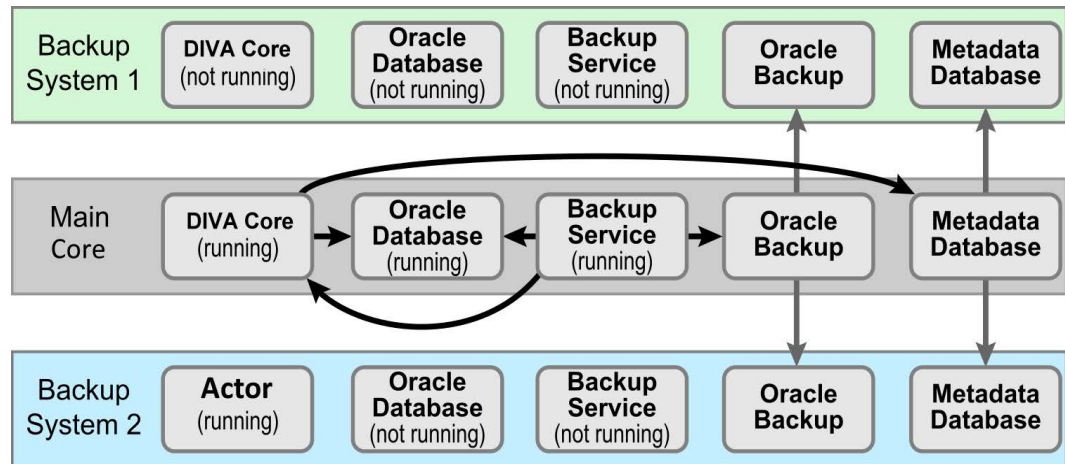
You use the following recovery sequence to complete the failover if any of the previous failures occur:

1. Failover to the Backup Core Manager computer.
2. Restore and recover the Oracle Database from an Oracle Secure Backup.
3. Discover if any Complex Objects are missing Metadata files.
4. Start the Core Manager.

Failover Procedures

You use the following procedure to recover the DIVA Core system if a failure occurs. The first figure is a typical DIVA Core System configuration showing the connections between the different modules, the second displays a failover case, and the third depicts a recovered, operational system. The Main Manager and Backup System 1 are configured identically. However, the Backup Service, Manager, and Core Database are not running until they are started (see the third figure). The Backup Service creates the backups on the Main Manager computer and then pushes copies of them to the Backup System 1, Backup System 2, and Backup System N. The N represents additional

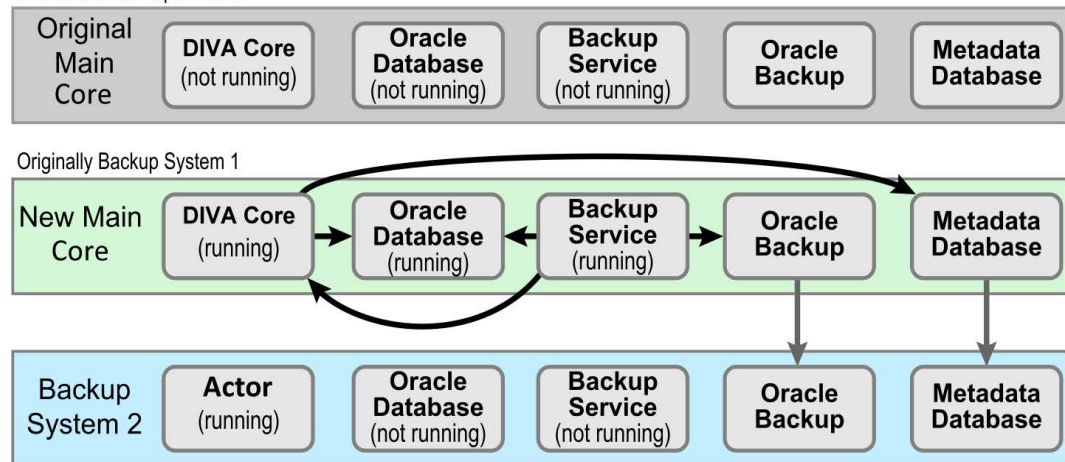
system numbering if applicable, for example Backup System 3, Backup System 4, and so on.



DIVA_026

For this example, assume the Main Manager computer failed and is offline. You are effectively switching the Original Backup Manager to be the New Main Manager and the Original Main Manager will be the New Backup Manager (they are trading places), resulting in the least amount of time the system is offline.

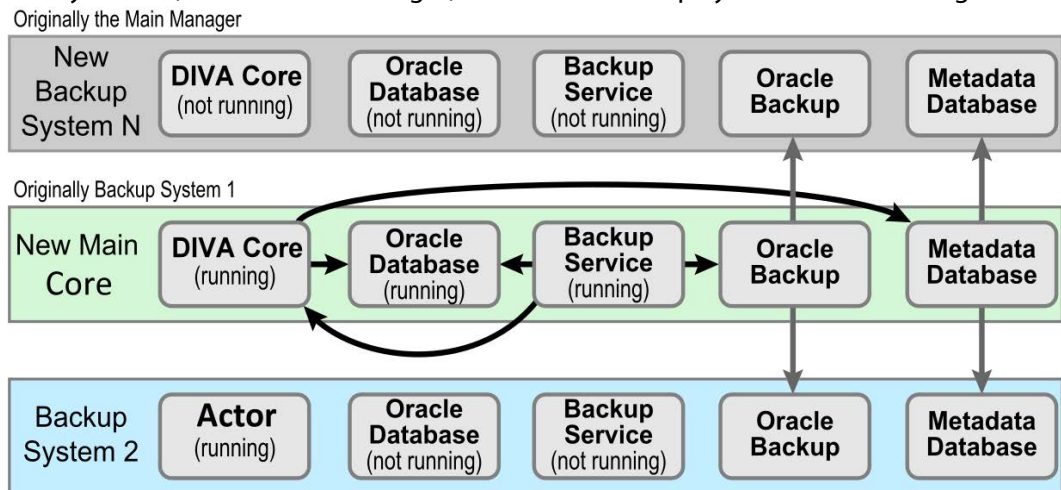
Offline and Non-Operational



DIVA_027

1. Restore the Core Database on the New Main Manager from the latest Oracle Database backup.
2. On the New Main Manager, adjust the Manager configuration file and Backup Service configuration file to point to the Core Database that has just been restored (see the previous step).
3. Update the Metadata Database Location to the location where the Metadata Database files were backed up on New Main Manager system (the Original Backup System 1). You update the parameter under the Manager Setting panel in the System Management App on the New Main Manager computer.

4. Run the Backup Service command on the New Main Manager system. This command lists all of the Complex Objects that are missing the Metadata file in the Metadata Database.
 If a Complex Object is missing the Metadata file, it must be restored from the Original Main Manager, or Backup System 2. Complex Objects are unusable without the associated Metadata file.
5. Start the Manager and Backup Service on the New Main Manager.
 After the Original Main Manager system is restored, recovered from its failure, and is operational, it is converted to the New Backup System N with no downtime.
6. Update the DB_BACKUP_REMOTE_DESTINATIONS and FBM_BACKUP_REMOTE_DESTINATIONS parameters in the Backup Service configuration file on the New Main Manager system by adding the New Backup System N (the Original Main Manager) as the additional remote backup location.
7. Restart the Backup Service on the New Main Manager for your configuration changes to take effect.
8. Copy the existing Core Database backups and Metadata files from the Backup System 2 (or New Main Manager) to the New Backup System N in the background.



DIVA_028

Metadata Database Failure Scenarios

This section describes possible Metadata Database failures and resolutions.

The typical Core Metadata Database backup configuration backs up the database and transfers the backup files to remote systems (as defined in the configuration) every 15 minutes. Technical Support recommends having at least two remote backup systems for redundancy.

Identifying Failure Scenarios, Causes, and Resolutions

The following are examples of possible failure scenarios. Each scenario includes the method of detection, the cause of the failure, a description of the failure, and recovery procedures. Contact Technical Support if you require additional assistance to resolve any of these issues.

Scenario 1: Metadata Database Storage Disk Failure

A disk failure is identified on the Main Manager because no more Complex Objects can be archived into the DIVA Core system. Only Delete requests are possible on existing Complex Objects. DIVA Core is still operational for archiving non-complex objects.

New Metadata files created for Complex Objects archived since the last successful backup, up until the disk failure, are not available immediately. However, they can be recovered from the AXF file.

A disk failure is identified on one of the backup systems because the Metadata Database files created by a new Archive request since the disk failure are backed up only to one backup system, instead of all identified backup systems.

The method of detection for this failure is that a Complex Object request fails with the error Internal error: metadata database error. Metadata Database Backup Failure events are logged in the Manager Event Log.

The possible causes of this failure include the following:

- RAID controller failures
- Power surges
- External process errors
- Disk volume reconstruction error if the RAID was previously rebuilt

Even though Technical Support recommends storing the Metadata Database on a RAID disk, disk failure scenarios cannot be totally eradicated, and the unlikely chance of Disk Failure still exist.

Use the following procedure to attempt recovery from disk failure on the Main Manager:

1. Stop the Manager and Backup Service.
2. Replace the failed disk.

3. Navigate to the Manager Setting page in the System Management App and confirm that the Metadata Database Location setting is pointing to the replaced disk.
4. Start the Manager and Backup Service.
5. Copy all of the Metadata files from a backup system to the Metadata Database Location on the replaced disk.
6. Confirm no Complex Objects are lost.
7. The Metadata files of Complex Objects archived since the last successful backup, and before the disk failure, are not immediately available. However, they are recoverable from the AXF file. Recovery from AXF files is not supported in this DIVA Core release; contact Technical Support for assistance.

Use the following procedure to attempt recovery from disk failure on one of the backup systems. The system can be operational if the backups made to other backup systems were successful.

1. Replace the failed disk.
2. Copy all Metadata files from the second Backup System and Main Manager System to the folder identified in the Metadata Database Location on the replaced disk.

Scenario 2: Metadata Database File Corruption

No operations or requests are possible on Complex Objects whose Metadata files are corrupted, except Delete Object requests, until it is restored. A Metadata file modified by any external source (other than DIVA Core) after it is backed up will not affect its backup copies in the backup systems.

You can identify when a Metadata Database file becomes corrupted because Complex Object requests fail with the following error:

```
Internal error: metadata database error:  
Message: Metadata file read error.
```

The possible causes of this failure include the following:

- External process errors
- The file is modified manually by mistake

Use the following procedure to attempt recovery from a corrupt Metadata Database file. If the corruption occurred after the Metadata file is backed up, the Metadata file can be restored from one of the backups servers.

1. Execute the *FindMetadataFile.bat* utility located in the %DIVA_HOME%/programs/utilities/bin folder on the Main Manager System.

This utility prints out the location of the Metadata file with its file name inside the specified Metadata Database Location, and accepts the database connection parameters and the Complex Object name and Collection as parameters.

2. Locate the file with the file name and path printed from the utility in the Metadata Database backup location on one of the backup servers.

3. Replace the Metadata file on the Main Manager System in the configured Metadata Database Location with the copy from the backup server.

If the corruption occurred before the Metadata file was backed up, the Metadata file is not immediately available. However, it is recoverable from the AXF file. Recovery from AXF files is not supported in this DIVA Core release; contact Technical Support for assistance.

Scenario 3: Lost or Manually Deleted Metadata Database File

Metadata deleted by any external source other than DIVA Core after it is successfully backed up does not affect its backup copies on the backup systems.

You cannot perform any operations or requests on Complex Objects whose Metadata file is corrupt, except Delete Object, until the Metadata file is restored.

You can identify when a Metadata Database file is lost or deleted because Complex Object requests fail with the following error message:

```
Internal error: metadata database error:  
Message: get: Error opening metadata for virtualobjectname/  
collection, db error=Error file not found.
```

The possible causes of this failure include the following:

- External process errors
- The file was manually deleted by mistake

If the file is lost after the Metadata File is backed up, the Metadata File can be restored from one of the Backup Servers. Use the following process to attempt recovery from a lost or deleted Metadata Database file:

1. Execute the *FindMetadataFile.bat* utility located in the %DIVA_HOME%/programs/utilities/bin folder on the Main Manager system.
This utility prints out the location of the Metadata file with its file name inside the specified Metadata Database Location, and accepts the database connection parameters and the Complex Object name and collection as parameters.
2. Locate the file with the file name and path printed from the utility in the Metadata Database backup location on one of the backup servers.
3. Replace the Metadata file on the Main Manager System in the configured Metadata Database Location with the copy from the backup server.

If the file was lost before the Metadata file was backed up, the Metadata file is not immediately available. However, it is recoverable from the AXF file. Recovery from AXF files is not supported in this DIVA Core release; contact Technical Support for assistance.

Scenario 4: Failure to Backup Metadata Database to All Backup Systems

Failure to back up the Metadata Database to all backup systems results in all Complex Objects archived after this failure not being backed up. You must resolve this failure as soon as possible because the DIVA Core system is at risk of data loss.

You can identify this error when a Metadata Database Backup Failure is logged in the Manager Event Log.

The possible causes of this error are as follows:

- Network errors
- The backup systems are offline
- The Backup Service has failed

Use the following referenced resolutions to attempt correction of this issue:

Network Errors

Resolve the network error.

Backup System Offline

Start, or restart, the Backup System.

Backup Service Failure

Restart the Backup Service and collect the logs for investigation.

After the problem is resolved, all of the Backup Systems sync automatically, and the missing Metadata files are backed up during the process. There is no data recovery required for this scenario.

Scenario 5: Failure of the Metadata Database Backup to One Backup System

In this scenario, the Metadata Database fails to back up to (only) one of the Backup Systems. However, the back ups to other Backup Systems continue successfully.

You can identify this error when a Metadata Database Backup Failure is logged in the Manager Event Log.

The possible causes of this error are as follows:

- Network errors
- The Backup System where the error occurred is offline

Use the following referenced resolutions to attempt correction of this issue:

Network Errors

Resolve the network error.

Backup System Offline

Start, or restart, the Backup System.

After the problem is resolved, all of the Backup Systems sync automatically, and the missing Metadata files are backed up during the process. There is no data recovery required for this scenario.

Core Manager Will Not Start

When the Manager starts it checks the following parameters. The Manager will not start if any combination of these parameters is incorrect. Confirm the Enable Metadata Database parameter is configured correctly, and the Metadata Database Path is a valid path that is not empty.

DIVA Core Backup Service Will Not Start

The DIVA Core Backup Service is designed to terminate execution immediately after attempting to start if it is configured incorrectly. This behavior can be caused by any of the following reasons:

- The configuration file is missing.
- The database connection information is incorrect, or the database is not running.
- The `BACKUP_SERVICE_MANAGE_METADATA_BACKUPS` parameter is set to Y (Yes, or enabled) in the Configuration file, but not enabled under the Manager Settings panel in the System Management App.
- The `BACKUP_SERVICE_MANAGE_METADATA_BACKUPS` parameter is set to Y (Yes, or enabled) in the Configuration file, but the Metadata Database Location is not set, or set to an invalid directory under the Manager Settings panel in the System Management App.
- The `BACKUP_SERVICE_MANAGE_METADATA_BACKUPS` parameter is set to Y (Yes, or enabled) in the Configuration file, and the Metadata Database Backup is enabled under the Manager Settings panel in the System Management App, but the Metadata Database Location is not set, or set to an invalid directory.
- `BACKUP_SERVICE_MANAGE_DATABASE_BACKUPS` and `BACKUP_SERVICE_MANAGE_METADATA_BACKUPS` parameters are set to N (No, or disabled) in the Configuration file.
- `RMANRecoverWindow.bat` is not in the bin folder for the Backup Service.

Cluster Manager Installation

This chapter provides general guidelines for the installation of MSCS (Microsoft Cluster Server) software and Oracle Fail Safe software combined with DIVA Core software, to achieve high availability for DIVA Core components by building a two node cluster.

This guide describes only MSCS and Oracle Fail Safe installation steps required for DIVA Core cluster installation.

The Active Directory installation and management is not documented, although it is mandatory for the two DIVA Core Cluster Node servers to be part of a Windows domain.

Note: DIVA Core supports the Oracle Linux 7 x86_64 and later environment. However, the Cluster Manager support is only applicable to Windows-based systems. The minimum server operating system for using Complex Objects is Windows Server 2016.

Topics:

- [Overview](#)
- [Installation Requirements](#)
- [Microsoft Cluster Configuration](#)
- [DIVA Core and Oracle Fail Safe Configuration](#)
- [Configuring Oracle Fail Safe](#)
- [Testing the Configuration](#)
- [Maintenance](#)

Overview

This section describes an overview of the MSCS (Microsoft Cluster Server), Oracle Fail Safe, DIVA Core integration, and tested releases.

Related Documentation

For more information, see the DIVA Core documentation set for this release located at <https://www.telestream.net/telestream-support/diva/support.htm>, and the following recommended Microsoft and Oracle documentation set:

- Oracle Fail Safe Installation Guide 4.2.1
https://docs.oracle.com/cd/E67869_01/OFSIG/installing-ofs.htm#OFSIG118
- What's New in Failover Clustering in Windows Server
<https://docs.microsoft.com/en-us/windows-server/failover-clustering/whats-new-in-failover-clustering>
- Configure and Manage the Quorum in a Windows Server Failover Cluster
<https://docs.microsoft.com/en-us/windows-server/failover-clustering/create-failover-cluster>
- NIC Teaming Overview
<https://docs.microsoft.com/en-us/windows-server/networking/technologies/nic-teaming/create-a-new-nic-team-on-a-host-computer-or-vm>
or
<https://docs.microsoft.com/en-us/windows-server/networking>
- Failover Clusters Cmdlets in Windows PowerShell
<https://docs.microsoft.com/en-us/powershell/module/failoverclusters/?view=windowsserver2022-ps>
- Microsoft Windows Firewall with Advanced Security
<https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-firewall/windows-firewall-with-advanced-security>
- Microsoft Cluster-Aware Updating
<https://docs.microsoft.com/en-us/windows-server/failover-clustering/cluster-aware-updating>
- Microsoft Cluster-Aware Updating Best Practice
<https://docs.microsoft.com/en-us/windows-server/failover-clustering/cluster-aware-updating-requirements>

Prerequisites

The prerequisites for installation must be met before the arrival of the Installation and Delivery Team at your location.

You are responsible for installing the Microsoft Cluster in your environment and providing a dedicated domain user with specific permissions. See [Domain Account Requirements](#) for the required user permissions.

During installation, you must make four shared partitions available so Installation and Delivery personnel can configure DIVA Core in your environment. The drive letters E:, F:, and H: must be used for the shared partitions as follows:

- E:\ - Oracle Database
- F:\ - Oracle Database Transaction Logs
- H: - Oracle Database Backups

When the Installation and Delivery team arrives they will install and configure the Oracle Fail Safe and DIVA Core software for you.

Oracle Fail Safe Integration with Windows

Oracle Fail Safe enables configuring and managing the Oracle Database and other Oracle and third-party applications for high availability on Windows clusters. An instance runs on only one node at a time.

A cluster is a group of independent computing systems that operate as a single virtual system. This type of configuration eliminates individual host systems as single points of failure. Oracle Fail Safe works with Microsoft Cluster Server to ensure that if a failure occurs on one cluster system, workloads running on that system fail over to a surviving system. Oracle Database combined with Oracle Fail Safe on a Windows cluster protects the system from both hardware and software failures.

When properly configured, Oracle Fail Safe ensures a surviving system becomes operational in less than one minute, even for heavily-used databases.

Real Application Clusters Integration with Windows

Real Application Clusters integrate with Microsoft Cluster Server clusters deployed on all Windows operating systems supporting clustering. This enhances high availability by offering:

- Optional automatic restarts of a failed instance or listener in a cluster.
- Detection and resolution of instance cluster hangs.
- Elimination of connect time failover TCP/IP timeout delays for new connection requests.
- Use of user written scripts after database state changes (from online to offline or vice versa).

DIVA Core Cluster Solution

DIVA Core Cluster uses Oracle Fail Safe. An external disk hosts the Oracle data file and backups. The disk serves the nodes through a SAS (Serial Attached SCSI) connection. Two Windows Standard nodes connect to the disk and host Oracle Fail Safe and DIVA Core software.

All software components on each node must have the same release. Release discrepancies may cause cluster failure. For example, if Node-1 has DIVA Core 8.3 installed, Node-2 must also have DIVA Core 8.3 installed, not a different release.

The following software releases are currently supported:

DIVA Core

Release 7.2 or later

Oracle Fail Safe

Release 4.2.1 or later

Microsoft Failover Cluster Manager

Windows Server 2016 or Windows Server 2019

Installation Requirements

In this section, you will identify and confirm that your systems have the proper installation requirements, and set permissions for the domain user and cluster.

Hardware Requirements

- Server requirements for DIVA Core Clustered Managers (two identical servers):
 - Rack-mount chassis
 - One CPU Xeon E5-2420 (six cores - 1.9GHz) minimum
Embedded Oracle license is restricted to one CPU (processor card).
 - 16 GB RAM
 - Two 300GB HDD (Hard Disk Drive) 10,000 RPM (configured in RAID 1) system disks

If you use DIVA Core to archive Complex Objects (for example, DPX), the best course of action is to request specific recommendations based on the estimated traffic (in terms of size and number of objects to be archived per day). In general, Technical Support recommends using a minimum of two 900GB HDD with 10,000 RPM if Complex Objects need to be archived.

This recommendation is also valid for the backup Core Manager or an Core Actor if an Actor server is used for the backup Manager.

For more information and assistance on setting up your RAID refer to Microsoft's Enable Support for Clustered Windows Servers using Clustered RAID Controllers: <https://support.microsoft.com/en-us/kb/2839292>.

- Redundant power supply and fans
- Two on-board Gigabit Ethernet interfaces (copper RJ45 interfaces)
- One SAS or Fiber Channel HBA (Host Bus Adapter) for the shared disk bay connection.

A shared disk bay with dual RAID controller (SAS or Fiber Channel interface) and seven 300 GB SAS disks connected to both servers for the Core Database.

- One Fiber Channel HBA for the tape library control. The Fiber Channel HBA is not required in the following cases:

With SONY Petasite Managed Storage (controlled through the PCS software and a network API).

With StorageTek Managed Storage if the ACSLS software with network ACSAPI interface is used in the configuration.


IMPORTANT: If ACSLS virtual Managed Storage are used, an HBA will be required (consult with Technical Support for more information).

If the library control is based on SCSI LVD interface but some legacy Managed Storage still use SCSI HVD interfaces which are no longer supported, contact

Technical Support in case the library control is based on a SCSI physical interface rather than Fiber Channel.

Windows Server 2016 or Windows 2019.

- Shared disk array requirements are:
 - One direct-attached shared disk array with dual controllers, dual power and dual fans.
 - Six 146 GB disk drives (6 Gb/sec 10,000) RAID 5 virtual disks
 - Two spare physical disks
- Two HBAs for direct attachment of servers to the shared storage



Software Requirements

The following software is required for successful MSCS installation, configuration, and operation:

- Windows Server 2016 or Windows 2019
- Core Database installation package
- Oracle Fail Safe 4.2.1 installation package
- Shared disk array drivers and management software
- All servers must be fully patched with important updates, recommended updates, and Microsoft updates - they must all be the same patch level.

Network Requirements

The following connectivity and parameters are required for successful MSCS installation, configuration, and operation:

- For cluster management, one IP address and host name (DIVA-CL-MSCS) from the public network with corresponding DNS (Domain Name Service) and Active Directory entries on the DNS and domain controllers.
- For the Oracle Cluster Group, one IP address and host name (DIVA-CL-ORC) from the public network with corresponding DNS and Active directory entries on the DNS and domain controllers.

- For the cluster node's public network, two IP addresses - one per node (internal access only).
- For the cluster node's private network, two IP addresses - one per node.
 - The private network is reserved for cluster communications and is commonly referred to as the heartbeat network.
- When configuring the network interfaces:
 - Do not specify a default gateway or DNS servers.
 - On the DNS Settings page, deselect the Register this connection's address in the DNS check box.
 - On the WINS Settings page, deselect the Enable LMHosts Lookup check box.
 - On the WINS Settings tab, select the Disable NetBIOS over TCP/IP check box.
 - Label the network interfaces as Public and Private respectively.
- The two server nodes must be members of a Windows domain.
- If NIC Teaming is in use, it must be configured before you create the cluster.

Example IP Addresses and Host Names

The following are examples of valid IP addresses and associated host name combinations:

- 172.20.128.129 DIVA-CL-MSCS
- 172.20.128.130 DIVA-CL-ORC (DIVA MANAGER VIP)
- 172.20.128.125 RD-MC1 (Public)
- 10.10.10.125 RD-MC1 (Private)
- 172.20.128.127 RD-MC2 (Public)
- 10.10.10.127 RD-MC2 (Private)

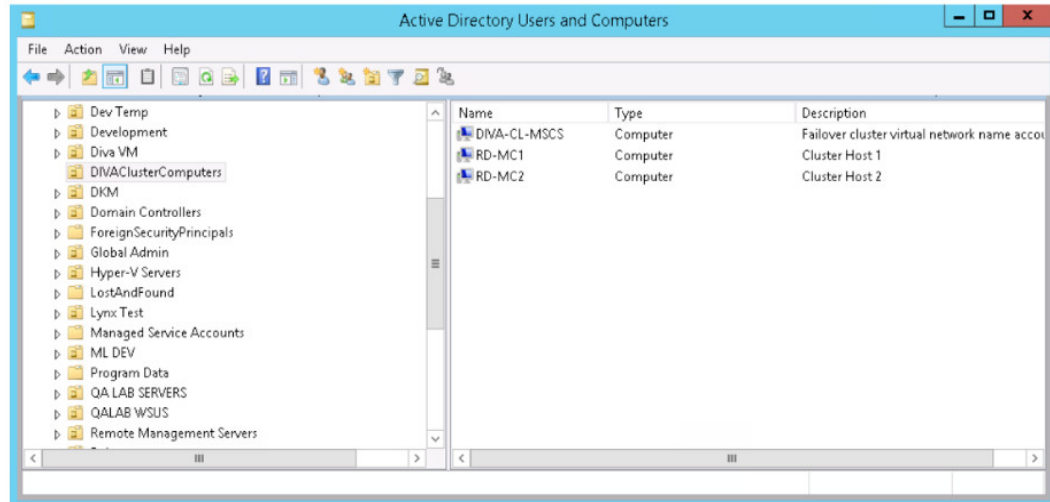
Domain Account Requirements

You must have a dedicated domain account to install and manage the DIVA Core Cluster Manager. You must set the following local permissions on each domain account cluster node:

- Local Administrator
- Logon as batch Request
 - Should be included with Local Administrator permissions.
- Logon as service mode
 - Should be included with Local Administrator permissions.

For example purposes, this book uses a domain account named DIVAClusterAdmin that is a member of the Domain Users Tape Group.

For organizational purposes, Technical Support recommends using a DIVAClusterComputers Active Directory OU (Organizational Unit). You use the Active Directory Users and Computers screen for managing the OU. Active Directory Users and Computers is an MMC snap-in that is a standard part of Microsoft Windows Server operating systems.



Granting Domain User Permissions to Create the Cluster

To successfully create a Cluster, you must ensure the Domain User has permission to Create Computer Objects in the Cluster Container and All Descendant Objects. Alternately, the domain administrator can create a computer object for each node and Cluster Name Objects in advance.

If the domain administrator created an existing computer object, ensure that it is in a disabled state. You must also ensure that the user creating the Cluster has Full Control permission to that computer object using the Active Directory Users and Computers tool before creating the cluster. After you create the Cluster, repeat the steps below to give the Cluster Name Object the same Full Control permissions as the domain user.

To find out more about Cluster Permissions visit:

<https://docs.microsoft.com/en-us/windows-server/failover-clustering/create-failover-cluster#verify-the-prerequisites>

Use the following procedure to add Full Control permissions to the OU for the domain user:

1. Open the Active Directory Users and Computers snap-in from the Windows Server Management console.
2. Right-click the DIVAClusterComputers computer object and click Properties on the context menu to display the Properties dialog box.
3. Click the Security page, and then select the Domain User (DIVAClusterAdmin in the examples) in the Group or user names area at the top of the screen.
4. Click the Advanced button on the bottom right side of the screen to open the Advanced Security Settings screen.

5. On the Permissions page, locate the domain user and click the listing one time to highlight the domain user.
6. Click Edit just under the Permission entries area to open the Permission Entry screen.
7. On the top of the screen, verify that the Type option is set to Allow, and the Applies to option is set to This object and all descendant objects.
8. Select all of the check boxes in the Permissions area.
9. Click OK on the bottom of the screen to apply the permissions.

Granting Microsoft Cluster Object Permissions to Create the Cluster Role

Use the following procedure to grant cluster object permissions to create the cluster role:

1. Open the Active Directory Users and Computers snap-in from the Windows Server Management console.
2. At the top of Active Directory Users and Computers, click View > Advanced Features.
3. Create a new computer object called DIVA-CL-ORC.
4. Right-click the newly created Computer object and select Properties.
5. Click the security tab (this tab is only displayed if Advanced Features are enabled).
6. Click Add.
7. Change the Object Type to include *Computers*.
8. Add *DIVA-CL-MSCS*.
9. Give DIVA-CL-MSCS full control of *DIVA-CL-ORC*.
10. Right-click DIVA-CL-ORC and disable it.

Microsoft Cluster Configuration

This section describes the steps to configure the Microsoft Cluster for use with DIVA Core.

Configuring the Operating System

Now that all disks have been created and configured, you need to configure the operating system on both Cluster Node Servers. First, you will join both server nodes to a single, common domain.

After the server nodes are both in a common domain, add the DIVAClusterAdmin domain account to the local Administrator's group. This must be completed on both Cluster Node Servers.

Now that the Cluster Administrator has been added to both nodes you must configure the MSCS Cluster.

Configuring the Microsoft Cluster Server Cluster

The following procedures for configuring the MSCS cluster must be completed on both node servers.

Installing the Windows Failover Server Clustering Feature

Use the following procedure to install the clustering feature on each node:

1. Log on to the first node server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Open the Server Manager Console and using the menu on the top right side of the screen, navigate to Manage, and then Add Roles and Feature Wizard.
3. When the Add Roles and Features Wizard opens click Next.
4. Select the Role-based or feature-based installation option.
5. Click Next.
6. Click Select a server from the server pool option.
7. In the Server Pool listing area, select the server to use and click Next to connect to the local server.
8. Do not select anything on the Server Roles screen - just click Next.
This screen is only for installing Server Roles.
9. On the Features screen select the Failover Cluster check box.
10. Click Next. A dialog box will open asking to add the required features for failover clustering.
11. In the dialog box, select the Include management tools (if applicable) check box if not already selected.
12. Click Add Features.

13. You will be returned to the Features screen. Click Next.
14. On the Confirmation screen check that the options you selected in the steps above are present.
15. Deselect the Restart the destination server automatically if required check box if it is selected.
16. Click Install.
17. When the installation is complete, click Close.
18. Repeat all of these steps for the second node.

Next you will enable the remote registry service on both node servers.

Enabling the Remote Registry Service

Use the following procedure to enable the remote registry service on each node:

1. Log on to the first node server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Click Start, enter *services.msc* in the search area, and press Enter. This opens the Windows Computer Management utility on the Services page.
3. Double-click the Remote Registry Service to open the Properties dialog box.
4. Select Enable to enable the service.
5. Select Automatic to start the service automatically in the future.
6. Click Start to start the service now.
7. Click OK.
8. Repeat all of these steps for the second node.

Next you will register the host names with the DNS Manager.

Registering the Required Host Names to the DNS Manager

You, or your DNS Administrator, must add the entries for the Cluster Hostname and the DIVA Tape Group Name to the DNS as follows (respectively):

[DIVA-CL-MSCS](#)
[DIVA-CL-ORC](#)

Technical Support recommends also adding each Cluster Host Server public IP address. Use the following procedure to register the host names and IP addresses in the DNS Manager:

1. Open the Server Manager.
2. Select Tools > DNS from the menu on the top right side of the screen.
3. Right-click the DNS Zone and select New Host from the resulting menu.
4. Add the host name (*DIVA-CL-MSCS*) and IP address in the appropriate fields.
5. Select the Create associated pointer (PTR) record check box (if it is not already).
6. Click Add Host.

7. Right-click the DNS Zone again and select New Host from the resulting menu.
8. Add the DIVA Oracle Tape Group Name (*DIVA-CL-ORC*) and IP address in the appropriate fields.
9. Select the Create associated pointer (PTR) record check box (if not already).
10. Click Add Host.

The following steps must be completed on each node server.

1. Log on to the first node server as a local administrator.
2. Open the Windows Network and Sharing Center using System Management App > Network and Internet > Network and Sharing Center, or enter *ncpa.cpl* into run box.
3. Click Change Adapter Settings in the left menu.
4. Locate the network adapter card for the Private network connection and right-click the icon.
The private network is the cluster's heartbeat network only and should not be registered in the DNS.
5. Select Properties from the resulting menu.
6. Double-click Internet Protocol Version 4 (TCP/IPv4) in the protocols area.
7. In the displayed dialog box, click Advanced on the bottom right side of the screen.
8. Select the DNS page on the Advanced TCP/IP dialog.
9. Deselect the Register this connection's addresses in DNS check box.

Next you will create the Windows Server Core Cluster.

Creating the Windows Failover Cluster Resources

The following procedure should be completed on one cluster node only.

1. Log on to the first node server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Select Start > Administrative Tools > Failover Cluster Management Console.
3. In the Management area (in the middle of the screen), click Create a Cluster. This will start the Create a Cluster Wizard.
4. When the wizard opens, click Next.
5. Enter the FQDN (Fully Qualified Domain Name) of the first Cluster Node Server in the Enter server name field and click Add.
6. Enter the FQDN of the second Cluster Node Server in the Enter server name field and click Add.
7. Click Next.

Note: You must be a local administrator on each of the servers that you are validating.

8. On the Testing Options screen, select the Run all tests (recommended) option. This is the default selection.

9. Click Next.
10. On the Confirmation screen, click Next.
11. Monitor the validation tests and wait for them to complete. The Summary screen will be displayed when testing is done.
12. If warnings or exceptions are noted in the summary, click View Report to see the details.
13. Resolve any issues and rerun the Validate Configuration Wizard if configuration changes were made.

Note: Disable unused network adapter cards to prevent minor warnings. Some network adapter cards may have IP addresses on the same subnet. If they are not operational, this may not be an issue.

14. Continue rerunning the Validate Configuration Wizard and resolving any errors until the test all complete successfully.
15. When all tests complete successfully, select the Create the cluster now using the validated nodes check box, and then click Finish to create the Cluster.
When the Validate Configuration Wizard closes, you will be returned to the Create Cluster Wizard to continue with the configuration.
16. Click Next to advance to the Access Point for Administering the Cluster screen.
17. Enter the cluster name (*DIVA-CL-MSCS*) in the Cluster Name field.
18. Enter the Cluster IP address in the Address field.
19. Click Next.
20. On the Confirmation screen, verify that all entered information is correct.
21. Select the Add all eligible storage to the cluster check box.
22. Click Next to create the cluster.
23. When the cluster creation is complete, verify that all configurations were successful by clicking View Report.
24. When you have confirmed that the configuration was successful, click Finish.
Next you must configure the Cluster Quorum Storage.
25. In the Failover Cluster Management Console, expand the navigation tree on the left side of the screen so you can see the cluster.
26. Expand the Storage menu item and select Disks.
27. In the middle of the screen, you should be able to see all of your disks.
28. Select the main cluster item in the navigation tree on the left side of the screen.
29. On the right side of the screen (under Actions), click More Actions > Configure Cluster Quorum Settings. This will start the Cluster Quorum Wizard.
30. Select the Select quorum witness option.
31. Click Next.

32. In the displayed list of Cluster Disks, select the check box for the 100 MB dedicated Quorum Disk. You can identify the Quorum Disk either by the Location (it will show Available Storage), or by expanding the entry using the plus sign and confirming that it is a 100 MB disk (typically Q:\).
33. Click Next.
34. Verify that all selections are correct on the Confirmation screen and click Next.
35. When the configuration is complete, click View Report and verify that all configurations were successful.
36. When you have confirmed that the configuration was successful, click Finish.

Next you will validate the node configurations.

Validating the Nodes Configuration for MSCS Clustering

The following steps are to be completed on one cluster node only.

1. Log on to the first node server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Click Start > Administrative Tools > Failover Cluster Management Console.
3. Select the cluster name in the navigation tree on the left side of the screen.
4. Click Validate Cluster on the right side of the screen (under Actions).
You run the Validate Configuration Wizard again to confirm that there are no errors in your configuration.
5. When the first screen of the Validate Configuration Wizard is displayed, click Next.
6. On the Testing Options screen, select the Run all tests (recommended) option. This is the default selection.
7. Click Next.
8. Click Next on the Confirmation screen.
9. Monitor the validation tests and wait for them to complete. The Summary screen will be displayed when testing is done.
10. If warnings or exceptions are noted in the summary, click View Report to see the details.
11. Resolve any errors and rerun the tests until all tests complete successfully.
12. Click Finish to exit the wizard when all tests complete successfully.

Testing the Configuration

Now that the installation and configuration is complete, you need to test everything to verify proper operation before going to live production. First you will do a manual failover test.

Performing a Manual Cluster Failover Test from the Failover Cluster Manager

Use the following procedure to test manual failover configuration and operation:

1. Open Failover Cluster Manager.
2. If the cluster that you want to configure is not displayed in the navigation tree on the left side of the Failover Cluster Manager, right-click Failover Cluster Manager, click Manage a Cluster, and then select or specify the desired cluster.
3. Expand the cluster in the navigation tree on the left side of the screen.
4. Right-click the cluster you want to manage (DIVA-CL-MSCS).
5. Select More Actions > Move Core Cluster Resources > Best Possible Node.
6. In the middle section of the screen near the top in the summary section you should see the Current Host Server change from one node to the other.

DIVA Core and Oracle Fail Safe Configuration

This section describes configuring DIVA Core and Oracle Fail Safe in the Microsoft Cluster environment.

Installing DIVA Core Prerequisites

Install the DIVA Core Prerequisites on both Cluster Node Servers using the following procedure:

1. Right click the Prerequisites_x.x.exe and run as administrator.
2. Select all options and click Next (only Cygwin+rsync is available in later packages).
3. Enter credentials for the Rsync service.
4. Click Next.
5. Monitor the installation.
6. Click Finish.

Creating the Diva Role

Use the following procedure to create the DIVA role. This is only required on one node.

1. In the left column of Failover Cluster Manager expand the cluster name until you see Roles.
2. Right-click Roles and select Create Empty Role.
3. Double-click the New Role in the center section and rename it to DivaRole.
4. Right-click the DivaRole in the center section and select Add Storage.
5. Add all available disks (E, F, H, G, and so on).
6. Right-click the DivaRole in the center section and select Add > Resource > Client Access Point.
7. Type in the hostname on the first line and populate the IP address in the area provided.
For example, DIVA-CL-ORC / 172.20.128.130 is DIVA MANAGER VIP.
8. Click Next.

Note: This is where most people have problems. If you get an error confirm that the permissions (see [Granting Domain User Permissions to Create the Cluster](#)) and the DNS entries are correct.

9. Click Next when the summary is displayed.
10. Click Finish when the confirmation dialog is displayed.
11. Click the Resources Tab at the bottom of the center section in the Failover Cluster Manager (there are two tabs).

12. You should now see your storage and your new Server Name. Right-click your Server Name and select Bring Online.

Installing Oracle Database on Node1

Use the following procedure to install the Oracle Database on Node1:

1. Log on to Node1 as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Confirm the role is running on Node1:
 - a. Expand the Cluster Name (diva-cl-mscs.domain) in the left section of the Failover Cluster Manager.
 - b. Click Roles.
 - c. Your Role should be displayed in the center section, including a name and status. The fourth column will be Owner Node; confirm this displays the name of Node1.
3. Move the role to Node1 is required:
 - a. Open the Failover Cluster Manager.
 - b. Expand the cluster.
 - c. Expand the roles.
 - d. In the center section, right-click the rolename.
 - e. Select Move > Select Node.
 - f. Select Node1, then click OK.
4. Unzip the Oracle installer (OracleDivaDB_3-1-0_12_2_0_1_0_SE2_Windows_64-bit).
5. Execute *Install.bat* as administrator.
6. Follow the prompts on the installer (you can usually accept all defaults):
 - a. Select Database Size.
 - b. Select Database Home.
 - c. Select Database Mount Points.
 - d. Enter Database Backup Location.
 - e. Enter Database Memory.
 - f. Enter Number of Database Processes.
 - g. Create a Sys Password; you will need this so save it and confirm it is the same on both nodes.
7. Monitor the installer (30-40 minutes).
8. Create a database wallet password, the press Enter.
9. Press enter to close the installer after all scripts complete.
10. Open a new command line as administrator.
11. Stop the Oracle and Listener services, then delete the Lib5 service that was just installed.

```
SC Delete OracleServiceLIB5
SC Delete OracleOraDB12Home1TNSListener
```

12. Delete the E:\Oradata folder that was just created.
13. Delete the F:\Oradata folder that was just created.

Installing Oracle Database on Node2

Use the following procedure to install the Oracle Database on Node2:

1. Log on to Node2 as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Move the role to Node2 is required:
 - a. Open the Failover Cluster Manager.
 - b. Expand the cluster.
 - c. Expand the roles.
 - d. In the center section, right-click the rolename.
 - e. Select Move > Select Node.
 - f. Select Node2, then click OK.
3. Unzip the Oracle installer (OracleDivaDB_3-1-0_12_2_0_1_0_SE2_Windows_64-bit)
4. Execute *Install.bat* as administrator.
5. Follow the prompts on the installer (you can usually accept all defaults):
 - a. Select Database Size.
 - b. Select Database Home.
 - c. Select Database Mount Points.
 - d. Enter Database Backup Location.
 - e. Enter Database Memory.
 - f. Enter Number of Database Processes.
 - g. Create a Sys Password; confirm it is the same on both nodes.
6. Monitor the installer (30-40 minutes).
7. Create a database wallet password, the press Enter.
8. Press Enter to close the installer after all scripts complete.

Configuring Oracle Fail Safe

The procedures in this section will install and configure Oracle Fail Safe. When the installation is complete, you will verify that it was installed properly.

Installing Oracle Fail Safe

The steps in this section must be completed on both Cluster Node Servers.

Fail Safe requires Microsoft's .NET 3.5 SP1 to be installed on the computer before installing Fail Safe. The Fail Safe installation program will notify you if it cannot find .NET 3.5 SP1 on the computer.

Fail Safe also requires that the Cluster Object (DIVA-CL-MSCS) must have full control permissions on the Cluster OU before installation proceeds so the cluster can create a Cluster Group Object. Confirm the permissions are correct (see [Domain Account Requirements](#)).

Oracle Fail Safe 4.2.1 References:

Oracle Fail Safe 4.2.1 Installation Guide

https://docs.oracle.com/cd/E27731_01/doc.41/e24700.pdf

Oracle 4.2.1 Fail Safe Tutorial

https://docs.oracle.com/cd/E27731_01/doc.41/e24702.pdf

Oracle Fail Safe 4.2.1 Concepts and Administrator Guide

https://docs.oracle.com/cd/E27731_01/doc.41/e24699.pdf

1. Log on to both node servers as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Install Microsoft .NET 3.5 SP1 on the computer if not already installed. You can install .NET from the Server Manager Console.
3. Extract the Oracle Fail Safe 4.2.1 installation package into a temporary directory.
4. Execute the *temp_folder\install\setup.exe* file to begin installation.
5. Click Next on the first screen.
6. Select the Typical (243MB) installation.
7. Click Next.
8. Leave the Path as the default and click Next.

Note: The installation path must be the same on both nodes.

9. Enter the Domain Username (*qalab\DIVAClusterAdmin*) in the Username field.
10. Click Next.

11. Enter the Domain User's password in the Enter Password field, and then enter it again to confirm it in the Confirm Password field.
12. Click Next.
13. Review the Summary. If everything is correct click Install; otherwise click Back and resolve any issues.
14. Click Exit when the installation is complete.
15. Restart the node.
16. Repeat all of these steps for the second node.

Next you will verify the Fail Safe installation.

Verifying the Oracle Fail Safe Installation

The steps in this section must be completed on one Cluster Node Server only. Use the following procedure to verify the Fail Safe installation:

1. Log on to the node2 server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Launch the Oracle Fail Safe Manager.

Note: If you receive a message to finish the installation, click OK. It will run a quick validation, click OK and Close.

3. Connect to the new cluster using the cluster alias (DIVA-CL-MSCS) as follows:
 - a. Select the cluster alias in the navigation tree on the left side of the screen.

Note: If the cluster is not shown in the navigation tree, you must add it before proceeding - select Action, and then Add Cluster from the menu.

- b. Select Actions > Connect from the menu on the right side of the screen. This should automatically connect to the cluster.
 4. Select the cluster alias in the navigation tree on the left side of the screen.
 5. Click Actions > Validate from the menu on the right side of the screen. The cluster validation will begin.
 6. You must resolve any warnings or errors before proceeding.
 7. When issues are resolved, run the validation again.
 8. Repeat Step 4 through Step 7 until the validation completes successfully.

Preparing the Database Installation for the Cluster

Use the following procedure to install the database for the cluster:

1. Confirm Role is running on Node2; move the role to Node2 if necessary.
2. Create the folder *E:\OracleClusterDBFiles*.

3. Copy C:\app\oracle\product\12.1.0\dbhome_1\database\initlib5.ora to E:\OracleClusterDBFiles\initlib5.ora.
4. Add the following lines to the bottom of initlib5.ora in C:\app\oracle\product\12.1.0\dbhome_1\database\ (not on E):

```
Ifile='E:\OracleClusterDBFiles\initlib5.ora'  
SPFILE='E:\OracleClusterDBFiles\spfilelib5.ora'
```
5. Save your changes.
6. Delete C:\app\oracle\product\12.1.0\dbhome_1\database\spfilelib5.ora.
7. Open windows services and restart OracleServiceLib5.
8. Open a command line dialog.
9. Log into sql as sysdba with this command:

```
Sqlplus sys/Div1L2b5 as sysdba
```

It should say connected to an idle instance.
10. Create a new spfile with this command:

```
Create spfile='E:\OracleClusterDBFiles\spfilelib5.ora' from pfile
```
11. Start the database using the Startup command.
12. The database should mount and open.
13. Exit the command line window.

Configuring Oracle Fail Safe

The procedure in this section must be completed on one Cluster Node Server only. Oracle Fail Safe will automatically configure some parameters and others you must manually configure. Use the following procedure to manually configure the necessary parameters:

1. Open the Oracle Fail Safe Manager.
2. Expand the Cluster Object (DIVA-CL-MSCS) in the navigation tree on the left side of the screen.
3. Click the Oracle Resources menu item.
4. In the center section you should see Available Oracle Resources.
5. Right-click LIB5 resource and select Add Resource.
6. Click Next on the Name screen.
7. On the Group screen, select the group to add the resource to from the list, and then click Next.
8. On the Parameters screen, you will point to the new initLIB5.ora file you created earlier (E:\OracleClusterDBFiles\INITLIB5.ORA).
9. Click Next.
10. Enter the Sys password you set during the Oracle install on the Authentication Page.
11. Click Next.

12. Click finish on the confirmation page. This will start validating your install and oracle configuration.
13. After validation is complete, look for any errors and resolve. There will be Warnings about logs not being a shared disk, you can ignore these. You may receive errors similar to the following:

```
The database uses a nonclustered disk in one of the system parameters. Value of parameter is C:\APP\ORACLE\ADMIN\LIB5\ADUMP.
```

```
OPW-00029: Password complexity failed for SYS user: Password must contain at least 1 special character.
```

14. Close Failsafe Manager; the remainder of the configuration will be completed in Microsoft Failover Cluster Manager.
15. After validation is complete your Role will have been moved to Node1.

Additional Required Cluster Configurations

Use the following procedure to complete the additional required configuration items for DIVA to be operational in this cluster environment:

1. Log in to Node1.
2. Add the following lines to the bottom of E:\OracleClusterDBFiles\initlib5.ora:

```
spfile='E:\OracleClusterDBFiles\spfilelib5.ora'  
ifile='E:\OracleClusterDBFiles\initlib5.ora'
```
3. Open Failover Cluster Manager on Node1.
4. Expand the cluster and click Roles.
5. In the center section highlight your DivaRole.
6. Select Resources at the bottom of the page.
7. Locate the LIB5 resource.
8. Right-click LIB5 and select Bring Offline.
9. After it goes offline, right-click again and select Bring Online.
10. On both nodes check the C:\app\oracle\product\12.1.0\dbhome_1\database folder. There should be a PWDLIB5.ora file on both nodes. If one node is missing this file, copy it from the other node.
11. Right-click My Computer and go to properties on both nodes.
12. Select Advanced System Settings.
13. Select Environmental Variables.
14. Add a new System Variable as follows:

```
Name= ORACLE_SID  
Value = LIB5
```
15. Open a command line and log into sqlplus as sysdba: *Sqlplus sys/** as sysdba.*

16. After you are logged in, check the local_listener setting:

- Show parameter local_listener
- The value returned should be similar to
(ADDRESS=(PROTOCOL=IPC)(KEY=REGISTER_FsIDIVA_CL_ORC))
- If it is not working execute this command:

```
alter system set
local_listener='(ADDRESS=(PROTOCOL=IPC) (KEY=REGISTER_FsIDIVA_
CL_ORC))' scope=both;
```

Installing DIVA and Adding Services to the Role

Perform the following procedure on both nodes:

1. Launch DivaCore_x.x.x.x.exe installer on the active node.
2. Select C:\Diva as the install directory and click Next.
3. Select Install and click Next.
4. Deselect the Database Schema check box and click Next.
5. Allow the installation to complete, then click Next/Finish.
6. After the installer is done, navigate to C:\DIVA\OracleClient\network\admin.
7. Open tnsnames.ora with notepad.
8. Adjust the HOST parameters from localhost to the DIVA-CL-ORC IP. There may be multiple localhost entries in this file; confirm that all have been adjusted.
9. Install appropriate DIVA services. Always use the DIVA-CL-ORC IP for any service that requires a database connection.

Creating the DIVA Database User

Perform the following procedure only on the active node:

1. Open a command line as administrator.
2. Change directory to C:\DIVA\Program\Database\DBInstaller\bin.
3. Run the following command to create your database user:

```
DIVADBInstaller.bat --requesttype=installrequest --dbuser=diva
--dbpass=pwd --syspass=pwd --dbhost=nnn.nnn.nnn.nnn
```

Note: Replace the password (pwd) as needed. The dbhost parameter should point to the DIVA-CL-ORC IP.

4. Switch to the other node and issue a diva_create_user DIVA and SYS -orapwd to generate the password files.

Add DIVA Services to the Cluster

After DIVA services have been installed and the database user has been created, add them to the cluster role as follows:

1. Open Microsoft Failover Cluster Manager.
2. Expand the clusters.
3. Expand the roles.
4. Right-click the DivaRole and select Add Resource > Generic Service.
5. Select the service you want to add to the list.
6. You may want to add additional dependencies or policies after the service has been added.
7. Right-click the service and add any additional dependencies or restart polices. It is recommended to add Lib5 as a dependency for any service that requires a database connection.
8. Repeat steps 4 through 7 for every service you want to add to the Diva role.


Additional Notes

- Resolving an Issue with SPM Service Not Starting

Perform the following procedure if you experience an issue with the SPM Service not starting periodically:

- a. On both nodes, add a new line in the Environment Variable path and put the line on top as follows:

Environment Variables



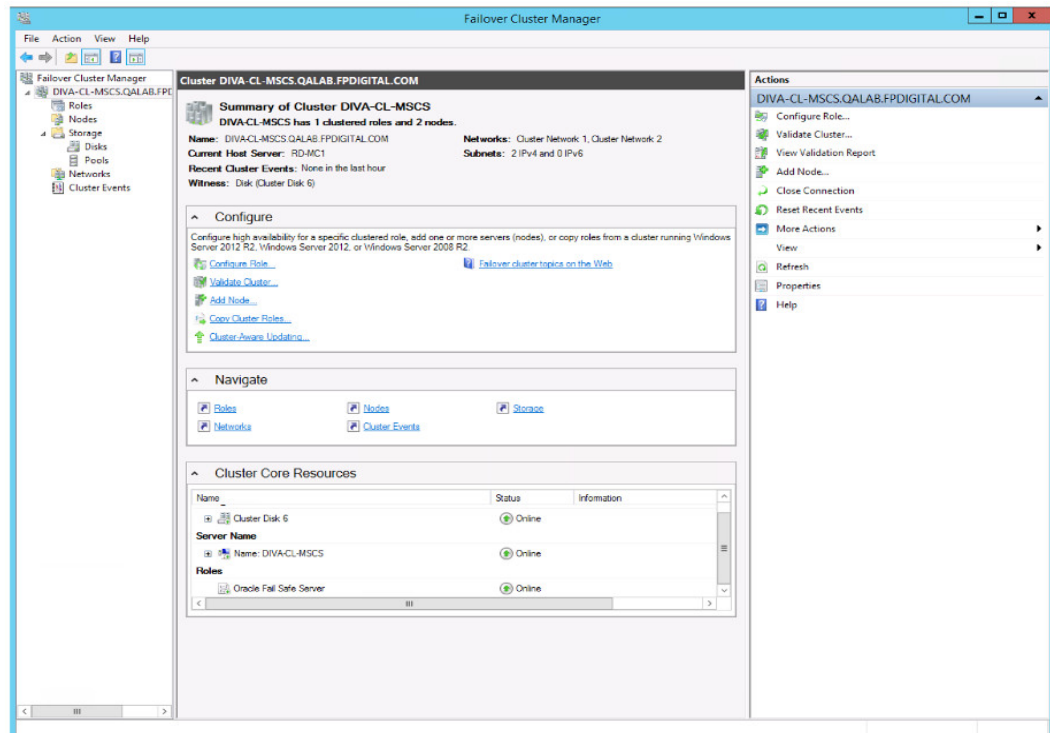
- b. Modify the SPM Service user as local admin or domain user. A resulting consequence is that sqlplus will always be the 11.2 version from C:\DIVA.
- For the REST API put the DIVA-CL-ORC IP in oracle.hostname parameter in C:\DIVA\Program\conf\restapi_dataservice\application.properties file, then test it with https://DIVA-CL-ORC_IP:8765/api-docs.
 - Also put the same IP in the backup_service conf file in DIVACoreAPISettings.


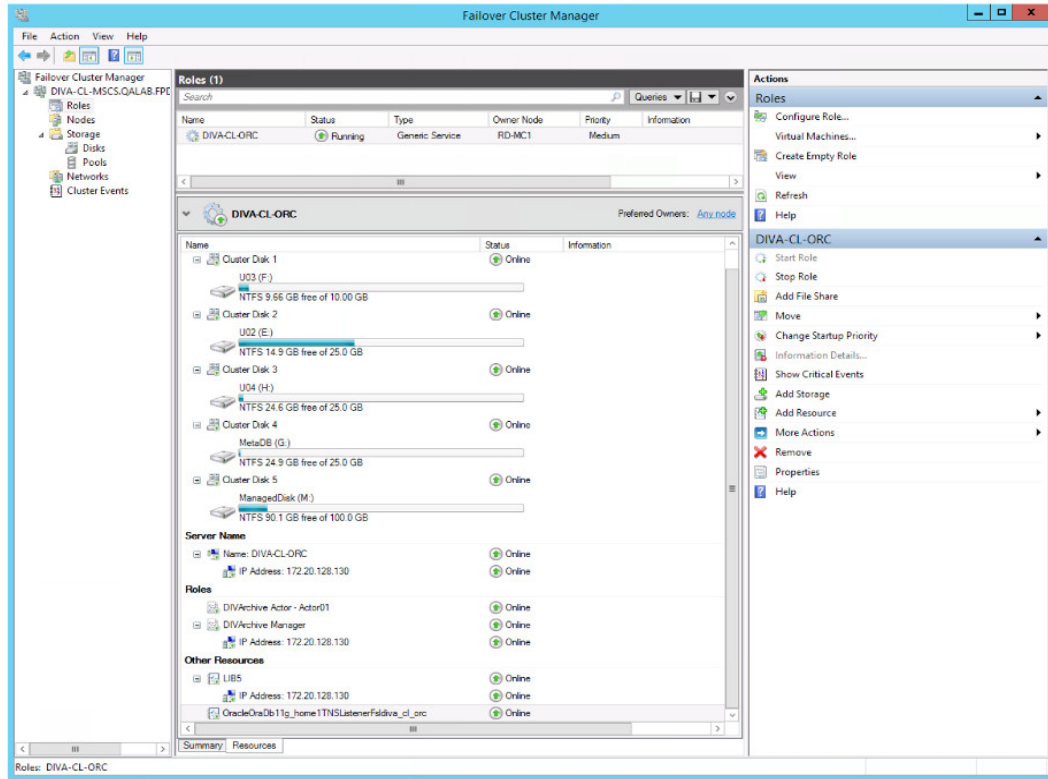
- For MongoDB installations:
 - a. Install MongoDB normally on the active node.
 - b. Switch to the other node.
 - c. Install MongoDB normally on the other node using a DIVA installer install or update.
 - d. Check if the service DIVA mds has been correctly installed; if not the do metadata_service install.
- If you have a disk for DIVA cache and/or storage on the shared RAID, you can configure it to be visible by both nodes at the same time using CVSFS protocol. Just remove it from any role, then right click and choose Add to Cluster Shared Volumes.



Name	Status	Assigned To	Owner Node
Cluster Disk 1	Online	DivADBRole	RUH-AVARCH-02
Cluster Disk 2	Online	Disk Witness in Quorum	RUH-AVARCH-02
Cluster Disk 3	Online	DivADBRole	RUH-AVARCH-02
Cluster Disk 4	Online	DivADBRole	RUH-AVARCH-02
Cluster Disk 5	Online	DivADBRole	RUH-AVARCH-02
Cluster Disk 6	Online	Cluster Shared Volume	RUH-AVARCH-02

Cluster Configuration Examples

This section only includes sample screen shots of successful cluster configuration and no instructional content.

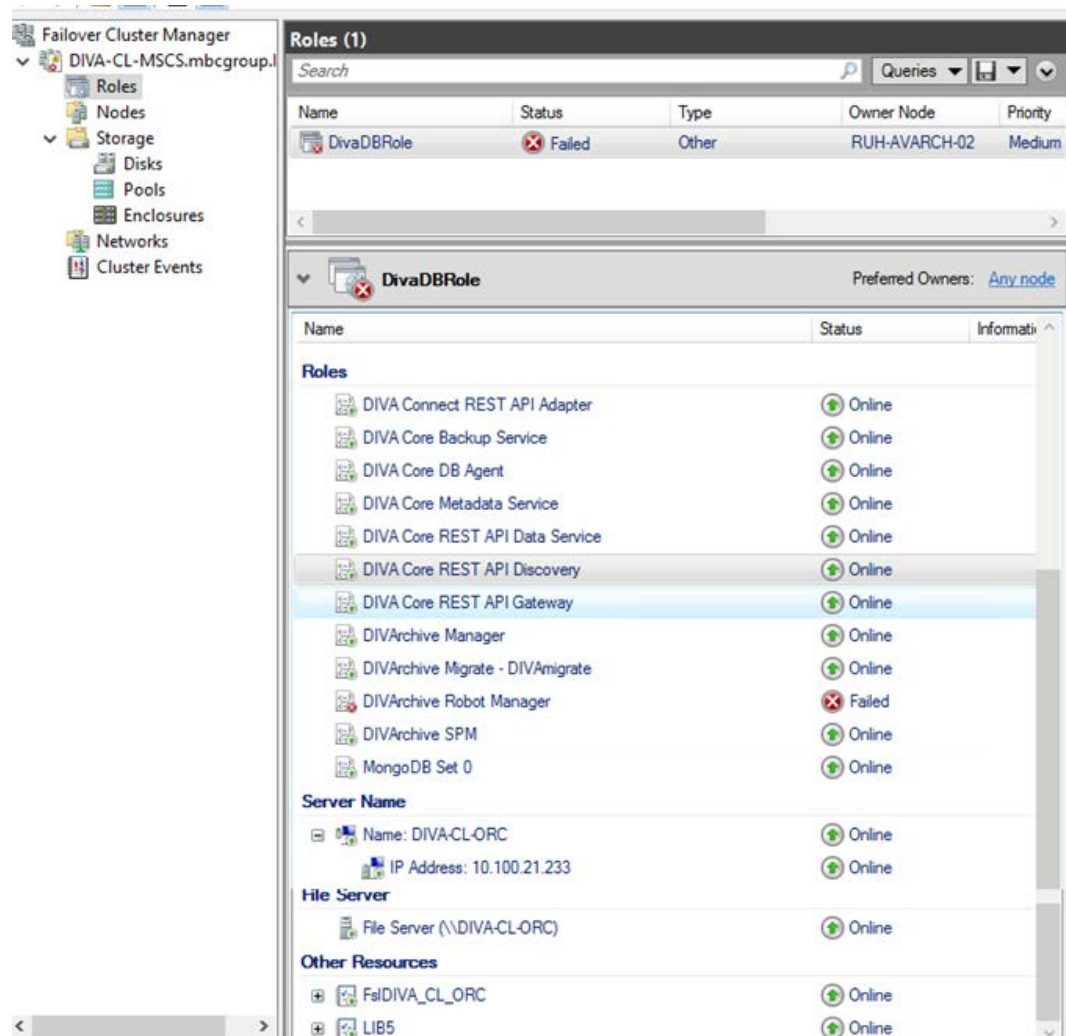






The next screenshot is with all DIVA Services on the cluster. The following are the dependencies configurations:

- The following depends on LIB5:
 - Manager
 - Backup
 - DB Agent
 - Migrate
 - SPM
 - REST API Dataservice
- The following depends on MongoDB:
 - Manager
 - Metadata Service



Testing the Configuration

Now that the installation and configuration is complete, you need to test everything to verify proper operation before going to live production. First you will do a manual failover test.

Performing a Manual Cluster Failover of the Core Resources

Use the following procedure to test manual failover configuration and operation of DIVA Core resources:

1. Open Failover Cluster Manager.
2. If the cluster you want to configure is not displayed in the navigation tree on the left side of the Failover Cluster Manager, right-click Failover Cluster Manager, click Manage a Cluster, and then select or specify the desired cluster.
3. Expand the cluster in the navigation tree on the left side of the screen.
4. Right click the cluster you want to manage (DIVA-CL-MSCS).
5. Select More Actions > Move Core Cluster Resources > Best Possible node.
6. In the middle section of the screen in the summary section you should see the Current Host Server change from one node to the other.

Next you will perform a restart failover test on the active node.

Performing a Cluster Failover Test by Restarting the Active Cluster Node

Use the following procedure to perform a restart failover test on the active node:

1. Connect to the DIVA Core System Management App using the virtual IP address (DIVA-CL-ORC) and confirm normal DIVA Core operation.
2. Disconnect the Public Network cable from the Active Cluster Node.
3. Confirm that the services move and start operation on the second Cluster Node.
4. Connect to the DIVA Core System Management App using the virtual IP address (DIVA-CL-ORC) and confirm normal DIVA Core operation.
5. Reconnect the Public Network cable to the Active Cluster Node.

Next you will test moving a configured role to another Cluster Node.

Moving a Configured Role to Another Cluster Node

Use the following procedure to move a configured role to another Cluster Node:

1. Open the Failover Cluster Manager (if not already open).
2. Expand the cluster in the navigation tree on the left side of the screen.
3. Select Roles.

4. Right-click the role to failover in the Roles area in the center of the screen.
5. Click Move > Select Node from the resulting menu.
6. In the Move Cluster Role dialog box, select the Cluster Node where you want to move the role
7. Click OK; the role will now move to the selected Cluster Node.
8. Verify the Owner Node in the Roles area in the center of the screen; it should now be the selected node.

If all tests have completed successfully, you are ready to place the system into live production.

Maintenance

This section describes routine maintenance and procedures necessary during normal operations. If you have an issue not covered here, refer to the appropriate [Related Documentation](#) or contact Technical Support.

Manually Placing a Service Offline

When a service is experiencing issues, Microsoft Cluster detects that it is offline and restarts the service on the active node. You can take the service offline for maintenance to avoid the service restart using the following procedures:

1. Open the Failover Cluster Manager.
2. Expand the Cluster Object (DIVA-CL-ORC) in the navigation tree on the left side of the screen.
3. Select Roles in the expanded tree on the left side of the screen.
4. Select the failing service in the Roles area in the middle of the screen.
5. Right-click the selected service, and then click Take Offline from the resulting menu.
6. The status of the selected service should now show Offline in the Roles area in the middle of the screen.

Adding a Network for Client Access

You can configure additional client access using the Failover Cluster Manager. This is useful when another subnet is configured for automation. Each node must have one static IP address on the same subnet as listed in the [Network Requirements](#). Use the following procedure to configure additional clients:

1. Configure the new interfaces and subnetwork on each node.
2. Click Start > Administrative Tools > Failover Cluster Management Console.
3. Expand the Cluster Object (DIVA-CL-ORC) in the navigation tree on the left side of the screen.
4. Select Networks in the expanded tree on the left side of the screen.
5. Select the new network to use for automation from the Networks list in the middle of the screen.
6. Click Properties under the listed network on the right side of the screen.
7. Enter a new name for the network used for automation in the Name field.
Using the name Automation for the network makes it easily identifiable.
8. Select the Allow clients to connect through this network check box.
9. Click Apply, and then click OK.
10. Right-click Roles in the navigation tree on the left side of the screen.

11. Click Add Resource from the resulting menu, and then click Client Access Point to open the Client Access Point Wizard.
12. On the Client Access Point screen enter an access point name (for example, DIVA-CL-AUTO) in the Name field.
13. Select the proper network and associated IP address in the Networks list.
You must add the FQDN to the DNS. Refer to the procedures in [Registering the Required Host Names to the DNS Manager](#) and [Creating the Windows Failover Cluster Resources](#) if necessary.
14. Click Next.
15. Verify the selected configuration on the Confirmation screen, and then click Next.
16. When the configuration is complete, verify that all configurations were successful by clicking View Report.
17. Click Finish after you have confirmed that the configuration was successful.

Rebuilding the Cluster after a Node Hardware Failure

Use this procedure when one node fails. The procedure requires downtime during Fail Safe configuration. To rebuild the cluster, complete the steps in the following sections:

1. [Evicting a Failed Node](#)
2. [Preparing New Hardware](#)
3. [Joining a New Node Server to a Cluster](#)
4. [Installing DIVA Core](#)
5. [Installing and Configuring Oracle Fail Safe](#)

Evicting a Failed Node

Do not perform this procedure as the primary troubleshooting method. Eviction should only be used when:

- Replacing a node with different hardware.
- Reinstalling the operating system.
- Permanently removing a node from a cluster.
- Renaming a node in a cluster.

Use the following procedure to evict a node:

1. Log in to the Active Node server as a dedicated cluster domain account user (DIVAClusterAdmin).
2. Click Start > Administrative Tools > Failover Cluster Management Console.
3. Expand the Cluster Object (DIVA-CL-ORC) in the navigation tree on the left side of the screen.
4. Right-click the failed node in the Nodes list in the middle of the screen.
5. Click More Actions from the resulting menu, and then click Evict.

6. A confirmation dialog box asks if you are sure you want to evict the node from the cluster - click Yes to evict the node (or No to leave the node in the cluster).

Preparing New Hardware

When the new hardware is ready, install Windows Server 2012 R2 Standard and all patches to match the Active Node.

Note: Both nodes must be at the same patch level.

Refer to the following procedures:

1. [Configuring the Operating System](#)
2. [Installing the Windows Failover Server Clustering Feature](#)
3. [Enabling the Remote Registry Service](#)

Joining a New Node Server to a Cluster

Use the following procedure to add a new server to the cluster:

1. Follow the procedure in [Validating the Nodes Configuration for MSCS Clustering](#).
2. Before connecting the external disk, ensure there are no local partitions using the D:, E:, F:, or H: drives.
Use the Windows Server Manager to view the disks and assigned drive letters.
3. Follow the procedure in [Replacing an HBA \(Host Bus Adapter\)](#).
4. Add the node to the cluster as follows:
 - a. Log in to the Active Node server as a dedicated cluster domain account user (DIVAClusterAdmin).
 - b. Click Start > Administrative Tools > Failover Cluster Management Console.
 - c. Expand the Cluster Object (DIVA-CL-ORC) in the navigation tree on the left side of the screen.
 - d. Right-click Nodes in the expanded tree on the left side of the screen.
 - e. Click Add Node in the resulting menu to open the Add Node Wizard.
 - f. Click Next on the first wizard screen.
 - g. Proceed through the wizard to add the new node to the cluster.

Installing DIVA Core

Refer to [Configuration Overview](#) to complete DIVA Core installation and configuration. Since the Core Database schema is already in place, do not reinstall the schema on the Active node.

Installing and Configuring Oracle Fail Safe

Use the following procedure to install and configure Oracle Fail Safe:

1. To install Oracle Fail Safe, refer to [Installing Oracle Fail Safe](#)

2. Complete the Oracle Fail Safe configuration as follows:
 - a. Confirm the Fail Safe service was created during the installation.
 - b. Confirm the LIB5 service instance was created during the installation.

Note: The initLIB5.ora file must be replicated on both nodes.

- c. Confirm the Oracle TNS Listener service was created during installation.
- d. Restart the new node and run the tests described in [Testing the Configuration](#).

Replacing an HBA (Host Bus Adapter)

The [SAS \(Serial Attached SCSI\)](#) HBA interfaces external disks dedicated for the database and quorum partitions. Use the following procedure if a SAS HBA fails, or if a node fails and you must rebuild the node using new hardware:

1. Replace the failed SAS HBA in the server following the manufacturer's installation and configuration instructions and recommendations.
2. Launch the Storage Manager software on the Active Node.
3. Locate the Host Mapping area of your Storage Manager.
4. Expand the DIVA Host Tape Group and select the host that contains the new HBA.
5. Right-click the host and click Manage Host Port Identifiers (your menu item listing may be different) from the resulting menu.
6. Select the failed port in the list, and then click Replace.
7. On the following screen, click the Replace by creating a new host port identifier option under Choose a method for replacing the host port identifier.
8. Enter the new host port identifier in the New host port identifier (16 characters required) field, and then click Replace.
9. When the replacement process completes, you should see the Cluster Volumes from the Active Node.

Configuring Windows Firewall with Advanced Security

Microsoft Best Practices recommend enabling the Windows Firewall, however it is not mandatory for DIVA Core. To use the Windows Firewall, use the `DIVACloud_Firewall_Exceptions_2012.ps1` PowerShell script to enable DIVA Core exceptions through the firewall. Use the following procedure to create and run the Firewall Exceptions script in PowerShell:

1. Open Notepad to create a text file.
2. Copy the following script content and paste it into the file you just created.

Note: You may (or may not) need to make adjustments to the line breaks, and so on due to formatting.

```
### Oracle DIVACloud Firewall Exception list. This will enable the
Windows Firewall for all profiles and exclude common DIVA ports.
###
### WINDOWS 2012 Only BELOW ###
Set-NetFirewallProfile -Profile Domain,Public,Private -Enabled
True
New-NetFirewallRule -DisplayName "DIVACloud SSH" -Description
"Oracle DIVACloud
(SSH Remote Access)" -Direction Inbound -LocalPort 22 -Protocol
TCP -Action
Allow
New-NetFirewallRule -DisplayName "DIVACloud DIVA View HTTP" -
Description
"Oracle DIVACloud (DIVA View HTTP)" -Direction Inbound -LocalPort
80
-Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud Remote Administration"
-Description
"Oracle DIVACloud (Remote Administration)" -Direction Inbound -
LocalPort 135
-Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud DIVA View HTTPS" -
Description
"Oracle DIVACloud (DIVA View HTTPS)" -Direction Inbound -LocalPort
443
-Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud CIFS" -Description
"Oracle
DIVACloud (Req. Collection Script)" -Direction Inbound -LocalPort
445
-Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud Oracle TNS Listener" -
Description
"Oracle DIVACloud (Oracle Database - Transparent Network
Substrate)"
-Direction Inbound -LocalPort 1521 -Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud VACP" -Description
"Oracle
DIVACloud (Automation (Harris) Control)" -Direction Inbound -
LocalPort 5010
-Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud DataExpedition" -
Description
"Oracle DIVACloud (ExpeDat - Accelerated File Transfer)" -
Direction Inbound
-LocalPort 8080 -Protocol UDP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud Core Robot Manager"
-Description "Oracle DIVACloud (Core Robot Manager)" -Direction
Inbound
-LocalPort 8500 -Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud Core Manager" -
Description
"Oracle DIVACloud (DIVA API Listener / Systems Monitoring)" -
Direction Inbound
-LocalPort 9000 -Protocol TCP -Action Allow
New-NetFirewallRule -DisplayName "DIVACloud DIVA Core Webservices"
```



```
-Description "Oracle DIVACloud (DIVA Systems Monitoring)" -  
Direction Inbound  
-LocalPort 9443,9763 -Protocol TCP -Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud DIVA Core  
AccessGateway"  
-Description "Oracle DIVACloud (DIVA Communications)" -Direction  
Inbound  
-LocalPort 9500 -Protocol TCP -Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud Core Actor" -  
Description  
"Oracle DIVACloud (DIVAactor)" -Direction Inbound -LocalPort 9900  
-Protocol TCP  
-Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud SNMP" -Description  
"Oracle  
DIVACloud (Systems Monitoring)" -Direction Inbound -LocalPort 161  
-Protocol  
UDP -Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud RDP" -Description  
"Oracle DIVACloud  
(Remote Desktop Protocol)" -Direction Inbound -LocalPort 3389 -  
Protocol TCP  
-Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud NRPE" -Description  
"Oracle  
DIVACloud (Icinga Systems Monitoring - Nagios NRPE)" -Direction  
Inbound  
-LocalPort 5666 -Protocol TCP -Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud NSClient++" -  
Description "Oracle  
DIVACloud (NSClient++ Monitoring w/Icinga)" -Direction Inbound -  
LocalPort  
12489 -Protocol TCP -Action Allow  
New-NetFirewallRule -DisplayName "DIVACloud ICMP" -Description  
"Oracle  
DIVACloud (Packet Internet Groper ICMPv4)" -Protocol ICMPv4 -  
IcmpType 8  
-Enabled True -Profile Any -Action Allow  
### OPTIONAL LOGRHYTHM ONLY### New-NetFirewallRule -DisplayName  
"DIVACloud  
LogRhythm TCP" -Description "Oracle DIVACloud (LogRhythm Log  
Collection TCP)"  
-Direction Inbound -LocalPort 135, 137, 138, 139, 445, 49153 -  
Protocol TCP  
-Action Allow  
### OPTIONAL LOGRHYTHM ONLY### New-NetFirewallRule -DisplayName  
"DIVACloud  
LogRhythm UDP" -Description "Oracle DIVACloud (LogRhythm Log  
Collection UDP)"  
-Direction Inbound -LocalPort 514 -Protocol UDP -Action Allow  
### OPTIONAL NEVERFAIL ONLY### New-NetFirewallRule -Program  
"C:\Program  
Files\Neverfail\R2\bin\nfgui.exe" -Action Allow -Profile Domain,  
Private,  
Public -DisplayName "DIVACloud Neverfail" -Description "Oracle  
DIVACloud  
(Neverfail)" -Direction Inbound
```

```
New-NetFirewallRule -Program "%SystemDrive%\Oracle\Ofs41_1\FailSafe\Server\FsSurrogate.exe" -Action Allow -Profile Domain, Private, Public -DisplayName "DIVACloud Oracle Fail Safe" -Description "Oracle DIVACloud (Fail Safe)" -Direction Inbound  
### WINDOWS 2012 Only ABOVE ###
```

3. Save the file with the file name *DIVACloud_Firewall_Exceptions_2012.ps1*.
4. Open a Windows PowerShell command prompt. You may have to open the PowerShell as a Windows Administrator to successfully execute the script.
5. Navigate to the folder where the script is located.
6. Execute the script by entering *DIVACloud_Firewall_Exceptions_2012.ps1* at the command prompt.
7. All necessary exceptions required for DIVA Core operations should now be included in the Windows Firewall configuration.

If you require additional information or assistance refer to the Microsoft TechNet document named Windows Firewall with Advanced Security located at <http://technet.microsoft.com/en-us/library/hh831365.aspx>.

Cluster-Aware Updating

Cluster-Aware updating automates the Microsoft software updating process on clustered servers while maintaining availability. It is a Microsoft best practice to perform regular Windows updates, however it is not mandatory for DIVA Core. Refer to the following Microsoft TechNet documentation for details on Cluster-Aware updating:

- Microsoft Cluster-Aware Updating
<http://technet.microsoft.com/en-us/library/hh831694.aspx>
- Microsoft Cluster-Aware Updating Best Practice
http://technet.microsoft.com/library/jj134234#BKMK_FW

DIVA Core Installation

This chapter describes DIVA Core software components and system installation.

Topics:

- [Software Component Relationships](#)
- [Installing the DIVA Core System](#)
- [System Management App Installation and Configuration](#)
- [Importing the DIVA Core License](#)
- [Using the MDDB \(Metadata Database\) Services](#)
- [Flashnet Migration Tool](#)
- [Avid AM \(Archive Manager\) Updater Tool](#)

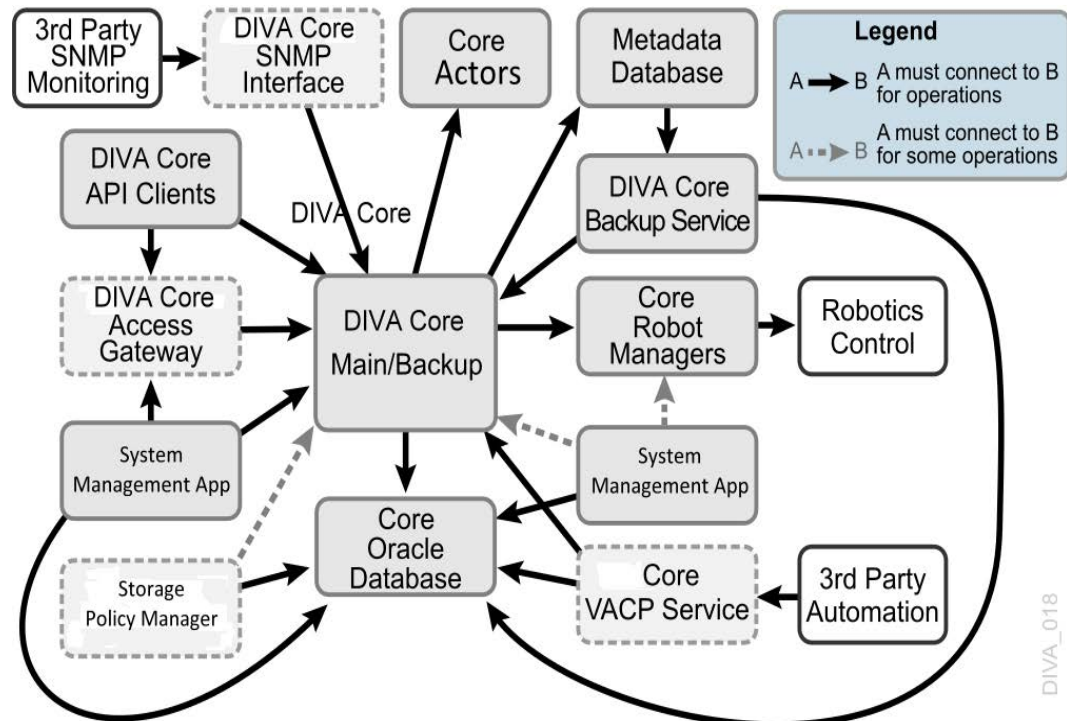
Software Component Relationships

The following figure displays the relationships and dependencies among the software components of a DIVA Core system. It specifically points out the client/server links between them.

A client/server link between two components does not necessarily mean that the server software must be started before the client. For example, the Core Manager to Actor connection. Each Actor acts as a server and the Manager initiates a client connection to the Actor. However, an Actor can be launched after the Manager is running since the Manager will attempt to reconnect to the Actor at periodic intervals.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Note: DIVA Core can run independently of the System Management App. It can be launched at any time after the Core Manager is running.



Software Component Distribution

The DIVA Core platform is flexible and scalable, so the installation of some software components can vary depending on the degree of storage and servers that are managed. Small installations can have all DIVA Core software components installed on a single computer. A very large installation will have these components distributed among several servers. All of these components run as system services.

The following list identifies where the components are typically installed:

Core Managers

Main and Backup Core Manager servers

Core Oracle Database

Main and Backup Core Manager servers

Core Metadata Database

Main and Backup Core Manager servers

DIVA Core Backup Service

Main and Backup Core Manager and Actor servers

Core Robot Managers

Main and Backup Core Manager servers. Robot Managers can also be installed on a separate server when the tape library is installed a substantial distance from the Core Manager servers.

DIVA Core Storage Policy Manager

Main and Backup Core Manager servers

DIVA Core VACP Services

Main and Backup Core Manager servers

DIVA Core SNMP Agent

Main and Backup Core Manager servers

DIVA Connect

Main and Backup Core Manager servers

Core Actors

Core Actor servers

DIVA Core Transfer Manager Communicator (TMC)

Core Actor servers

DIVA Core Archive Manager Communicator (AMC)

Core Actor servers

DIVA Core Watch Folder Monitor

Core Actor servers

Installing the DIVA Core System

The following sections describe installation of the DIVA Core system. Contact Technical Support if you need assistance.

Notes: The Oracle Database must be available for DIVA Core before installation. See [Database Installation and Configuration](#).

Before upgrading, you must confirm that the Array Name, Disk Name, and Cloud Account Name are all the same name. If you configured multiple arrays per cloud account it will not work because of database constraints. In this case, you will need to convert this manually after the upgrade is complete by using the Configuration Utility. Otherwise, the System Management App will not be able to display the cloud array settings correctly.

The basic process to get the System Management App to a functional state is as follows; detailed instructions are in the following subsections:

1. Configure the manager.conf file.
2. Setup the three REST API services (data service, discovery and gateway). This step installs the manager automatically.
3. You can now access the System Management App under the default URL address of <https://127.0.0.1:8765>.

Installation Overview

DIVA Core 8.3 installer includes an option to install the MDDB (Metadata Database). It is run in silent mode during the DIVA installation. This database is not included in the DIVA Linux Installer because there are too many flavors of Linux. DIVA users running under Linux must install and setup MDDB separately, like installing OracleDB separately from the DIVA installer.

The following procedure is a basic overview of the installation process and is common to both Windows and Linux installations. See the following operating system-specific sections for detailed instructions.

1. Install the Core Database user when running the DIVA installer. In Windows this is a check box; in Linux it is a command line question.
2. While installing the database user, make sure to import the license (otherwise the Manager Service will not start until the license is imported using the System Management App after installation).
3. Configure the basic, essential, Manager settings to get the Manager Service operational.
4. Configure the REST API.
5. Start the System Management App and log in under the sysadmin account; then create a System Management App user. This is done so the sysadmin account is not

being used to configure or view the DIVA Core system in the System Management App.

- a. Click the Add User button.
 - b. In the displayed dialog box enter the Username, Password, and select the user's role. In this case, the new user should be assigned an admin role. Sysadmin and admin have the same authority in the system with the exception that an admin cannot manage users.
 - c. Click Save to save the new user. The user will now appear in the Users list.
6. Log out of the System Management App and then log back in with the user account just created (not the sysadmin account).
 7. Configure the Network Servers, and so on until DIVA is fully installed and configured. See [DIVA Core Configuration](#) for configuration details.

Downloading the Software

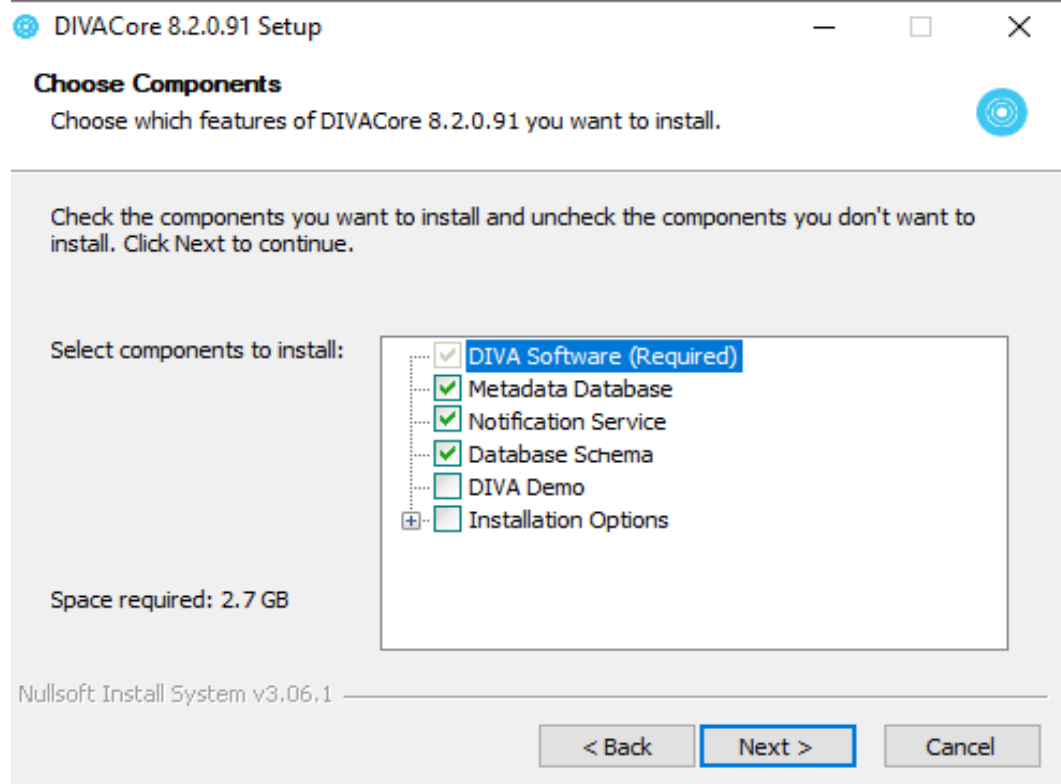
You must stay current with the release of DIVA Core that you install and operate. Current releases of the software are found on the Software Delivery Cloud.

Use the following procedure to obtain the DIVA Core software:

1. Log in to the Software Delivery Cloud and search for DIVA Core.
2. Select the licenses you require (for example, Core Actor, Core Manager, and so on). You must search each time after adding a new license to the list.
3. Select the operating system you run for each selected license using the Select Platform button.
4. Continue through the download wizard, accepting the terms, until the final download screen appears.
5. Confirm that all the licenses you require are listed.
6. Click Download All on the bottom right of the screen, or click the file name link, to download the software.
7. Save the download where it is easily accessible.

Installing DIVA Core for Windows

1. Double-click the executable file to begin the installation.
2. After installation begins, select the check box to install the Database Schema, the Mddb, and the Notification Service (RabbitMQ) then click Next to proceed.



3. The Mddb packaged in DIVA only supports Windows 10, Windows Server 2016 and later. It is selected by default and if you forgot to deselect it and run it on an unsupported operating system such as Windows Server 2016, you will see an error during installation. This error can be ignored because it will not affect the rest of the DIVA installation.
4. Enter the desired installation folder in the Destination Folder field. Technical Support highly recommends using the default installation folder (C:\DIVA). However, if another location is desired, click Browse to navigate the computer to locate the folder. Click Next.
5. Select Install, and then click Next.
6. Enter the required values for the database schema and record them for use later in the configuration process.
7. Enter the License File Path in the text box to import the license during installation.
8. Provide a port and data directory where the Mddb will store log and database files. The default port value is shown in the following figure. The data directory will default to H:\MddbData if H: drive exists. If H: drive does not exist, the default will be C:\MddbData. This directory must exist before proceeding to the next step.

9. Installation will continue in the specified destination folder using the selected components. The installation progress screen is displayed until installation is complete. Clicking Show Details will show the detailed progress (per file) of the installation.
10. The Close button will be highlighted when the wizard is finished.
11. Click Close to complete installation and close the program.

Manually Creating the Database User and Schema for 7.6.1 and earlier

The database user must be created using the DIVA operating system user account. Use the following procedure to create the database user:

1. Open a terminal console.
2. Change to the DIVA_HOME/Program/Database/Core/Install directory.
3. Execute *create_diva_user.bat* (Windows) or *create_diva_user.sh* (Linux), which creates the given DIVA database user and its associated tables

Usage:

```
create_diva_user syspasswd username userpasswd
oracle_connection [-useronly|-tableonly] [-custom_tablespaces
tables_tablespace indexes_tablespace temp_tablespace]
```

```
create_diva_user {DIVA|SYS} current_password new_password [-
orapwd]
```

Parameter Definitions:

- **syspasswd**
Password of the Oracle 'sys' account
- **username**
Username to create
- **userpasswd**
Associated user password
- **oracle_connection**
Oracle TNS service name or Oracle connection string (such as IP_ADDRESS:PORT/ORACLE_SERVICE_NAME)
- **DIVA|SYS**
Mention either DIVA or SYS to reset the respective password in the password file
- **new_password**
New password
- **current_password**
Current password. If there is no current database password, then enter the new password for the is parameter.

- **-useronly**
Only creates the database user and no database objects
- **-tablesonly**
Only creates the database objects for the given user.
- **-custom_tablespaces**
Use of custom tablespaces
 - **tables_tablespace**
tablespace for tables
 - **indexes_tablespace**
tablespaces for indexes
 - **temp_tablespace**
database temp tablespace
 - **-orapwd**
Option to reset/generate password file.

Installing DIVA Core for Linux

Installing DIVA Core in a Linux environment is a manual installation. The following sections describe installation procedures for DIVA Core 8.3 and above on a Linux host computer.

Note: DIVA Core 8.3 installer for Windows includes an option to install the MDDB (Metadata Database). This database is not included in the DIVA Linux Installer because there are too many flavors of Linux. DIVA Core users running under Linux must install and setup MDDB separately, like installing OracleDB separately from the DIVA installer.

Prerequisites and Initial Set-up

These instructions assume that Oracle Linux 7, x86_64 or later, is installed with sqlplus, and the Oracle client.

If you require a Linux environment in a language other than English, create a user and identify the desired language in the user profile. Oracle Linux 7 x86_64 and later has support for a variety of languages (other than English) and the language can be selected during Linux installation.

For more information on Oracle Linux 7 x86_64 and later see the documentation located at <https://docs.oracle.com/en/operating-systems/?tab=2>, or contact Technical Support for assistance.

Note: Linux paths and file names are case-sensitive.

Use the following procedure to prepare for installation:

1. Use the following command to create a directory on the host computer:
`mkdir /home/oracle/Downloads/DIVA_INSTALL`
2. Copy the installation packages to the directory.
3. Confirm you have the latest DIVA Core and DIVA Core API releases and copy them into the directory you created in Step 1. The file transfer can take a bit of time due to the large file size.
4. If the system will have the Oracle Database server, see [Database Installation and Configuration](#) to verify the system meets requirements.

Installing FTP Services

1. Open a terminal console on the host computer.
2. At the prompt enter `yum install vsftpd.x86_64` and press Enter.
3. When prompted if it is OK to install, enter `y` and press Enter.
4. When installation is complete, start the service and confirm that it starts on system startup using the following commands:

```
service vsftpd start
chkconfig vsftpd on
```

5. Create a directory in the `/home/diva` path for managed storage and then mount the / managed partition in this location as follows:

```
mkdir /home/diva/managed
mount --bind /managed /home/diva/managed
```

Installing DIVA Core 8.3 for Linux

The DIVA Core 8.3 installer has an option to create the database user schema. If you select this option you must provide the database information.

1. Open a terminal console.
2. Use the following command to change the permissions and make the installation script executable:

```
chmod +x DIVACore-8.3.{build_number}.sh
```

The `{build_number}` in this command will be the last two digits of the file name. For example, in `DIVACore-8.3.17.sh`, the `0.17` is the build number

3. Use the following command to execute the installation script:
`./DIVACore-8.3.{build_number}.sh`
4. When the Please specify diva user home directory [`/home/diva`] prompt is displayed, press Enter to accept the default directory.
5. The 8.3 installer can create the DB user schema using the installer. When prompted, enter the DB information.

For 7.6.1 and earlier, after installing DIVA Core you must create the database user and schema manually. See [Manually Creating the Database User and Schema for 7.6.1 and earlier](#).

Installing the DIVA Core Services

You control the DIVA Core services using the `divaservice` script.

1. Open a terminal console.
2. Change to the `/home/diva/DIVA/Program` directory.
3. Execute the `divaservice` script using the following options:

command ¹	Descriptions
<code>divaservice configure {SERVICE_NAME}</code>	Configures the specified (already installed) DIVA Core service. The first time you install a service you must use the <code>configure</code> option to include the configuration settings. It will generate a configuration file, install, and then start the service.
<code>divaservice install {SERVICE_NAME} {configuration_file_absolute_path}</code>	Installs the specified DIVA Core service using the specified configuration file. If you already have a fully configured configuration file, use the <code>install</code> option and include the absolute path to the configuration file for that service.
<code>divaservice {start-all stop-all restart-all}</code>	Starts, stops, or restarts all of the services at the same time.
<code>divaservice {start stop restart uninstall status} {SERVICE_NAME}</code>	Starts, stops, restarts, uninstalls, or gets the current status of a specific service.
<code>divaservice list</code>	Lists the names of all currently installed DIVA Core services.
<code>divaservice profile</code>	Displays the DIVA Core services profile.
<code>chkconfig {SERVICE_NAME} on</code>	Use this to start the DIVA Core services when Linux starts. For example, <code>chkconfig DIVAmanager_manager80 on</code> will cause the Manager service to start with Linux.

1. {SERVICE_NAME} can be one of the following: `manager`, `Actor`, `robotmanager`, `migrate`, `WFM`, `lynxlocaldelete`, or `spm`.

Example:

If you are upgrading, or want to install the services with preconfigured configuration files, you can use the `&&` command to do it consecutively (linking them together):

```
divaservice install manager '/home/diva/DIVA/Program/conf/manager/manager.conf' && divaservice install Actor '/home/diva/DIVA/Program/conf/actor/actor.conf' && divaservice install robotmanager '/home/diva/DIVA/Program/conf/robot_manager/robotmanager.conf'
```

Creating System Management App Shortcut

You can add the System Management App Shortcut to your desktop (for easy access) using the following procedure:

1. Open a terminal console.
2. Open the gedit program with root user permissions. If you are not logged in as the root user, use the following command:

```
sudo gedit
```

3. To create the System Management App shortcut, enter the following text and save the file as `/usr/share/applications/diva-control-gui.desktop`:

```
[Desktop Entry]
Version=8.3
Name=System Management App
Comment=DIVA Core CSM
Exec=sh -c "cd /home/diva/DIVA/Program/GUI/bin/ && ./gui.sh"
Icon=/home/diva/DIVA/Program/GUI/bin/gui.ico
Terminal=false
Type=Application
Categories=Application;DIVA Core;
```

4. Use the following command to copy the shortcut to the desktop after you have created it:

```
cp /usr/share/applications/{diva-system-management-app.desktop} /home/diva/Desktop
```

When you click each shortcut for the first time you may be asked if you trust the file. You must confirm them as being trusted files and they will be marked *trusted*.

Starting, Stopping, and Accessing DIVA Core in Linux

The following aliases become available after DIVA Core installation and are defined in `/home/diva/DIVA/Program/.divaenv`:

```
alias DIVAgui="CurrDIR=`pwd`; cd /home/diva/DIVA/Program/GUI/bin; ./gui.sh; cd ${CurrDIR}"
```

```
alias DIVAconf="CurrDIR=`pwd`; cd /home/diva/DIVA/Program/DIVACommand/bin; ./DIVACommand.sh; cd ${CurrDIR}"
```

Note: All Linux paths, file names and command are case-sensitive.

Use the following procedure to start DIVA Core when running in a Linux environment:

1. Open a terminal console.
2. Change to the proper directory as follows:

```
cd /home/diva/DIVA/std_linux
```

3. Start all DIVA Core services as follows:

```
./divaservice start-all
```

4. Open the System Management App as follows (or use the Desktop shortcut):

```
DIVAconf
```

Use the following connection parameters:

User Name

Enter the database user name that was created.

Password

Enter the database user's associated password.

SID

Enter *lib5*

Service Name

Leave this field blank.

IP Address

Enter the IP address of the database host computer.

Oracle Port

Enter *1521*

5. Open the System Management App as follows (or use the Desktop shortcut):

```
DIVAgui
```

When shutting DIVA Core down, close the System Management App and System Management App. When they have closed, use the following command to stop all DIVA Core services:

```
./divaservice stop-all
```

You can use the following command to restart the services (if necessary) for any reason when they are already running:

```
./divaservice restart-all
```

System Management App Installation and Configuration

DIVA Core System Management App is installed as part of the DIVA Core 8.3 installer (and later). It is hosted by the Manager Service. Installing the Manager and the REST API Data services will automatically set it up; see the [REST API Installation and Configuration](#) section for instructions.

Note: The REST API Discovery and Gateway services do not need to be installed to use DIVA System Management App. However, the Notification Service must be installed.

For System Management App to function properly the following must be installed:

- Notification Service: Required for the Running Requests page to display Request Status in real-time.
- RestAPI Data service: Required for login to DIVA System Management App.
- Manager: Required by everything else in DIVA System Management App.

By default, the Manager expects DIVA System Management App SPA files under the DIVA/Program/DIVAWebUI folder. You can modify it (although it is not necessary) in Manager's configuration file.

```
# Location of DIVA Command UI relative to the Program directory
api.command_ui.relativepath=./DIVACommand/

# Location of new Web UI relative to the Program directory
api.web_ui.relativepath=./DIVAWebUI/
```

Navigation Menu

The DIVA Core System Management App is accessible at <https://127.0.0.1:12443/DIVAWebUI>. Use the same login and password as the database the first time you log in.

The application navigation menu will be displayed on the left-hand side. DIVA Core section icons are permanently displayed, and full navigation details are revealed when you hover the mouse over them. Sub-items are revealed when the user clicks on a DIVA Core item. You can pin the full menu by clicking the hamburger button on the top banner.

Backend Support

DIVA System Management App only requires a running Manager and Data Service. You need to update `api.server.port` in the `manager.conf` file to update the System Management App server port. The default port is set to 12443. You must update the

api.dataservice.url in the manager.conf file If you change the data service address and/or port.

```
# URL of the secure API Data Service.  
#api.dataservice.url=https://localhost:13443  
  
# The secure REST API port.  
#api.server.port=12443
```

The manager only supports https, so port 12443 will point to a secure URL where you can access the System Management App interface at <https://xxx.xxx.xxx.xxx:12443/DIVAWebUI> where xxx.xxx.xxx.xxx is the System Management App server IP address. For example, https://127.0.0.1:12443/DIVAWebUI.

Importing the DIVA Core License

DIVA Core 8.3 requires a license. The Manager will not start without a valid license in the database. The license can be imported as part of the DIVA installer if you create the license before DIVA is installed. If DIVA is already installed, a license can be imported using the DIVA Configuration, Licensing, License History section in the System Management App. In addition to enabling the Manager, the license includes a set of options that are necessary to enable the associated features in DIVA. See [Appendix A: Core Options and Licensing](#) for detailed information.

Using the MDDB (Metadata Database) Services

After the MDDB (Metadata Database) is installed, it can be used by DIVA's MDDB Service. The MDDB service is DIVA's REST API micro-service that allows Manager and other services to access the database. The MDDB Service is installed in the `DIVA\Program\Metadataservice` folder. The configuration file is located in the `DIVA\Program\conf\metadata_service` folder, and log file is located in the `DIVA\Program\log\metadata_service` folder.

To install the service run `cmd.exe` as administrator, change to the `DIVA\Program\Metadataservice\bin` folder and type `metadata_service.bat install`. See `metadata_service.bat` help for more details of what other options this script supports.

Note: This command accepts parameter such as `-dburl`, `-certpath`, and so on, which will reset values configured in `appsettings.json` file. If you decide to modify the `appsettings.json` file directly, their values will be overwritten if the service is re-installed again.

```

Administrator: Command Prompt
ERROR: No valid command specified.

DIVA Core Metadata Service Command Line Interface
Usage: metadata_service.bat [command] [options]

where command is one of:
install          (or -i) To install the module as a system service
  options:
    -log          Path to log directory. Default: ..\..\log\metadata_service
    -conf         Path to configuration directory. Default: ..\..\conf\metadata_service
    -httpport    Port to listen for http connections. Default: 1776
    -httpsport   Port to listen for https connections. Default: 1777
    -dburl       Url for DB connection. Default: mongodb://127.0.0.1:27017/Core
    -certpath    Path to certificate located on disk.
                Must be used in conjunction with -certpass and
                cannot be used with -certname
    -certpass    Password to disk based certificate.

uninstall       (or -u) To remove the executable as a system service
start          Starts the module
stop          Stops the module if it is currently running
restart       Stops and subsequently starts the module
status       Determines whether or not the module is running
installdb     Installs MongoDB
  options:
    -datadir     Path to the data directory to store the MongoDB database.
                Default: C:\MongoData
    -port       Port for MongoDB to listen on. Default: 27017
uninstallDB   Uninstalls MongoDB if installed locally.
version       (or -v) Display the module version information and exits
help         (or -h) Displays this information and exits
             (or -?)

Command Failed

C:\DIVA\Program\MetadataService\bin>

```

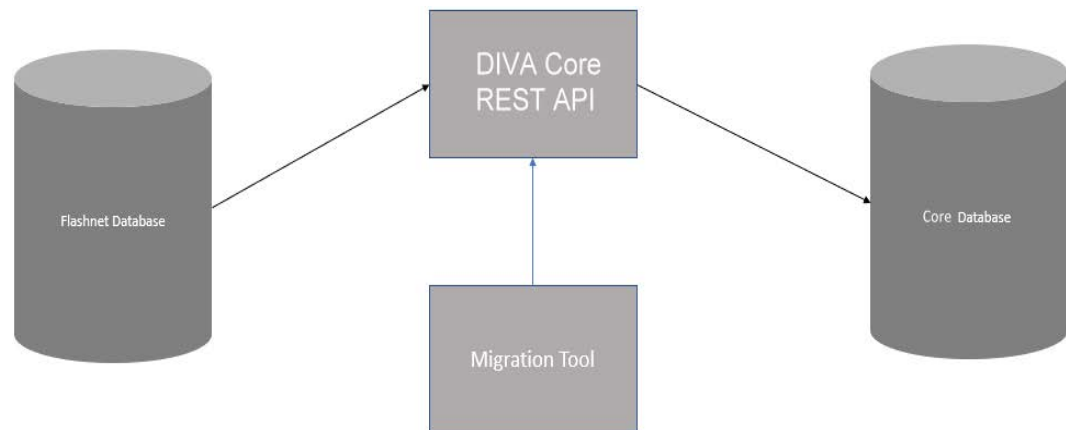
The MDDB service requires MDDB to work correctly; which must be configured in `DIVA\Program\conf\metadata_service\appsettings.json` file as a `ConnectionString`. The `metadata-service.bat` assumes `ConnectionString = "mongodb://127.0.0.1:27017/Core"` by default. However, this only works if MDDB is installed on the same server. If the MDDB Service is running on an operating system that MDDB does not support, you must manually update the connection string to point to correct server where MDDB is installed.

You can verify whether the MDDB Service is running correctly by navigating to <https://127.0.0.1:1777/index.html> which shows the Swagger documentation page for this service.

Flashnet Migration Tool

The Flashnet Migration Tool leverages a hidden REST API endpoint that reads from the Flashnet Database and writes to the DIVA Database.

Note: This initial tool release supports only Flashnet installs with SONY ODA media and objects archived by Avid or IPWS clients.



Caution: Incremental backups (disable full backups temporarily while performing the migration) must be running to clean-up redo logs. Your disk may fill-up during the migration if full backups are running.

The tool is a stand-alone application similar to the Import Tapes Tool to migrate the Flashnet database. The application is located in /Program/Utilities/bin/FlashnetMigration.bat.

The following steps must be completed (in the following order) before using the migration tool:

1. Synchronize the ODA library (media and drives). Make sure that Sync Tape/ Cartridge List check box is not selected.
2. Run a migration script to migrate the tape, groups and objects from the Flashnet Database to the DIVA Database.
3. Synchronize the tape list in the configuration to differentiate between online and offline cartridges.

Use the following format to run the application:

```
FlashnetMigration divaRESTAPIUser divaRESTAPIUserPassword flashnetDBUser
flashnetDBPassword flashnetDBName flashnetDBAddress flashnetDBPort
skipGroupMigration skipTapeMigration] [skipVirtualObjectMigration tapeType
resumeVirtualObjectMigrationOffset
```

For example:

```
FlashnetMigration sysadmin lib5 MSDB tempPassword flashnet
172.16.10.81 1433 false false false 600 0
```

Note: The DIVA Rest API user must have admin or sysadmin access. It is recommended to use the sysadmin / lib5 defaults for simplicity.

If you do not skip a migration and an error occurs during that migration, the current migration step will stop after it has attempted to migrate every resource of the current step.

For example, Tape Group migration will terminate with an error after attempting to migrate every Tape Group. If one of the Tape Groups cannot be migrated (for whatever reason) you can re-run the migration, skipping the Tape Group migration.

For the Tape migration, a user must specify the tape type to use for every migrated tape. This is only known after a database sync. Therefore, a database sync must be performed prior to tape migration.

Object (instance, component and element) migration is a long running step. If the step fails at some point, an error is reported with an offset that can be used to resume the object migration rather than having to start from the beginning.

Avid AM (Archive Manager) Updater Tool

The AM Updater is used to be able to migrate the AAF files into the AM (Archive Manager) Database. The AM Database is needed to be able to initiate restores from Avid using AMC.

The AMUpdater application is a migration/utility tool designed to import AAF metadata files into the Avid Interplay | Production Archive Engine database. The primary scope is to import the AAF files generated by Flashnet in a scenario without Archive Engine into the database of a new fresh installed Archive Engine.

Note: Be aware that in cases where the AAF has not been provided with a Flashnet IPWS archive, restore using Interplay is not possible. A destination located on the ISIS/NEXIS available to the media indexer must be added to the DIVA Core application. Manual restores of these archives/objects can then be restored to the folder, and the Media Indexer should pick up the media as if restored Via Avid.

Partial File Restore support for assets archived using Flashnet IPWS will not be available due to incomplete archive information presented in the .aaf files at the time of archiving (a limitation of IPWS). Partial File Restore of subsequently archived AMC-DIVA assets will be available.

System Requirements

The following are the minimum requirements to run the AM Updater Tool:

- Windows (64-bit)
- 64-bit desktop or server
- .NET framework release 4.5 or later

Installation and Configuration

Use the following procedure to install and configure the AM Updater Tool:

1. Copy the unzipped package to a folder on the local disk.
2. Configure the tool by updating the xml configuration file `AMUpdater.exe.config` with the actual values in the `appSettings` section.

In addition to Avid Interplay | Production Web Services settings (address, port, workgroup, user name and password), the following specific application settings must be configured:

AMUpdater.EnableAvidAMUpdate

This is a Boolean switch to choose between the effective import (true) and the simulated import (false). In the second case the AAF files are processed without updating the AE database; this option can be useful for detection of invalid AAF files prior import.

AMUpdater.AvidAMFolder

The destination folder in Avid AM as the target for the imported assets. In the case of multiple target folders it is the base name for all the target folders. The imported assets are evenly distributed between the target folders.

AMUpdater.AvidAMFolder.Count

The number of destination folders in Avid AM. The default is one (a single target folder with the given name). When the option is for multiple target folders the name of the folders are composed by adding a numeric suffix to the base name. If multiple target folders are created (the automatic creation is off), the manual creation of target folders must follow the same pattern for name.

AMUpdater.AvidAMFolder.CreatelfNotExists

This is a Boolean switch requesting the automatic creation of the target folders in the case they do not exist. If the switch is off (false) then each target folder must exist in Interplay and must be manually created before running the application; otherwise the import will fail. If the switch is on (true) then each target folder will be created by the application unless it already exists.

AMUpdater.AAF.SkipForAMA

This is a Boolean switch requesting to ignore the AMA clip assets. This option should be used for analysis purpose only; it is **not** recommended in a production migration as it can lead to ignoring valid clips.

AMUpdater.AAF.SkipForSequence

This is a Boolean switch requesting to ignore the sequence assets. It can be used in a multiple phase import scenario (see also the complement [AMUpdater.AAF.SkipNotSequence](#)). Note when this option is active it implies and overrides `AMUpdater.AAF.File.SkipForSequence`.

AMUpdater.AAF.SkipNotSequence

This is a Boolean switch requesting to ignore all but sequence assets. It can be used in a multiple phase import scenario (see also the complement [AMUpdater.AAF.SkipForSequence](#)). Note when this option is active it implies and overrides `AMUpdater.AAF.SkipForAMA`.

AMUpdater.AAF.File.SkipForSequence

This is a Boolean switch requesting to ignore the explicit set of the media files belonging to a sequence. It can be used to minimize the overall import duration. Most parts of these files are assumed to be imported using the referred master clips. When this option is active the computed clips for the rendered effects are not imported; the effects must be re-rendered again after restore. Note that this option can be used in a multiple phase import scenario (see also [AMUpdater.AAF.SkipNotSequence](#)).

AMUpdater.AAF.File.SkipForWellKnownInvalidPath

This is a Boolean switch (that applies to any kind of asset) requesting to ignore the specific media files having an invalid file path in AAF. These kinds of files cannot be set into Avid AM and trying to do so will always generate errors for the files in question.

1. After updating the configuration file with the Interplay | Production and application specific settings, run `AMUpdater AAF-Directory` at the command prompt where `AAF-Directory` is a folder containing (in its sub-folders) the hierarchy the AAF files to be imported
2. In the `system.diagnostics` section of the configuration file:
 - The application generates logs. These logs trace the steps of the `AMUpdater` execution. The file path of the log can be set by updating the `initializeData` value. The path must exist and must be accessible; otherwise the log is off.
 - The log granularity respects the switch setting. Full logging details can be obtained by changing the `appTraceSwitch` switch value to 4 (Verbose) while the value of 0 (Off) disables the log.

After the migration process completes, it is **strongly recommended** to uninstall the application by removing it from disk. The process of importing AAF files should not be used on a regular basis due to the risk of affecting the database integrity.

DIVA Core Configuration

Use the System Management App to perform DIVA Core configuration.

Topics:

- [Configuration Overview](#)
- [Prerequisites](#)
- [DIVA Core System Management App and Configuration Utility](#)
- [System Management App and Configuration Utility Tabs Overview](#)
- [Cloud Storage Configuration](#)
- [Cloud Replicated Bucket Scanning](#)
- [Object Auto-Indexing](#)
- [Disk Storage Configuration](#)
- [Object Storage Servers](#)
- [Media Storage Configuration](#)
- [DIVA Core General Settings](#)
- [Licensing Configuration](#)
- [Synchronizing Media and Drive Compatibility with the Database](#)

Configuration Overview

DIVA Core 8.3.x includes the following changes:

- Deprecates and removes DIVA Command completely.
- Begins the deprecation of the original Configuration Utility.
- Begins the deprecation of the original Control Panel.
- Begins migrating configuration and operational functionality to the web-based System Management App.

The following sections identify functions in both the Configuration Utility that have not been migrated yet, and the functions now available in the System Management App. Some functionality is still available in both places, but it is recommended to use the System Management App as much as possible at this stage of the migration to become familiar with the app and its functionality.

Module Configuration Files

Each DIVA Core software module has its own static configuration text file with parameters needed to launch that particular application. The files are typically denoted with the .conf file name extension. There are some DIVA Core modules that use an XML based file rather than a text file for their configuration and those will be noted where applicable.

Unlike older releases of DIVA Core that stored these configuration files in the same folder as the application itself, DIVA Core 8.0 centralizes them to a dedicated conf subfolder under the DIVA Core Program Group.

The configuration files are typically updated with additional or changed settings in newer releases of the software. A new or patch release of DIVA Core will have the new releases of the .conf files appended with a .ini extension. For example, the new release of the Core Manager Configuration file will be named manager.conf.ini. You must remove the .ini extension after the installation is complete and the new configuration file updated.

Each configuration file can be opened and edited with any plain text editor (for example, Windows Notepad).

Any changes made to the configuration file of a DIVA Core software component requires that the component be shut down and then restarted for the changes to take effect. The exceptions to this are the Manager and DIVA Connect options, both of which allow configuration changes to be reloaded while they are still running. There are code dependencies between some applications in the DIVA Core platform, so other components may also need to be restarted for changes to take effect.

DIVA Core Databases

At the system level, settings that relate to the overall operation of each DIVA Core component and their interaction are configured and retained by an Oracle Database. This is commonly known (and will be referred to in this document) as the Core Database (or just simply as the database).

User modification of this database is performed through the DIVA Core System Management App.

The DIVA Core System Management App connects only to the database and does not necessarily require the Core Manager to be running. It is only intended for experienced users and caution should be exercised when altering settings using the utility. An incorrect setting can impede DIVA Core operations or prevent the Core Manager from starting successfully. Contact Telestream Support for assistance if you are unsure about making a particular change.

When launched, the Core Manager obtains the DIVA Core system configuration from the database. However, it does not poll the database for changes made through the System Management App. Therefore, the database must be notified of any changes made. This is performed using the Notify Manager menu item in the System Management App.

You can accomplish most changes to the configuration while the Core Manager is running. There are a small number of configuration changes that require a restart of the Core Manager to become effective. A full list of changes that can be made to the system configuration dynamically while the Manager is running is listed in [Appendix D: Dynamic Configuration Changes](#).

The System Management App also does not dynamically poll the database for changes that are made when the Manager is running. In such cases, you click the Update button in the utility to refresh the information displayed from the database.

You can install the System Management App on any computer that has TCP/IP connectivity to the database and a supported Java Runtime Environment installed. DIVA Core release 8.0 requires the Java Runtime Environment 64-bit (build 1.8.0_45-b14), to be installed to launch the System Management App successfully.

In some cases, a network firewall between the two can prevent a connection. For complete operation and functionality of the System Management App, the Oracle Listener Port (typically 1521) and the Core Robot Manager Ports (typically 8500 and higher) must be opened in the firewall. Full functionality of the System Management App also requires that the Core Manager Port (typically 9000) is open.

Metadata Database

The Core Metadata Database has very high performance and almost unlimited scalability. The Metadata Database should be treated with the same caution as the Oracle Database. It should be backed up at regular intervals through the DIVA Core Backup Service.

Telestream highly recommends that the Metadata Database files are stored on a RAID disk array. The Metadata Database should not be on a standard disk due to decreased performance and the real-time backup functionality that a RAID array offers the system.

Metadata Database files stored on a standard disk are vulnerable to data loss if a single disk failure until the information is replicated through the DIVA Core Backup Service. Storing the Metadata Database files on a RAID array isolates the data from this type of failure.

The information stored in the Oracle Database is already stored on a RAID-1 array and is not subject to data loss if a single disk fails.

Metadata Database Sizing

The following formula can be used as a rough guide to determine the minimum amount of disk space required to support the Complex Object Metadata Database:

$$(100 + \text{average_path_filename_size}) * 1.15 * \text{avg_num_component_files} * \text{num_objs}$$

The following is a general example using the equation:

average_path_filename_size = 60

For example, this/nested/subdir01/As_The_World_Turns_24fps_scenes1-10.avi

avg_num_component_files = 200,000

The Average Number of Files and Folders within the Complex Object.

num_objs = 50,000

The number of Complex Objects to be archived.

In this example, minimum budgeting for a Metadata Database size of approximately 1.67 TB would be recommended.

When planning the system, you must allocate enough Metadata Database disk space to ensure for expected, or unexpected, growth of the environment. The same amount of disk space must also be allocated for the Metadata Database in all of the backup systems.

Environment Variables

Some DIVA Core software components may require defining one or more Windows operating system environment variables for those components to launch successfully.

An environmental variable allows the configured variable to be available to all programs rather than requiring it to be configured from the application each time it is executed. This makes the variable independent of the application and therefore you do not need to manually insert or update the value when the application software is updated or modified.

A User Environmental Variable only applies to an individual Windows User Profile. A System Environmental Variable is applicable to all Windows User Profiles.

The following example illustrates how to configure the DIVA_JAVA_HOME environment variable on a Windows system.

Note: This is simply an example and not required for DIVA_JAVA_HOME as it is already pointing to a valid JRE after installation.

This variable defines the path of the Java Runtime Environment for DIVA Core applications on the Windows host. This particular parameter is required on any Windows computer that will run either the DIVA Core System Management App or the DIVA Core System Management App.

Use the following procedure to configure an environment variable:

1. Open the Windows Control Panel.
2. Double-click the System icon.
3. Click the Advanced tab.
4. Click the Environment Variables button.
5. Click the New button.
6. Enter the variable name in the Variable name field. In this example the name is DIVA_JAVA_HOME.
7. Enter the variable value in the Variable value field. This is the path (or other value) to use for the named variable. In this example the value is C:\DIVA\java.
8. Click OK to complete the process.

You have now defined the variable and it is displayed in the System variables list. The DIVA_JAVA_HOME environment variable is now accessible to all users (and applications) on the system and does not need to be defined each time an applications is executed.

SSL Authentication and Security

DIVA Core 8.0 includes [SSL Certificate Authentication](#) for authentication of services, and securing the internal and API communications in DIVA Core. Certificate authentication provides unique identification and secure communications for each DIVA Core service in a network.

Certificate authentication functions similar to identification cards like passports and drivers licenses. For example, passports and drivers licenses are issued by recognized government authorities. [SSL \(Secure Sockets Layer\)](#) certificates are signed by a recognized [CA \(Certificate Authority\)](#). An SSL certificate verifies the identity of its owner. When the SSL certificate is presented to others, it helps verify the identity of its owner based on the quality of the contents of the certificate.

DIVA Core it comes with Default Root Certificate Authority called DIVA_CA. The DIVA_CA is a self signed CA created for this purpose, and signs all SSL certificates for the DIVA Core services. Every DIVA Core service now has its own password protected private keys and an SSL certificate which will be signed by the DIVA_CA.

See [SSL \(Secure Sockets Layer\) and Authentication](#) for configuration information.

External Certificate Authorities

You can use external third party CAs (for example, VeriSign, Comodo, and so on) with DIVA Core. The external CA must create a [CSR \(Certificate Signing Request\)](#) for DIVA_CA, signed by the third party CA, and the third party certificate must be added to the [Trust Store](#) to satisfy the certificate chain.

See [SSL \(Secure Sockets Layer\) and Authentication](#) for instructions to create the CSR, and installing the thirty party CA in your DIVA Core installation.

Prerequisites

This section describes the prerequisites necessary for successful DIVA Core configuration.

Database Installation and Configuration

See [Database Installation and Configuration](#) for Database installation and configuration details.

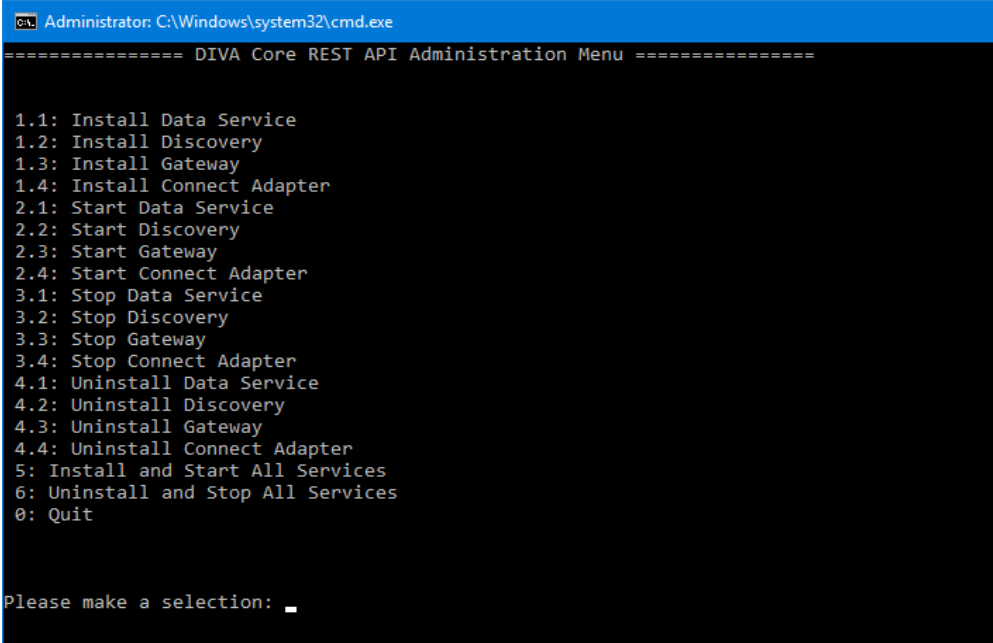
DIVA Core Installer and Database Schemas

See [DIVA Core Installation](#) for details on the DIVA Core installer, Database Schema, and License Importing.

REST API Installation and Configuration

Use the following procedure to install, configure and start the REST API Service:

1. Open the %DIVA_HOME%\Program\conf\restapi_dataservice\application.properties file and confirm that the Data Source parameter settings match the database settings used in the installer.
2. Save the changes and close the file.
3. Navigate to %DIVA_HOME%\Program\RestApi and run the menu.bat file to install the REST API Service.



```
Administrator: C:\Windows\system32\cmd.exe
===== DIVA Core REST API Administration Menu =====

1.1: Install Data Service
1.2: Install Discovery
1.3: Install Gateway
1.4: Install Connect Adapter
2.1: Start Data Service
2.2: Start Discovery
2.3: Start Gateway
2.4: Start Connect Adapter
3.1: Stop Data Service
3.2: Stop Discovery
3.3: Stop Gateway
3.4: Stop Connect Adapter
4.1: Uninstall Data Service
4.2: Uninstall Discovery
4.3: Uninstall Gateway
4.4: Uninstall Connect Adapter
5: Install and Start All Services
6: Uninstall and Stop All Services
0: Quit

Please make a selection: _
```

4. The listed services can be selected individually, or select option 5 to install and start all services (option 5 is recommended).

5. Open the Windows Services panel (services.msc) and confirm that the services were installed and are running.

Note: If the REST API Data Service is not running, change the IP Address in the application.properties file to 127.0.0.1, save the file, then restart the service.

Open your browser and navigate to the Manager's IP address using secure port 8765 (https://nnn.nnn.nnn.nnn:8765). You may receive a notice that your connection is not private; the connection is private and safe. The System Management App uses a self-signing certificate which results in the notice being displayed. Click Advanced, then click Proceed to nnn.nnn.nnn.nnn:8764 (unsafe) to proceed to the System Management App interface. You log in the first time to the System Management App using sysadmin and lib5 (the same password as the database).

IMPORTANT: Log in to the System Management App system the first time and immediately create a new user so that the sysadmin account is not being used for configuration and viewing DIVA Core functions.

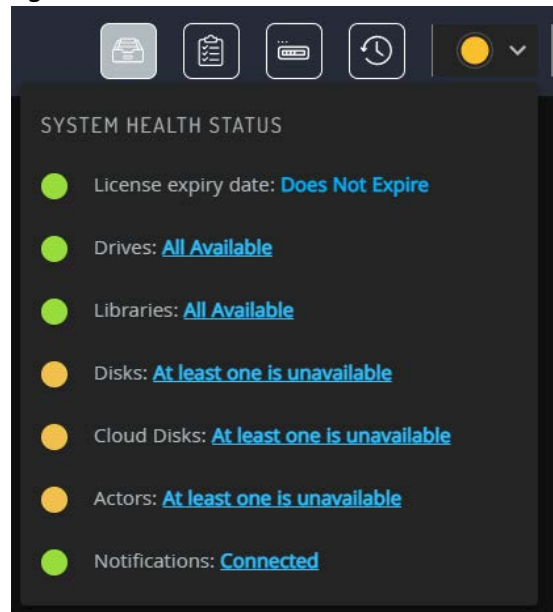
Core Manager Configuration

Copy the %DIVA_HOME%\Program\conf\manager\manager.conf.ini to %DIVA_HOME%\Program\conf\manager\manager.conf and set the DIVAMANAGER_DBHOST parameter to the actual IP of the server hosting the database. Confirm the other parameters in the Database Settings section of the configuration file are identical to the settings entered during installation.

DIVA Core System Management App and Configuration Utility

Caution: The System Management App and Configuration Utility are intended only for experienced users. Incorrect or incomplete changes can adversely affect DIVA Core operations (and possibly even delete data from the archive), or prevent the Core Manager from running. Contact Telestream Support for assistance if you are unsure about making desired changes.

The DIVA Core System Management App primarily connects to the DIVA Core REST API. The app is browser-based and can be run on any computer with TCP/IP connectivity. The status of devices in the system can be viewed from the app using the System Health Status pull down menu on the top right of the screen as shown on the following figure.



When a user logs into the System Management App, it automatically connects and checks the system status. Contact Telestream Support if you still cannot connect after attempting to resolve the error.

Disconnect the System Management App when not in use by logging out of the app and closing the browser being used.

Configuration Utility Overview

The original DIVA Core Configuration Utility is currently used to configure some settings and options in the system. Eventually the System Management App will replace the Configuration Utility, but currently they run side-by-side. However, at this time, there are still items that can only be accessed using the Configuration Utility.

The following sections describe the functionality and use of both the System Management App and the original Configuration Utility.

Configuration Utility Frame Buttons

Each frame in the System Management App includes a set of buttons that perform various functions as follows:

Plus (+)

This button is a plus sign (+). Clicking the button launches a dialog box to add an entry to the frame.

Edit

Highlighting a frame entry and clicking this button enables editing of the entry's properties.

Minus (-)

This button is a minus sign (-). Highlighting a frame entry and clicking this button will remove the entry from the frame. Entries with child dependencies cannot be removed.

Update

Clicking this button refreshes the associated frame content listing from the database.

DIVA Core System Management App Profiles and Passwords

The DIVA Core System Management App provides four fixed user profiles (Administrator, Operator, Advanced Operator, and User) to provide varying levels of access. The profiles require a password that can be changed when logged into the app as a System Administrator (sysadmin account).

To add users navigate to Configuration > User Management and click the +Add User button. Complete all necessary entries on the form and click Save.

The difference between the Operator and Advanced Operator profiles are the Insert and Eject commands, which are only accessible from the Advanced Operator profile. You use the Operator profile during normal operations unless you are inserting or ejecting a tape.

There is no default password to log in to the System Management App as an Administrator or Operator. You must assign an Administrator and Operator password in the System Management App after DIVA Core installation is complete. Without an assigned password you are not permitted to switch to the respective profile in the System Management App. If you attempt to switch to Administrator or Operator mode without first assigning a password to the profile, an error message is displayed notifying you that you must set a password. After you set the profile password in the System Management App the first time, it no longer matters what you use for the old password when changing passwords.

Setting Profile Passwords in the System Management App

Profile passwords are initially created when the new user is created. Use the following procedure to change the profile password:

1. Open the System Management App.
2. Navigate to Configuration > User Management.
3. Locate the user profile to edit and click the edit button (the pencil icon).
4. Enter the following information in the appropriate fields on the edit form:

Old Password

Enter the old password in the Old Password field. You must leave this field blank the first time you set the Administrator password.

New Password

Enter the new password in the New Password field.

Re-type New Password

Enter the new password again in the Re-type New Password field.

6. Click OK to save the changes.

Changing the Database Logging Level in the Configuration Utility

Use the following procedure to change the Database Logging Level:

1. Open the System Management App.
2. Connect to the Core Database.
3. Click the Tools menu item.
4. Click the Change DB Logging Levels menu item (or use F12).
This option is only available when the System Management App is connected to the database.
5. Use the menu lists to select the desired logging level for each package listed in the Change DB Logging Levels dialog box. Available levels are:

FATAL

When selected, this level displays very severe errors that may cause DIVA Core to terminate.

ERROR

When selected, this level displays errors that still allow DIVA Core to operate.

WARN

When selected, this level displays warning messages that are potentially harmful to operations.

INFO

When selected, this level displays informational messages about the progress of the operations.

TRACE

When selected, this level displays messages used to help debug the system.

6. Click OK to save your changes.

System Management App and Configuration Utility Tabs Overview

The following sections describe a general overview of each tab in the Configuration Utility and System Management App. Each tab includes multiple frames where you configure different aspects of the system. The term tabs relate to the Configuration Utility. The System Management App does not have tabs, but rather a menu on the left side of the screen that is expandable using the Start Orb on the top of the screen. Some settings and options under each tab in the Configuration Utility are now available in the System Management App under the Configuration, Resources Management, and Troubleshooting menus as noted in each section.

To notify the Actors of any changes in the Actor configuration in the Configuration Utility, click Notification > Notify Actors while connected to the Manager. In the System Management App the Notify Manager button is located between the username and the System Health Status pull down on the top right of the screen. The Actors must be running and connected to the Manager to receive the notifications.

System Tab

The System tab defines key parameters for your DIVA Core installation and is the starting point for creating your DIVA Core configuration. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

You should create a drawing of the system components before entering details. The drawing includes the data and control paths between components, how they interact with each other, established naming conventions for resources such as disks and tapes, and the workflow of the platform. Some parameters are difficult to change once they have dependencies from other configuration parameters in the database.

System Tab Frames

The System tab includes the following frames:

Production System Definitions

All DIVA Core installations have at least one production system. Additional production systems allow dedication of a particular Actor for specific destinations.

Sites

All installations have at least one site. Additional sites are optional and may be considered by the Core Manager for optimal resource allocation.

Note: Site Support must be enabled in the Core Manager configuration, otherwise all sites will be considered equally.

Sources and Destinations

These define where DIVA Core archives from (Sources) and restores to (Destinations).

Actor Settings

These define Actor host definitions and logical functions. All installations have at least one Actor.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Transcoders

These define DIVA Core transcoders and analyzers. DIVA Core automatically selects the Actors either attached to a BitStream Flip Factory Transcoder installation or integrated with the DIVAnalyze Harris QuiC compressed file analysis software. DIVA Core allows a single transcoder to perform multiple transcodings. DIVA Core assigns additional ports as needed from the base port specified in the configuration. Therefore, a gap of one hundred between individual transcoder port settings is recommended to avoid port conflicts.

Actor Configuration in the Database

Except for the Service Name and Port, all Actor configuration settings are located in the System Management App under the Actor Advanced and Partial Restore Settings tabs in the Actor frame of the Systems tab. Some settings are only available In Engineering Mode and are labeled with an X in the Engineering Mode column of the following tables.

Name	Type	Minimum	Maximum	Engineering Mode	Default
DISABLE_DISK_PREALLOCATION	Boolean			X	Yes
TAPE_TEST_UNIT_READY_TIMEOUT (S)	Integer	60	1200		180
DO_NOT_CHECK_VIRTUAL_OBJECT_NAME	Boolean				No
DO_NOT_CHECK_COLLECTION	Boolean				No
SIMULATION	String			X	
SIMULATION_READING_ERROR_RATE (%)	Integer	0	100	X	0
SIMULATION_WRITING_ERROR_RATE (%)	Integer	0	100	X	0
SIMULATION_TAPE_SIZE (MB)	Integer	20	500000	X	300000
SEACHANGE_CHECK_DELAY (MS)	Integer	0	10000	X	1000
PROFILE_READ_BLOCK_SIZE (B)	Integer	1500	262144		1500
PROFILE_WRITE_BLOCK_SIZE (B)	Integer	1500	262144		32768

Name	Type	Minimum	Maximum	Engineering Mode	Default
QUANTEL_RENAME_CLIPS	Boolean				No
QT_SELF-CONTAINED_THRESHOLD (MB)	Integer	10	100		50
RENAME_TRANSCODED_CLIPS	Boolean			X	
DIRECTORY_SERVER_ENABLED	Boolean			X	Yes
DISK_FTP_PASSIVE_MODE	Boolean				No
DISK_FTP_BLOCK_SIZE	Integer	1024	524288		32
DISK_FTP_SOCKET_WINDOW_SIZE (B)	Integer	65536	10485760		65536
QT_IGNORE_START_TIMECODE	Boolean				No
QT_OMNEON_FIRST_FRAME_HANDLING	String				Reset
AVI_IGNORE_START_TIMECODE	Boolean				No
EVS_MXF_IGNORE_START_TIMECODE	Boolean				No
GXF_TIMECODE_REFERENCE	Integer	0	2		1
GXF_PROGRESSIVE_TIMECODE_TRANSLATION	Boolean				No
LXF_IGNORE_START_TIMECODE	String				No
MXF_PARTIAL_RESTORE_DICTIONARY_FILE	String				
MXF_TIMECODE_FROM_SOURCE_PACKAGE	Boolean				No
MXF_TIMECODE_VALUE_TO_SWITCH_PACKAGE	String				-1
MXF_ENFORCE_CLOSED_HEADER	String				Yes
MXF_RUN_IN_PROCESSOR	String				
MXF_IGNORE_START_TIMECODE	Boolean				No
MXF_USE_BMX_LIBRARY	Boolean				No
MXF_USE_OMNEON_DARK_METADATA	Boolean				No
MXF_SERIALIZE_DEPTH_FIRST	Boolean				No
MXF_GENERATE_RANDOM_INDEX_PACK	Boolean				Yes
MXF_NUMBER_FRAMES_PER_BODY_PARTITION	Integer	50	500		250

Name	Type	Minimum	Maximum	Engineering Mode	Default
MXF_UPDATE_TCTRACK_ORIGIN	Boolean				No
MXF_TOLERANCE_ON_TCOUT	Integer	0	250		0
MXF_DURATION_FROM_FOOTER	Boolean				Yes
MXF_MAX_QUEUE_SIZE	Integer	100	1000		200
SEACHANGE_IGNORE_START_TIMECODE	Boolean				No
MPEG2_TRANSPORT_STREAM_IGNORE_START_TIMECODE	Boolean				No
MPEG2_PROGRAM_STREAM_IGNORE_START_TIMECODE	Boolean				No

Robots Tab

All DIVA Core installations include the Robots tab, although not every installation necessarily has a library. This tab defines basic associations with the robotics software and hardware components.

In the System Management App these options are located under either the Configuration menu or the Resources Management menu. The Resources Management menu contains items for day-to-day operations. The Configuration menu item contains the actual definition and configuration of the resources.

Robots Tab Frames

The Robots tab includes the following frames:

Robot Managers

This frame defines to DIVA Core the connection parameters to each host running a Core Robot Manager Instance.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Managed Storage

Displays the tape or DVD Managed Storage currently configured through one or more Core Robot Managers and their online status.

Media Compatibility

This frame maps the Tape Media Type defined in the Tapes tab, to the Drive Types defined in the Drives tab. Although you can manually remove entries in this frame, they can only be added or updated during a database synchronization with a Robot Manager.

Robot Managers-ACS

This frame associates each Robot Manager with an ACS (Automated Cartridge System) number. Although you can manually remove entries in this frame, they can only be added by performing a database synchronization with the specific Robot Manager.

Disks Tab

The Disks tab defines the physical disks that are to be used by DIVA Core, how they are grouped together for either permanent or cache storage, and how each disk is logically accessed by the Actors. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Disks Tab Frames

The Disks tab includes the following frames:

Arrays

An Array defines a logical association of disks in which one or more physical disks are assigned for use by DIVA Core. The Array Name is equivalent to the Tape Group Name for tapes.

Disks

The symbolic name and location for each disk in your system, whether confined to a single host or shared between hosts. These disks are then assigned to Arrays.

Actor-Disk Connections

Configures how each disk is logically connected to each Core Actor, and how it is to be used. For shared disks accessible by more than one Actor, the disk connection must be declared for all Actors.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Object Storage Accounts

Configures how each object storage account is logically connected to each Core Actor, and how it is to be used.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Drives Tab

The Drives tab is where the drives in your tape managed storage are identified and configured for use with DIVA Core and its Actors. In some installations, a tape library and its drives may be shared with other applications. The configuration options enable you to disable any of the identified drives from DIVA Core use. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Drives Tab Frames

The Drives tab includes the following frames:

Drives

Displays the drives currently identified to DIVA Core in a database synchronization and their current status.

The Drive Edit dialog box enables editing the Serial Number, Status (online or offline), Enabled Operations (Archive, Restore, and so on), and Used (yes or no), information for a drive. This is useful if this information was not retrieved, or was entered improperly, during a SyncDB process. The firmware release for the drive is also displayed in a non-editable field. The firmware information is obtained from the Actors when they scan for tape drive devices. Other additional non-editable information is also displayed in this dialog box, and all of the information is displayed in the Drives frame.

Drive Properties

This displays the drive models currently configured for use with DIVA Core.

Although you can manually remove entries in this frame, they can only be added by performing a database synchronization with a Robot Manager.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Actors-Drives

Indicates to DIVA Core which Actors have access to the drives configured in the Drives frame. In this area associations can be added, edited, or deleted.

Clicking the + button adds an existing Actor-Drive association. The Add New Row in Actors-Drives dialog box is displayed. Use the menu list to select the Actor, and then use the check boxes next to each drive to associate with the selected Actor.

Clicking the Edit button enables edit of an existing Actors-Drives association. The Edit Drives Entry dialog box is displayed. Make the required or desired updates and click OK to save the changes.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Tapes Tab

The Tapes tab defines each Tape Media Type capacity in DIVA Core, and each individual tape's write, repack or to be cleared status. Tapes that do not contain any DIVA Core Objects (that is, they are empty or are from another archive application in a shared library environment) and have been ejected from a DIVA Core managed library can also be deleted from the Core Database from this tab. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Tapes Tab Frames

The Tapes tab includes the following frames:

Tape Properties

Displays the Tape Media Types and configuration parameters currently configured in DIVA Core after a library database synchronization.

In the Tape Properties frame, you can highlight an existing tape and click Edit to open the Tape Properties dialog box.

Empty Ejected Tapes

Displays the tapes that no longer have any DIVA Core content and have been ejected from an attached library. Clicking the - button permanently removes the selected tape from the Core Database.

Inserted Protected Tapes

When a tape is externalized, it is set to Protected Mode by DIVA Core. You must manually remove this state after reinsertion into the library if the tape is to have content written to it.

Tape States

A tape will appear in this frame if either the Enable for Writing or the Enable for Repack states is set to N. The Enable for Writing state can be automatically disabled by DIVA Core if it encounters an error during a read, write, or repack operation.

Click the + button in the Tape States frame to add a tape to the Tape States. Select the tape to add from the list in the displayed dialog box, and then click OK to add the tape.

The Tape States frame gives an overall indication of the reliability of your tape drives. Tapes appearing in this frame (if not manually inserted) indicates that either a read or write error occurred on that tape during DIVA Core operations. If you have many tapes present here this may indicate an issue with one or more of your tape drives and should be promptly investigated.

Sets, Tape Groups & Media Mapping Tab

Use the Sets, Tape Groups & Media Mapping tab to allocate tapes into pools for use by DIVA Core. The Set ID represents each media pool. The Set ID is typically used to distinguish different types of tape media. However, it may also be used to dedicate a specific set of tapes to specific Tape Groups.

A Tape Group is a logical name for the storage of DIVA Core Objects. Each Tape Group is assigned a Set ID of tapes to draw upon. Each Tape Group can only be assigned one Set ID. Several Tape Groups can share the same Set ID.

You can use the System Management App to define the format of an array or Tape Group. The format is configured in the Disks and Sets, Tape Groups & Media Mapping tabs for arrays and Tape Groups respectively. Alternatively, you can use the addTape Group API call define a Tape Group or array and its format. The default format is AXF. This can also be achieved by selecting Legacy in the System Management App, or specify the corresponding value for the format using the API call.

Changing the format of an array is performed through the Edit Array Entry dialog box. Changing the format of a Tape Group is performed through the Edit Tape Groups Entry dialog box. In either case, highlight the desired array or Tape Group and then click Edit to open the associated dialog box. Use the menu list to select the format (Legacy or AXF).

In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Use the following procedure to change the format of an array or Tape Group:

1. Navigate to either the Disks (for arrays) or the Sets, Tape Groups & Media Mapping (for Tape Groups) tab.
2. Highlight the desired array or Tape Group in the displayed list.
3. Click the Edit button at the top of the frame.
4. Use the Tape Format menu list to select the format for the selected array or Tape Group.
5. Click OK to complete the change.

Sets, Tape Groups & Media Mapping Tab Frames

The Sets, Tape Groups & Media Mapping tab includes the following frames:

Unused Tape Sets

Displays empty tapes that are recognized by DIVA Core and the library module where they are located. The Set ID of each tape can also be defined in this frame.

You can highlight an existing tape and click Edit to display the Edit Unused Tapes Sets Entry dialog box. When done editing the Tape Set click the Refresh button to refresh the list.

The Unused Tape Sets frame includes the following columns:

Barcode

These are tapes not currently in use by a Tape Group.

ACS

The ACS number is the specific library where the tape is located.

LSM

The LSM number is the specific library where the tape is located.

Media Type

This is the set's Media Type.

Set ID

This is the tape's Set ID. Click the desired tape to edit and then click the Edit button.

You can select multiple tapes by holding the CTRL key and clicking each tape. You select a range of tapes by holding the SHIFT key, click the first tape in the range, and then click the last tape in the range. Click Edit to open the Edit Multiple Rows

dialog box.

Select the Set ID for the tape from the Set ID list. Only Set IDs that have already been created in the Tape Groups window will be listed.

Setting the Set ID to 99 indicates that the DIVA Core is not to use the tape. This particularly applies to cleaning tapes installed in the library if they are reported to DIVA Core after a library audit. For example, a cleaning tape's typical barcode is CLNn-nnn.

This also applies to some installations where DIVA Core shares its Managed Storage with other applications. Tapes in use by other applications should also have their Set ID set to 99 to prevent DIVA Core from using them.

Tape Groups

Adds, removes, or edits existing Tape Groups and each Tape Group's association with the tape pools defined in the Unused Tapes Sets frame. A Tape Group can only be removed when it no longer contains any DIVA Core Objects.

Additional Set IDs for the Unused Tape Sets frame are only available after they are first created in a Tape Group. Tapes that should not be used by DIVA Core (for example, cleaning tapes) must be configured with a Set ID of 99.

The Tape Groups frame includes the following columns:

Id

This is the library ID the Tape Group belongs to. This is automatically generated by the system and not editable.

Tape Group Name

The name assigned to the Tape Group. These names will appear in the MEDIA list of an Archive request in the System Management App.

Set ID

Default Set ID of each Tape Group is 1. You cannot assign tapes to additional Set IDs until after they are included in a Tape Group.

Description

This is an arbitrary description of the Tape Group.

Media Types

This is the tape Media Types currently in use by this Tape Group. This is updated automatically when a tape is assigned to the Tape Groups Set ID in the Unused Tapes Sets frame.

Tape Format

This is the format of the tape (Legacy or AXF).

Encryption

Identifies whether tape Tape Group encryption is enabled or disabled.

Compression Enable

Identifies whether Tape Group compression is enabled or disabled.

Worse Fit Enable

By default, DIVA Core attempts to fill any tapes already assigned to a Tape Group before assigning an unused tape. The Worst Fit option attempts to span objects on as many tapes as possible.

Repack Reservation

This only applies if the Worst Fit option is enabled. It sets the number of unused tapes in the pool to reserve for tape repacking. All other Tape Groups that also use this Tape Group's Set ID must also have identical values.

VW

This column identifies whether Verify Write is on or off for the Tape Group.

Media Mapping

Media Mapping enables DIVA Core to automatically alter the specified media in an Archive request to another Disk Array, Tape Group or Storage Plan. In this way, you can alter the storage for Archive requests without requiring any changes in the archive initiator (automation or MAM system).

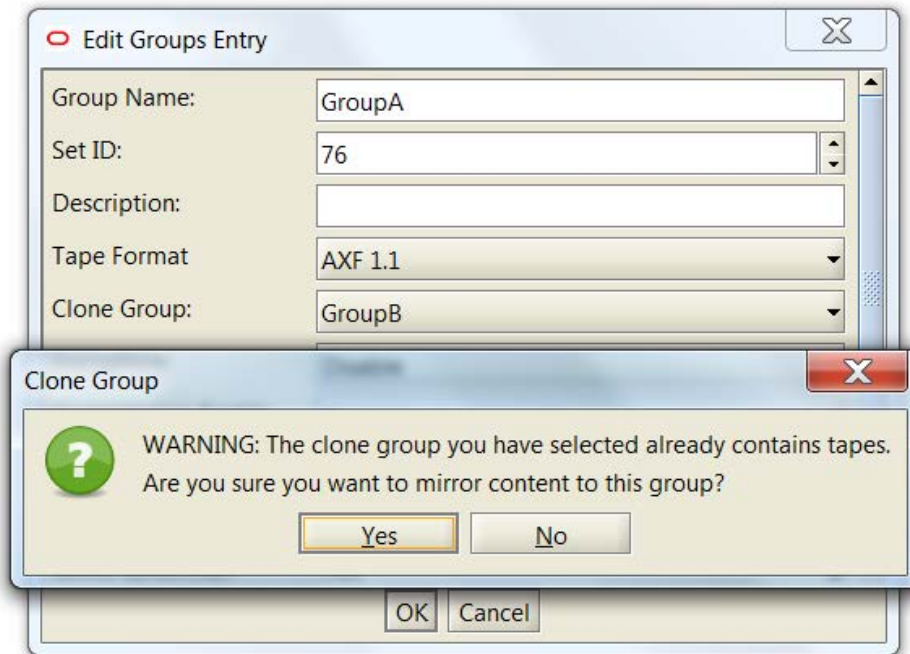
To edit an existing Media Map, highlight the desired mapping and then click Edit to edit the entry. Click OK to save your changes.

Click the + button in the Media Mapping frame to add a Media Mapping entry. Enter the Name for the mapping in the Name field. Use the menu lists to select the From (source), Media to Map to and Storage Plan to Map to, and then click OK to save the entry.

Configuring Clone Tape Groups

A unidirectional Storage Link from a Source Tape Group to a Clone Tape Group can be established by selecting Clone Tape Group in the System Management App's Tape Group configuration window. If the Clone Tape Group already contains tapes, a

warning dialog box is displayed to alert the user that this Tape Group will also contain tapes that are not clones of tapes in the Source Group.



Only Tape Groups that are not already part of a unidirectional Storage Link can be selected as a clone Tape Group.

Groups					
Group Name	Set ID	Description	Media Type	Tape Format	Clone Group
GroupA	76		LT04 , LT03	AXF 1.1	GroupB
GroupB	77		LT04 , LT03	AXF 1.1	
GroupC	78		SAIT2	AXF 1.1	GroupD
GroupD	1		LT05 , SAIT27 , LT04 , SAIT28 , SAXF 1.1		
GroupE	1		LT05 , SAIT27 , LT04 , SAIT28 , SAXF 1.1		
GroupF	1		LT05 , SAIT27 , LT04 , SAIT28 , SAXF 1.1		
GroupG					
GroupH					
GroupI					
GroupJ					
GroupK					

It is not possible to chain Clone Tape Groups, so no Clone Tape Groups are available for selection when a user attempts to edit a Tape Group that was already designated as a Clone Group.

Group Name	Set ID	Description	Media Type	Tape Format	Clone Group
GroupA	76		LT04 , LT03	AXF 1.1	GroupB
GroupB	77		LT04 , LT03	AXF 1.1	
GroupC	78		SAIT2	AXF 1.1	GroupD
GroupD	1		LT05 , SAIT27 , LT04 , SAIT28 , SAXF 1.1		
GroupE	1		LT05 , SAIT27 , LT04 , SAIT28 , SAXF 1.1		
GroupF					
GroupG					
GroupH					
GroupI					
GroupJ					
GroupK					

Edit Groups Entry

Group Name:

Set ID:

Description:

Tape Format:

Clone Group:

Encryption:

Compression Enable:

Worse Fit Enable:

Repack Reservation:

Verifv Write(VW):

OK Cancel

Note: A Clone Tape will remain linked to its Source Tape even after the Tape Group Storage Link is removed. The only way to remove a tape Storage Link is through the Modify Clone Storage Link action.

Analytics App Tab

The Analytics App settings are identified in the Configuration Utility’s Analytics App tab as described in the following sections.

Configuration Frame

Set the main parameters in the Configuration frame as follows:

DB: Maximum possible history of Events in Months

Enter the number of months to retain Analytics App event history.

DB: Maximum possible number of Metrics

Enter the maximum number of Analytics App Metrics stored in the system. Analytics App will remove the oldest entries after this number is exceeded. This is completed through an automated database Request that executes once per day, every day.

Manager: Size triggering Event Queue DB flush (nb events)

Enter the number of events collected in memory before saving them to the database.

Manager: Time delay triggering Event Queue DB flush (secs)

Enter the maximum interval for saving events to the database. If this interval is reached before the size triggering parameter is reached, the events will be saved to the database regardless of how many have been collected.

Event Definitions Frame

The Event Definitions panel displays the list of Event Definitions available for use in the metrics. Double-clicking an Event Definition or clicking Open will display a dialog box listing its associated parameters

Event Definitions are factory set and cannot be modified. Built-in metrics (Analytics App* metrics) cannot be edited and therefore do not appear in the Metric Definitions frame.

Metrics Definitions Frame

Double-click a Metric Definition to display an editing dialog box where the metric can be examined or modified. This has the same effect as selecting a metric in the list and clicking the Edit button.

The + and - buttons at the top of the frame enable adding or deleting a metric.

When adding a metric definition or editing an existing one, the Metric Definitions Properties dialog box is displayed. You can now enter or edit the following information as necessary:

Name

This is the name of the Metric Definition.

Description

This field enables you to enter a description of the Metric Definition that is displayed next to the Metric Name in the Metric Definitions panel. This description also appears in the System Management App when pausing your mouse over an entry of the Metric Definition list.

Enabled

Select this check box to enable the metric. Deselect it to disable the metric.

Collection Type

The Collection Type fields specify which event parameter (for example, Transfer Size) is collected as the data and the statistical computation performed on it (for example, Sum). Available statistics are as follows:

- Average
- Count

- Maximum
- Minimum
- Sum
- Weight-based Average

Weighted By

The Weighted By field specifies the divider parameter for Weight-Based Average collection (for example, Duration).

Collected Event

The Collected Event list specifies the events from which the collected event parameter is retrieved. The list will only display event types suitable for the parameter specified in the Collection Type second field. Event types that have no such parameter attached are absent from the listing.

Resource Type

The Resource Type field specifies which resource breaks down the data. For example, if you select Drive Serial Number, separate metrics will be generated for each drive. Use the menu list to select the resource type for the metric.

Interval

The Interval specifies the interval for metric calculation. For example, selecting 1 Day will generate a metric each day (if corresponding data is available). The metric calculation is based on the associated events that occurred in the last 24 hours. Use the menu list to select the desired interval.

Default Events and Metrics Configuration

The following table identifies the default events and metrics that are internal to DIVA Core:

Event Field ID	Displayed Name	Is Aggregatable? Is Resource?	Is Collectible?	Date or Number	Quantifier
1	Event ID	No	Yes	Number	
2	Event Type	Yes	No		
3	Tape Type	Yes	No		
4	Tape Barcode	Yes	No		
5	Drive Type	Yes	No		
6	Drive Name	Yes	No		
7	Drive Serial Number	Yes	No		
8	Actor Name	Yes	No		

Event Field ID	Displayed Name	Is Aggregatable? Is Resource?	Is Collectible?	Date or Number	Quantifier
9	Object Name	Yes	No		
10	Object Collection	Yes	No		
11	Object Instance	No	No		
12	Media	Yes	No		
13	Request ID	No	No		
14	Event End Time	No	No		
15	Event Duration	No	Yes	Number	Seconds
16	Transfer Size	No	Yes	Number	Bytes
17	Transfer Rate	No	Yes	Number	MB/ Second
18	Transfer Error Rate	No	Yes	Number	Errors/GB
19	Error Code	Yes	No		
20	Error Message	No	No		
21	Disk Name	Yes	No		
22	Library Serial Number	Yes	No		
23	SD Name	Yes	No		
24	Transcoder Name / Analyzer Name	Yes	No		
25	Local DIVA Core System	Yes	No		
26	Number of Operations	No	Yes	Number	

Sample Metrics Definition

The following is a sample use case scenario:

You want to create your own metric for average duration of read and write operations on a tape in a DIVA Core system. Use the following procedure to create this metric:

1. Open the System Management App.
2. Click the Analytics App tab.
3. Click the + button on the Metrics Definitions frame to open the Metric Definition dialog box.
4. Enter a unique name for the metric in the Name field.
5. Enter a description in the Description field.

6. Select the Enabled check box to enable the metric.
7. Set the Collection Type and Weighted By fields as appropriate using the menu lists.
For example, if you select Weight-Based Average as the Collection Type, the Weighted By field is enabled. Because the Weighted By field is active, you are required to select a value to use to weigh the metric definition. In this case, the values for the Weighted By field are identical to the second Collection Type field.
8. Use the check boxes in the Collected Event area to select the events for collection.
9. Use the menu list to select the aggregation Resource Type.
10. Use the menu list to select the aggregation Interval.
11. Click OK to save your metric definition and complete the process.

Media Tab

The Media tab displays information (properties) of the media identified in the DIVA Core system. The display is for informational purposes and read only. You click the Refresh button to refresh the displayed list. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

The Source Media Priority determines which source instance is preferred during the instance selection process of a Restore, Partial Restore, and Copy To Tape Group request, by the media on which the instance resides. Instances on media with a higher priority are preferred. If two instances reside on media with the same priority, DIVA Core will select an instance based on its internal algorithm (same algorithm used in earlier versions of DIVA Core).

Note: Cloud instances are only copied or restored if all local instances are offline or no local instances exist. In other words, this condition is an absolute condition independent of the Source Media Priority.

Storage Plans Tab

Caution: Misconfiguration of the DIVA Core Storage Policy Manager may lead to unexpected and disastrous results! Minor changes can lead to catastrophic consequences. For example, the deletion of hundreds of thousands of instances on tape or database corruption. Without special training and familiarity with the product, it is recommended to contact Telestream Support before making any changes to SPM. Failure to do so may result in severe damage to the DIVA Core system or even permanent data loss.

The Storage Plans tab enables creation of simple and advanced rules for automated management and movement of content within the archive.

For detailed configuration information, see the DIVA Core Storage Policy Manager (SPM) User's Guide. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

See [Appendix A: Core Options and Licensing](#) for detailed information.

Storage Plans Tab Frames

The Storage Plans tab includes the following frames:

Storage Plans

Displays the Storage Plan names and definitions. Use the following procedure to add a Storage Plan:

1. Click the + button in the Storage Plans frame.
2. Enter the Storage Plan name in the Storage Plan Name field.
3. Use the menu lists to select whether to Allow Last Instance Deletion, Please Specify Origin (Internal/External) (this is typically Internal), and the Tape Group/Array Name to associate with the Storage Plan.
4. Click OK to save the changes.

Use the following procedure to edit an existing Storage Plan:

1. Highlight a desired Storage Plan
2. Click Edit to edit the selected Storage Plan.
3. Click OK to save your changes.

Media Tape Groups

Defines the tape groups or disk arrays to be allocated to slots, and if content deletion will be managed by the Storage Policy Manager.

Use the following procedure to add a Media Tape Group:

1. Click the + button in the Media Tape Groups frame.
2. Enter the Medium Name and Storage Name in the designated fields.
3. Use the menu lists to select the Tape Group/Array Name, Watermarked, and the Disk Cleaning Strategy.
4. Click OK to save the changes.

Use the following procedure to edit an existing Media Tape Group:

1. Highlight a desired Media Tape Group.
2. Click Edit to edit the selected Media Tape Group.
3. Click OK to save the changes.

Storage Plans

Displays the Storage Plan names and definitions.

Use the following procedure to add a new Storage Plan:

1. Click the + button in the Storage Plans frame.
2. Enter the Storage Plan name in the Storage Plan Name field.
3. Use the menu lists to select whether to Allow Last Instance Deletion, Please Specify Origin (Internal/External) (this is typically Internal), and the Tape Group/Array Name to associate with the Storage Plan.
4. Click OK to save the changes.

Use the following procedure to edit an existing Storage Plan:

1. Highlight a desired Storage Plan.
2. Click Edit to edit the selected Storage Plan.
3. Click OK to save your changes.

After you click OK, A warning dialog box appears asking whether you want to continue saving the changes.

Click Yes to save the changes, or No to cancel.

Filters

This frame displays filter definitions related to the Storage Plan Objects. It enables performing actions on all or specific objects (based on object filters).

Slots Tab

This tab defines the Slots associated with the Storage Plans for the Storage Policy Manager. Slots define which tape groups or disk arrays are related to each storage plan, and the parameters for storage plan execution. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

For detailed configuration information, see the DIVA Core Storage Policy Manager User Guide.

Slots Tab Frame

The Slots tab includes only one frame named Slots. The Slot Configuration screen serves two purposes; new slot configuration and editing an existing slot configuration. Both functions use the same dialog box. However, the information displayed in the dialog box is determined by whether a slot is being added, or an existing slot is being edited.

Use the following procedure to add a new Slot:

1. Click the + button in the Slots frame.
2. Configure the Slot's parameters by entering the information desired for this slot, or using the menu lists to select the options.
3. Click OK to save the changes.

Use the following procedure to edit an existing Slot:

1. Highlight the desired Slot.

2. Click Edit to edit the selected Slot.
3. Click OK to save your changes.

Manager Setting Tab

Use the Manager Setting tab to set several parameters related to the Media and the Metadata Database in the system. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Media Configuration

There are two settings to configure for the Media in the Manager Setting tab:

Media/Storage Plan Submission Delimiter

The object is assigned to a specific Storage Plan and saved to the specified media. The Media Name and the SP Name must be separated by the “&” delimiter.

Maximum Number of Records in DP_OPERATIONS Table

The maximum number of records maintained in the DP_OPERATIONS table in the database.

Metadata Database Configuration

The following three parameters must be configured in the Manager Settings tab to enable Complex Objects processing:

Complex Objects Metadata Database Location

Enter the full path to the Metadata Database files in this field.

Database Backup Notification

Select the check box to enable the Metadata Database backup notifications, or deselect the check box to disable notifications.

Notifications must be enabled to receive DIVA Core Backup Service messages to the System Management App. If this parameter remains disabled, there will be no notification of errors or warnings displayed in the System Management App. The default is enabled (selected).

Enable Metadata Database Feature

Select the check box to enable the Metadata Database Backup feature. The default is disabled (deselected).

Keystore Configuration

The Keystore password is set in the System Management App in the Manager Configuration view. The Keystore password is entered in the Export: Tape Encryption Keystore Password field. The password must be at least eight characters and contain at least one digit, at least one lowercase alphabetic character, at least one upper case

alphabetic character, and at least one special character within a set of special chars (! @ # % \$ ^).

Exporting encryption keys are enabled by selecting the Export: Enable Export of Encryption Keys check box. Exporting encryption keys is disabled by default. Engineering mode must be used to view or edit both settings.

You can verify the integrity of the Keystore file using the Java keytool. See <https://docs.oracle.com/javase/6/docs/technotes/tools/solaris/keytool.html> for details on Java's keytool.

License Tab

This tab is used to import or export your DIVA Core licensing information into the Core Database. New licenses can be added using this tab without restarting the system.

Click the Import button to open the Import License window. Enter all of the required information and click Submit to import the license.

The following figure displays the License Tab in the System Management App. The Expiration Date column shown is used for temporary licenses.

License Status:

- Green: Valid more than 60 days
- Yellow: Valid less than 60 days
- Red: Valid less than 30 days

Cloud Storage Configuration

The Cloud Storage Configuration screen in the System Management App enables configuring and viewing the DIVA Core Cloud Arrays.

Use the following procedures to configure the listed Cloud Storage components:

1. Click the DIVA Configuration, Cloud Storage Configuration menu item on the left of the System Management App screen.
2. Click the + icon in the Cloud Arrays area of the screen to display a dialog box where you will configure the various DIVA Core Cloud Storage Arrays.

There are four sections to the configuration; Cloud Account, Array Settings, Disk Settings, and Connected Actors. Each parameter is filled in using default values; change the values as applicable for your system.

3. Enter the configuration for this specific array, then click Save to add the array to the Cloud Arrays list.

Repeat this step for all Cloud Storage Arrays that should be accessible in DIVA Core.

The following subsections describe the available parameters for cloud storage accounts.

Common Parameters

The following is a list of common parameters:

Account Parameters

Name

Enter a name for the cloud array in the Array Name field. This is symbolic and typically represents the purpose of the stored objects. This name is used as the Media parameter in the archive request.

Type

The Type field defines the type of protocol to be used to access the object storage server. For example, S3, OCI, Azure, and so on.

URL

The URL field has the following format: [<protocol>://]<ipaddress or host-name>[:<port>], where the following applies:

- The protocol can be http or https. If this parameter is omitted the default protocol is https.
- IP address or hostname is mandatory.
- Port is optional (it is defaulted to 80 with http and 443 with https)

Proxy

If the object storage can only be accessed through a HTTP proxy server, the host-name and port (optional) of a proxy server can be specified. For example, `http://proxy:80`.

Threads Per Transfer

This value sets the number of uploading threads a request can send to the object storage at a time. The default value is 5.

Part Size (MB)

When uploading using multiple threads, large files must be broken down into smaller parts to optimize upload speeds. The default value is 50 MB and the maximum is 200 MB. Small part sizes (10 to approximately 20 MB) are more suitable for slower networks. Medium part sizes (20 to approximately 50 MB) are good for high speed and low latency networks. Larger part sizes (50 MB and greater) could be needed to upload large files (< 1 TB).

Note: The performance of a transfer to an HTTP based object storage server depends on multiple factors: latency, bandwidth, CPU/memory on the client machine and these two parameters; Threads per Transfer and Part Size (MB). There is no default optimal configuration. A high number of threads per transfer will consume more CPU. A large part size will consume more memory. On the client side, the system will under perform if the CPU is saturated or if there is no available memory. The recommended methodology to get optimal settings is to set the part size to a low value; for example, 10 to approximately 20 MB. Then run some performance tests with a large file by increasing the number of threads per transfer until the CPU get saturated, or until there is no more gain in terms of performance. When the best number of threads has been identified, increase the Part Size until you get the best performance without consuming too much memory per transfer.

Unique System ID

When Bucket Name is not specified, this identifier is used by DIVA Core in the naming of buckets/containers. By default System Management App will populate this field using the system UUID that can be found under DIVA Core settings. However if multiple arrays were created that depend on this ID, the user should generate a unique UUID for each account instead of using the default value populated by System Management App. This is because the identifier should be unique to each storage account.

Array Parameters**Format**

This is the same format field that is used for Disk Arrays. LEGACY format is not supported for object Storage.

Storage Class

The Storage Class defines the type of backend storage to be used by the server to store objects. The default class is STANDARD for all the different types of cloud accounts. Depending on the implementation, more storage classes are supported.

Verify Write

Select whether to enable Verify Write from the list. Verify Write is not compatible with Complex Objects.

Default Checksum Type

The MD5 algorithm is the Default Checksum Type. This field is not editable in this dialog box. However, it can be changed in DIVA Core Settings.

Max Instances Per Bucket

It is possible to assign a maximum number of instances per bucket. Previously, this value was hard coded to 1000, but it is now configurable and simply defaults to 1000 instances per bucket for an Oracle OCI / OPC account and 100,000 for an Amazon S3 account. It can be configured to 1,000, 10,000 or 100,000 instances per bucket. When Bucket Name is specified, this parameter is ignored and DIVA stores an infinite number of instances under the name specified in the Bucket Name parameter.

Bucket Name

With AXF and AXF native files and folders formats, it may be useful to specify a custom bucket name by specifying it in the storage options. When a custom Bucket Name is specified all instances will be written into the same bucket. The maximum number of instances per bucket is unlimited when a custom Bucket Name is specified. If DIVA is allowed to generate the Bucket Name, then DIVA will only put the configured maximum number of instances per bucket before creating a new bucket. The bucket name parameter may also contain an optional sub-directory. If there is a directory specified, DIVA will prefix all the files with this directory name (for example, my-bucket/directory). This option applies to object storage arrays only (OCI, S3, Azure, GCS, and so on).

Priority

During restore, an instance on the media with higher priority is used first. However, for cloud arrays this field is ignored unless `DIVAMANAGER_ALLOW_MEDIA_PRIORITY_CONTROL_OF_CLOUD_INSTANCES` is set to `TRUE` in `manager.conf`. If the value is set to `FALSE` (or commented out), then all cloud instances are considered to have the least priority regardless of the value set here.

Description

Enter an optional description for the Cloud Array.

Disk Parameters

Site

Similar to Disk Arrays, this parameter defines the location of a cloud array.

Status

Set the current logical status of the cloud array to ONLINE or OFFLINE.

Min Free Space MB

This is set to the minimum free space before DIVA can consider the disk as full. With most object storage implementations with unlimited space, this field will not be considered.

Protocol Specific Configuration

The following subsections describe protocol-specific parameters.

Configuring S3 Storage Accounts

Storage accounts allow users to configure programmatic access to a their S3 account. The configuration data in a DIVA Core storage account is exclusively used by Core 's Actors to query S3 storage and transfer content to and from S3 buckets.

The following are S3 Storage Account parameters:

Type

Select S3 as the storage type.

Access Key/Secret Key

Specify the Access Key and Secret Key provided by the S3 server to the user during creation of the S3 account.

Region

Enter your local region in the text box.

Restore Tier

Users can specify an additional configuration setting called the Restore Tier for Glacier and Deep Archive storage classes. The Restore Tier is used to specify the rate of retrieval. For Glacier, it can be either EXPEDITED (1-5 minutes), STANDARD (default - 3-5 hours), or BULK (5-12 hours).

For Deep Archive it can only be either STANDARD (within 12 hours) or BULK (within 24 hours).

It is possible to override both the Storage Class and Tier using Request Options. Users can specify the options as `-storageClass=<STORAGE CLASS>` or `-storage_class=<STORAGE CLASS>` and `-restoreTier=<RESTORE_TIER>` or `-restore_tier=<RESTORE_TIER>`.

Use Virtual Hosted Style

Some S3 implementations provide two styles for accessing S3 Objects in a bucket; virtual-hosted-style and path-style. Virtual-hosted-style can be used when buckets are part of the global domain in the URL. For example, <https://bucketname.s3.amazonaws.com>.

Configuring Azure Blob Storage Accounts

Azure Blob Storage is the object storage service offered by Microsoft to Azure account owners. Like any other object storage service, Azure Blob Storage offers an interface to create blobs (objects) under buckets.

The following are Azure Blob Storage Account parameters:

URL

The URL is usually set to <https://blob.core.windows.net> or <http://127.0.0.1:10000> when using the Azure Storage Emulator. The URL can also be set to https://<IP_Address> when using a dedicated Azure server. If this format is used, DIVA will internally map the IP address to blob.core.windows.net.

Login

This is set to the account name. The account name is *devstoreaccount1* when using the storage emulator.

Password

This is the account key that is used to sign HTTP requests. This is a Base64 encoded string. The account key is *Eby8vdM02xNOcqFlqUwJPLlmEtlCDXJ1OUzFT50uSRZ6IF-suFq2UVERCz4l6tq/K1SZFPTOtr/KBHBeksoGMGw==* when using the storage emulator.

Bucket Name

This field works as previously described. Bucket limitations are listed at <https://docs.microsoft.com/en-us/rest/api/storageservices/naming-and-referencing-containers--blobs--and-metadata>.

A Custom Bucket Name is required if you are using Azure Lifecycle Policy Manager (<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts?tabs=azure-portal>) to automatically change the tiers of your objects after N (number) of days. You also need to specify the Custom Bucket Name to use in the Filter section of the lifecycle management tool.

Note: You cannot change the Bucket Name after data is written to it because DIVA is looking for the specific Bucket Name. Changing the Bucket Name is equivalent to changing the disk mount point on a local disk; DIVA will not find it.

Storage Class

This parameter specifies the target storage tier associated with the buckets created by DIVA. Azure supports multiple storage tiers as follows: HOT, COOL and ARCHIVE as described on the following pages:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types#premium-ssd>

<https://docs.microsoft.com/en-ca/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal>

When the Storage Class is set to STANDARD, the actual storage class used is the Default Storage Tier configured on the Azure portal. If Standard Storage Class is specified at the DIVA array level, then no storage tier is specified with archive (put) request. The Default Blob Access Tier applies to objects created with this account.

Rehydration from Archive Tiers is always specifically set to HOT by the Actor. This is hardcoded and cannot be changed. After it is rehydrated and successfully restored, it is deleted from the rehydration area. While an object is being rehydrated it is temporarily put into a .diva folder.

Default Blob Access Tier also applies to the Metadata Bucket that DIVA Core creates. DIVA Core does not specify HOT or COOL when creating the Metadata bucket.

Note: If you do not specify a name, your metadata bucket may get moved to the Archive Tier.

Best Practices

The following are recommended best practices when using Azure Blob accounts:

- Use a Custom Container Name (-bucket_name= in Array Options).
- Set the Default Tier to HOT.
- Use Azure Lifecycle Policies on Data and Metadata containers if you want to change tiers at some point.

This will be cheaper and easier than using SPM. You must use filters to separate Data and Metadata containers.

- Data Container - move to the Archive Tier after N number of days (this number is up to you)
- Metadata Container - keep on HOT tier for three days, then move to COOL after three days. This is in case you archive something by accident and then delete it within a couple days. Deleting early from COOL will have an extra charge. Metadata files must be on COOL or HOT, they cannot be in ARCHIVE.

Configuring Google Cloud Storage Accounts

DIVA Core must be associated with a service account and a specific role/permission. This account is created from the console (<https://console.cloud.google.com>) under IAM & Admin > Service Accounts.

To be able to Archive, the role needs the following authorizations:

- storage.virtualobjects.create
- storage.virtualobjects.delete
- storage.virtualobjects.get
- storage.virtualobjects.list
- storage.virtualobjects.update
- storage.multipartUploads.abort
- storage.multipartUploads.create
- storage.multipartUploads.listParts

The following permissions are required to only be able to Restore objects:

- storage.virtualobjects.get
- storage.virtualobjects.list

When creating the required Service Account, you will generate and save a JSON. The JSON key contains information that is required to configure a GCS account in the DIVA configuration. For example:

```
{
  "type": "service_account",
  "project_id": "gcs-integration-to-diva",
  "private_key_id": "513bd217d88b38f8f6a7edcalfadd79bb62826e0",
  "private_key": "-----BEGIN PRIVATE KEY-----\nMIIEvgIBAD.....hhpRt5\n-----END PRIVATE KEY-----\n",
  "client_email": "donald-s-dev-actor@gcs-integration-to-diva.iam.gserviceaccount.com",
  "client_id": "117535311288780186075",
  "auth_uri": "https://accounts.google.com/o/oauth2/auth",
  "token_uri": "https://oauth2.googleapis.com/token",
  "auth_provider_x509_cert_url": "https://www.googleapis.com/oauth2/v1/certs",
  "client_x509_cert_url":
  "https://www.googleapis.com/robot/v1/metadata/x509/gregory-s-dev-actor%40gcs-integration-to-diva.iam.gserviceaccount.com"
}
```

The following are Google Cloud Storage Account parameters:

URL

The object storage URL is normally set to <https://www.googleapis.com>.

Client Email

This is set to the client email address found in the JSON key.

Region

This is the region where buckets and objects are stored. Available locations are listed at <https://cloud.google.com/storage/docs/locations>.

Private Key

This private key can be found in the JSON key of the service account. The private key must have new lines instead of \n characters.

Private Key ID

The Private Key ID can be found in the JSON key.

Project ID

This is the Google Cloud Project ID and the JSON key.

Storage Class

This parameter specifies the target Storage Tier associated with the buckets created by DIVA. Google Cloud Storage supports multiple Storage Tiers as described on the page at <https://cloud.google.com/storage/docs/storage-classes>.

Bucket Name

This is the optional target Bucket Name. Bucket limitations are listed at <https://cloud.google.com/storage/docs/naming>.

Configuring Oracle Cloud Infrastructure Accounts

DIVA Core supports Archives and Restores to the OCI (Oracle Cloud Infrastructure) object storage. The authorization access between OCI and its clients (like DIVA) is based on a Public/Private Key pair.

The information required to generate and retrieve all the connection parameters is described at <https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/apisigningkey.htm>.

The following are Oracle Cloud Infrastructure Account parameters:

Login

This is set to the User ID of the user to which the public key was associated. For example, ocid1.user.oc1..aaaaaaaap4x4no7ona25j5tn4dk635a6rq3-jun2mwj3eej65c565t3owxsoq.

Tenant ID

This is set to the Tenant ID of the OCI account. This information can be found on the OCI console. For example, ocid1.tenancy.oc1..a152aaaaykq43t6e22bati7hb-wrddb1tmhwhft54m3hrxiskgqtn3sgsula.

Private Key

Set this parameter to the RSA Private Key (PEM format).

Key Password

Set this parameter to the passphrase of the Private Key.

Fingerprint

This field must be set to the fingerprints of the Private Key.

Namespace

This is the namespace used by the OCI tenant. You can find this information on the OCI Console.

Compartment ID

This parameter is the ID of the compartment to be used by DIVA when archiving content. You can find this information on the OCI Console. For example, `ocid1.compartment.oc1..aaaaaaaeeobvubqa73c2kap5hed5vgfyjia7lejkccegj4pd-cfp5uetr3u4q`.

Storage Class

OCI supports both Archive and Standard storage classes.

Configuring EMC-ECS SWIFT Accounts

DIVA Core supports local object storage arrays that include disks with Swift interfaces. For example, an EMC ECS Object Store.

The following are EMC-ECS SWIFT Account parameters:

Type

This parameter should be set to *Local* because instances stored on EMC Elastic Cloud Storage are local instances whose priority is lower than other types of local disk instances, but a higher priority than tape storage instances.

In DIVA Core 8.0 and later you can define Oracle Storage Class and Storage Location separately. If you require new cloud or local arrays in the future, you can specify all of these parameters as options.

In DIVA Core 8.0 and later, both SWIFT and S3 are supported for interfacing with EMC ECS. However, you cannot change the existing configuration after the Array is configured.

You can set the Media Priority of a source instance for a Restore, Partial File Restore, and Copy to Tape Group requests. This enables restoring an instance stored on a local non-EMC ECS array with a higher priority than an instance on an EMC ECS array. The Manager decides which source instance is preferred during these requests if the media priorities are all the same.

Login

This must be set to the Auth V1.0 Swift username.

Password

This is the Swift password.

Service Name

The Service Name is typically empty with EMC-ECS.

Identity Domain

Set this field to *swift_namespace* for EMC ECS.

Cloud Replicated Bucket Scanning

The purpose of this feature is to keep scanning a cloud bucket containing AXF instances (AXF and AXF_RF). This bucket is populated by a third party software:

- Can be a bucket replication software
- Can be another DIVA Core system

In both cases, DIVA Core will scan the bucket for new objects and populate its database when new objects are detected.

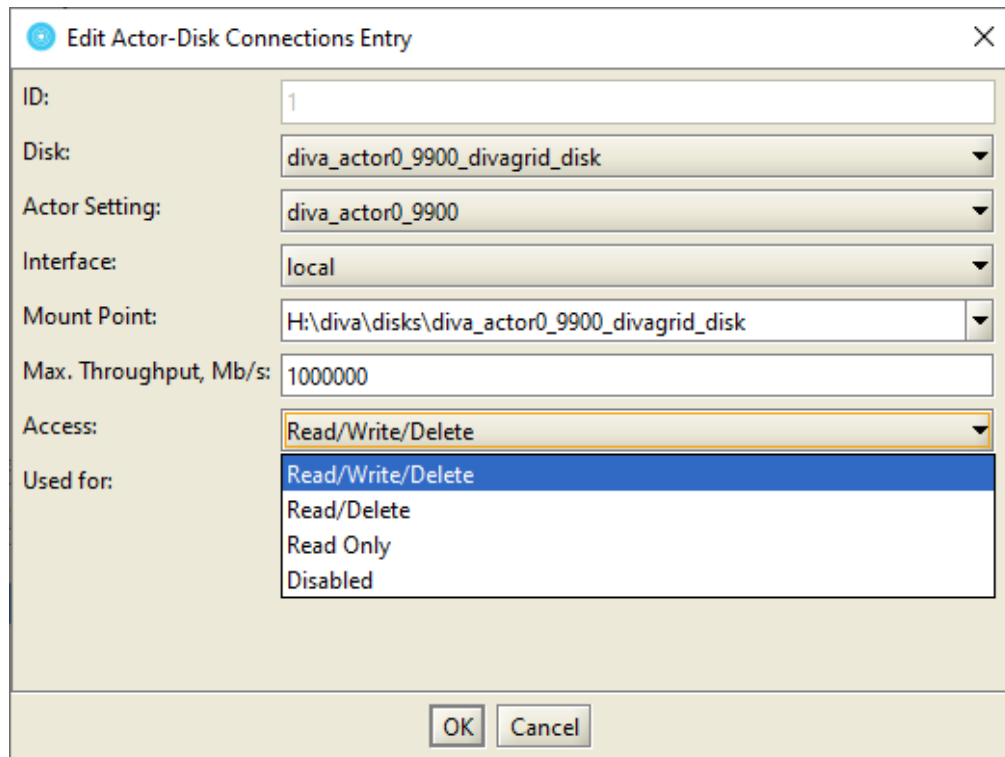
Automatic Scan Restart Configuration

At the start of a scan, the Core Manager will set all Actor-Disk connections of the disk being scanned to Read-Only access to prevent writes to that disk. If any request attempts to write to the disk, the request attempting the write fails and displays the "Warning: One or more resources not available for source to disk transfer: No usable disk." warning message.

Note: The reverse is not true. When a scan is terminated, the associated Actor-Disk connections access does not revert to its original setting. The user will need to manually reset the access of each Actor-Disk connection if needed.

New Access Type for Actor-Disk Connections

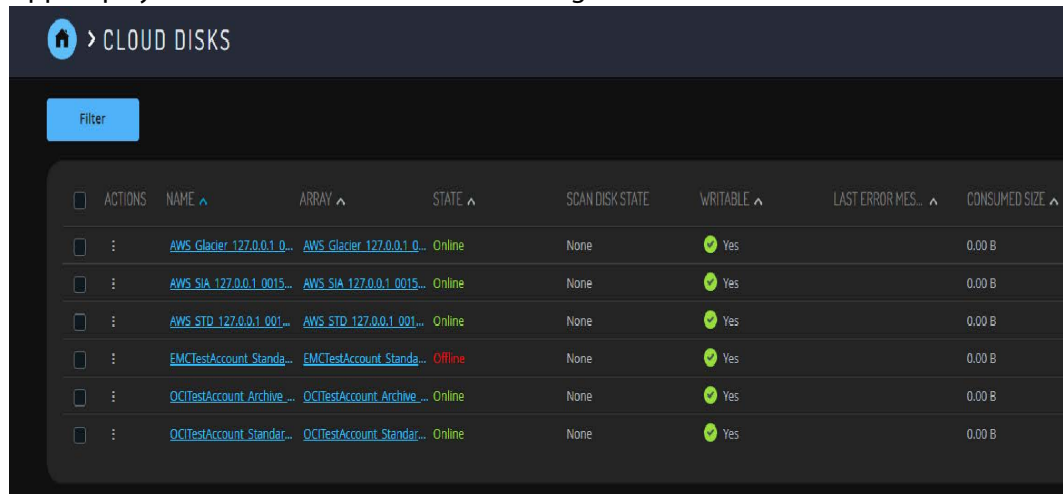
DIVA Core 8.3 supports true READ-ONLY connections. Previously, READ-ONLY access allowed deletes. DIVA Core now supports a new and separate access type called READ/DELETE to allow both operations. Additionally, READ/WRITE has been relabeled as READ/WRITE/DELETE because it supports all three options.



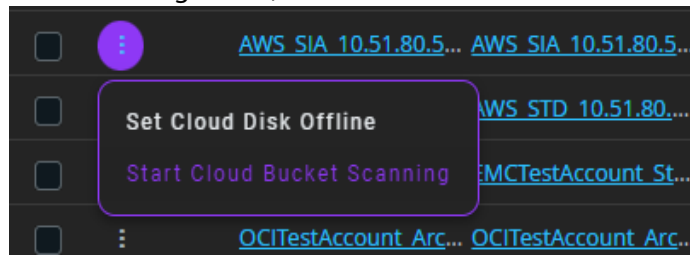
Note: If a disk is currently being scanned, changing the access type of all the associated connections to Read/Write/Delete, will still not allow writes to that disk. DIVA Core will only allow writes after the scan is stopped if the access type for that disk is configured to permit it.

System Management App Cloud Bucket Scanning Support

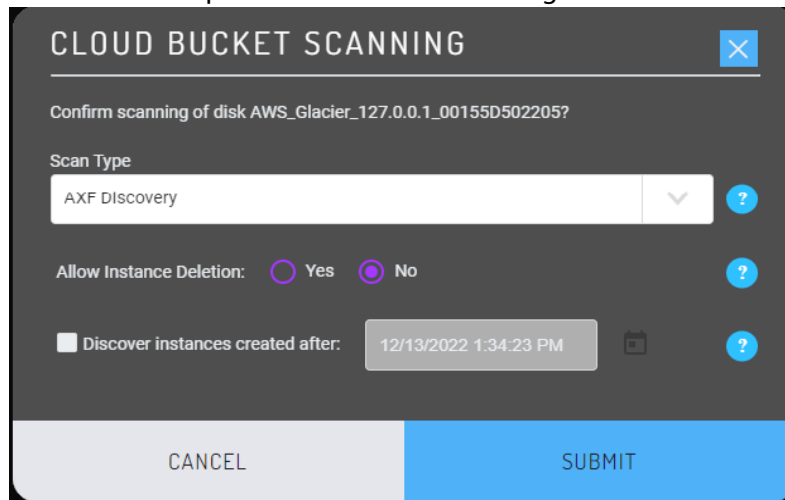
The Cloud Disks page in Resource Management for the DIVA Core System Management App displays the cloud disks and their scanning status in the main table.



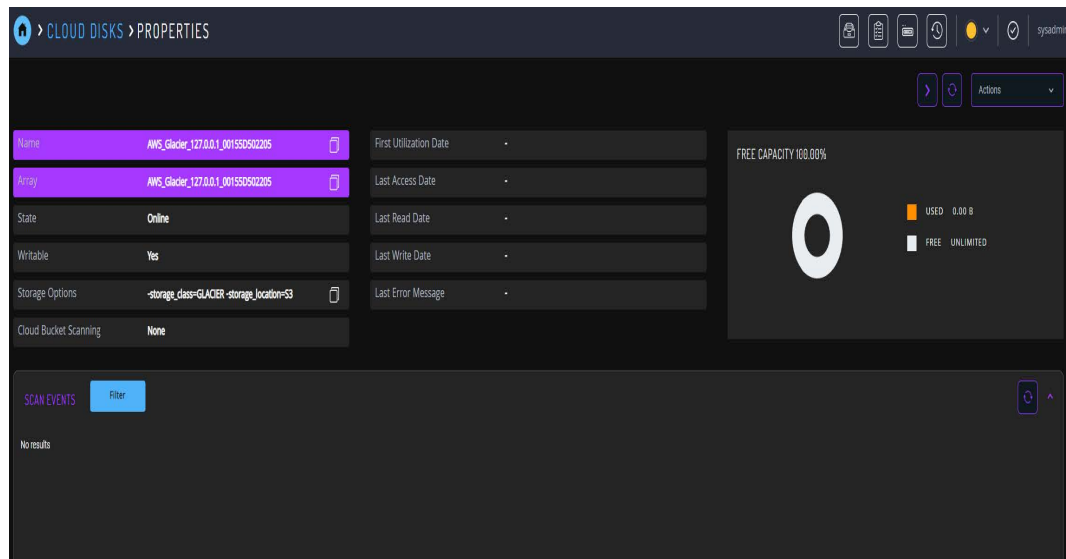
Start Cloud Bucket Scanning from the Action menu for an individual disk. When a Cloud Disk is running a scan, the Action menu also allows the user to stop the scan.



The user will be presented with the following form to enter the scan settings:



The Cloud Disk Properties page displays the scan events associated with the disk in a table at the bottom of the page. Events can be filtered by Severity, Type, Actor, Object, or Description. It is also possible to control the scan state from the property page in the Actions menu.



Object Auto-Indexing

This feature detects files from a cloud bucket or a bucket subdirectory, creates an AXF reference-file metadata file, and indexes them as new objects in the Core Database. The data files are not transferred by DIVA, they are simply indexed as they are on the managed cloud bucket.

The auto-indexing service works like the bucket scanning. It is not a separated component or a Windows service, it is a DIVA Core function.

Currently Supported Disks and Functionality

The following disks and functions are currently supported:

- Managed disks using S3, Azure and Google Cloud are supported. Others types of disks are not yet supported (for example, local, NAS, and so on).
- The service can index a file from any type of storage class.
- DIVA creates an AXF_RF_1.1 reference-file metadata for each file detected.
- One file per object in this release.
- Empty files are excluded.
- Overwriting existing file is prohibited.
- Direct file deletions are not supported; deletes must be sent to a DIVA API.
- A trial mode is available to be able to test the behavior of auto-indexing.
- Azure blob tags are only supported by v2 storage accounts for general purposes; object auto-indexing does not work with v1 storage accounts.

Duplicate Objects

If an object already exists in DIVA with a different UUID than is discovered during the auto-indexing of a cloud disk, DIVA will skip the object and post an ERROR event.


The event can be retrieved from the API in the normal way using `GET /disks/scans/events`.

Duplicated objects are not common but it could happen. This would leave the file associated with an AXF metadata file that is not referenced in the database; kind of an orphaned AXF instance.

Here is how to clean this up:

From the data bucket, the file that was supposedly indexed will have a tag containing the UUID of the object that auto-indexing tried to create.

For example:

Tags (2)
Track storage cost of other criteria by tagging your objects. [Learn more](#) 

Key	Value
diva_status	persisted
diva_uuid	82aa75b1-fb55-ed11-b6bf-1b65982d7a18

1. Navigate to the AXF metadata location and delete the associated `<uuid>.axf` file. (for example, `82aa75b1-fb55-ed11-b6bf-1b65982d7a18.axf`).
2. Upload the same file under a different name so that auto-indexing will generate a different objectname or different collection.
3. Delete the file that the auto-indexing failed on.

Trial Mode

Trial Mode can be specified when starting a scan. This mode of operation will not save objects to the database and will not update the objects in the cloud.

What is different when running in Trial Mode:

- Because there is no persistence in Trial Mode, duplicated objects (same name, same collection) are not detected.
- Auto-indexing attempts to detect files created or modified since a specific date and time. The date keeps updating and auto-indexing will attempt to detect new files created or modified 12 hours before this date and time. The overlap between each listing iteration makes it possible to see the same object multiple times in Trial Mode.

Configuration

The auto-indexing offers some options when the service is started as follows:

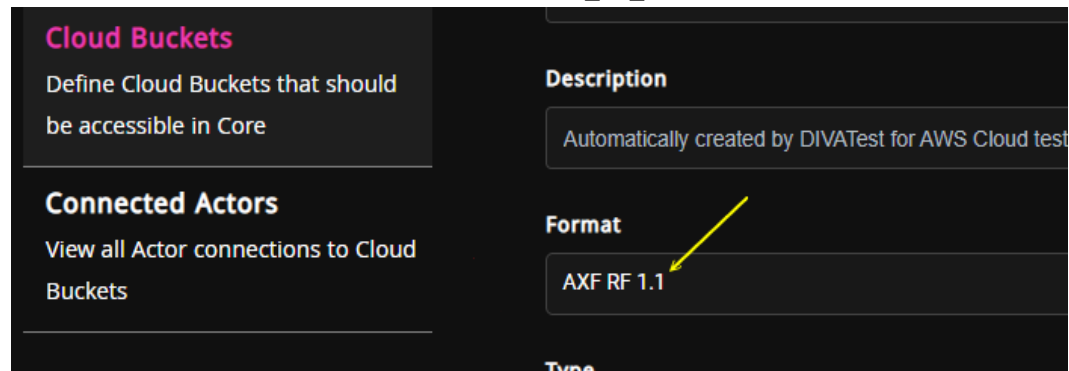
Option	Description
Index files created after [date/time]	The service will only consider the files created after the specified data/time and ignore the other ones.
Default Category	Depending on the value of Object Path Prefix, the service will use a default collection for all the objects. If a default collection is needed, this parameter must be set, otherwise nothing can be indexed.
SubPathFilter	<p>An optional filter can be specified if the auto-indexer should only detect the file path starting from SubPathFilter inside of the specified Bucket Name (see cloud disk configuration below).</p> <p>Example with these parameters set:</p> <ul style="list-style-type: none"> • Whatever the configuration of Object Path Prefix • Bucket name is set to <code>MyBucket/MyVideoFiles</code> • SubPathFilter is set to <code>ABC</code> <p>Auto-indexing will only consider the cloud objects starting from <code>MyVideoFiles/ABC</code> inside <code>MyBucket</code>.</p>
Trial Mode	<p>Trial Mode can be specified when starting a scan. This mode of operation will not save objects to the database and will not update the objects in the cloud.</p> <p>This mode of operation is useful to test some settings and anticipate the behavior of auto-indexing. When AXF metadata files and database entries are created it is difficult to revert the changes.</p>

Cloud Array Parameters used by Object Auto-Indexing

The behavior of the service also depends on the configuration of the cloud disk associated with the service.

- Archive Format

The auto-indexing feature will generate AXF_RF_1.1 instances. It only works when the format of the cloud bucket is set to AXF_RF_1.1.




- Bucket Name

DIVA will scan and index files under `<BucketName>[/Directory]`. The bucket name may contain an optional subdirectory; if so it will only scan the contents of this subdirectory.

- Object Path Prefix

The behavior of the auto-indexer depends on the selected prefix.







The following table describes how auto-indexing will set Objectname, collection, and uuid depending on the object path prefix selected:

Object Path Prefix	Objectname	Collection (Category)	UUID	Comment
None	Name of the file without the extension.	Default collection that is set in the configuration of the service.	Generated by the Actor.	
Name&Category Auto-indexing will only process the files starting from two subfolders after BucketName[/directory].	subfolder2	subfolder1	Generated by the Actor.	<p>ObjectName and Category are extracted from the path using a regular expression containing three patterns in parenthesis: $([^\wedge/]) / ([^\wedge/]) / ? (.*)$</p> <ul style="list-style-type: none"> • The first is the collection. • The second is the Objectname. • The third is the relative path of the file within the DIVA object. <p>The file paths not matching the regular expression are ignored.</p>
UUID Auto-indexing will only process the files starting from a subfolder matching a UUID.	Name of the file without the extension.	Default collection that is set in the configuration of the service.	subfolder1	<p>The UUID is extracted from the path using a regular expression: $([0-9a-f]+-[0-9a-f]+-[0-9a-f]+-[0-9a-f]+-[0-9a-f]+) / ? (.*)$</p> <p>The file paths not matching the regular expression are ignored.</p>

Preventing Auto-Indexing and AXF Archiving on the same Array

The auto-indexing service can generate AXF_RF objects and quickly populate the database with new objects. Starting the service on a cloud bucket that has already been used by DIVA to store AXF instances could be disruptive. For example, if auto-indexing is started with a cloud array that already contains AXF_RF objects, the auto-indexing may create additional AXF_RF objects pointing to the same content. If one of the objects or instances is deleted, the content of the second object/instance will also be deleted. The same issue could happen if you archive or copy to a cloud array that is already auto-indexing.

To prevent that situation from happening, DIVA will create a small file on the cloud array describing the utilization of the cloud array (AXF-ARCHIVING, AXF_RF-ARCHIVING, AUTO-INDEXTING). The name of this file is the name of the array with a .usage extension. For example:

<input type="checkbox"/>	 f0d91dc5-fb55-ed11-b6bf-1b65982d7a18.axf	axf
<input type="checkbox"/>	 fb1688bc-fb55-ed11-b6bf-1b65982d7a18.axf	axf
<input type="checkbox"/>	 GREG-BUCKET-BadInstances.json	json
<input type="checkbox"/>	 GREG-BUCKET.usage	usage

If DIVA is trying to use the same array for different purposes, the operation will fail.

If someone starts the auto-indexing service but the cloud array has already been use for archives or copies, actor will terminate auto-indexing with the following error in the event log:

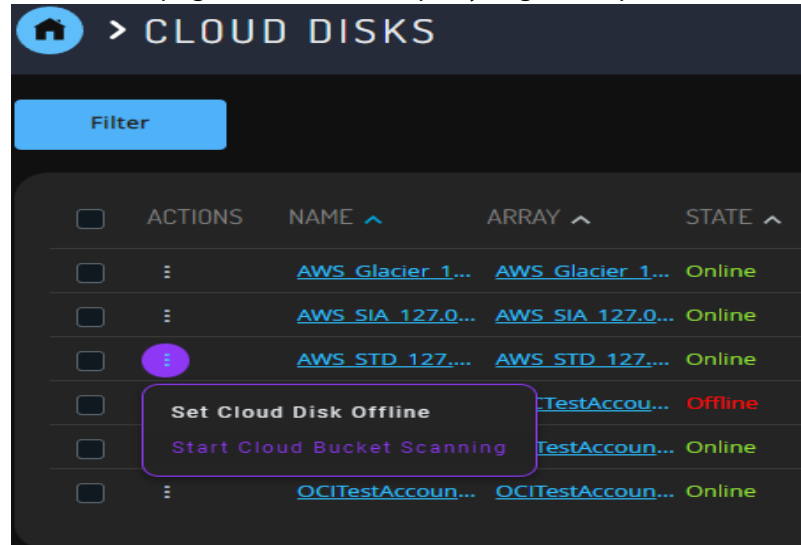
```
Cannot use this disk location for AUTO-INDEXTING purpose because it's already in use for AXF_RF-ARCHIVING.
```

If an attempt is made to archive or copy to a cloud array that has already been used for auto-indexing, the request will abort with the following error message:

```
Cannot use this disk location for AXF_RF-ARCHIVING purpose because it's already in use for AUTO-INDEXTING.
```

System Management App Object Auto-Indexing Support

Object Auto-Indexing support has been combined with Cloud Bucket Scanning support that was introduced in DIVA Core 8.2. Both types of scan can be triggered in the same manner as before by using the context menu for the relevant disk from the Cloud Disks page, or from the Property Page of a specific disk.



When Start Cloud Bucket Scanning is selected, a form is presented to the user. This is the same form for starting a cloud bucket scan as before, but a Scan Type selector has been added to choose the type of scan to perform. AXF Discovery is the type of scan that was supported in DIVA 8.2.

CLOUD BUCKET SCANNING

Confirm scanning of disk AWS_STD_127.0.0.1_00155D502205?

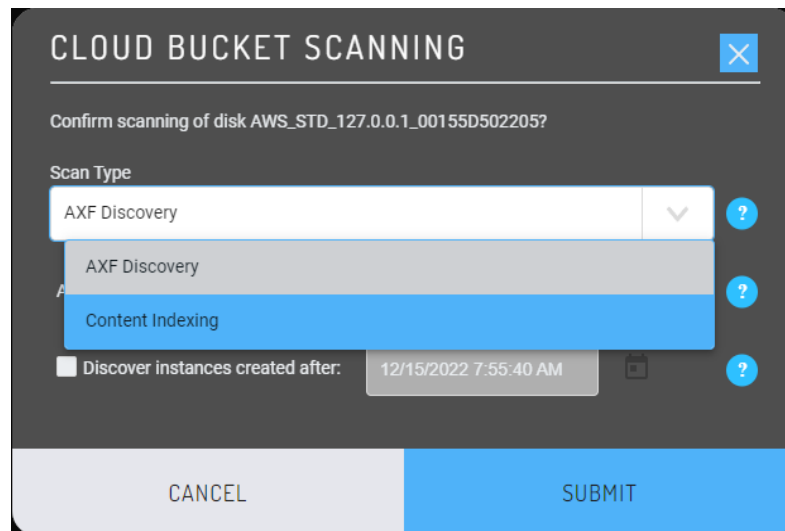
Scan Type: AXF Discovery

Allow Instance Deletion: Yes No

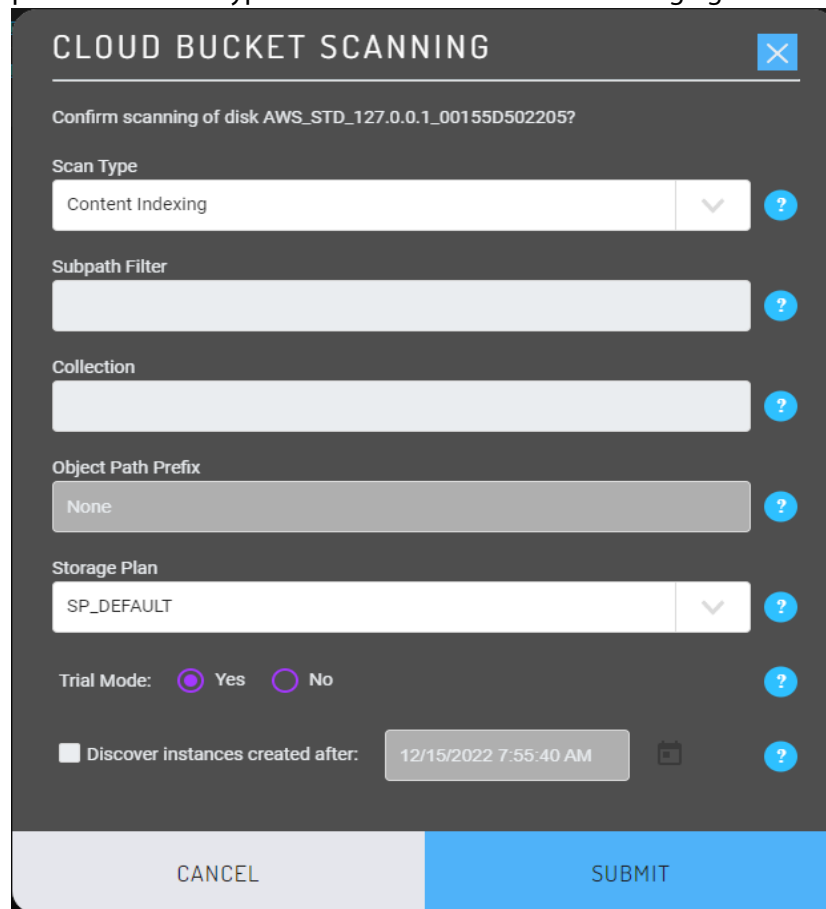
Discover instances created after: 12/15/2022 7:55:40 AM

CANCEL SUBMIT

Content Indexing is the name for the Object Auto-Indexing scan. This option is only available if the Cloud Bucket format for the disk is set to AXF RF 1.1.



Selecting Content Indexing as the scan type displays the dialog for the options pertinent to that type of scan as shown in the following figure:



Viewing the status of the indexing is the same as previously done, and scan events will appear in the Scan Events table on the Property page of the cloud disk.

The screenshot shows the 'PROPERTIES' page for a cloud disk. The top section displays various attributes: Name (AWS_STD_127.0.0.1_00155050...), Array (AWS_STD_127.0.0.1_00155050...), State (Online), Writable (Yes), Storage Options (-storage_class=STANDARD -stor...), and Cloud Bucket Scanni... (None). A 'FREE CAPACITY 100.00%' gauge is also visible. Below this is the 'SCAN EVENTS' table, which is filtered. The table has columns for ID, ACTOR NAME, TYPE, SEVERITY, EVENT TYPE, EVENT DATE, and DESCRIPTION. It contains three rows of scan events related to content indexing.

ID	ACTOR NAME	TYPE	SEVERITY	EVENT TYPE	EVENT DATE	DESCRIPTION
7	disk_actor1_9901	Content Indexing	Information	Start	11/21/2022 11:00:34 AM	[28453471 - Starting scan of type AXF_INDEXING on disk AWS_STD_10.51.80.54_001550502204 with actor disk_actor1_9901.
8	disk_actor1_9901	Content Indexing	Information	Stop	11/21/2022 11:00:37 AM	[28453471 - Stopping scan of type AXF_INDEXING on disk AWS_STD_10.51.80.54_001550502204 with actor disk_actor1_9901.
9	disk_actor1_9901	Content Indexing	Error	Service	11/21/2022 11:00:37 AM	[28453471 - Stopping service as error encountered in scan for disk AWS_STD_10.51.80.54_001550502204 and actor disk_actor1_9901: Cannot use this disk location for AUTO-INDEXING.
10	disk_actor1_9901	Content Indexing	Information	Resource Selection	11/21/2022 11:00:37 AM	[28453471 - Releasing scan service for disk AWS_STD_10.51.80.54_001550502204 and actor disk_actor1_9901.

It is not possible to run both types of scan against a bucket. The default view of the properties table is to show all scan events. However, the Scan Events table can be filtered by scan type.

The screenshot shows a 'Filter' dialog box for the Scan Events table. It contains several filter criteria: Severity (set to '*'), Event Type (set to '*'), Type (set to '*'), Actor Name (set to '*'), and Collection (set to '*'). The 'Type' dropdown menu is expanded, showing options for 'AXF Discovery' and 'Content Indexing'. At the bottom of the dialog are 'Reset Filters' and 'Apply Filters' buttons.

Disk Storage Configuration

DIVA Core's Disk Management defines each physical disk, how it is attached (or mounted) to the system, and then groups the disks together to perform specific roles within the archive. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

The Disks section defines the physical disks that are to be used by DIVA Core, how they are grouped together for either permanent or cache storage, and how each disk is logically accessed by the Actors.

The first step to disk management is to define an array. In DIVA Core an array is a logical grouping of one or more disks for the storage of DIVA Core Objects. You define arrays in the Arrays area of the Disk Storage page in the System Management App.

1. Click the DIVA Configuration > Disk Storage Configuration menu item on the left of the System Management App screen. Expand all of the sections using the arrow icons so you can see all of the configuration areas.
2. Click the + icon in the Array Settings area of the screen to display a new tab where you will configure the new Disk Array. Each parameter is filled in using default values. Enter the following information into the Array fields:

Name

Enter a name for the array in the Array Name text box. This is symbolic and typically represents the purpose of the stored objects.

Format

Select the Array Format from the pull down menu. The AXF format is required for Complex Objects. The following list identifies the available formats:

LTFS_AXF_1.1

This new format offers the same features as AXF_1.1 on an LTFS formatted tape and is only supported on tape. This format is not recommended for complex objects because it would generate very large LTFS indexes.

The tape block size must be set to 524288 Bytes or greater; LTFS does not support lower block sizes.

If an LTFS tape is loaded by LTFS software on a standalone drive, the contents of that tape can be accessed using Windows Explorer.

WARNING: Accessing LTFS tape content in Windows Explorer should only be used for recovery purposes, and the LTFS software must not be running concurrently with DIVA on the same drives.

AXF_RF_1.1

This format uses the AXF 1.1 structures, but AXF files will not contain any overhead.

AXF_1.1

This format is compliant with AXF 1.1 standards. Technical Support does not rec-

ommend using the AXF_RF_1.1 format with Complex Objects.

AXF_1.0

This format is compliant with AXF 1.0 standards.

AXF

This is redirected to AXF_1.1.

LEGACY

This is the formal archive format used by DIVA Core (index.txt, 00000001, 00000002, and so on).

Max Allowed Disk Percent For Repack

Enter the percentage of disk space available for use by repack requests.

Max Allowed Disk Percent For Migrate

Enter the percentage of disk space available for use by migration requests.

Verify Write

Select whether to enable Verify Write from the list. Verify Write is not compatible with Complex Objects.

Default Checksum Type

The MD5 algorithm is the Default Checksum Type. This field is not editable in this view.

Priority

To be completed.

Description

Enter a description for the array in the Description field. This is arbitrary and typically denotes the function of the array.

UseRandomSMBAddress

Used if Actor-Disk MountPoint contains a CIFS hostname (for example, \\host-name\share\folder). When enabled the Actor will resolve the list of IP addresses associated with the hostname and choose one of them randomly. If this parameter is disabled, the Actor will let the operating system choose the IP address because the connection is re-used between multiple requests. Therefore, only one IP is used until the connection is re-established if this parameter is disabled.

3. Click Save to record the array configuration.
4. Click the + icon in the Disk Settings area of the screen to display a new tab where you will configure the disks to be included in this array. Next you define the physical disks that are going to be used by DIVA Core, and assign them to arrays based on their intended function. You configure disks by selecting the array containing them. Each configured disk represents a distinct physical volume. Logical associations of disks to DIVA Core are performed in the Connected Actors section.

Note: If you intend to share a physical disk between two or more arrays, you can declare the same disk multiple times, but each declaration must have a unique name.

5. When you click the Disk Settings + icon a section for Connected Actors is also displayed.
6. Each parameter is filled in using default values. Enter the following information in the fields in the Disk area:

Disk Name

Enter a symbolic name for the disk in the Disk Name field to describe its function or its location.

Site

Select the Site that defines the location of this disk. This parameter is taken into consideration by DIVA Core for optimum allocation of disk resources in the array if the Site Selection parameter is enabled in the Core Manager configuration file.

Status

Set the current status of the disk using the pull down menu (ONLINE or OFFLINE). OFFLINE indicates that the disk is offline and not to be used. During DIVA Core operations, the status may be set Offline by DIVA Core if an unexpected disk I/O error occurs.

Min Free Space MB

You set the minimum free space of the disk in this field. DIVA Core considers the disk full when the remaining free space reaches this amount. Use this setting on disks that are shared with other applications, or with file systems that suffer poor or degraded performance when approaching 100% capacity.

Verify Write

Select whether to enable Verify Write from the pull down menu. This parameter overrides the same parameter set at the array level. Verify Write is not compatible with Complex Objects.

Default Checksum Type

The MD5 algorithm is the Default Checksum Type. This field is not editable in this dialog box.

7. Click Save to save the new connection and all Disk Storage parameters for this array.

Defining Actor to Disk Connections

After you have configured the Actor definitions, you must define the logical connections (mount points) of the physical disks previously identified during the initial DIVA Core configuration. This configures how each disk is logically connected to each Actor, and how it is to be used. For shared disks accessible by more than one Actor, the disk connection must be declared for all Actors.

If the same resource on a physical disk is to be shared between multiple Actors, and file sharing software has been installed, Technical Support recommends that the drive letter or volume of the disk connection configured in each Actor host is identical (for simplicity). Actors retrieve these mount point definitions when the Core Manager first connects to each of them. Any modifications performed here require the relevant Actor to be restarted.

To edit the parameters, double-click the Actor Name in the Actor-Disk Connections area to open the Add new row in Actor-Disk Connections dialog box. Click the + button on the top of the area to add an Actor-Disk connection. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Note: Multiple selections are available in Add mode, but not in Edit mode. Nearline storage is used for disk instances created during a Restore or N-Restore request with a Nearline QOS when no other disk instances are available.

The following list describes the options on the Add new row in Actor-Disk Connections dialog box:

Disk

Select a physical disk in the drop-down menu for this Actor association. Only entries previously defined in the Disks area will be displayed. Multiple disks may be selected using the check box next to each disk name.

Actor

Select the Actor for this disk association. Only Actors declared in the Actors area of the System page will be listed. The selected disk must be directly accessible by this Actor.

Max. Throughput, MB/s

This allows bandwidth throttling of the transfers performed by the Actor. Typically used to load balance transfers with other Actors or non-DIVA Core applications.

When DIVA Core has multiple disks to choose from for object storage, this parameter is the first criteria for disk allocation (that is, the disk with the highest throughput will be used first). The second criterion is the percentage of used capacity of each disk considered.

Used for

This defines how the associated disk is to be used by this Actor as follows:

Cache Only

DIVA Core will only use this disk for caching operations.

Storage Only

DIVA Core will only use this disk for object storage.

Cache and Storage

DIVA Core will use this disk for both cache and object storage.

Storage and Nearline

DIVA Core will use this disk for both object and Nearline storage.

Cache and Storage and Nearline

DIVA Core will use the disk for cache, object, and Nearline storage.

Access

This defines this Actor's read/write access to the associated logical disk. This allows further granularity in load balancing with other Actors.

Mount Point

The Mount Point is used with the Interface selection.

Interface

Select the access method the Actor will use to connect to the disk.

Actor to Disk Interfaces and Mount Points

The disk interface method and the corresponding mount point in an Actor-Disk connection are determined by how the drive is logically connected and presented to the Actor host computer's operating system. The following sections describe different interface methods and mount point configurations. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Local Interface

This option specifies that unbuffered I/O will be used with the disk to maximize transfer performance. Disks that use this option can reside within the Actor host itself (for example, disks to be used for cache purposes), disks connected to the host through either SCSI or Fiber Channel HBAs (for example, in a SAN), or those specified with a UNC (Universal Naming Convention) mount point. Some network drives may actually suffer with this type of interface. In these instances use the Remote option instead.

Note: Windows-based Actors do not support network drives mapped to a Windows drive letter (this is a Microsoft security restriction). Networked disks in Windows must use the Remote option instead.

The Mount Point is the drive letter or volume of the drive as it appears to the host operating system, plus any additional directory path.

Remote Interface

This interface specifies that buffered I/O will be used with the disk, and allows access to disks hosted by another computer using CIFS protocol. This option must be used for networked disks with the Windows Actor Service.

The mount point for a CIFS connection is a UNC path. For example, `cifs://192.168.56.26\shared` or `cifs://user\domain:password@//192.168.56.26\shared`.

Appropriate permissions for any CIFS-based disk must be enabled for the Actor to access the network drive. Otherwise, the disk will be set Offline.

Linux-based Actors can use CIFS by automatically mounting them. Linux-based Actors can also automatically mount to a specified Vantage Transcoder Cache, or use a fixed user-created mount point to transfer content to and from SMB network shares.

Note: Linux-based Actors support UNC paths for CIFS Source and Destination Servers by automatically mounting/dismounting from the SMB share. However, you can define a local path to a mounted SMB share. UNC paths are supported for SMB Servers and Managed Disks if the UNC path is directly mounted on the Windows Actor.

BML Interface

This interface enables the Actor to use a SeaChange BML (non-Infiniband Media Managed Storage) as disk storage. For regular disks, DIVA Core stores objects under multiple subdirectories. The BML however uses a flat file system (that is, no directory structure). DIVA Core automatically incorporates a directory structure into the file name when it is archived to the BML, and removes this addition from the file name as it is restored.

FTP Interface

This interface enables DIVA Core to use FTP servers as disk storage using the FTP protocol. Only Linux-based FTP servers are supported when operating in a Linux environment; not Windows-based FileZilla and IIS FTP servers. This is because Windows FTP servers cannot handle the large numbers of files. The mount point must be in the format `ftp://login:password@host/rootdir`.

MetaSAN Interface

Select this interface when MetaSAN manages the disk volume. By default, Core Actors preallocate storage on disks to prevent disk fragmentation. MetaSAN implements its own anti-fragmentation mechanisms. Selecting this option will disable preallocation when dealing with this volume.

Simulation Interface

Use this interface when setting up a DIVA Core Simulator. See the DIVA Core Simulator Operations Guide for details. This book is only available to OPN partners.

The mount point must be a real path name to a directory on a local disk. When used to store objects, only the file size is recorded to the disk (no content is actually saved). You cannot use a simulated disk as cache for a repack request.

Object Storage Servers

DIVA Core can interface with various Object Storage Servers as a Source or Target Server. The supported types of object storage are as follows:

- Simple Storage Service: S3
- Azure Blob Storage: AZCS
- OpenStack: SWIFT
- Google Cloud Storage: GCS
- Oracle Cloud Infrastructure: OCI

Users must manually configure a Server resource linked to a storage account. The user must select the storage account to use as a Server. Defaults are pre-filled for fields such as the IP Address. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

The other parameters can be configured through object Storage specific connect options as noted in the following table:

Attribute	Value	Example
IP Address	Name of the cloud account	diva_account
Source Type	Type of Object Storage	S3, AZCS, SWIFT, GCS, OCI
Connect Options	-storage_class {Target Storage Class} -restore_tier {AWS S3 restore tier} -max_virtualobject_size -proxy {HTTP Proxy server} -threads_per_transfer {Number of threads} -part_size {Part size in MB}	-storage_class Standard -restore_tier Expedited -max_virtualobject_size -proxy http://proxy.domain.com:8080 -threads_per_transfer 10 -part_size 20

-storage_class

The option can be used to specify the target storage tier when DIVA is restoring content to an Object Storage Server. The default value is Standard. Depending on the type of object storage, some specific storage classes may be available as follows:

- S3: REDUCED_REDUNDANCY, STANDARD_IA, ONEZONE_IA, INTELLIGENT_TIERING, GLACIER, DEEP_ARCHIVE, OUTPOSTS
- AZCS: Hot, Cool, Archive
- OCI: Archive
- GCS: NEARLINE, COLDLINE, ARCHIVE
- SWIFT: Archive

-restore_tier

This option is specific to AWS S3. It can be used to specify the retrieval option when restoring an archived object in AWS S3 (GLACIER or DEEP ARCHIVE). The possible values are Expedited, Standard and Bulk. The default value is Standard.

-max_virtualobject_size

When restoring files to SWIFT, this option can be used to specify the maximum segment size when creating static large objects. The value is in megabytes and can be set between 5 MB and 4096 MB. The default: 512 MB.

-proxy

Use this option if the upcoming request to the HTTP server must go through a HTTP or HTTPS proxy.

-threads_per_transfer

Use this option to modify the network concurrency when uploading or downloading large files using multipart upload/download. The default value is 1.

-part_size

This is the size in MB of each part to be used when uploading or downloading large files using multipart upload/download. The value can be set between 5 MB and 200 MB. The default is 50 MB.

Media Storage Configuration

The Media Storage Configuration screen in the System Management App enables configuring and viewing the following DIVA Core system components: RobotManager, Media Compatibility, Tape/Drive Properties, Sets and Tape Groups. In the System Management App these options are located under either the Configuration menu or the Resources Management menu.

Two views are available and can be selected using the toggle button on the top right of the interface:

- A structural view that shows all the elements of media storage in a hierarchical view. For example, a RobotManager can control one or multiple ACSs. One ACS contains one or more LSMs.
- A list view that presents the same information using tables of items. This view looks like what was presented by the legacy System Management App.

Adding a Robot Manager

The purpose of adding a Robot Manager to the Media Storage setting is to declare it in the DIVA configuration. To be able to use it, a Robot Manager service connected to a library must be configured (see [Component Configuration](#)).

At the system level, each instance of the Core Robot Manager must be declared to the Core Manager in the Robot Managers area of the Robots page in the System Management App.

Use the buttons to add (+), edit (Edit), or delete (-) a Robot Manager. The Update button refreshes the displayed Robot Manager information from the database.

Clicking the + button adds a Robot Manager to the configuration. The Add New Row in Robot Managers dialog box is displayed. Enter the following information in the appropriate fields and then click OK to add a Robot Manager:

Name

The name of the Core Robot Manager attached to this DIVA Core system.

Address

The IP address of the host running the Core Robot Manager installation.

Port

The Robot Manager TCP port. This must match the RM_PORT parameter specified in robotmanager.conf.

Site

The Core Manager uses this parameter to determine optimal use of resources in resource allocation. Use the menu to select the appropriate site for this Robot Manager. Site Selection must be enabled in the Core Manager configuration file or all sites are considered equally.

Database Configuration Synchronization

To facilitate the configuration, DIVA can retrieve most of the information from the library through a RobotManager: RobotManager-ACS association List of Drives, List of Media, Drive Types, Media Types, Media-Drive Compatibility. A RobotManager service must be configured and connected to a library before you can use it to synchronize the database.

Use the SynchronizeDb button to synchronize all the information from a Robot Manager. The database can also be synchronized by using the (+) button for each Collection.

Robot Manager-ACS Association

Each Core Robot Manager is logically referred to by the Core Manager using its ACS (Automatic Cartridge System) number. This value should be unique among all Core Robot Managers. Individual Managed Storage (or frames) are typically referred to by their LSM (Library Storage Module) number.

Use the following procedure to associate a Robot Manager with an ACS:

1. Open System Management App.
2. Select the Synchronize DB option from the Tools menu and acknowledge the warning message.
3. Select the individual Robot Manager to synchronize from the menu in the Database Synchronization dialog box, or select ALL to synchronize all Robot Managers.
4. Only select the Synchronize Robot Manager ACS Associations check box. Confirm that all other check boxes are deselected.
5. Click Go to update the selected associations.
6. Confirm correct, successful, operation in the Status area at the bottom of the screen.
7. Enter the details for each library in the Library Data Entry dialog box when prompted, and then click OK to continue.
8. Click Close to exit the Database Synchronization dialog box.
9. Confirm the association in the Robot Managers-ACS area.

Defining Tape Capacity and Block Sizes

The values in the following two tables must be used when entering adding a Drive Type or Media Type in the Core Database. The values have been tuned by Technical Support to avoid tape spanning, and therefore may be lower than the theoretical capacity.

The following table identifies tape capacities to use when entering a Drive Type or Media Type in the database:

Media Type	Drive Type	Capacity
9840	STK 9840A	19 531 008
	STK 9840B	19 531 008
	STK 9840C	39 062 272
9940	STK T9940A	58 593 536
	STK T9940B	195 312 384
T10000T1	STK T10000A	488 281 008
	STK T10000B	976 562 176
T10000TS	STK T10000A	117 187 072
	STK T10000B	234 374 656
	STK T10000C	5 243 000 000
	STK T10000D	7 812 500 480
	STK T10000D (maximum capacity enabled)	8 300 781 056
DTF-2	GY-8240	195 312 448
SAIT1	S-AIT 1	488 281 088
SAIT2	S-AIT 2	781 249 536
AIT3	AIT 3	97 656 192
DLT-IV	Quantum DLT7000	34 179 648
LTO-100G	IBM, HP, Seagate LTO-1	97 656 192
LTO-200G	IBM LTO-2	195 312 128
LTO-400G LTO-400W	IBM or HP LTO-3	390 624 768
LTO-800G LTO-800W	IBM or HP LTO-4	781 249 536
LTO-1.5T LTO-1.5W	IBM or HP LTO-5	1 464 843 264
LTO-2.4T LTO-2.4W	IBM or HP LTO-6	2 441 405 952
LTO-6.4T LTO-6.4W	IBM LTO-7	5 859 374 592
LTO-9.0T	IBM LTO-8 (M8)	8 789 062 500

Media Type	Drive Type	Capacity
LTO-12.8T LTO-12.8W	IBM LTO-8	11 718 750 000
3592-JA 3592-JW	3592-J1A TS1120 TS1130	292 968 750 488 281 250 625 000 000
3592-JB 3592-JX	TS1120 TS1130 TS1140	683 593 750 976 562 500 1 562 500 000
3592-JK	TS1140 TS1150	488 281 250 878 906 250
3592-JC 3592-JY	TS1140 TS1150	3 906 250 000 6 835 937 500
3592-JL	TS1150 TS1155	1 953 125 000 2 929 687 500
3592-JD 3592-JZ	TS1150 TS1155	9 765 625 000 14 648 437 500

The following table identifies tape block sizes to use when entering a Drive Type or Media Type in the database:

Manufacturer	Tape Drive Type	Block Size in Bytes
HP	LTO Ultrium 1	65536
	LTO Ultrium 2	524288
	LTO Ultrium 3	524288
	LTO Ultrium 4	524288
IBM	LTO Ultrium 1	65536
	LTO Ultrium 2	524288
	LTO Ultrium 3	524288
	LTO Ultrium 4	524288
	LTO-5	524288
	LTO-6	524288
	LTO-7	524288
	LTO-8	524288
Oracle StorageTek	T9840A	262144
	T9840B	262144
	T9840C	262144
	T9940A	262144
	T9940B	262144
	T10000A	524288
	T10000B	524288
	T10000C	524288
	T10000D	524288
Quantum	DLT 7000	65536
Seagate	LTO Ultrium 1	65536
Sony	GY-8240 (DTF-2)	65536
	AIT-3	65536
	S-AIT 1	524288
	S-AIT 2	262144

DIVA Core General Settings

The following DIVA Core General Settings are configured using the System Management App under the Configuration menu item.

Checksums

Configure the following Checksum Support basic settings in the first section of the DIVA Core General Settings screen. When changes are complete, (at the top of the screen) click Save to save your changes, Cancel to cancel your changes, or Refresh to refresh the screen.

Manager: Checksum feature is enabled

Select the check box to enable Manager Checksum Support; deselect the check box to disable this function.

Manager: Default Checksum type

Click the Edit icon (looks like a pencil) on the right of this setting to enable the pull down list. Select the desired default checksum type from the list.

Manager: Number of retries following failed checksum

Click the Edit icon (looks like a pencil) and enter the desired number of retries to attempt after a failed checksum verification is processed.

Manager: Select different drive per retry on failed checksum

Select the check box to enable using a different drive after each failed checksum verification; deselect the check box to disable this function.

Media

Configure the following Media settings in the first section of the DIVA Core General Settings screen. When changes are complete, (at the top of the screen) click Save to save your changes, Cancel to cancel your changes, or Refresh to refresh the screen.

Manager: Frequency of Automated Clones (every X hours)

Click the Edit icon (looks like a pencil) and enter the desired number of hours between automated cloning process executions.

Manager: Maximum Simultaneous Automated Clones

Click the Edit icon (looks like a pencil) and enter the desired number of automated clones to execute simultaneously.

Objects

Configure the following object settings in the second section of the DIVA Core General Settings screen. When changes are complete, (at the top of the screen) click Save to save your changes, Cancel to cancel your changes, or Refresh to refresh the screen.

Manager: Allow Manager to Restore Objects when Complex Object Metadata is not available

Select the check box to enable allowing Manager to restore Complex Object that have no metadata available; deselect the check box to force restoring only objects that include metadata.

Manager: Enable UUID Preservation

Select the check box to enable preserving the UUID; deselect the check box to disable this function.

Auto-Discovery Publisher Endpoint

Click the Edit icon (looks like a pencil) and enter the URL for the Publisher endpoint.

DB: Maximum retry for locking a DIVA Core Object

Click the Edit icon (looks like a pencil) and enter the desired maximum number of retries to lock an object.

DB: Retry interval in seconds for locking a DIVA Core Object

Click the Edit icon (looks like a pencil) and enter the desired interval (in seconds) to retry locking an object.

UUID to uniquely identify a DIVA Core System

Click the Edit icon (looks like a pencil) and enter the desired UUID to identify a DIVA Core system.

Complex Objects Metadata Service URL

Click the Edit icon (looks like a pencil) and enter the Object Metadata Service URL.

SMTP Notifications

Configure the following SMTP Notification settings in the first section of the DIVA Core General Settings screen. When changes are complete, (at the top of the screen) click Save to save your changes, Cancel to cancel your changes, or Refresh to refresh the screen.

Enable email Notification

Select the check box to enable email notifications; deselect the check box to disable this function. When disabled the remaining settings do not need to be set.

Database Backup Notification

Click the Edit icon (looks like a pencil) and use the pull down menu to select the type of notifications to receive concerning database backups. The options are ERRORS and WARNINGS, ERRORS, or DISABLED.

Manager: Set the default DIVA Core Backup Service monitor timeout (Minutes)

Click the Edit icon (looks like a pencil) and enter the desired default number of minutes for the backup service timeout.

(SMTP) Outgoing Mail Host

Click the Edit icon (looks like a pencil) and enter the SMTP Mail Host domain.

(SMTP) Outgoing Mail Port

Click the Edit icon (looks like a pencil) and enter the SMTP outgoing mail port; typically port 25 (unsecure), 465 (secure) or 587 (secure).

(SMTP) Outgoing Mail Required Authentication

Select the check box if the SMTP server requires authentication; deselect the check box if authentication is not required.

Account name (Full email Address)

Click the Edit icon (looks like a pencil) and enter the full email address to use for sending email notifications.

Account Password

Click the Edit icon (looks like a pencil) and enter the email account password. Clicking the Eye icon toggles between displaying and hiding the password. Show the password allows the user to check the entry. However, for security purposes the password should always be hidden before leaving the screen.

DIVA Core System Administrator email Address

Click the Edit icon (looks like a pencil) and enter the DIVA Core System Administrator email address.

email Subject

Click the Edit icon (looks like a pencil) and enter the desired subject for the email.

Notification email Recipients (comma delimited)

Click the Edit icon (looks like a pencil) and enter all email addresses that will receive the notification emails. This is a comma separated list when entering more than one address.

Number of hours between e-mail notifications

Click the Edit icon (looks like a pencil) and enter the desired number of hours between sending email notifications.

Number of minutes before first e-mail notification

Click the Edit icon (looks like a pencil) and enter the desired number of minutes before sending the first email notification.

Determines whether to send an e-mail notification when an actor goes offline

Select the check box to send an email notification if a Actor goes offline; deselect the check box to disable this function.

Determines whether to send an e-mail notification when a drive goes offline

Select the check box to send an email notification if a drive goes offline; deselect the check box to disable this function.

Determines whether to send an e-mail notification when a disk goes offline

Select the check box to send an email notification if a disk goes offline; deselect the check box to disable this function.

Determines whether to send an e-mail notification when an actor-drive connection goes offline

Select the check box to send an email notification if an Actor-Drive connection goes offline; deselect the check box to disable this function.

Determines whether to send an e-mail notification when an actor-disk connection goes offline

Select the check box to send an email notification if an Actor-Disk connection goes offline; deselect the check box to disable this function.

Minimum disk space in MB at or below which an e-mail notification will be sent

Click the Edit icon (looks like a pencil) and enter the desired minimum disk space (in MB) to trigger an email notification.

Minimum empty tapes at or below which an e-mail notification will be sent

Click the Edit icon (looks like a pencil) and enter the desired minimum number of empty tapes to trigger an email notification.

Maximum number of aborted requests, at or above which an e-mail notification will be sent

Click the Edit icon (looks like a pencil) and enter the desired number of aborted requests to trigger an email notification.

Security

Configure the following Security settings in the first section of the DIVA Core General Settings screen. When changes are complete, (at the top of the screen) click Save to save your changes, Cancel to cancel your changes, or Refresh to refresh the screen.

Export: Tape Encryption Key Store Password

Click the Edit icon (looks like a pencil) and enter the Tape Encryption Key Store Password.

Export: Enable Export of Encryption Keys

Select the check box to enable exporting the encryption keys; deselect the check box to disable this function.

Licensing Configuration

The Licensing screen in the System Management App enables configuring and viewing the following DIVA Core system components:

- Active License Information
- License Histories

The Active License Information section only allows the user to view the current active DIVA Core licenses. See [Appendix A: Core Options and Licensing](#) for detailed information.

When the System Management App was installed the DIVA Core license should have been imported during installation. If the installer skipped importing the license during installation, or if you need to add a new license, you can use the following procedures to import a DIVA Core system license:

1. Click Import in the License Histories section.
2. A dialog box appears requesting the required information for importing a license.
3. Enter all of the required information and click Choose File to find the license file to import.
4. Select the check box if the Core Manager should be notified of the new license; otherwise leave the check box deselected.
5. Click Import to import the new license. The new license will appear in the License Histories section with a green check mark in the Import Status column.

Click the icon under the Actions column to export and download an existing license to the local machine in plain text format.

Synchronizing Media and Drive Compatibility with the Database

Media-Drive compatibilities are created during database synchronization

Synchronizing Media Types with the Database

You must import the values that have been uncommented in the Tape_Types configuration files into the Core Database. Each Core Robot Manager to be queried must be online to complete this procedure successfully.

Use the following procedure to import and synchronize the values from the Tape_Types files in the database:

Caution: Only perform this operation if you are adding Media Types to the library.

1. Open the System Management App and connect to the database.
2. Select the Synchronize DB option from the Tools menu and acknowledge the warning message.
3. Select the individual Robot Manager to synchronize from the menu list in the Database Synchronization dialog box, or select ALL to synchronize all Robot Managers.
4. Only select the Synchronize media types list check box. Confirm that all other check boxes are deselected.
5. Click Go to update the selected associations.
The System Management App will connect to the Core Robot Manager. The Robot Manager parses the SCSI_Tape_Types (or ACSLS_Tape_Types if used) configuration file.
6. If a Tape Type is not currently in the database, you will be prompted to enter it. Click No for any Tape Types not currently in use.

Note: If you report cleaning tapes in the following two steps, you must enter a Tape Size and Block Size of 1 KB for each cleaning tape added so they do not interfere with the total available size computation of all tapes in the System Management App.

7. Enter the Total Size for this Media Type and click OK.
8. Enter the Block Size for this Media Type. Ensure you enter the Block Size correctly before clicking OK because you cannot change it later.
9. Click Close to exit the Database Synchronization dialog box.
10. Confirm the Tape Type has been correctly entered in the Tape Properties area of the System Management App Tapes page.

Synchronizing Drive Types with the Database

You must also import the uncommented values in the Drive_Types configuration files into the Core Database. Each Core Robot Manager to be queried must be online to complete this procedure successfully.

Use the following procedure to import and synchronize the values from the Drive_Types files in the database:

Caution: Only perform this operation if you are adding Drive Types to the library.

1. Open the System Management App.
2. Select the Synchronize DB option from the Tools menu and acknowledge the warning message.
3. Select the individual Robot Manager to synchronize from the menu list in the Database Synchronization dialog box, or select ALL to synchronize all Robot Managers.
4. Only select the Synchronize drive types list check box. Confirm that all other check boxes are deselected.
5. Click Go to update the selected associations.

The System Management App will connect to the Core Robot Manager. The Robot Manager parses the SCSI_Drive_Types (or ACSLS_Drive_Types if used) configuration file.

If a Drive Type is not currently in the database, you will be prompted to enter it. Click No for any Drive Types not currently in use.

6. Enter the Block Size for this Drive Type. Ensure you enter the Block Size correctly before clicking OK because you cannot change it later.
7. Confirm there are no errors in the status window. If errors appear, recheck the Tape_Types and Drive_Types definition files.
8. Click Close to exit the Database Synchronization dialog box.
9. Confirm the Drive Type has been correctly entered in the Drive Properties area of the System Management App Drives page.

Synchronizing the Library Drive List with the Database

If you add Drive Types or additional drives to a DIVA Core Managed Library, you must declare them in the Core Database. Drives that are added are initially set Offline, and therefore disabled. Before they can be used, you must set them Online and notify the Manager (if running). During DIVA Core operations, the Manager may automatically set a drive Offline if it encounters a problem with it.

When the Used field in the Drives area is set to N, DIVA Core ignores the drive and it is not displayed in the System Management App Drives tab. If you subsequently set a

drive to Y, DIVA Core will not use it until you notify the Manager. This field restricts using drives in Managed Storage that are shared with other backup or archive applications.

The Operations field in the Drives area defines which operations are permitted on each drive. Operations can be one of the following:

R

The drive is dedicated to only Repack operations.

S

The drive will perform all standard operations only. That is, all operations except Repack.

A

The drive can perform all operations including Repack.

N

The drive will not be used for any operations. However, it can be enabled later without a Manager restart.

Use the following procedure to add the drives to the database:

1. Open the System Management App and connect to the database.
2. Select the Synchronize DB option from the Tools menu and acknowledge the warning message.
3. Select the individual Robot Manager to synchronize from the menu list in the Database Synchronization dialog box, or select ALL to synchronize all Robot Managers.
4. Only select the Synchronize drive list check box. Confirm that all other check boxes are deselected.
5. Click Go to update the selected associations.
The System Management App will connect to the Core Robot Manager. The Robot Manager obtains the current drive list and drive location for each drive from the library.
6. Confirm there are no errors in the status window. If the drives reported from the library do not match those declared in the Drive Properties area, an error is displayed and no drives are entered into the database.
7. Click Close to exit the Database Synchronization dialog box.
8. Confirm the drives have been correctly entered in the Drives area of the System Management App Drives page.

Manually Identifying Drive Serial Numbers

When using a tape library with DIVA Core there are two logical connections to each drive in that library. The first is the Robotics Control (managed by the Core Robot Manager) for mounting and dismounting the tapes from specific drives, and the Data Interface to the drive from the Actors.

Tape Managed Storage identify their drives by the Drive ID (typically 0, 1, 2, and so on). DIVA Core needs to know the corresponding data path to that drive from each Actor when the Robot Manager instructs the library to mount a tape to a specific Drive ID. If the Actor-Library mapping is incorrect, DIVA Core attempts to read or write to the incorrect drive, resulting in possible data loss or corruption.

The host computer operating system presents each drive to applications using their SCSI ID. The SCSI ID for a drive can vary as hardware is added or removed. This is particularly true when shared among multiple hosts in a SAN based environment. This configuration requires statically configured SCSI IDs using persistent bindings. This configuration dramatically complicates drive replacement.

To simplify configuration and streamline future drive replacements, the data path mapping to each drive (for its physical location in the library) is achieved by using its unique serial number rather than its SCSI ID. When a Core Actor is launched it interrogates each drive's serial number and compares it to the values in the database. Then the Actor establishes the correct data path to the drive, irrespective of its SCSI ID.

Each drive's serial number is automatically identified by library synchronization with the database during initial installation or drive replacement. Some cases may require you to manually determine the serial number and enter it into, or verify it against, the database.

You can manually identify the drive serial number either using the library's front panel display, or using the Scandrive Utility and the Robot Manager Client or GUI.

The latter method involves mounting a tape into a specific drive number in the tape library, establishing which drive the Actor is reporting that has that tape mounted, and then recording its serial number and entering, or verifying, it with the corresponding library Drive ID in the database. You must only complete this process one time for each drive in the library.

Caution: The Robot Manager Client GUI utility issues direct commands to the Robot Managers and will interfere with DIVA Core operations. It interacts directly with both the Robot Managers and the Tape Drives in the Managed Storage. You must not use it while the Core Manager is running.

You can use the Robot Manager Client GUI utility to send manual mount commands to a Core Robot Manager. See the [Robot Manager Client GUI](#) section for information.

The serial number of each drive can be discovered by using the scandrive.exe utility located in the %DIVA_HOME%\Program\Actor\bin folder. This utility automatically reports all SCSI devices installed in the host computer, and their corresponding port,

bus, target and logical unit numbers. For tape devices, the utility also indicates the drive's firmware, serial number, and whether a tape is loaded into the drive.

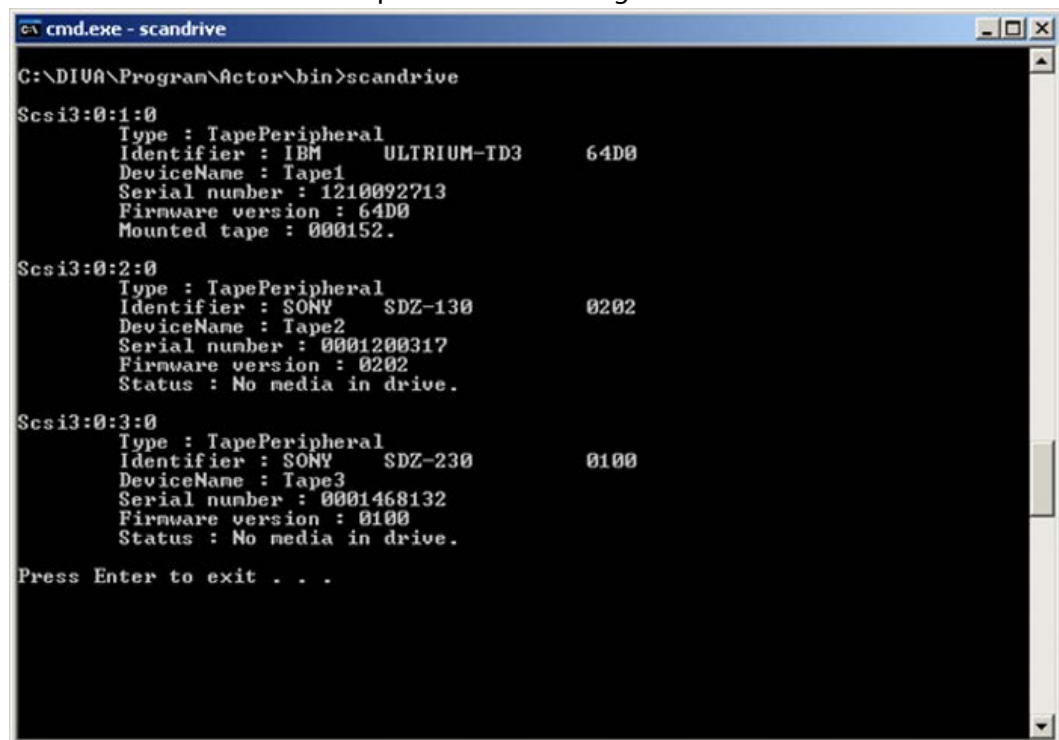
After a tape is mounted in a drive (using the Robot Manager Client GUI), run the scandrive.exe utility on an Actor host (that will use the selected drive) to determine which drive has the tape mounted and its corresponding serial number.

See [Determining the SCSI Library Connection](#) for information on using the Scandrive utility.

In the following figure the Type section refers to that peripheral's class (HDD, CDROM and so on). Each tape drive will be reported as a Tape Peripheral, and the Identifier for each corresponding device should match the model number of the drive itself (for example, IBM Ultrium TD2).

Confirm the tape barcode is the correct one loaded through the Robot Manager Client GUI. You must then enter the serial number for the appropriate drive by highlighting it in the Drives section of the System Management App, and then selecting Edit. Repeat the process by mounting a tape into the next library drive.

Remember to dismount the tape after determining the drive's serial number.



```
cmd.exe - scandrive
C:\DIVA\Program\Actor\bin>scandrive
Scsi3:0:1:0
  Type : TapePeripheral
  Identifier : IBM    ULTRIUM-TD3    64D0
  DeviceName : Tape1
  Serial number : 1210092713
  Firmware version : 64D0
  Mounted tape : 000152.

Scsi3:0:2:0
  Type : TapePeripheral
  Identifier : SONY    SDZ-130    0202
  DeviceName : Tape2
  Serial number : 0001200317
  Firmware version : 0202
  Status : No media in drive.

Scsi3:0:3:0
  Type : TapePeripheral
  Identifier : SONY    SDZ-230    0100
  DeviceName : Tape3
  Serial number : 0001468132
  Firmware version : 0100
  Status : No media in drive.

Press Enter to exit . . .
```

Synchronizing the Library Tapes with the Database

Each tape inserted into a library is initially identified by its barcode label. DIVA Core keeps track of tapes currently in the library and that have been externalized in its database.

The labels and status (whether internalized or externalized) are updated in the database by Insert Tape or Eject Tape commands issued to DIVA Core. The database can become out of sync with a library's contents when tapes are added or removed directly in the library rather than through DIVA Core.

Use the following procedure to re-synchronize the tape list in the database with the library contents:

Tip: This procedure is a quick way to populate the database with tapes from the library when tapes are initially loaded.

1. Open the System Management App and connect to the database.
2. Select the Synchronize DB option from the Tools menu and acknowledge the warning message.
3. Select the individual Robot Manager to synchronize from the menu list in the Database Synchronization dialog box, or select ALL to synchronize all Robot Managers.
4. Only select the Synchronize tape list (can be very long) check box. Confirm that all other check boxes are deselected.
5. Click Go to update the selected associations.

The System Management App will connect to the Core Robot Manager. The Robot Manager obtains the current tape list from the library.

Tapes in the library are compared to the tape tables in the Core Database. New tapes are inserted into the table and existing tapes have their status updated (internalized or externalized).

6. If a Tape Type is reported that does not match the types configured in the Tape Properties area, an error is reported and no update of the database occurs. This type of error can also occur if a library cannot correctly read a tape's barcode label. You must carefully check the Robot Manager logs in this case.
7. Click Close to exit the Database Synchronization dialog box.
8. New tapes discovered during the audit are added to the Unused Tape Sets area in the Sets, Tape Groups, & Media Mapping page of the System Management App, and assigned Set ID 1. Tapes currently tracked by DIVA Core that are missing from the audit will have their status updated to externalized. You can examine the status of all tapes in the DIVA Core System Management App.

Clearing Unused Tapes

To view the To Be Cleared setting of Tapes from the regular (structured) view, go to Robot Managers, Library Settings, LSM. Select an LSM and scroll down to Unused To Be Cleared Tapes. Alternatively, switch to the List View at the top of the Media Storage page, and from the navigation bar, select Unused To Be Cleared Tapes.

You can limit the number of tapes displayed by using the Barcode filter, or disable an individual tape's To Be Cleared setting by clicking the corresponding row's Edit button under Actions. You can also disable the To Be Cleared setting for multiple tapes. Click the Multiple Edit icon at the top of the Unused To Be Cleared Tapes dialog box and select the tapes to disable the To Be Cleared setting on. Then click Add Barcodes and then click Save.

You can enable the To Be Cleared setting on a set of tapes by clicking the Add button and selecting the tapes you want to mark to be cleared. Click Add Barcodes and then click Save. This will set all selected tapes, with To Be Cleared enabled.

Creating Tape Groups

You use the System Management App to define Tape Groups within the archive. Tape Groups segment material within the tape library, or associate content with a particular Media Type. The default Tape Group is present in all installations and cannot be removed. However, you can specify your own Tape Group Names and not use the default Tape Group. Generally, the Tape Group Name is descriptive of the function or content that is being stored.

A Tape Group is associated with a Set ID defining the pool of tapes it can draw upon to store DIVA Core Objects. When DIVA Core writes an object to a tape from the pool, the tape is assigned to a Tape Group. It is released from the Tape Group when all objects have been deleted or the tape has been repacked.

The Tape Group concept in combination with the Set ID enables optimal use of tape resources. Some tape drives and media are extremely fast but typically have less storage than their larger capacity (and slower) counterparts. Content that is small, or required very quickly, should be archived to this Tape Group and should use the faster drives.

For example, the 9840C tape drive is small in capacity, but it provides extremely fast access times (approximately fifteen seconds from mount to data retrieval), and is better suited to storing large numbers of relatively small data files. This is particularly true related to tape repacking.

For example, if the Commercials Tape Group is allocated Set ID 3 and all 9840C tapes are assigned to that set. Short form commercial material written to tape will exclusively use the 9840C media. Longer (and larger) material, such as movies and interstitial programs are better suited to the larger capacity tape sets.

Tape Group Encryption

DIVA Core 8.3 tape drive encryption securely supports bulk tape migration between DIVA Core systems. You enable, disable, or update tape group encryption in the System Management App. Tape group encryption is disabled by default.

After enabling encryption on a tape group, all additional tapes added to the group will also be encrypted. However, any existing tapes in the Tape Group remain unencrypted if encryption was previously disabled.

Enabling encryption on a tape group generates an encryption key, which is also encrypted. You can change the encryption key at any time. Use the following procedure to enable, disable, or update the encryption key:

1. Navigate to the Tape Groups view in the System Management App.
2. Double-click the tape group from the list on the Tape Groups view to display the Edit Tape Groups Entry screen.
3. Select Enable, Disable, or Update from the Encryption options list.

When you enable encryption any new tape added to the Tape Group will be encrypted. However, any tape already in the Tape Group at the time of this assignment is unaffected and remains unencrypted if encryption was previously disabled on the Tape Group. You will receive a warning that you are about to enable encryption on the Tape Group when you click OK.

Disabling encryption (after it is already enabled) only affects additional tapes added to the Tape Group, and the existing tapes remain encrypted.

Updating the encryption generates a new key. You will receive a warning notifying you that a new encryption key will be assigned to the Tape Group, and that any new tapes added will use the new encryption key. The existing tapes that were already encrypted will continue to use the original key. Therefore, tapes in the same tape group can have different encryption keys. You must notify the Manager of the change when updating the encryption key.

4. Click OK to save your changes.

DIVA Core generates an encoded 256-bit encryption key. For security reasons, the encryption key is also encrypted. If you disables and re-enable encryption on a Tape Group, the same encryption key is used.

You can view the encryption status of the tape on the Home, Tapes screen in the System Management App.

See the DIVA Core Operations Guide for detailed information.

Creating Tape Sets

When a new tape is entered into a library, or DIVA Core clears a tape of its objects (whether all objects on that tape have been deleted, migrated to another tape, or moved to another tape after a tape repack), the tape is released back to the Unused Tapes Sets pool.

New tapes are automatically assigned a Set ID1, which is the default in all DIVA Core installations. Other Set ID numbers are typically used to distinguish between different types of media, but could be used to create restricted pools of tapes for particular applications. If this is the case in your installation, the Set ID must be updated for these tapes after they are inserted into the library.

Remapping Media

You can put transformation rules in place for the specified Tape Groups on Archive requests on the Media Mapping area in the System Management App. The remapped destination media can be either a disk array, tape group, or a storage plan. This is not typically used during initial installation, but rather at a later time in the object's life cycle.

Transformation rules allow transparent redirection of objects from one media type to another without needing to alter the archive initiator. Some examples are migration of an existing Tape Group to a new drive or tape generation, or migration from tape to disk.

Note: You must use a migration Request to change a tape format from Legacy to AXF. Repacking a tape will not change the tape format. Repacking of existing Legacy format objects retains the format of the tape even if the tape group format was updated in the configuration from Legacy to AXF.

The following events appear in the request details when an object's media is remapped to another media, a storage plan, or both:

- Media Name Translation has changed the Destination Media to media.
- Media Name Translation has changed the Destination Media to storageplan.
- Media Name Translation has changed the Destination Media to media & storage-plan.

Component Configuration

This chapter describes DIVA Core component configuration.

Topics:

- [Configuration Overview](#)
- [Module Configuration Files](#)
- [Manager Configuration](#)
- [Actor Configuration](#)
- [Robot Manager Configuration](#)
- [DIVAmigrate Installation and Configuration](#)

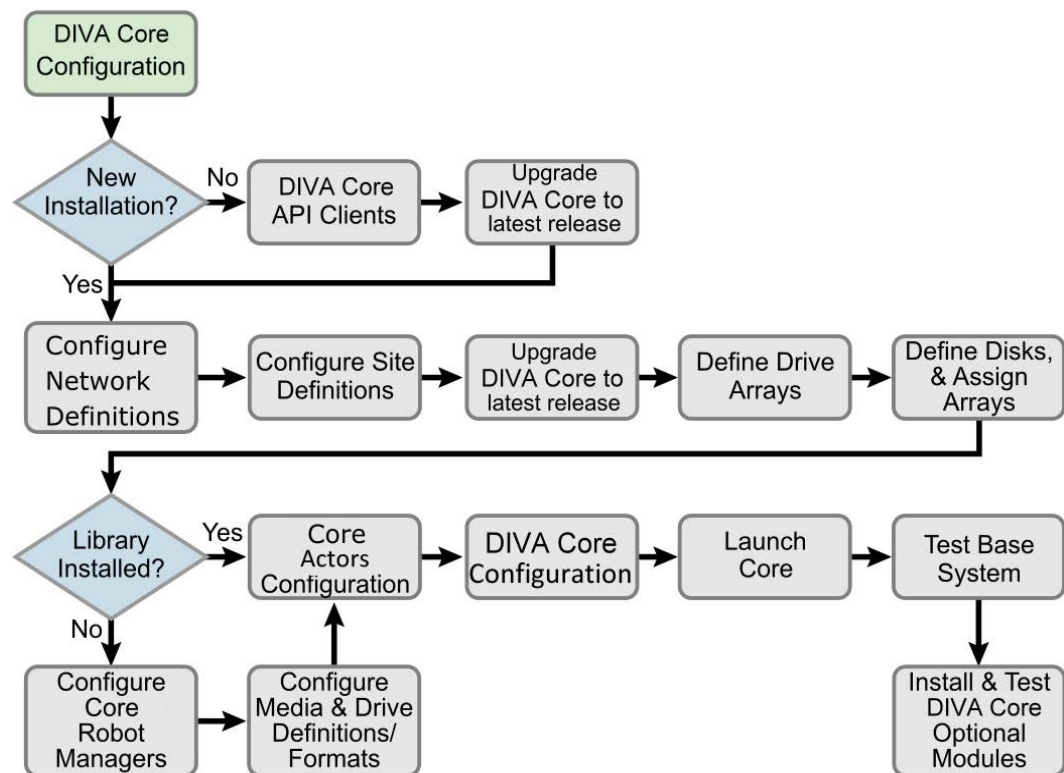
Configuration Overview

There are many interrelated components in a DIVA Core System. The following figure shows the basic configuration workflow.

The configuration of DIVA Core is hierarchical and top-level parameters such as Networks, Sites, Arrays, and Disks need to be configured before configuring other components such as Core Actors.

If you intend to modify an existing DIVA Core system, you must always start by backing up the existing DIVA Core installation, configuration files, and especially the DIVA Core and Metadata Databases.

Contact Technical Support before making any modifications to your DIVA Core platform if you are unsure about any steps in the procedures, or require clarification.



DIVA 031

Module Configuration Files

Each DIVA Core software module has its own static configuration text file with parameters needed to launch that particular application. The files are typically denoted with the .conf file name extension. There are some DIVA Core modules that use an XML based file rather than a text file for their configuration and those will be noted where applicable.

Unlike older releases of DIVA Core that stored these configuration files in the same folder as the application itself, DIVA Core 8.3 centralizes them to a dedicated conf subfolder under the DIVA Core Program Group.

The configuration files are typically updated with additional or changed settings in newer releases of the software. A new or patch release of DIVA Core will have the new releases of the .conf files appended with a .ini extension. For example, the new release of the Core Manager Configuration file will be named manager.conf.ini. You must remove the .ini extension after the installation is complete and the new configuration file updated.

Each configuration file can be opened and edited with any plain text editor (for example, Windows Notepad or Notepad++).

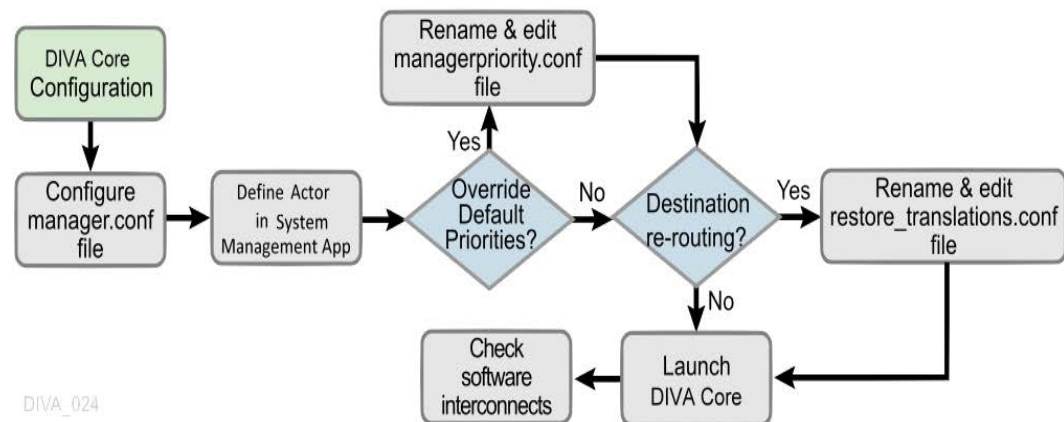
Any changes made to the configuration file of a DIVA Core software component requires that the component be shut down and then restarted for the changes to take effect. The exceptions to this are the Manager and DIVA Connect options, both of which allow configuration changes to be reloaded while they are still running. There are code dependencies between some applications in the DIVA Core platform, so other components may also need to be restarted for changes to take effect.

Manager Configuration

The Manager module is located in %DIVA_HOME%\Programs\Manager\bin in Windows, and in /home/diva/DIVA/Program/Manager/bin in Linux, and runs as a service. The static configuration file for the Manager is manager.conf. You can typically leave most settings in this file left at the default values. The settings that would normally require updating are highlighted in bold type.

See [Appendix A: Core Options and Licensing](#) for detailed information.

The following figure is the workflow for installing a Core Manager:



Configuring the Local Manager

The static configuration file in new installations is initially named manager.conf.ini. You must remove the .ini extension for it to be recognized by the Core Manager.

The configuration file is divided into five distinct groups; Basic, Database, Advanced, Logging, and Service settings. You must not modify the Service settings section, and therefore, not covered in this manual. Values defined in this section must only be altered with instruction from Technical Support.

Each parameter section in the configuration file contains information on defining that parameter. The information lines are commented out (start with #) and ignored by the Manager. Any parameter definition that is missing the equal sign is also ignored.

Spaces in the parameter settings are significant. Do not put extra spaces before or after the parameter names or their values. If you have trouble running the Manager after configuring the manager.conf file, confirm that spaces are not present in any of the parameter values you have defined.

Restarting the Manager can disrupt a live Network. You can make most of the customizations in the configuration file effective immediately using the restart command line switch.

If you intend to update your existing DIVA Core system with a newer software release, you must use the manager.conf.ini from the new release. You must update the Basic

and Database settings with the values from the old configuration file. The new release configuration file may have additional settings or updates included; this applies to all DIVA Core software modules when installing a release updated.

Basic Settings

Except for the SERVICE_NAME, these parameters are always required and must be defined for the Manager to start successfully. These settings define how other DIVA Core software components and DIVA Core API clients connect to the Manager.

Note: These settings are not reloadable while the Manager is running. You must restart the Manager for them to take effect.

The following table describes the Basic settings in the manager.conf file:

Parameter	Parameter Type	Description	Default
SERVICE_NAME	Name	You can use this parameter to specify the name of the service. If not defined, the Service Name defaults to Core Manager.	
DIVAMANAGER_NAME	Name	This is the name this Manager instance uses to identify itself to other Core Managers sharing its resources. Otherwise, this is arbitrary. It must be unique in a system running multiple Managers except for Main and Backup Managers (configured as a cold standby). In this instance, the names should be identical.	DIVA
DIVAMANAGER_PORT	TCP Port Number: unsecure connections	This is the name this Manager instance uses to identify itself to other Core Managers sharing its resources. Otherwise, this is arbitrary. It must be unique in a system running multiple Managers except for Main and Backup Managers (configured as a cold standby). In this instance, the names should be identical.	DIVA
DIVAMANAGER_SECURE_PORT	TCP Port Number: secure connections	The secure TCP port used by DIVA Core Services and the DIVA Core API.	8000

Database Settings

These parameters define the location and instance of the Core Database. Except for the DIVAMANAGER_TNSNAME parameter, you must define all settings in this section for the Core Manager to launch successfully.

The following table describes the Database settings in the manager.conf file:

Parameter	Parameter Type	Description	Default
DIVAMANAGER_TNSNAME	Name	<p>The TNS Name of the DIVA Core Schema within the Oracle database. DIVA Core ignores this setting if the DIVAMANAGER_DBHOST and DIVAMANAGER_DBPORT settings are defined.</p> <p>This feature requires Oracle 11g or higher installed on the host running the Manager. If this setting is defined, the location of the Oracle OCI driver (for example, ocijdbc11.dll) must be added to the wrapper.java.library.path setting (located in Service settings section of the file); otherwise, the Manager will not start as a service.</p> <p>Example: wrapper.java.library.path=.;C:\app\oracle\product\11.1.0\BIN</p>	
DIVAMANAGER_DBHOST	IP Address or Host Name	This specifies the Host Name or IP Address of the computer containing the Core Database. If using a host name, this must be present in the hosts file on the computer where the Core Manager is installed.	
DIVAMANAGER_DBPORT	TCP Port Number	The Oracle Listener port configured during the Core Database installation.	1521
DIVAMANAGER_DBSID	Name	The Core Database SID (Instance System Identifier) in Oracle where Core Manager connects.	

Parameter	Parameter Type	Description	Default
DIVAMANAGER_DBUSER	Name	The user name the Core Manager uses to connect to the Core Database. This is typically diva (case sensitive).	diva
DIVAMANAGER_DBSERVICENAME	Name	Oracle ServiceName setting. Either this value or DIVAMANAGER_DBSID must be set. If both are set, this takes precedence over the SID.	No default value, but lib5.world is recommended.
DIVAMANAGER_DBSID	Name	Oracle ServiceName setting. Either this value or DIVAMANAGER_DBSERVICENAME must be set. If both are set, DIVAMANAGER_DBSERVICENAME takes precedence over SID.	No default value, but lib5.world is recommended.

Advanced Settings

You typically leave the parameters in this section are typically left at their defaults. They customize DIVA Core's default behavior for task execution, resource allocation, and the number of connections it will accept from DIVA Core Applications and DIVA Core API Clients. These parameters are normally adjusted or fine-tuned after completing the initial installation of DIVA Core.

Most (but not all) of these settings can be altered while the Manager is running by using the reload option.

The following table describes the Advanced settings in the manager.conf file:

Parameter	Parameter Type	Description	Default
DIVAMANAGER_TO_LOWER	true or false	Sets case sensitivity for DIVA Core. If set to true, then all object names, categories and tape groups will be set to lowercase.	false
DIVAMANAGER_REQUEST_SCHEDULING_QUEUE_SIZE	Number of requests	The maximum number of requests that can be queued for processing by DIVAMANAGER_MAX_CONCURRENT_REQUESTS processors of the Request Scheduler.	500
DIVAMANAGER_MAX_CONNECTIONS	Number of Connections	Specifies the maximum number of simultaneous client connections the Manager will accept. This includes Core Actors, System Management Apps, API connections, and support tools.	200
DIVAMANAGER_MAX_SIMULTANEOUS_REQUESTS	Number of Requests	The maximum number of requests processed by the Core Manager. When this limit is reached, any further requests will be rejected. The maximum tested value for this setting is 2000.	500
DIVAMANAGER_API_TASK_QUEUE_SIZE	Number of tasks	The number of tasks that will be accepted to the API command processing queue. If this queue is full, subsequent commands will be rejected. The maximum tested value is 2000.	
DIVAMANAGER_MAX_INACTIVE_REQUESTS	Number of Requests	Maximum number of inactive requests that cannot find resources examined by the Request Scheduler each time it is activated.	0
DIVAMANAGER_TYPICAL_VIRTUALOBJECT_SIZE	Percentage	<p>During operation a Core Actor retrieves the file size of an object before an archive transfer. This value determines the best location on the tape for the file.</p> <p>Some servers do not indicate the file size of an object before a Direct Archive. Therefore, DIVA Core will use this value as an estimate for tape selection.</p> <p>You must define this setting so that most objects to be archived in the DIVA Core system are below this size.</p>	10 (percent)
DIVAMANAGER_MAX_CONCURRENT_REQUESTS	Number of Requests	The maximum number of concurrent requests executed by the Core Manager. The maximum tested value for this setting is 16.	8

Parameter	Parameter Type	Description	Default
DIVAMANAGER_MAX_SPAN_SEGMENTS	Number	DIVA Core will attempt to span the file across two or more tapes if no more writable tapes with enough free space are available to archive a file. This setting defines the maximum number of tapes across which the object will be spanned.	2 (segments)
DIVAMANAGER_INITIAL_DB_CONNECTION_LIMIT	Number of Connections	The initial number of database connections available to the Core Manager.	1
DIVAMANAGER_MIN_DB_CONNECTION_LIMIT	Number of Connections	The minimum number of database connections available to the Core Manager.	1
DIVAMANAGER_MAX_DB_CONNECTION_LIMIT	Number of Connections	The maximum number of database connections available to the Core Manager.	10
DIVAMANAGER_CAPACITY_LOW_WATER_MARK	Percentage	When the percentage of the total used capacity reaches this amount, periodic warning messages are issued in the System Management App.	90 (percent)
DIVAMANAGER_ENABLE_SPANNING_LARGE_VIRTUALOBJECTS	true or false	Enables spanning of large objects. This parameter overrides SPAN_SEGMENTS if any object in the system is known to be too large.	true
DIVAMANAGER_INACTIVITY_TIMEOUT	Time in Seconds	The maximum time a physical connection can remain idle in a connection cache before it is terminated (in seconds).	3600
DIVAMANAGER_MAX_VIRTUALOBJECTS_FOR_REPACK	Number	Repacking a tape with many objects can consume resources for a lengthy period without reclaiming a great deal of unused space in the process. This setting prevents this by limiting the selection of tapes in manual and automatic repacks based on the number of objects.	500
DIVAMANAGER_SIZE_OF_STATEMENT_CACHE	MB	The size of the database statement cache.	10
DIVAMANAGER_STOP_IMMEDIATELY_FOR_REPACK	true or false	This setting specifies whether to complete any repack requests still running or to terminate them after the Automatic Tape Repack period. If this is set to true then repack requests still in progress after the Automatic Repack period will be stopped.	true

Parameter	Parameter Type	Description	Default
DIVAMANAGER_DEFAULT_ROW_PREFETCH	Number of Rows	The default number of rows to prefetch from the database per query.	1000
DIVAMANAGER_DISMOUNT_AFTER	Time in Milliseconds	This specifies the time in milliseconds to automatically dismount a mounted tape no longer needed by any other request.	120000 (two minutes)
DIVAMANAGER_FAILOVER_ENABLED	Boolean	Whether to enable Fast Connection Failover. This feature introduces a slight performance penalty.	false
DIVAMANAGER_UPDATE_PRIORITIES_PERIOD	Time in Milliseconds	DIVA Core periodically examines all requests in its request queue and increments the request priority. This prevents a condition where low priority requests might be continually superseded by higher priority requests. This setting specifies the period between updates of the queue by the Manager. You set this value to 0 to disable priority updates.	60000 (one minute)
DIVAMANAGER_NUM_REQUEST_SOLUTIONS_TO_EVALUATE	Boolean	The number of immediate solutions to evaluate per invocation of the Best Solution Finder during resource selection. Values are 0 (disabled) or 1 (enabled).	0 (disabled)
DIVAMANAGER_MAX_DELAY_BETWEEN_SCHEDULER	Time in Milliseconds	The maximum number of milliseconds between two Request Scheduler activations when the Manager is constantly busy.	5000 (five seconds)
DIVAMANAGER_SCHEDULER_AFTER_INACTIVITY	Time in Milliseconds	The number of milliseconds after which a requested Request Scheduler activation can be launched if the Manager is idle. This duration should be significantly lower than DIVAMANAGER_MAX_DELAY_BETWEEN_SCHEDULER. You should not need to modify this value.	500
DIVAMANAGER_PING_INTERVAL	Time in Milliseconds	This defines the interval in milliseconds between Manager checks to see if the connections to its clients and services are still active (Actors, SPMs, System Management Apps, and so on).	600000 (ten minutes)

Parameter	Parameter Type	Description	Default
DIVAMANAGER_EXPORT_ROOT_DIR	Directory Path	The Export Tapes command enables the sharing of tapes between two or more separate DIVA Core platforms. This setting defines the root folder for the exported tape's Metadata files. The folder must exist and have write permissions enabled on the host computer where the Core Manager is running.	Exported
DIVAMANAGER_MAX_RESTORE_SERVERS	Number between 2 and 200	The maximum number of servers allowed in an N-Restore request by a Core Actor.	5
TAPE_FULL_ON_SPAN_REJECTED	true or false	If true, and spanning is disabled, the Manager marks a tape full when spanning occurs.	false
DIVAMANAGER_MAX_EXPORT_TAPES	Number between 1 and 100	The maximum number of tapes allowed in an Export Tapes request.	10
DIVAMANAGER_MAX_EXPORT_ELEMENTS	Number between 1 and 10,000,000	The maximum number of elements that can be exported using the Export command.	1000000
DIVAMANAGER_MAX_FILES_IN_ARCHIVE	Number between 1 and 1,000,000	The maximum number of files allowed in an Archive request.	1000000
DIVAMANAGER_MAX_FILES_IN_PARTIAL_RESTORE	Number between 1 and 1,000,000	The maximum number of files allowed in a Partial File Restore request.	1000000
USE_IMPROVED_BEST_WORST_FIT_ALGORITHM	true or false	When a file was archived to tape in earlier DIVA Core releases, the Best/Worst Fit algorithm selected the tape with the largest remaining free size. This could result (over time) in a low number of blank tapes for tape repacking, and so on. The current algorithm selects the tape based on the smallest free space and then fills all tapes before using more free tapes.	true

Parameter	Parameter Type	Description	Default
DIVAMANAGER_SITE_SUPPORT_ENABLED	true or false	Resources within DIVA Core can be defined by their location. If you set this parameter to true, the Manager first tries to perform the request from the sites identified as MAIN. If unsuccessful, it retries the request with resources from all other sites. If you set this parameter to false, DIVA Core ignores site identification and all site resources are considered equally.	false
DIVAMANAGER_CACHE_QOS_USE_DISK	true or false	In earlier DIVA Core releases, a Restore request with a Quality of Service of CACHE or CACHE and DIRECT resulted in the tape instance being used as first priority, even if a disk instance existed. This setting instructs DIVA Core to use the disk instance regardless of the QOS method specified.	true
DIVAMANAGER_PRIORITY_TIER	Number between 0 and 100	<p>DIVA Core bases the execution of requests in its request queue by the request priority number. However, there are instances where a request in the queue with lower priority uses a tape that is already mounted. Giving this request priority over others lower in the queue can save a substantial amount of time in tape mount and dismount operations, and help reduce wear and tear on the tape drives.</p> <p>If this setting is enabled, DIVA Core examines the request queue for lower priority requests involving a tape that is already mounted in a drive and adds the number specified here to the request priority.</p> <p>For example, if the request priority is 25, and the Priority Tier value is 50, the total request priority is 75.</p> <hr/> <p>Note: This feature applies only to Restore and Copy Requests that read from tape. Archive and Copy requests that write to tape are not supported by this feature.</p> <hr/>	0 (disabled)

Parameter	Parameter Type	Description	Default
DIVAMANAGER_ETC_FEATURE	true or false	This parameter enables the Estimated Time to Complete feature. This function gathers statistics (over time) on the time for completion of all execution states of each DIVA Core request. Setting this value to true enables this feature.	false
DIVAMANAGER_ETC_CONFIDENCE_LEVEL	Number	The percentage of Slope Confidence Interval for the simple regression statistical function used in the Estimated Time to Complete feature. DIVA Core ignores this setting if the DIVAMANAGER_ETC_FEATURE is disabled.	50
DIVAMANAGER_OVERWRITE_POLICY	Number between 0 and 2	This value determines how DIVA Core handles files that already exist on a Destination Server when executing a Restore, Partial File Restore, or N-Restore request as follows: 0 - If the file to be restored to the Destination Server already exists no overwrite will occur. 1 - The Actor does not verify if the files with the same names exist before attempting to overwrite these files. If files with the same names do exist, a backup of the existing files is made before overwriting them. 2 - The Actor attempts to delete and then write to files with the same names.	1
DIVAMANAGER_OVERWRITE_OVERRIDE	true or false	Overrides the policy sent by the external application through a request with the policy set in DIVAMANAGER_OVERWRITE_POLICY.	false
LICENSE_EXPIRATION_NOTIFICATION_PERIOD	Number of Days	Number of days before a temporary license is to expire that a notification message will be displayed on the GUI. The range of possible values is 1 to 99.	15
LICENSE_EXPIRATION_TIME_OF_DAY	Time of Day	The time of day the Manager will shut down if the license has expired. The Manager will stop at the designated time on the day after the license validity date. (00-23:00-59)	8:00

Parameter	Parameter Type	Description	Default
ATTEMPT_ACCESS_TO_OFFLINE_DISK	true or false	If a disk is offline or not visible to all available Actors, the Manager will automatically terminate a transfer request for objects residing on that disk. If this is set to true, the Manager attempts the transfer irrespective of disk status.	false
CHANGE_DISK_STATE_ON_ERROR	true or false	Defines whether the Manager will automatically vary a disk's status to Offline if a transfer error occurs.	true
MANAGER_ACTOR_DISK_RETRY_NUMBER	Number	If a disk I/O error occurs during a transfer, this sets the maximum number of transfer retry attempts with alternate Actors that also have access to the disk. Values are 0 to 7.	3
DISK_STATUS_POLLING_RATE	Number	This defines the rate in milliseconds in which each disk in the system is polled to obtain its total and remaining free space.	60000 (one minute)
DISK_BUFFER_SPACE	Number	This defines the percentage of the overall space of a disk to keep free.	0.05 (percent)
DISK_CONNECTION_STATE_RESET_DELAY	Time in Minutes	A disk connection will be reset from the Out of Order state when a successful access is completed and this amount of time has passed since the connection was set to Out of Order.	1.0 (minute)
COMPONENT_SIZE_CONVERSION_TO_KB_RULE	Number	When an element is successfully transferred to tape or disk, the Actor reports the size of the element in bytes. This value is then converted to KB before it is saved to the database. The conversion may be one of three possible values: 1 - $KB = (bytes / 1024) + 1$ 2 - $KB = bytes/1024$, but if $(KB < 1)$ then $KB = 1$ 3 - $KB = \text{Math.ceil}(bytes/1024)$	3
DIVAMANAGER_MAX_EXCLUDED_INSTANCES	Number	The maximum number of instances excluded from a request that are logged as an event.	3

Parameter	Parameter Type	Description	Default
LOGGING_TRACE_LEVEL	DEBUG, INFO, WARN, ERROR, FATAL	<p>Defines the level of information written to the respective log files as follows:</p> <ul style="list-style-type: none"> • DEBUG - All messages within the Manager are logged. Log files grow rapidly. • INFO - Information, Warning, Error, and Fatal messages are logged. • WARN - Warning, Error, and Fatal messages are logged. • ERROR - Error and Fatal messages are logged. • FATAL - No messages are logged unless the Manager stops unexpectedly. 	INFO
DIVAMANAGER_MAX_SPAN_SEGMENTS	Number	<p>DIVA Core will attempt to span the file across 2 or more tapes if no more writable tapes with enough free space are available to archive a file.</p> <p>This setting defines the maximum number of tapes that the object will span. This setting will completely disable spanning if set to 1 or below. If a span case arises, the Manager retries the request with a new tape using the old Worst Fit algorithm, and the first tape in the attempted span will be marked full. If the second attempt fails, the request will terminate.</p>	0 (segments)
DIVAMANAGER_MAX_DATABASE_CONNECTION_ATTEMPTS	Number	The maximum number of allowable attempts to connect to the database.	10000
DIVAMANAGER_MIN_DATABASE_CONNECTION_PERIOD	Number	The minimum period (in milliseconds) between connection attempts.	1000 (milliseconds)
DIVAMANAGER_MAX_FOLDERS_IN_ARCHIVE	Number	The maximum number of folders allowed in an Archive request. Performance degradation can occur for values greater than 10000. The maximum value is 10000.	10000
DIVAMANAGER_COMPLEX_VIRTUALOBJECT_THRESHOLD	Number	The maximum number of files allowed before an object is classified as a Complex Object. The maximum value is 10000.	1000

Parameter	Parameter Type	Description	Default
COPY_ONLY_FROM_DISK_INSTANCE_WHEN_POSSIBLE	Boolean	Controls instance selection for Copy and CopyAs requests when the Destination Server is tape. Copy requests always check if a disk instance can be used as the Source Server of a copy. If the required resources for a disk to tape transfer are not available, a tape to tape transfer may be used if this parameter is set to false. When set to true the request will wait for the resources to use the disk instance as the Source Server. This parameter is reloadable in SERVICE mode.	true
COMPONENT_SIZE_CONVERSION_TO_KB_RULE	Number	This is the Object Size Conversion Rule. Use one of the following rules to convert an object component size from Bytes to Kilobytes: 1 - KB = (bytes/1024) + 1 2 - KB = bytes/1024, but if (KB < 1) then KB = 1 3 - KB = Math.ceil(bytes/1024)	3
COPY_ONLY_FROM_DISK_INSTANCE_TIMEOUT	Time in Minutes	Tape instance is available for a Tape to Tape transfer. After this time, either a disk or tape instance may be selected as the Source Server of a copy to tape.	15 (minutes)
DIVAMANAGER_RESTORE_QOS	CACHE_ONLY, DIRECT_ONLY, DIRECT_AND_CACHE, CACHE_AND_DIRECT, NEARLINE_ONLY, NEARLINE_AND_DIRECT	This identifies the default Quality of Service for Restore requests.	NEARLINE_AND_DIRECT
NTH_PROGRESS_MESSAGE	Number	The number of progress messages sent to GUIs. Every Nth progress message will be sent. The N=100 progress message will always be sent.	5 - implies send every fifth progress message to all GUIs.

Parameter	Parameter Type	Description	Default
DIVAMANAGER_TIME_T O_WAIT_FOR_GRACEFU L_SHUTDOWN	Minutes	The time to allow for a graceful shutdown to complete.	1440 (one day)
ABORT_ARCHIVES_ON_ EMPTY_FILES	true or false	If true the Manager terminates an Archive request if it contains an empty file or folder.	false
TAPE_FULL_ON_SPAN_R EJECTED	Boolean	If true, the Manager will mark a tape full when a span occurs but spanning is disabled.	false
DIVAMANAGER_RETRY_ ON_SPAN_REJECTED_AL GORITHM	1 = Prefer empty tapes 2 = Prefer used tapes with less remaining space 3 = Prefer tapes with more remaining space	The tape selection retry algorithm to use when a span is rejected. The Manager enables configuring the retry logic when spanning is disabled, but an object is too large to fit on the selected tape. By default, the Manager retries with an empty tape, but you can alternatively retry with a used tape with most or less remaining space.	1

Logging Settings

The following table describes the Logging settings in the manager.conf file:

Parameter	Parameter Type	Description	Default
LOGGING_TRACE_LEVEL	DEBUG, INFO, WARN, ERROR, FATAL	Defines the level of information written to the respective log files as follows: <ul style="list-style-type: none"> • DEBUG - All messages within the Manager are logged. Log files grow rapidly. • INFO - Information, Warning, Error, and Fatal messages are logged. • WARN - Warning, Error, and Fatal messages are logged. • ERROR - Error and Fatal messages are logged. • FATAL - No messages are logged unless the Manager stops unexpectedly. 	INFO
LOGGING_MAXFILESIZE	Kilobytes or Megabytes	When the log file reaches this size, a new file is generated and the old one renamed with appropriate time and date stamps. Older log files are subsequently compressed automatically into zip files at one hour intervals.	10 MB
LOGGING_LIFETIME	Hours	This setting defines how long to maintain trace service and zipped log files before deleting them.	50

Configuring Request Priorities

Each request submitted to the Core Manager is placed in the Manager transfer queue. Request priorities enable DIVA Core to differentiate between important requests, such as Restore requests, over less important events. For example, tape repacks, and so on.

The request priority is a number from zero to one hundred with zero being the lowest priority and one hundred being the highest. The request priority is typically specified when you submit the request (either from the System Management App or the DIVA Core Client API). You can also alter the priority after you submit the request using the Change Priority command.

The default request priority for each request type is preset within DIVA Core. You can override the default priorities (at your discretion) using the following procedure:

1. Navigate to the %DIVA_HOME%\Program\conf\manager folder.
2. Rename the managerpriority.conf.ini file to managerpriority.conf.
3. Edit the managerpriority.conf file using a plain text editor (for example, Notepad or Notepad++) to set the desired values for each request type.

4. You must reload the Manager configuration using the reload option or restart the Manager for the new settings to take effect.

Regardless of the configured request priority, the Manager will (by default) periodically increment the priority of every request already the request queue. This prevents a condition where a low request priority can be continually overridden by higher priority requests and never executed.

You can disable this feature by setting the `DIVAMANAGER_UPDATE_PRIORITIES_PERIOD` parameter in the Manager configuration file to 0. You must then reload the Manager configuration or restart the Manager.

Rerouting Destinations (`restore_translations.conf`)

To simplify production workflows, you can configure DIVA Core to automatically override the original Destination Server specified in a Restore, Partial File Restore, or N-restore request based on the Object Collection and original Destination Server. This is called Destination Rerouting. Typically, you use this function to enable selective transcoding based on an object Collection.

You configure Destination Rerouting by editing the `restore_translations.conf` file. The file is located in the `%DIVA_HOME%\Program\conf\manager` folder with the Manager configuration file.

The `restore_translations.conf` file is delivered with a `.ini` extension. You must remove the `.ini` extension for this file to be considered by the Manager.

All re-routing entries must be in the following format:

```
DT_Number=Destination_1;Collection_1;TranslatedDestination_1
```

The following list describes these parameters:

DT_Number

This must be the first string in the line and start with `DT_Number`. The Number can be any value unique among all entries. For example, `DT_0`, `DT_1`, `DT_2`, and so on. Up to three hundred entries are supported.

Destination_1

The Destination Server in a Restore request for this rule to apply.

Collection_1

If the Object Collection of the request also matches the Destination Server will be re-routed.

TranslatedDestination_1

This is the new Destination Server for the Restore request.

The following example describes how to configure rerouting a destination:

- A video server accepts clips with Format1
- The archive contains clips with both Format1 and Format2
- Format 1 Objects are in Collection 1 (Cat1)
- Format 2 Objects are in Collection 2 (Cat2)

You configure this example as follows:

1. Define a Source Server (Source1) that points to the video server with no restore transcode options.
2. Define another Source Server (Source2) that points to the video server with options to transcode to Format1.
3. Create a `restore_translations.conf` file containing the following line:

```
DT_0=Source1;Cat2;Source2
```

When an object with the Collection Cat2 is restored to Destination Server Source1, re-route it to Destination Server Source2 instead. In this manner, the automation can always use Source1 as the Destination Server in the request.

Objects having a format of Format1, which are directly compatible with the video server, will be restored to Source1 without transcoding.

Objects having a format of Format2 and a Collection of Cat2 match the configuration line and are rerouted to Source2. Source2 has options to transcode them to Format1 when restoring.

Controlling the Manager

Core Manager control and management functions are performed from a command prompt on Windows platforms using the `manager.bat` batch file, and from a terminal window using the `manager.sh` script file on Linux platforms. The executable is located in the `%DIVA_HOME%\Program\Manager\bin` folder in Windows, and in the `/home/diva/DIVA/Program/Manager/bin` directory in Linux.

Installing and Removing the Manager Service in Windows

You must first install the Core Manager as a system service on new systems. You can accomplish this using the `install` (or `-i`) and `uninstall` (or `-u`) command line switches as follows:

manager install

This (or `manager -i`) installs the Core Manager service set by the `SERVICE_NAME` parameter defined in `manager.conf`. If this parameter is undefined, the service is installed as Core Manager.

manager uninstall

This (or `manager -u`) removes the Core Manager service set by the `SERVICE_NAME` parameter defined in `manager.conf`.

In the Windows Services applet, confirm that the Core Manager service is installed correctly. If you must change the service name, uninstall the existing service before editing the `manager.conf` file. Then reinstall the service after changing the service name.

The default path to the `manager.conf` file is `%DIVA_HOME%\Program\conf\manager`.

You can identify a specific configuration in the command line if you require using an alternate file using the `-conf` or `-f` switch as follows:

```
manager install -conf [configuration file]
manager uninstall -conf [configuration file]
```

Installing and Removing the Manager Service in Linux

The `divaservice` executable in the Manager `/home/diva/DIVA/Program` directory installs (or uninstalls) the Core Manager as a service from a Linux terminal. See [Installing the DIVA Core Services](#) for more information about using the `divaservice` command.

Use the following command sequence to install the Manager service:

```
cd/home/diva/DIVA/Program

./divaservice install manager /home/diva/DIVA/Program/conf/
Manager/manager.conf
```

Use the following command sequence to uninstall the Manager service:

```
cd/home/diva/DIVA/Program

./divaservice uninstall manager /home/diva/DIVA/Program/conf/
Manager/manager.conf
```

Managing the Manager Service

You can manage the Manager Service using the following command line switches after the service is installed:

manager start

This switch starts the Manager service (if stopped).

manager stop

This switch stops the Manager service (if running).

manager shutdown

This switch finishes currently jobs and stops accepting new requests, then it stops the Core Manager service (if running).

manager restart

This switch stops and subsequently starts the Manager service.

manager reload

Some changes in the Manager configuration files take effect after reloading the Manager. This switch reloads the `manager.conf`, `managerpriority.conf`, and

restore_translations.conf files from the default path (%DIVA_HOME%\Program\conf\manager).

Use the following command to reload the Manager using a different configuration file:

```
manager reload -conf [configuration file]
```

manager status

This switch displays the current status of the Manager service (running or not running).

manager dump

This switch requests a system dump from the Manager service.

manager version

This switch (or `manager -v`) displays the Manager service release information and then exits.

manager help

This switch (or `manager -h`) display all command line options and then exits.

Logging Manager Activity

The Core Manager keeps detailed logs of its operations and stores them in the %DIVA_HOME%\Program\log\manager folder. The logs are used for troubleshooting and diagnostics purposes, and may be requested by Technical Support.

The logging settings in `manager.conf` determine the level and quantity of information captured in each log file. If you must alter the settings, you can make the changes effective immediately using the `manager reload` command, or (in DIVA Core 8.3) change them dynamically from the System Management App. See the DIVA Core Operations Guide for more details.

Class-level logging is supported through the `manager.classLog.properties` file. Any class set to one of the following values will log at the specified logging level:

- TRACE
- DEBUG
- INFO
- WARN
- ERROR
- FATAL

New statical data is generated every five minutes that lists various Manager performance related metrics, and collected in a statistics folder.

After logs have reached the size defined by `LOGGING_MAXFILESIZE` in `manager.conf` they are renamed with date and timestamps, compressed (zipped), and a new file is

started (named manager.trace). The manager.trace file is the log file currently being written to by the Manager.

Confirming System Connectivity

After the Core Manager has been successfully configured and launched you must check that the Manager can successfully be connected to by other DIVA Core clients (for example, the System Management App). Also, the Manager itself must be able to connect to the configured Actors and, if installed, Robot Managers.

Confirming Remote Client to Manager Connectivity

This short test establishes whether the Manager is configured correctly and accepting remote connections from clients:

1. Launch the DIVA Core System Management App from a remote client (that is, not on the same host computer as the Core Manager).
2. Click the Menu Orb on the top left of the System Management App.
3. Click Connect.
4. Enter the IP Address and TCP Port of the Manager in the Connect to the Manager dialog box.
5. Click Connect.
6. A successful connection will be indicated by a Connected status in the System Management App notification area (at the bottom of the screen).

Confirming Manager to Actors Connectivity

This short test establishes whether the Manager can connect to all Actors in the system. This test assumes all Actors have been configured correctly and are online.

With the System Management App still open, click the Actors icon in the Home tab on the icon bar to display the Actors view.

Confirm that the Manager has established a connection to all configured Actors, and troubleshoot if necessary.

Confirming Manager to Robot Manager Connectivity

This short test establishes whether the Core Manager can be connected to each configured Core Robot Manager. This test assumes the following:

- All Core Robot Managers are configured correctly.
- Each Core Robot Manager is running.
- All Managed Storage are loaded with tapes.
- Any library management software (for example, ACSLS) is running, and the library is set to Online.

- Manual operation has been confirmed successfully with the Core Robot Manager Client Tools.

Use the following procedure to confirm connectivity:

1. Click the Tapes icon on the Home tab to display the Tapes view.
2. Take note of the ACS and LSM number for each tape to test each particular library.
3. Right-click a tape for each ACS and LSM to test and click Eject Tape from the resulting menu.
4. Click the Manager icon on the Home tab to display the Manager Current Requests view.
5. Double-click the Eject Tape request entry to check if an error was encountered during request execution.

Manager Failover Procedures

Caution: The procedures in this section are critical and sensitive. They should only be performed under the control of Technical Support.

The following steps are required to failover a Core Manager to the Backup when the database is still accessible on the original Manager:

1. Ensure all contents of the DIVA folder from main Manager exist in the Backup Manager (particularly the correct .conf files). If they do not exist move the .conf files to the Backup Manager.

Caution: Make sure to confirm the Backup Manager has the correct DIVA binary files including major/minor version, patches, and proper database version. Always keep a backup of the original DIVA folders if making any file changes.

2. Confirm all services are installed, for example WFM, Manager, Backups, SPM, Oracle, and so on, on the Backup Manager machine. If not, the services must be installed before proceeding. Ensure the services are at the same version and patch level as the main Manager.
3. Stop all services and export the database from the original Manager. Contact Telestream Support if the database is not accessible due to failure.
4. Create a new DIVA user on the Backup Manager using the -notable option, then import the database to the Backup Manager and verify the count of archived objects is correct from the Original Manager to the Backup Manager. This can be done with the following query in SQL;

```
SELECT COUNT(*) AO_VIRTUALOBJECT_NAME from  
DP_ARCHIVED_VIRTUALOBJECTS;
```

Contact Telestream Support if you need assistance exporting and importing the database.

5. Change the Backup Manager IP to the Original Manager IP by first applying a placeholder IP on the Original Manager.
6. Confirm the configuration is valid in the `manager.conf`, `robotmanager.conf`, `spm.conf`, and all disk and file paths in the configuration are accessible from the Backup Manager machine.
7. Enable and start all services and confirm Backup Manager is running as anticipated; monitor activities.

Cluster Failovers

Use the following procedures if a cluster fails to initiate:

1. Check that the backups are synced on the Active Node and Backup Manager.
2. Stop all DIVA Services from the Microsoft Cluster on the Active Node.
3. On the Active Node, run `SELECT COUNT(*) AO_VIRTUALOBJECT_NAME from DP_ARCHIVED_VIRTUALOBJECTS;`
4. Create an export of the current database from the Active Node.
5. Stop the DB Services on the Active Node from the Microsoft Cluster Manager.
6. Start the DB Services on the Backup Manager server.
7. Recover the database from the backups. Contact Telestream Support if you require assistance recovering the database.
8. Start the DIVA Services on the Backup Manager and run some tests to confirm functionality.
9. When all testing is successful, stop all Backup Manager services, and restart all services from the Microsoft Cluster Manager. Verify all operations are functioning normally.

Actor Configuration

This section describes Core Actor configuration and operations.

Configuration Overview

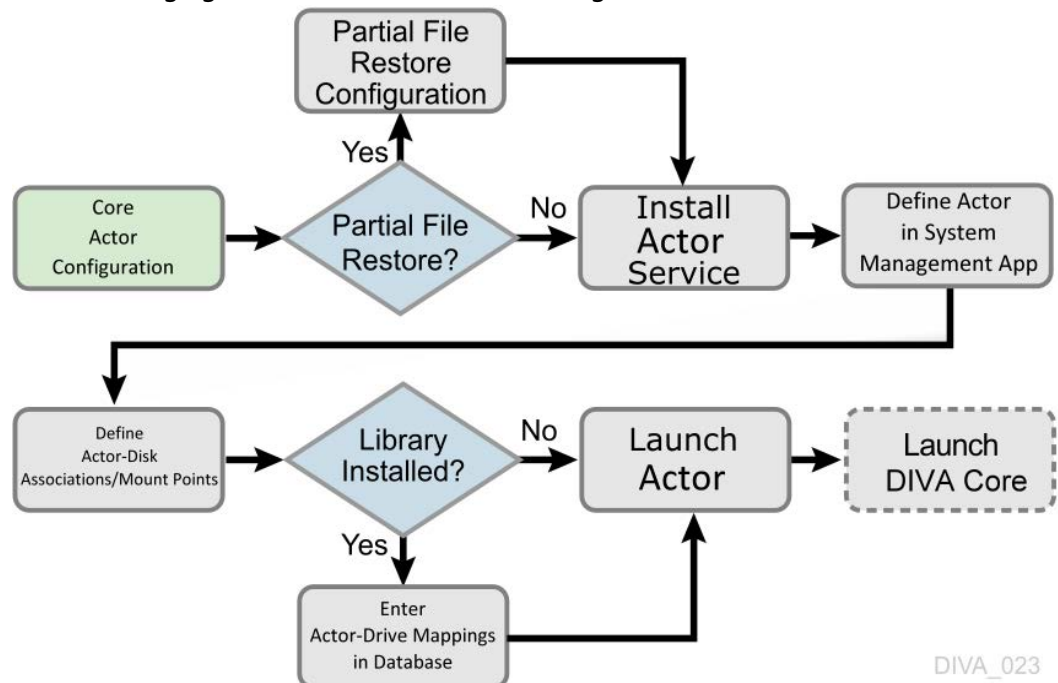
The Core Actor runs on both Windows and Linux platforms. Windows Actors no longer start automatically with Windows; the Actor runs as a standalone server application. The Core Manager connects to each Actor as a client application.

The Actor is installed in the %DIVA_HOME%\Program\Actor\bin\ folder in Windows, and in the /home/diva/DIVA/Program/actor/bin/ directory in Linux. The Actor's configuration files are located separately in the %DIVA_HOME%\Program\conf\Actor\ folder in Windows, and in the /home/diva/DIVA/Program/conf/actor/ directory in Linux. At the system level, the location and capabilities of each Core Actor are defined in the System Management App.

The Actor configuration parameters are located the System Management App, except for the Service Name and Port. These settings are located under Actor Advanced and Partial Restore Settings pages of the Actor area of the System page. Some settings are only available In Engineering Mode.

You must notify the Actors of any changes to the configuration by clicking on Notification, Notify Actors while connected to the Manager. The Actors must be running and connected to the Manager to receive the notifications.

The following figure is the workflow for installing a Core Actor:



DIVA_023

Configuring the Local Actor (actor.conf)

The Actor configuration file contains the Service Name and Port parameters. Remove the .ini extension from the actor.conf.ini file and edit the file with a plain text editor (for example, Notepad or Notepad++) to insert the Service Name and Port number as described in the following table.

Parameter	Parameter Type	Description	Default
DIVAActor_PORT	TCP Port Number	The TCP Port Number for the Actor to listen on for incoming requests. If running more than one Actor on the host, the TCP Port Number must be unique for each Actor.	9900
SERVICE_NAME	Name	The DIVAActor_SERVICE_NAME parameter specifies the name of the Actor and the service during installation. This is required if you install two or more Actors on a single Windows host computer because both cannot have the same Actor Service Name. If this parameter is not defined or commented out, the Service Name defaults to the Host Name of the Actor computer and will be DivaAct Host_Name.	

Configuring DIVA Core Partial File Restore

The Partial File Restore parameters are located on the Partial Restore Settings page in the System Management App Actor area. These options provide additional parameters to the Actor for specific partial file restore formats.

To edit the parameters, double-click the Actor Name in the Partial Restore Settings page to open the Edit Partial Restore Settings dialog box. The Partial File Restore options are defined on the Partial Restore Settings tab of the dialog box.

DIVA Core 7.5 and later MPEG2 Transport Stream supports HD MPEG video essences with AES3 audio tracks.

The following table describes the Partial File Restore parameters available on the Edit Partial Restore Settings Entry dialog box. There is a request option available as indicated in the table that can be used when creating the request.

Parameter	Value or Type	Request Option	Description	Default
Name	String		This is the name of the Actor associated with these Partial File Restore options. This value is automatically filled in from the Actor settings. If you modify the name here, or in the Actor settings screen, it will be modified in both places.	
QT Ignore Start Timecode	N (disabled) Y (enabled)	-PfrQtIgnoreStartTimecode	If this setting is enabled, Partial File Restore will ignore the SOM value of the original clip and process TCIN and TCOUT as if it starts from 00:00:00:00.	N
QT Omneon First Frame Handling	IGNORE RESET UPDATE	-PfrQtOmneonFistfrmHandling	This setting identifies how the Actor will handle the first frame of a QuickTime clip: <ul style="list-style-type: none"> • IGNORE: Partial Files Restore will ignore this field. The value found in the original clip will remain unchanged in the restored clip. • RESET: Partial File Restore will reset the value of this field to zero. • UPDATE: Partial File Restore will increment this value using the frame count from which the partially restored file begins. 	RESET

Parameter	Value or Type	Request Option	Description	Default
AVI Ignore Start Timecode	N (disabled) Y (enabled)	-PfrAvilgnoreStartTimecode	If this setting is enabled, Partial File Restore will ignore the SOM value of the original clip and process TCIN and TCOUT as if it starts from 00:00:00:00.	N
EVS MXF Ignore Start Timecode	N (disabled) Y (enabled)	-PfrEvsMxflgnStartTimecode	If this setting is enabled, Partial File Restore will ignore the SOM value of the original clip and process TCIN and TCOUT as if it starts from 00:00:00:00.	N
GXF Timecode Reference	Integer	-PfrGxfTimecodeRef	<p>This setting specifies how the time code SOM reference is to be derived for a GXF Partial File Restore request. The options are defined by the following values:</p> <ul style="list-style-type: none"> • The objects start time codes are ignored. TCIN and TCOUT must be relative to 00:00:00:00. • SOM is derived from the first field number of the MAP packet (default). • SOM is derived from the time code at Mark In from the UMF packet. 	1

Parameter	Value or Type	Request Option	Description	Default
GXF Progressive Timecode Translation	N (disabled) Y (enabled)	-PfrGxfProgTimecodeTrans	Partial File Restore is expecting TCIN and TCOUT to be in conformance with the frame rate of the archived clip by default. For example, if the frame rate of the clip is 29.97fps NTSC (or 25fps for PAL), the frame count of TCIN and TCOUT can be comprised between 0 and 29 (25 if it is PAL). HD formats have progressive frame rates (23.976, 24, 29.97, 30, 59.94, 60). For automations, the actual frame rate of the clip can be unknown. When this parameter is set to Y (enabled), DIVA Core considers that TCIN and TCOUT are PAL or NTSC timecodes and translates these timecodes according to the actual frame rate of the archived clip.	N
LXF Ignore Start Timecode	N (disabled) Y (enabled)	-PfrLxfIgnoreStartTimecode	If this setting is enabled, Partial File Restore will ignore the SOM value of the original clip and process TCIN and TCOUT as if it starts from 00:00:00:00.	N

Parameter	Value or Type	Request Option	Description	Default
MXF Partial Restore Dictionary File	Path and File Name	-PfrMxfPrDictFile	<p>This parameter must point to the name and location of the MXF dictionary file. The dictionary is normally distributed with the Core Actor installation in the %DIVA_HOME%\Program\Actor\bin folder. The default dictionary file name is mxf_file.bin.</p> <p>Set this parameter to %DIVA_HOME%\Program\Actor\bin\mxf_file.bin.</p> <p>Where %DIVA_HOME% is the root path of your DIVA Core installation for the Actor (typically C:\Diva in Windows and home/diva/DIVA in Linux).</p>	
MXF Timecode From Source Package	N (disabled) Y (enabled)	-PfrMxfTimecodeFrmSrcPkg	<p>If you set this parameter Y (enabled), the time code track used to locate the in and out points will be the one from the source package. Otherwise, timecode will be sourced from the Material Package.</p>	N

Parameter	Value or Type	Request Option	Description	Default
MXF Timecode Value To Switch Package	-1 (no switch) 0 (switch)	-PfrMxfTCValuetoSwitchPkg	If the SOM value found in the MXF package specified by the parameter MXF Timecode From Source Package is equal to this value, the Actor will automatically look for the SOM in the other MXF Package. The default value of -1 avoids switching from one package to the other.	-1
MXF Enforce Closed Header	N (disabled) Y (enabled)	-PfrMxfEnforceClosedHeader	If this parameter is set to Y (enabled) the extraction will fail if the metadata in the header is not closed. If set to N (disabled), the Actor will attempt to find closed metadata in the footer partition.	Y
MXF Run In Processor	File Name	-PfrMxfRunInProcessor	If this parameter is defined it must contain the name of the RunInProcessor.dll. In this case, the run-in processor will be used to read and create run-ins. For example: RUN_IN_PROCESSOR=RunInProcessor.dll.	

Parameter	Value or Type	Request Option	Description	Default
MXF Ignore Start Timecode	N (disabled) Y (enabled)	-PfrMxfIgnoreStartTimeCode	If this parameter is set to Y (enabled), MXF Partial File Restore will ignore all start time code values of the original clip and TCIN and TCOU (SOM and EOM) is processed as if the original clip starts at 00:00:00:00. This option overrides the MXF TIMECODE FROM SOURCE PACKAGE parameter.	N
MXF Use Omneon Dark Meta	N (disabled) Y (enabled)	-PfrMxfUseOmneonDarkMeta	Certain Omneon MXF clips have their start time code located in a Dark Metadata Set. By default the MXF Partial File Restore does not pay attention to this field. Set this parameter to Y if you want the MXF Partial File Restore to manage this field.	N
MXF Use BMX Library (instead of MOG SDK)	N (disabled) Y (enabled)	-PfrMxfUseBMXLibrary	The MOG SDK library has been replaced by BMX under Linux. Under Windows, the use of either MOG SDK or BMX can be selected from the Config Utility under Advanced Actor Settings, by setting the Use BMX Library parameter to Y. Under Linux, BMX will always be used.	N

Parameter	Value or Type	Request Option	Description	Default
MXF Serialize Depth First	N (disabled) Y (enabled)	-PfrMxfSerializeDepthFirst	If this parameter is set to Y (enabled) the MXF Partial File Restore serializes the Metadata Sets of the partially restored clip using a depth-first approach. This option is recommended when the Destination Server is a QUANTEL ISA gateway. If it is set to N (disabled), the MXF Partial File Restore serializes the Metadata Sets with no ordering.	N
MXF Generate Random Index Pack	N (disabled) Y (enabled)	-PfrMxfGenerateRip	RIP (Random Index Pack) is an optional small structure located after an MXF file that contains file offset information for each partition in the file (when present). You can set this parameter to N (disabled), for incompatible servers (for example, SONY XDCAM).	Y

Parameter	Value or Type	Request Option	Description	Default
MXF Number of Frames Per Body Partition	Integer between 50 and 250.	- PfrMxfFramesPerBodyPartition	This parameter defines the number of frames per partition in the output file. Only values between 50 and 250 are valid. If a value greater than 250 is entered, the MXF Partial File Restore will use 250. If the entered value is less than 50, it will use 50. This parameter is rounded automatically by the Actor to align body partitions on GOP boundaries.	250
MXF Update TC Track Origin	N (disabled) Y (enabled)	-PfrMxfUpdateTctrackOrigin	When the video essence is MPEG2 LGOP, Partial File Restore will use the origin field of each track to be frame accurate. The origin specifies GOP precharge frames. Your video server may use a different implementation or interpretation of this field. If this parameter is set to Y (enabled), the Origin field is modified in all tracks. If this parameter is set to N (disabled), the Origin field is modified in all tracks except the timecode track.	N

Parameter	Value or Type	Request Option	Description	Default
MXF Tolerance on TCOU	Integer between 0 and 250.	-PfrMxfTcoutTolerance	This parameter can be set to indicate a tolerance on the TCOU supplied to a Partial File Restore request. This tolerance value is 0 by default, but you can set it to a specific number of frames. If the supplied TCOU is beyond the end of the clip, but not too far out (within the tolerance), DIVA Core will perform the Partial File Restore until the end of the clip instead of reporting and invalid TCOU.	0
MXF Duration From Footer	N (disabled) Y (enabled)	-PfrMxfDurationFromFooter	When the duration of the input clip is -1 in the header partition, the MXF Partial File Restore loads the footer partition in to obtain the correct value. Some older clips may not have a correct RIP after the file, and the footer partition may not be accessible. If you set this value to N (disabled), the MXF Partial File Restore does not load the footer partition and performs a blind Partial File Restore, if TCIN and TCOU are valid.	Y

Parameter	Value or Type	Request Option	Description	Default
MXF Maximum Queue Size	Integer between 0 and 200.	-PfrMxfMaxQueueSize	The maximum size (in MB) that the extractor can queue before producing an error (to avoid running out of memory).	200
Seachange Ignore Start Timecode	N (disabled) Y (enabled)	-PfrSealgnoreStartTimeCode	If you set this parameter to Y (enabled), SeaChange Partial File Restore ignores the start time code value of the original clip and processes TCIN and TCOU as if it starts from 00:00:00:00. The configuration of the MXF parser is also required for MXF. However, because this is a SeaChange clip, it ignores the MXF Ignore Start Timecode in this workflow.	N
MPEG2 Transport Stream Ignore Start Timecode	N (disabled) Y (enabled)	-PfrTslgnoreStartTimeCode	If you set this parameter to Y (enabled), the MPEG2 transport stream Partial File Restore ignores the start time code value of the original clip, and processes TCIN and TCOU as if it starts from 00:00:00:00.	N
MPEG2 Program Stream Ignore Start Timecode	N (disabled) Y (enabled)	-PfrPSlgnoreStartTimeCode	If you set this parameter to Y (enabled), MPEG2 transport stream Partial File Restore ignores the start timecode value of the original clip and processes TCIN and TCOU as if it starts from 00:00:00:00.	N

Defining and Declaring Actors

Each Core Actor must be declared in the Core Database. You declare the Actors in the Actors area in the System Management App. The Actors area has three tabs:

Actor Settings

This tab includes general Actor definition settings such as Actor name, IP address, port, Network, and so on.

Actor Advanced Settings

This tab includes advanced settings such as read and write block sizes, tape unit timeout, Quantel, QuickTime and FTP settings.

Partial Restore Settings

This tab includes Partial File Restore settings previously in the Partial File Restore configuration file.

Actor and Partial File Restore settings are configured and edited on the Actor Settings Entry screen. Click + on the top right of the Actor Settings area to create and configure an Actor, or double-click the Actor you want to edit to access the settings screen.

The following list describes the maximum operations parameters on the Actor Settings Entry screen.

Name

This is the name of the Actor associated with the Partial File Restore options. This value is automatically filled in from the Actor settings. If you modify the name here, or in the Actor Settings Screen, it will be modified in both places.

IP Address

This is the IP address of the Actor.

Port

This is the port number the Actor listens on for commands.

Prod. System

This parameter identifies the Network where the Actor is in use.

Site

This parameter identifies the physical location of the Network.

Max Drive Operations

This is the maximum number of simultaneous requests to and from drives that this Actor can perform. You can use this parameter to distribute requests and bandwidth among all Actors.

Max Server Operations

This is the maximum number of simultaneous requests to and from servers from the Servers configuration that this Actor can perform. You can use this parameter to distribute requests and bandwidth among all Actors.

Max Disk Operations

This is the maximum number of simultaneous transfers to and from disks (both read and write) that this Actor can perform. You can use this parameter to distribute requests and bandwidth among all Actors.

Max Stage Operations

This is the maximum number of staging request that an Actor is allowed to run at the same time.

Max Bridge Operations

This is the maximum number of concurrent requests using DIVA Bridge that an Actor is allowed to run at the same time.

Verify Tape

This parameter defines whether tapes are verified.

Direct Restore

This parameter defines whether this Actor can be used for direct restores to a Source or Destination Server.

Cache Restore

The Actor is permitted to perform cache restores to a Source or Destination Server. You must disable this option if this Actor has no local cache storage for the temporary storage of the DIVA Core Object during a transfer.

Copy To Tape Group

This parameter defines whether this Actor can be used for Copy To Tape Group requests. You can use this option to isolate specific Actors involved in critical operations from mass Copy To Tape Group requests, such as those from the DIVA Core SPM option.

Associative Copy

This parameter defines whether this Actor can be used for Associative Copy requests.

Repack

This parameter defines whether this Actor can be used for tape repack requests. You must set this to N if the Actor has no local cache for temporary storage during the repack operation. Because tape repacking is a lengthy operation, you can also use this setting to dedicate an Actor solely to repack requests by disabling the other options (except Delete) and disabling repack on the other Actors.

Delete

This parameter defines whether this Actor can be used for requests that involve deleting DIVA Core Objects from a disk. You can use this option to isolate an Actor from mass deletion requests (for example, requests issued from the SPM option).

Direct Archive

This parameter defines whether this Actor can be used for direct Archive requests.

Cache Archive

This parameter defines whether this Actor can be used for cache Archive requests. You must disable this option if this Actor has no local cache storage for the temporary storage of the DIVA Core Object during a transfer.

First Utilization Date

This is the date the Actor was first put into use.

Advanced Actor Settings

Advanced Actor parameters are displayed, configured and edited on the Actor Advanced Setting page in the Actors Panel of the System Management App. To configure or edit advanced Actor parameters, double-click the Actor you want to edit to access the settings screen.

The following list describes the parameters on the Actor Advanced Settings Entry screen:

Name

This is the name of the Actor associated with the Partial File Restore options. This value is automatically filled in from the Actor settings. If you modify the name here, or in the Actor Settings Screen, it will be modified in both places.

Tape Test Unit Ready Timeout (s)

The time in seconds to wait for a drive to become ready after a tape is mounted. If the drive is not ready within this period, the drive is considered to be not responding.

Linux SMB Mount Point

The root path for a Linux-based Actor to create a mount point to SMB shares. The default value is /mnt. When a Linux-based Actor connects to a CIFS Unmanaged Storage Repository Server, it will mount the share to a directory within the specified mount point.

For example, if the root path is /mnt, a Linux-based Actor connecting to a CIFS share of \\hostname\share_folder, will result in the share being accessible from the /mnt/hostname/share_folder.

Profile Read Block Size (B)

The FTP block size used for transfers on profile video servers when reading. The default value (1500) is the best block size to use with GVG profile servers. This value

may be different when using other servers. Possible values are between 1500 and 262,144 bytes.

Profile Write Block Size (B)

The FTP block size used for transfers on profile video servers when writing. The default value (32,768) is the best block size to use with GVG profile servers. This value may be different when using other servers. Possible values are between 1500 and 262,144 bytes.

Quantel Rename Clips

Automatically rename clips when restoring them to Quantel.

- Setting this to N disables this feature. This is the default setting.
- Setting this to Y renames files using the first part of the object name (before the comma) truncated. This is Omnibus renaming.

QT Self-contained Threshold (MB)

When performing a QuickTime Partial File Restore, the Actor must determine if a clip is self-contained, or not based on the size of the input file. This parameter is a limit in MB. When this limit is exceeded, the Actor considers the clip to be self-contained. The unique objective of this parameter is to prevent the Actor from loading a large self-contained clip into memory. Values range from 10 MB through 100 MB.

Disk FTP Passive Mode

FTP data connections are, by default, created in Active mode. The DivaFTP client connects from a random unprivileged port (greater than port 1023). Then it immediately starts listening to the port and sends a PORT command to the FTP server.

When you set this parameter to Y, data connections are created in Passive mode rather than Active mode. In Passive mode the DivaFTP client sends a PASV command to the FTP server and the server creates socket, not the client.

Disk FTP Block Size (KB)

This parameter defines how much data the Actor attempts to send and receive using a single system call during FTP transfers.

For example, if the Actor internal buffer size is set to 2 MB, and this parameter is set to 32768 bytes, 64 system calls are required to write a single buffer to a data socket.

Disk FTP Socket Window Size (B)

This parameter adjusts the normal buffer size allocated for output and input buffers. This parameter is internally used to set the send and receive buffers for FTP-managed disk types.

Configuring Actor to Drive Connections

The Data Transfer component of the drives must be configured for use with the Actors separate from the Tape Drive Control configuration for the Robot Manager. You must logically configure of each drive in the Actor-Drive configuration in the database.

The Actors-Drives area is located on the Drives page. The area displays the current Actor-Drive associations including the Actor Name, Drive Number, and Library location. If a drive is connected to multiple Actors through a SAN, the Actor-Drive mapping must be repeated for each Actor accessing this drive.

You can combine the Drive Operations settings and the Actor Capability settings to dedicate a drive to a particular set of Actors for specific operations. For example, tape repacking.

To edit the parameters, double-click the Actor Name in the Actors-Drives area to open the Add new row in Actors-Drives Connections dialog box. Click the + button on the top of the area to add a Actors-Drives connection.

Two options are available on the Add new row in Actors-Drives Connections dialog box as follows:

Actor

Select the Actor the drive is connected to from the list. Only Actors already defined in the Actors area of the System page are listed.

Drives

Select the logical drive in the relevant library for this mapping. Only drives defined in the Drives area of the Drives page are listed. You can select one or more drives using the check boxes. Multiple selections are only available when adding an association, not while editing an existing one.

When you select a different Actor, the drives available for configuration are displayed. If all drives have already been configured for the selected Actor, the Drives list is not available and indicates there are no drives available for the selected Actor.

Defining Core Proxy Actors

Note: This feature is only supported for disk and Server based requests.

The user must first define an Actor with a UDP port to configure a Proxy Actor. The UDP port allows a regular Actor to message a Proxy Actor using the connection-less protocol. In the following figure Actor diva8024_actor1_9901 is configured as a Proxy Actor with UDP port 10001. The TCP port is irrelevant for a Proxy Actor.

You must configure the link between the Actor and Proxy Actor to notify DIVA Core that this Actor is a Proxy by adding an Data-Proxy Actor Connection.

After configuration, DIVA Core is now aware that Actor `diva8024_actor0_9900` can see Proxy `diva8024_actor1_9901`. This means that any remote resources only visible to the Proxy Actor can now be accessed using the regular Actor.

The Actor configuration file corresponding to the proxy must also be updated with the UDP port. In this example, the Actor configuration file for `diva8024_actor1_9901` (the Proxy Actor) only requires a UDP port.

```
DIVAActor_PORT=UDP/10001
```

If you want to specify both a TCP and UDP port, then you must use `DIVAActor_PORT2` as shown here:

```
DIVAActor_PORT=9901
```

```
DIVAActor_PORT2=UDP/10001
```

You can now configure a remote disk that is not connected to a regular Actor and still archive to that disk if a Proxy Actor is connected to that disk.

Note: The Manager does not directly connect to a Proxy Actor. It can only directly communicate with a regular Actor. A Proxy Actor exclusively communicates with a regular Actor.

Resource Selection and Manager-Actor Communication

The Manager selects what regular Actor to use to satisfy a request based on the resources that Actor can directly or indirectly (via a proxy) access. If multiple proxies are configured for a single Actor, the decision of which proxy to use is based primarily on the load on that Actor.

The Manager does NOT directly connect to a proxy. It can only directly communicate with a regular Actor. A proxy exclusively communicates with a regular Actor.

Cloning Actors and Tapes

In addition to configuring Clone Tape Groups, Actors and Source Tapes must be enabled for cloning. By default, all Source Tapes are enabled for cloning. However, a Source Tape will be disabled for automatic cloning if a read failure occurs during a clone request. The user will have to manually re-enable the Source Tape for automatic cloning by setting the corresponding Tape State in the System Management App.

If a write error occurs during a clone request, the Source Tape is unaffected and can still be used for writing content. If the Clone Tape is bad and cannot be used, the existing clone link must be removed, and then either manually invoke the clone or use the automated clone scheduler to invoke it. On invocation, the clone request will select a new tape from the Clone Tape Group.

See the DIVA Core Operations Guide for details on tape selection, manual cloning, and automatic cloning processes.

Logging Actor Activity

Core Actors log all activities during normal operations. The log files are named `actor.log`, or `actor_SERVICE_NAME.log`. The files are stored in the `%DIVA_HOME%\Program\log\actor` folder.

Each Core Actor also provides additional logging functions for some specific server protocols (for example, the Quantel QCP interface, FTP servers, and Partial File Restore). Core enables logs by default, and they are unique for each server type. They provide detailed logging information from that protocol to the standard Actor log file.

These files are useful in diagnosing transfer errors with either drives or servers, and particularly for debugging the configuration when a Source or Destination Server has been added. Technical Support may request these logs when providing assistance.

Installing and Uninstalling Actor Services in Windows

You can use the `actorservice.exe` executable in the Actor bin directory to install (or uninstall) the Core Actor as a service from a Windows command-line prompt.

By default, the Actor Service uses the `actor.conf` file located in `%DIVA_HOME%\Program\conf\actor` folder to define the Service Name. If you are installing multiple Actors on a single host, you must create additional Actor configuration files and specify them to the service to create unique instances for each Actor (see [Actor Service Management Functions](#) for more information).

See [Appendix A: Core Options and Licensing](#) for detailed information.

Use the following commands to install or uninstall the Actor Service from the Windows command line:

actorservice -i

Installs the Actor Service using the `SERVICE_NAME` parameter defined in `actor.conf`. If this parameter is undefined, then the service is installed as Core Actor - Host_Name.

actorservice-u

Removes the Actor Service using the `SERVICE_NAME` parameter defined in `actor.conf`. If this parameter is undefined, then the service to be removed is Core Actor - Host_Name.

Installing and Uninstalling Actor Services in Linux

The `divaservice` executable in the Actor bin directory installs (or uninstalls) the Core Actor as a service from a Linux terminal.

Use the following command sequence to install Actor services:

```
cd/home/diva/DIVA/Program
```

```
./divaservice install actor /home/diva/DIVA/Program/conf/actor/  
actor.conf
```

Use the following command sequence to uninstall Actor services:

```
cd/home/diva/DIVA/Program
```

```
./divaservice uninstall actor /home/diva/DIVA/Program/conf/actor/  
actor.conf
```

See [Installing the DIVA Core Services](#) for more information on using the `divaservice` command.

Actor Service Management Functions

When installing or uninstalling additional Actor Services on the same host, you must specify the path to each Actor's configuration file for each instance. You add the `-conf` (or `-f`) command switches when installing the service as follows:

```
actorservice {-i|-u} {-conf|-f} {Path and file name}
```

The command syntax is the same for Windows and Linux. However the path and file name will be different. The following examples install the Actor services for two different Actors on the same host computer. You use the `-u` command switch (instead of `-i` to install) to uninstall these same Actor services.

Check the services applet after installation to verify that each Actor Service was installed correctly.

For example, use the following command in Windows to install the Actor defined by the `SERVICE_NAME` in the `actor1.conf` configuration file:

```
actorservice -i -conf C:\DIVA\Program\conf\actor\actor1.conf
```

Use the following command in Windows to install the Actor defined by the `SERVICE_NAME` in the `actor2.conf` configuration file:

```
actorservice -i -conf C:\DIVA\Program\conf\actor\actor2.conf
```

Use the following command in Linux to install the Actor defined by the `SERVICE_NAME` in the `actor1.conf` configuration file:

```
actorservice -i -conf ../../conf/actor/actor1.conf
```

Use the following command in Linux to install the Actor defined by the `SERVICE_NAME` in the `actor2.conf` configuration file:

```
actorservice -i -conf ../../conf/actor/actor2.conf
```

The following additional command options are also available for the Actor Service:

actorservice debug

Starts the Actor Service in console mode. This is used for troubleshooting.

actorservice version

Displays the Core Actor Service software release information. You can also use the "-v" switch instead of "version".

actorservice help

- Displays all command line options.

Launching the Actors

Windows Core Actors no longer start automatically with Windows. You can manage the Actor Services through the Windows Services applet, from a Windows command line, or from Linux terminal.

In Windows, you can locate the Actor Service in the Windows Services applet, right-click the name, and then select the desired management function (Start, Stop, Restart, and so on) from the context menu.

Note: The quotation marks in the following commands must be used when specifying a Windows service with spaces in the name.

You can restart an Actor from a Windows command line or Linux terminal using the following command sequence:

```
net stop "Core Actor"  
net start "Core Actor"
```

If a SERVICE_NAME is specified in the actor.conf file (for multiple Actors on a single computer), then you can restart an Actor from a Windows command line or Linux terminal using the following command sequence:

```
net stop "Core Actor -SERVICE_NAME"  
net start "Core Actor -SERVICE_NAME"
```

Tip: Create a Windows batch file containing these commands and place it on the desktop for easy access.

Robot Manager Configuration

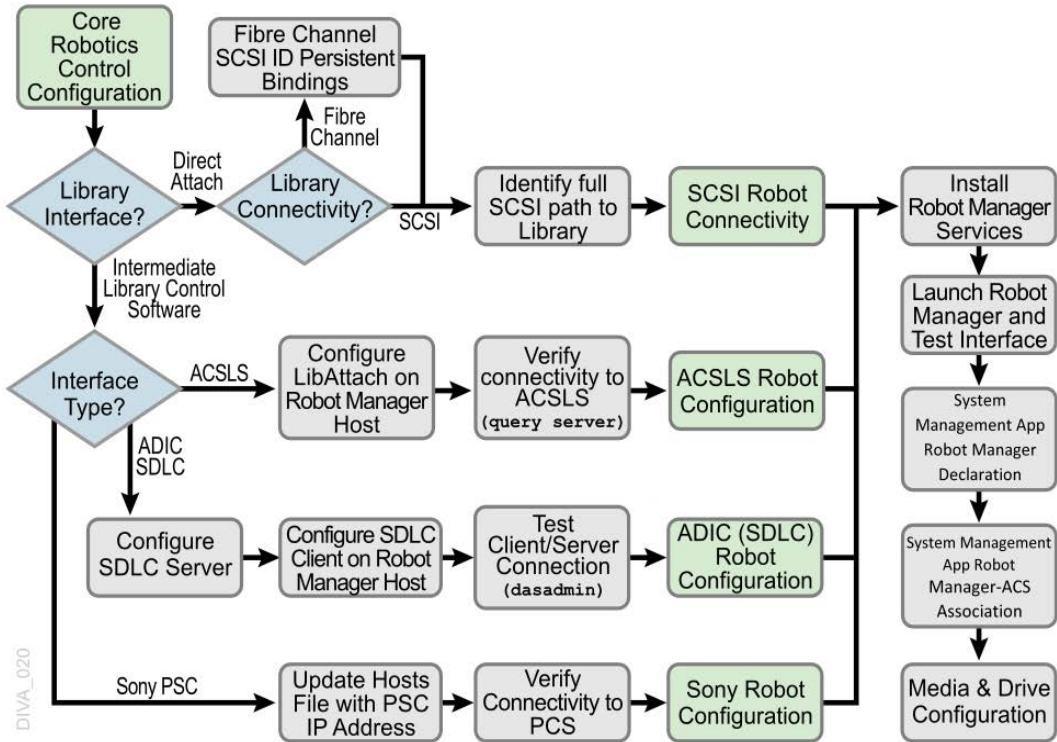
Configuration Overview

The Core Robot Manager on Windows platforms runs as a Windows service and is launched automatically with Windows. See [Appendix A: Core Options and Licensing](#) for detailed information.

You configure the type of interface a specific library in a static configuration file. The file name is robotmanager.conf and is located in the DIVA_HOME\Program\conf\robot_manager folder on the computer where the Core Robot Manager is installed. In a new installation (or upgrade) the file is provided with a .ini extension. You must copy the file, remove the .ini extension, and then edit the new file.

Since many different types of Managed Storage and connections are supported, not all sections of the configuration file will be relevant to your particular installation. Also, some parameters are specific to the operating system where the Robot Manager is installed. Therefore, some settings in the configuration file are initially commented out (that is, they have # in front of the parameter). This indicates to the Robot Manager to ignore the setting. For the setting to be taken into account the # must be removed.

The following figure outlines the steps for configuring the robotics to be controlled by DIVA Core:



SCSI Connected Managed Storage

For directly attached SCSI controlled Managed Storage, you must configure and correctly identify the SCSI ID controlling the library, and enter this value into the `RM_SCSI_DEVICE_LSM` parameter in the Robot Manager configuration file. Before changing the configuration, you must understand several concepts as described in the following sections.

The `robotmanager.conf` configuration file includes the following main parameters:

RM_SCSI_MOVEMEDIUM_TIMEOUT

Robot SCSI uses the `MOVE MEDIUM` SCSI command during mount, dismount, enter, and eject requests. The value of `RM_SCSI_MOVEMEDIUM_TIMEOUT` is indicated in minutes, and the default timeout is fifteen minutes for the communication between the library and the robot manager.

Some Managed Storage, like Spectra T950, may require more time to be able to complete a `MOVE MEDIUM` request and you should set this parameter value accordingly.

RM_SCSI_EJECT_USEGLOBALLOCK

You must set this parameter to one if you want the SCSI Robot Manager Eject calls obtain the lock number of the LSM and hold that lock until all associated tapes to be ejected have completed the ejection process. When all tape ejections are complete, the call unlocks the drive and proceeds on to the next drive. The default setting is zero.

Fiber Channel HBA (Host Bus Adapter) and SCSI Persistent Binding

Most installations use FC (Fiber Channel) rather than native SCSI to interface to the library (typically over a SAN). In these instances, the FC HBA in the Core Robot Manager host presents the World Wide Name of the library interface as a SCSI ID. By default, most HBAs automatically map these to a SCSI ID for the host operating system to access. This presents a problem if a device is added or removed on the SAN because it could alter the SCSI ID of the library by the HBA, and automatically remap the existing devices. Disable the Automap feature to avoid this issue and use Persistent Bindings instead. This feature allows the SCSI mapping of the library to remain consistent between host restarts, and from the advent of any addition or removal of devices on the SAN.

If the library controller or the HBA in the Core Robot Manager host is changed, this might alter the library's SCSI Persistent Bindings to the host operating system. This requires the Persistent Binding for the library to be reconfigured in the HBA configuration software on the Core Robot Manager computer.

Determining the SCSI Library Connection

For the SCSI interface Managed Storage the Core Robot Manager communicates with the library directly over the SCSI hardware layer and does not require the Windows driver interface.

For all other Managed Storage it is essential that no library driver be loaded for the library interface. If a driver is loaded, the Core Robot Manager will be unable to communicate with the library. In this case, if your library does not appear in Windows Device Manager as an Unknown Medium Changer, the Robot Manager will be unable to communicate correctly with the robotics.

If you cannot locate a specific library in the Scandrive Utility (see the following), but that library is visible in your HBA, then the library has likely been disabled in the Windows Device Manager (denoted by an X over the device icon). You must re-enable the device for it to appear in the Scandrive Utility.

For Windows, you can determine the RM_SCSCI_DEVICE_LSM(n) settings for the Core Robot Manager using the scandrive.exe utility. The utility is located in the %DIVA_HOME%\Actor\bin directory. The utility automatically reports all devices located in the Windows SCSI hardware tree in the registry and their corresponding Port, Bus, Target, and LUN (Logical Unit Numbers).

```

C:\Diva\Program\Actor\bin\scandrive.exe
Scsi5:0:0:0
  Type : MediumChangerPeripheral
  Identifier : STK    SL500    1026

Scsi5:0:1:0
  Type : TapePeripheral
  Identifier : IBM    Ultrium-TD2  4770
  DeviceName : Tape0
  Status : OK -> The media may have changed.

Scsi5:0:2:0
  Type : TapePeripheral
  Identifier : IBM    Ultrium-TD2  4770
  DeviceName : Tape0
  Status : KO -> No media in drive.

Scsi5:0:3:0
  Type : TapePeripheral
  Identifier : IBM    Ultrium-TD2  4770
  DeviceName : Tape0
  Status : KO -> No media in drive.
  
```

The utility reports the SCSI Device ID of the library in the format ScsiP:B:T:L (see the previous figure), where P is the port number, B is the bus number, T is the target number, and L is the Logical Unit Number.

The Type section of the utility's output refers to that peripheral's class (HDD, CDROM, and so on). A tape library will be reported as a Medium Changer Peripheral, and the Identifier for each corresponding device reported should match the model number of the library itself (for example, SL500). You can then enter the full SCSI path reported for each library into the RM_SCSCI_DEVICE_LSM(n) settings in the robotmanager.conf file.

Sony ODA Drives

DIVA Core supports the Sony new generation of ODA drives; the ODS-280F (Fiber Channel) and ODS-280U (USB). DIVA Core has only been tested with the Fiber Channel type. The drives are twice as fast as the Gen1 drives. The ODS-280U has not been qualified for use with DIVA Core.

A new cartridge type is also available for this drive, the ODC3300R. This is a WORM drive with a 3.3 TB capacity.

Gen2 drives can read content written on Gen1 media with Gen1 drives. DIVA Core does not support the READ-ONLY media-drive compatibility. Technical Support recommends isolating Gen1 media from Gen2 media in the configuration (because there is no cross-generation compatibility) and there must be at least one Gen1 drive in a library containing Gen1 cartridges.

DIVA Core supports Sony ODA ODS-D55U and ODS-D77F drives only in the Windows environment. These are Blu-ray Optical Drives and the media is WORM media using a UDF format. Only AXF formatted objects can be written to the discs. The drives are controlled by the Robot Manager and the media is viewed as a Tape Cartridge.

In the Windows Device Manager these drives will be shown as Unknown Medium Changer under the Medium Changer section because there are no device drivers for them. The drive itself will also appear as an Optical SCSI Device with the make and model number under the Disk Drives section.

Sony ODA Gen 3 is supported. The drive type is ODS-D380F and uses the following cartridge:

Cartridge Type

ODC5500R

Capacity

5.5 TB

Block Size

64 KB

Drive Type

WORM

Note: The drive is still R/W compatible with ODC3300R and read-only compatible with older cartridge types.

There are seven different types of disc media available for use with the Sony Optical Drives as follows:

SONY-ODC300R

293,265,408 KB capacity

SONY-ODC300RE

293,265,408 KB capacity

SONY-ODC600R

586,530,816 KB capacity

SONY-ODC600RE

586,530,816 KB capacity

SONY-ODC1200RE

1,173,086,208 KB capacity

SONY-ODC1500R

1,500,020,736 KB capacity

SONY-ODC3300R

3,222,717,696 KB capacity

SONY-ODS-D380F

5,372,184,576 KB capacity

The disc types are identified in the `scsi_tape_types.ini` file (described in the following section).

Note: You must configure the drive settings before configuring DIVA Core. The recommended parity setting is PARITY ON.

You can view the drive specifics using the Optical Disc Archive Utility. This utility enables viewing of device logs, and viewing and changing drive settings.

To change the drive settings, click the Setting page in the Optical Disc Archive Utility. Technical Support recommends leaving the Default Volume Type set to PARITY ON, and to use the default settings for the remaining items.

Click the Media item under the Drive navigation tree to view information about the media in a drive.

You click the Write-protect button to write-protect a drive. Once an Optical Disc is write-protected, you can no longer write objects to the device. However they are still retrievable.

Configuration File Adjustments

You must change several parameters in the `scsi_drive_types.ini` configuration file to use these optical drives.

In the `robotmanager.conf` configuration file, under the SCSI module specific options, the serial number must be identified. You can find the serial number in the `RM_SCSI_DEVICE_LSM(n)` parameter line. For example, `RM_SCSI_DEVICE_LSM(0)=00001003`, where (0) is the LSM number, and 00001003 is the serial number. You must identify the serial number for all listed devices (LSM(0), LSM(1), LSM(2), and so on).

In the `scsi_drive_types.ini` file, the drive types must be uncommented (remove the #). For example, remove the # from in front of the line that reads `#601 0x00 0x00 SONY-ODS-D77F 600 601 602 603 604 605` to use your D77F drive as shown. The `TransportDomain` and `TransportType` are obtained automatically and not used in the configuration, so you must leave these set to `0x00` as shown in this example.

```
#-----
# If the SCSI Robot Manager is connected to a SONY ODA library
# UNCOMMENT ALL LINES IN THE FOLLOWING PART
#-----
#TypeID TransportDomain TransportType TypeName CompatibleTapeTypes
#-----
--
#600 0x00 0x00 SONY-ODS-D55U 600 601 602 603 604 605
601 0x00 0x00 SONY-ODS-D77F 600 601 602 603 604 605
```

Also, in the `scsi_tape_types.ini` file, uncomment all of the disc types listed as shown in the following example. The R or RE after the disc number indicates whether the disc is Write Once (R) or Rewritable (RE). This indicator is used because the barcode does not contain the video type as in normal tape barcodes.

```
#-----
# If the SCSI Robot Manager is connected to SONY ODA library,
# UNCOMMENT ALL LINES IN THE FOLLOWING PART
#-----
#TypeID TransportDomain TransportType TypeName
CompatibleDriveTypes
#-----
600 0x00 0x00 SONY-ODC300R 600 601
601 0x00 0x00 SONY-ODC300RE 600 601
602 0x00 0x00 SONY-ODC600R 600 601
603 0x00 0x00 SONY-ODC600RE 600 601
604 0x00 0x00 SONY-ODC1200RE 600 601
605 0x00 0x00 SONY-ODC1500R 600 601
```

System Management App Settings and Information

You must configure the following settings in the DIVA Core System Management App:

Drives Page

Set the Drive Properties to 64 KB. The serial number comes from the Robot Manager and the firmware release number comes from the drive.

Tapes Page

The Tape Properties area displays all of the enabled Tape Types from the `scsi_tape_types.ini` file.

System Management App Settings and Information

The Optical Drives and Discs are displayed in the DIVA Core System Management App on the Drives page as Tape Drives and Tapes respectively.

Repack of the discs and deletion of objects is available. However, the space is not recoverable. When trying to repack the disc, the normal Repack dialog box is displayed, but there is a warning that the space is non-recoverable. Due to this limitation of the discs, auto-repack has been disabled for these drives and discs.

Additional Information

Additional information related to the use of the Optical Drives and Discs includes the following:

- Because Write-Once media must be finalized, zero remaining space will be reported to the Manager.
- Objects are spanned when there is 100 MB of space remaining. This is so that there is space left for the disc to be finalized. Once an object is spanned, the disc is considered full and is automatically finalized.
- The Actor will auto-finalize the discs when there is 500 MB of space remaining unless an object was spanned. However you can manually finalize the disc through the Optical Disc Archive Utility.
- If a drive is manually mounted and viewed in the Windows Explorer, the display will show the individual files on the disc. Each file name will begin with a numeric value at the beginning that identifies the object's location on the tape.

Configuring Direct Attached SCSI Managed Storage

A Direct Attached Library is directly connected to the Core Robot Manager host computer either through a native SCSI interface and SCSI HBA, or through a SCSI over Fiber Channel connection and Fiber Channel HBA (either directly or through a SAN).

In either case, the Core Robot Manager uses its own DIVA Core provided driver (SCSI_Robot.dll in Windows or libSCSI_Robot.so in Linux) to directly interface with the library without the need for intermediate library management software. For this type of SCSI attached library, you must uncomment the entries (in the following sections) and configure them in the robotmanager.conf file. Library Drive Models and Tape Types parameters are located in other configuration files.

Common Settings for SCSI-based Managed Storage

The following are typical settings for the SCSI-based Managed Storage:

Robot Manager Common Options

Uncomment only the `RM_MODULE=SCSI_Robot.dll` in the Windows environment.

SCSI Device Parameters

The following table identifies common SCSI device parameters.

Parameter	Parameter Type	Description	Default
SERVICE_NAME	Name	The display name of the Robot Manager Windows Service. You must set this variable if multiple Robot Managers are installed on the same server. If this variable is used, the Service Name will be DivaRbt-SERVICE_NAME. If this variable is not set, the Service Name will revert to just DivaRbt.	Uncommented
RM_PORT	TCP port number	The TCP port that the Core Robot Manager listens on for incoming requests. This value must be unique if there are multiple Core Robot Managers running on a single host computer. This is typically, TCP port 8500 and greater.	8500
RM_ACS	Number	The ACS (Automated Cartridge System) controlled by the Core Robot Manager module. This value will appear in the Robot Manager/ACS Association List in the System Management App for this Robot Manager after database synchronization	0

SCSI Module Parameters

The following table identifies common SCSI module parameters. See [Determining the SCSI Library Connection](#) for parameter details.

Module Parameter	Operating System	Description	Values
RM_SCSI_DEVICE_LSM0	Windows	This specifies the SCSI target of the library as it identified by the host operating system.	ScsiP:B:T:L
RM_SCSI_DEVICE_LSM1	Windows	This setting is specific to a StorageTek dual L1400M library with a PTP (Pass Through Port), and specifies the SCSI target of the 2nd frame (LSM). Although this type of library configuration can be addressed using only the RM_SCSI_DEVICE_LSM0 connection, DIVA Core manages this type of library more effectively when both frames are specified. DIVA Core also manages the PTP in this case.	ScsiP:B:T:L

Additional Settings for Media Type Detection

The following table identifies an additional parameter that can be set to enable media type detection from the barcode.

Parameter	Parameter Type	Description
RM_SCSI_ENABLE_MEDIA_TYPE_DETECTION_LAYOUT	String pattern	The purpose of this parameter is to detect the type of a media from the volume tag returned by the library. The layout is a string of 8-10 characters indicating where the label and the mediatype are. It must contain these three characters only: L: The character at this position is part of the tape label/barcode considered into Core Database. T: The character at this position will be used for media type detection X: The character at this position will be ignored Example: for a given volume tag ABC003L6, if the layout is set to LLLLLLTT, RobotManager will detect an LTO6 tape and report ABC003 to DIVA Core.

Configuring ACSLS Attached Managed Storage

DIVA Core can directly interface to most Oracle StorageTek Managed Storage using the Robot_SCSI driver. Some library configurations require the use of the Oracle StorageTek ACSLS library management software for the Robot Manager to control the library.

You can only install ACSLS (Automated Cartridge System Library Software) on Solaris platforms. The Solaris host and ACSLS are sold and supported by Oracle. See the Oracle ACSLS documentation at <http://docs.oracle.com> for detailed information.

Technical Support does not support DIVA Core installations under the Solaris operating system.

Configuring LibAttach

LibAttach is an intermediate Windows driver providing connectivity to the ACSLS host. LibAttach runs as a Windows service and is typically installed on the same computer running the Core Robot Manager. The Core Robot Manager communicates to the ACSLS host using the LibAttach driver.

You must enter the following settings on the LibAttach Configurator dialog box (part of the ACSLS software):

Library server host name

Host name or IP address of the ACSLS server. If you use a host name, it must be resolvable by the Core Robot Manager host.

Firewall support

These settings are only required if a firewall is installed between the Robot Manager host and the ACSLS server. If no firewall is present leave these parameters set to 0.

Testing the LibAttach Connectivity to ACSLS

You can verify connectivity from the Robot Manager host to the ACSLS server with the `query_server.exe` utility located in the LibAttach installation directory. When you launch the utility a Windows command prompt opens. Statistics from the library will be returned if the connection is successful.

Firewall Support

You must have a TCP or UDP port open (to allow communication) if there is a network firewall between your Robot Manager host and ACSLS server. If there is a firewall, enter the open port numbers into the Firewall Support settings in the LibAttach Configurator.

Early implementation of firewall support for LibAttach did not work correctly with the Core Robot Manager, even though the `query_server` utility returned a successful connection. Ensure that you have the latest release of LibAttach that incorporates the patch released to address this issue. Contact Technical Support for additional information.

robotmanager.conf Common Options

The following table identifies common robotmanager.conf options:

Parameter	Parameter Type	Description	Default
RM_MODULE=ACSL5_Robot.dll		Uncomment only this line	Commented
RM_PORT	TCP Port Number	The TCP Port the Robot Manager will listen on for incoming requests. This value must be unique if there are multiple Robot Managers running on a single host. The assigned port is typically TCP Port 8500 and higher.	8500
RM_ACS	Number	ACSL5 configurations ignore this value because the ACS number is supplied from ACSLS.	Ignored
SERVICE_NAME	Name	This is the display name of the Robot Manager Windows service. This variable must be set if multiple Robot Managers are installed on the same server. If this variable is used, the Service Name will be DivaRbt-SERVICE_NAME. The Service Name will revert to DivaRbtif this variable is not set.	Uncommented

The following table identifies the ACSLS parameters:

Parameter	Parameter Type	Description	Default
RM_ACSLS_SERVER	IP Address or Host Name	ACSLs ignores this parameter and it can be left blank.	
RM_ACSLS_SSI_SOCKET	TCP Port Number	ACSLs SSI socket is the UNIX domain socket used by SSI. If this value is left undefined, it defaults to TCP port 50004.	50004
RM_ACSLS_TIMEOUT	Time in milliseconds	This sets the timeout period for queries to ACSLS through LibAttach. If you leave this value set to 0, the timeout period used by the Robot Manager is 10 minutes. If you must alter this timeout period, replace 0 with your own value (in milliseconds).	0
RM_ACSLS_IE_TIMEOUT	Time in milliseconds	When an Insert or Eject tape command is issued you must open the CAP and insert or eject tapes within this timeout period. If you leave this value set to 0, the timeout period used by the Robot Manager is 10 minutes. If you must alter this timeout period, replace 0 with your own value (in milliseconds).	0
RM_ACSLS_MAX_DISMOUNT_RETRIES	Number of retries	The maximum number of retries when the dismounted drive is still in use. If the setting is 5, the initial delay is five seconds and then doubled after each retry.	5
RM_ACSLS_DISMOUNT_FORCE	0 (disabled) 1 (enabled)	Under normal circumstances, you must unload a tape first (using an Actor) before issuing a dismount command to the library. A forced dismount instructs the library to issue the unload command to the drive directly. This option is not recommended because this may interfere with operations on the Actors.	0

Configuring Sony PetaServe Managed Storage

Control of Sony PetaServe Managed Storage from the Core Robot Manager is directed through the Sony PSC controller over an Ethernet connection. The PSC controller parameters for the Robot Manager configuration file must match those on the PetaSite Controller.

robotmanager.conf Common Options

The following table identifies common robotmanager.conf options:

Parameter	Parameter Type	Description	Default
RM_MODULE=SONY_Robot.dll		Uncomment only this line	Commented
RM_PORT	TCP Port Number	The TCP Port the Robot Manager will listen on for incoming requests. This value must be unique if there are multiple Robot Managers running on a single host. The assigned port is typically TCP Port 8500 and higher.	8500
RM_ACS	Number	ACS (Automated Cartridge System) controlled by the Robot Manager module.	0

The following table identifies common Sony PetaSite options:

Parameter	Parameter Type	Description	Default
RM_SONY_ENABLE_MEDIA_TYPE_TRIMMING	Number	<p>This parameter must not be modified during production. The database may need to be patched if it is changed during production.</p> <p>Some tape labels contain and additional two or three characters identifying the type of media. For example, 004452L2 is an LTO2 tape and S1000052 is a SAIT1 tape.</p> <p>If this parameter is set to 1, the Sony Robot detects the tape using the label and filters out the two or three additional characters from the label.</p>	1
RM_SONY_MEDIA_TYPE_TRIMMING_LEFT	Number	<p>This parameter must not be modified during production. The database may need to be patched if it is changed during production.</p> <p>Depending on the label, the two characters may be on the right or on the left of the label. Set this parameter to 1 if the Media Type information is on the left, otherwise set it to 0.</p>	0
RM_SONY_PSCSERVERNAME	IP Address or Host Name	<p>This parameter specifies the Host Name or IP Address of the Sony PSC (PetaSite controller). If you specify a Host Name, this must be defined in the operating system's hosts file.</p>	

Parameter	Parameter Type	Description	Default
RM_SONY_PSCUSERID	Number	This specifies the User ID that the Robot Manager uses when it connects to the Sony PetaSite Controller.	1
RM_SONY_PSCTIMEOUT	Time in milliseconds	Command time out to the PSC in milliseconds. This is only used for mount operations.	900000
RM_SONY_PSCDISMOUNTRETRIES	Number of retries	The maximum number of retries when the dismantled drive is still in use. If the setting is 5, the initial delay is five seconds. The delay is then doubled after each retry.	5

Configuring ADIC Managed Storage with SDLC

This interface is available on both Windows and Linux platforms. Refer to [Appendix E: ADIC SDLC Installation and Configuration](#) for setting up the SDLC server and client components for the Core Robot Manager interface.

robotmanager.conf Common Options

The following table identifies common robotmanager.conf options:

Parameter	Parameter Type	Description	Default
RM_MODULE=ADIC_Robot.dll		Uncomment only this line	Commented
RM_PORT	TCP Port Number	The TCP Port the Robot Manager will listen on for incoming requests. This value must be unique if there are multiple Robot Managers running on a single host. The assigned port is typically TCP Port 8500 and higher.	8500
RM_ACS	Number	ACS (Automated Cartridge System) controlled by the Robot Manager module.	0

The following table identifies common ADIC parameters:

Parameter	Parameter Type	Description	Default
RM_ADIC_DAS_CLIENT	Host Name	Host Name of the computer running the ADIC DAS client.	
RM_ADIC_EJECT_AREA_NAME	Name	Symbolic name of the Cartridge Access Port.	E01
RM_ADIC_TIME_INSERT	Time in milliseconds	Number of milliseconds to wait to put away the tape after closing the CAP.	5000
RM_ADIC_MAX_DISMOUNT_RETRIES	Number of retries	Maximum number of retries when the dismounted drive is still in use. If the setting is 5, the initial delay is five seconds. The delay is then doubled after each retry.	5

Configuring Simulated Managed Storage (for DIVA Core Simulators)

Simulated robots are available on Windows and Linux platforms. The settings are shown here for reference only. Refer to the DIVA Core Simulator Operations Guide (available to OPN partners only) for more information on installing and configuring a DIVA Core Simulator platform.

robotmanager.conf Common Options

The following table identifies common robotmanager.conf options:

Parameter	Parameter Type	Description	Default
RM_MODULE=SIMULATOR_Robot.dll		Uncomment only this line	Commented
RM_PORT	TCP Port Number	The TCP Port the Robot Manager will listen on for incoming requests. This value must be unique if there are multiple Robot Managers running on a single host. The assigned port is typically TCP Port 8500 and higher.	8500
RM_ACS	Number	ACS (Automated Cartridge System) controlled by the Robot Manager module.	0

The following table identifies the DIVA Core Simulator parameters:

Parameter	Parameter Type	Description	Default
RM_SIMU_BASEDIR	Directory Path	The DIVA Core simulation files base directory path. This is typically C:\Diva\Simulation.	
RM_SIMU_OPERATION_SHORT_DELAY	Time in milliseconds	This setting simulates physical delays in mount, dismount, enter, and eject operations. The recommended setting is 10000 msec.	0
RM_SIMU_OPERATION_LONG_DELAY	Time in milliseconds	You can use this setting to simulate an operation that takes more time than expected for execution. The recommended setting is 120000 msec.	0
RM_SIMU_OPERATION_LONG_DELAY_FREQUENCY	Number	This setting specifies how often a long delay should occur. The recommended setting is 50.	0

Parameter	Parameter Type	Description	Default
RM_SIMU_LIST_SHORT_DELAY	Time in milliseconds	This setting introduces a simulated physical delay in list operations. The recommended setting is 500.	0
RM_SIMU_LIST_LONG_DELAY	Time in milliseconds	You can use this setting to simulate a list operation that takes more time than expected for execution. The recommended setting is 60000 msec.	0
RM_SIMU_LIST_LONG_DELAY_FREQUENCY	Number	This setting specifies how often a long delay should occur in list operations. The recommended setting is 100.	0

Robot Manager Command Options

You perform Core Robot Manager control and management functions using `robotmanager.exe` from a command prompt. On Windows servers the executable is located in the `%DIVA_HOME%\Program\Robotmanager\bin` folder. On Linux servers, `robotmanager.sh` is located in the `/home/diva/DIVA/Program/RobotManager/bin` directory.

Installing and Uninstalling the Robot Manager Services in Windows

Use the following command line options to install or uninstall the Core Robot Manager from a Windows command prompt:

robotmanager -i

Installs the Robot Manager Service as set by the SERVICE_NAME parameter defined in robotmanager.conf. If this parameter is undefined, the service is installed as Core Robot Manager - host_name.

robotmanager -u

Removes the Robot Manager Service set by the SERVICE_NAME parameter in robotmanager.conf. If this parameter is undefined the service to be removed is Core Robot Manager - host_name.

These Robot Manager command options default to the robotmanager.conf file located in the %DIVA_HOME%\Program\conf\robot_manager folder to define the Service Name (if any). If you are installing multiple Robot Managers on a single host (see [Appendix A: Core Options and Licensing](#) for DIVA Core options and licensing information), additional Robot Manager configuration files must be created and specified to the service during installation to create unique instances for each Robot Manager.

You can create additional configuration files for each Robot Manager by copying and renaming the original robotmanager.conf file. For example, robotmanager1.conf, robotmanager2.conf, and so on. Each configuration file must contain unique SERVICE_NAME, RM_PORT, and RM_ACS entries.

For example, robotmanager1.conf might have the following parameters for a SCSI interface:

```
RM_MODULE=SCSI_Robot.dll
SERVICE_NAME=Robot1
RM_PORT=8500
RM_ACS=0
```

While robotmanager2.conf might have the following parameters for an ACSLS interface:

```
RM_MODULE=ACSLs_Robot.dll
SERVICE_NAME=Robot2
RM_PORT=8501
RM_ACS=1
```

The path to each Robot Manager configuration file must be specified for each instance when installing additional Robot Manager Services on the same host. Identify the path by adding the -conf (or -f) command switches when installing the service.

For example, robotmanager -i -conf ..\..\conf\robot_manager\robotmanager2.conf installs the Robot Manager service as defined by the SERVICE_NAME parameter from the robotmanager2.conf configuration file.

If one or more Robot Manager Services must be uninstalled, the configuration file path must also be specified. For example, robotmanager -u -conf ..\..\conf\robot_manager\robotmanager2.conf removes the Robot Manager Service as defined by the SERVICE_NAME parameter in robotmanager2.conf configuration file.

After installing the services check the Windows Services applet to confirm that the Robot Manager Services were installed correctly. To change the SERVICE_NAME, uninstall the existing service before editing the robotmanager.conf file. Then reinstall the service after changing the SERVICE_NAME parameter.

Installing and Uninstalling the Robot Manager Services in Linux

Use the following command line options to install or uninstall the Core Robot Manager from a Linux terminal.

Use the following command sequence to install the Robot Manager service:

```
cd /home/diva/DIVA/Program
```

```
./divaservice install robotmanager /home/diva/DIVA/Program/conf/  
robot_manager/robotmanager.conf
```

Use the following command sequence to uninstall the Robot Manager service:

```
cd /home/diva/DIVA/Program
```

```
./divaservice uninstall robotmanager /home/diva/DIVA/Program/conf/  
robot_manager/robotmanager.conf
```

See [Installing the DIVA Core Services](#) for information on the divaservice command.

Robot Manager Service Management Functions

The following command options are also available for the Robot Manager Service:

robotmanager debug

Starts the Core Robot Manager in console mode. Console mode displays diagnostic messages and other information from the library in the console window.

robotmanager version

Displays the Core Robot Manager software release information.

You can also use “-v” instead of “version”.

robotmanager help

This displays all command line options.

Testing the Robot Manager Library Interface

After configuring the Robot Manager configuration file, launch the Core Robot Manager and confirm that the library itself can be controlled.

Library interfaces that use ACSLS, SDLC, or PSC intermediate control software must be running before launching the Core Robot Manager. ACSLS controlled Managed Storage should also be varied online (for example, vary lsm0 online).

Starting, Stopping, and Restarting the Robot Manager

Windows Core Robot Managers start automatically with Windows. You manage (start, stop, restart, and so on) the service through the Windows Services applet.

Note: If the library is offline when the service is started, the Robot Manager does not automatically reconnect after the library comes online. You must restart the service to connect.

A Robot Manager can also be stopped and then started (restart) from a command window. The quotation marks in the commands must be used when specifying a service with spaces in the name. Use the following command sequence to stop and then start the service:

```
net stop "Core RobotManager"  
net start "Core RobotManager"
```

You use the following command sequence if the *SERVICE_NAME* is specified in the robotmanager.conf file:

```
net stop "Core RobotManager SERVICE_NAME"  
net start "Core RobotManager SERVICE_NAME"
```

Testing the Robot Manager Library Control

Caution: These utilities must not be used in a live DIVA Core system. You must not send commands to a Robot Manager using either of these utilities under any circumstances when the Core Manager is running. Technical Support is not responsible for any complications arising from inappropriate use of these utilities.

Either the Robot Manager Client (a command-line interface) or GUI can be used to establish basic control functionality of a Robot Manager to its controlled Managed Storage. You can use either of these utilities to send manual commands to a Core Robot Manager to initiate simple operations, for example, drive mounting, dismounting, enter or eject operations from the CAP (Cartridge Access Port). Both utilities connect to a Robot Manager through TCP/IP and can be run from a remote computer. This feature enables the Robot Manager GUI to be used from a remote computer.

If you mount a tape with either of these utilities, you must first unload the tape before it can be dismounted, unless the library supports Forced Dismount commands and they are enabled in the Core Robot Manager configuration file.

Robot Manager Client

This command-line client is typically located with the Robot Manager executable files in the %DIVA_HOME%\Program\RobotManager\bin folder.

You must specify the IP address of the Robot Manager and its TCP port when launching the client as follows:

```
RobotManagerClient {IP_Address} {TCP_Port}
```

The IP_Address is the IP address of the Robot Manager computer, and the TCP_Port is the Robot Manager listening port. You can hard-code these two parameters in the Robot Manager Client batch file if there is only a single Robot Manager requiring access.

All of the client commands are self-explanatory after you start the program.

Robot Manager Client GUI

The Robot Manager Client GUI is typically located with the Robot Manager executable files located in the %DIVA_HOME%\Program\RobotManager\bin folder. The GUI provides the same functionality as the command line client. You execute RobotManagerGUI.bat to open the GUI interface.

The GUI interface includes the following buttons and functionality:

Connect Button

Click this button to connect to the Core Robot Manager. You must enter the IP address and TCP port of the Core Robot Manager to be tested in the **Connect** prompt.

Tape List Button

Click this button to load the tape list from the library.

Reload Config. Button

Click this button to reload the configuration.

Exit Button

Click this button to exit the program.

Tape List

To manually mount a tape, select a Barcode ID and drag and drop it on to one of the drives displayed in the LSM list.

LSM List

This area lists all of the available drives and the tapes in the drive. You right-click a tape and select **Dismount** from the menu to dismount a tape.

CAP List

To manually eject a tape from the library, select a Barcode ID and drag and drop it to one of the listed CAPs.

Status Area

The Status area is at the bottom of the screen and displays status messages from the Robot Manager.

Configuring the Robot Manager at the System Level

At the system level, each instance of the Core Robot Manager must be declared to the Core Manager in the Robot Managers area of the Robots page in the System Management App.

You use area buttons to add (+), edit (Edit), or delete (-) a Robot Manager. The Update area button refreshes the displayed Robot Manager information from the database.

Clicking the + button adds a Robot Manager to the configuration. The Add new row in Robot Managers dialog box is displayed. Enter the following information in the appropriate fields and then click OK to add a Robot Manager:

Name

The name of the Core Robot Manager attached to this DIVA Core system.

Address

The IP address of the host running the Core Robot Manager installation.

Port

The Robot Manager TCP port. This must match the RM_PORT parameter specified in robotmanager.conf.

Site

The Core Manager uses this parameter to determine optimal use of resources in resource allocation. Use the menu list to select the appropriate site for this Robot Manager. Site Selection must be enabled in the Core Manager configuration file or all sites are considered equally.

Logging Robot Manager Activity

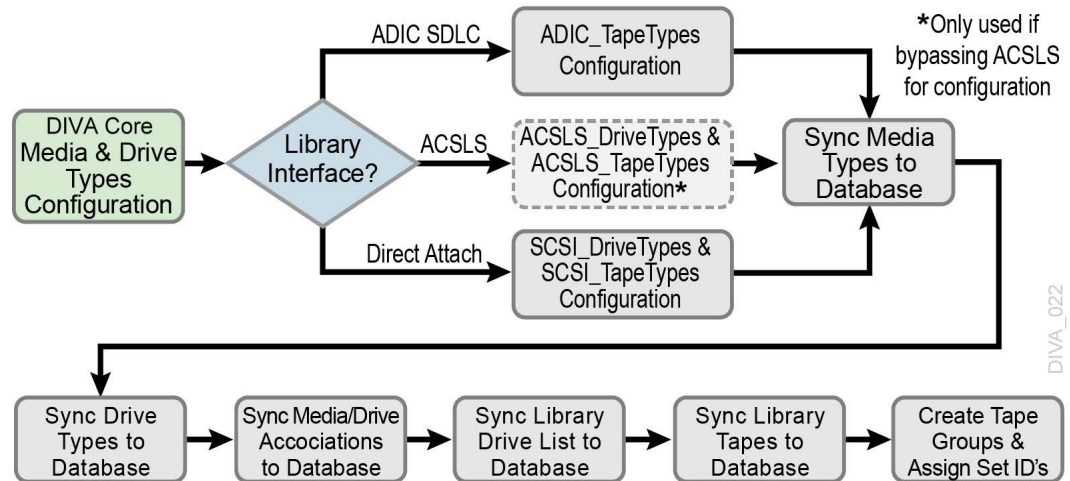
During normal operation, each Core Robot Manager logs its communications with the library and stores them in the %DIVA_HOME%\log\robot_manager folder. These logs are useful for troubleshooting issues. You may be asked to provide the log files when contacting Technical Support.

The most recent log file is named robot_manager.log or robot_manager_SERVICE_NAME.log and is located in the ..\log\robot_manager folder. Older logs are renamed with the time it was saved as its file name and moved to dated subfolders under the name of each Robot Manager.

Configuring Media and Drive Types

After you have successfully configured the Core Robot Manager for your Managed Storage, and the appropriate details for all Core Robot Managers entered into the Robots page section of the System Management App, the Tape Media, Drive Models, and the Drive Locations currently installed in each library must be entered.

The following flowchart lists the workflow of this portion of the configuration. All of the Core Robot Managers configured must be running and successfully connected to each library before commencing this portion of the configuration.



Tape Drives and their associated media types that are installed in a particular library are initially configured in the Core Database using static configuration files. The files are located in the %DIVA_HOME%\Program\conf\robot_manager folder. The Core Robot Manager selects the appropriate files according to the RM_MODULE setting configured in robotmanager.conf.

The following list identifies the configuration file names and use:

scsi_drive_types.ini and scsi_tape_types.ini

Used for direct attached SCSI Managed Storage. These files are only considered if the .ini extension is removed.

acsls_drive_types.ini and acsls_tape_types.ini

Used for Managed Storage managed by ACSLS. Normally, tape and drive types are derived from ACSLS during library synchronization with the database. However, you can use these files to override the values returned from ACSLS. These files are only considered if the .ini extension is removed.

adic_media_types.ini

Used with ADIC Managed Storage controlled by SDLC. Drive Types for this library are directly returned from the SDLC server. These files are only considered if the .ini extension is removed.

When a hardware audit is initiated on the specific library by the System Management App (through the Core Robot Managers, either directly or through intermediate library management software), hexadecimal codes are returned to identify the model and order of the tape drives currently installed, and the media types present in the library.

These library hardware codes are mapped to drive and media IDs within the Core Database using the Tape Types and Drive Types configuration files.

It is only necessary to modify these files when Drive Types or Media Types are added to the library.

SCSI_drive_types and ACSLS_drive_types

You can edit these files using any plain text editor (for example, Notepad or Notepad++). No modification of these files is required other than to remove comment fields for the appropriate library and drive types for your installation.

Remove the # at the beginning of the line in the appropriate library section for the drives to be recognized in a Synchronize Drive Types List in the System Management App. You must leave drive types in Managed Storage not installed commented out.

The Compatible Drive Types column cross-references the Tape Type ID in SCSI_Tape_Types (or ACSLS_Tape_Types if used). These values are examined in the Synchronize Media/Drive Compatibility List procedure in the System Management App.

SCSI_tape_types and ACSLS_tape_types

You can edit these files using any plain text editor (for example, Notepad or Notepad++). No modification of these files is required other than to remove comment fields for the tape types for your specific library.

Remove the # at the beginning of the line in the appropriate library section for the tapes or DVDs to be recognized in a Synchronize Media Types List in the System Management App. You must leave tape types (or DVDs) in Managed Storage not installed commented out.

The Compatible Drive Types column cross-references the Drive Type ID in SCSI_Drive_Types (or ACSLS_Drive_Types if used). These values are examined in the Synchronize Media/Drive Compatibility List procedure in the System Management App.

ADIC_media_types

You can edit these files using any plain text editor (for example, Notepad or Notepad++). No modification of these files is required other than to remove comment fields for the tape types for your specific library.

Remove the # at the beginning of the line in the appropriate library section for the tapes to be recognized in a Synchronize Media Types List in the System Management App. You must leave tape types in Managed Storage not installed commented out.

The Compatible Drive Types column cross-references the Drive Type ID returned from the SDLC controller. These values are examined in the Synchronize Media/Drive Compatibility List procedure in the System Management App.

DIVAmigrate Installation and Configuration

This section describes an overview of the DIVA Core DIVAmigrate Embedded Utility, and installation and configuration of the tool.

DIVAmigrate Embedded Utility Overview

DIVAmigrate is installed as part of the DIVA Core Suite's standard installation. It is located in the %DIVA_HOME%\Program\ folder, and runs as a Windows Service.

You create Migration Requests through the DIVA Core System Management App connected to the Core Manager, or using the command-line interface through the client.bat file located in the %DIVA_HOME%\Program\Migrate\bin folder.

You control the utility using the migrate.bat file, also located in %DIVA_HOME%\Program\Migrate\bin folder. See the DIVA Core Operations Guide in the DIVA Core Library for details.

Migration Requests are stored in the Core Manager Database. The DIVAmigrate Service monitors the Core Manager Database, runs new Migration Requests, and also updates the status of existing Migration Requests in the database so that the System Management App displays the status to users.

Installing DIVAmigrate

The DIVAmigrate Utility is part of the standard DIVA Core installation, and is placed in the %DIVA_HOME%\Program\ folder. You can install DIVAmigrate on the Core Manager computer, or any other computer capable of communicating with the Manager using the TCP/IP protocol. You can confirm connectivity by successfully pinging the Manager from the client computer.

Windows Files and Folders

You will find the following new files and folders after you complete the initial DIVA Core installation in Windows:

```
%DIVA_HOME%\Program
  Migrate
    bin
      client.bat
      migrate.bat
    lib
      migrate.jar
  conf
    migrate
      migrate.conf.ini
  log
    migrate
```

Linux Files and Directories

You will find the following new files and directories after you complete the initial DIVA Core installation in Linux:

```
%DIVA_HOME%/Program
  Migrate
    bin
      client.sh
      migrate.sh
    lib
      migrate.jar
  conf
    migrate
      migrate.conf.ini
  log
    migrate
```

Configuring the DIVAmigrate Service

The DIVAmigrate Service requires a valid configuration file during install and start procedures. The default DIVAmigrate configuration file is named `migrate.conf` and is located in the `%DIVA_HOME%\Program\conf\migrate\` folder.

The configuration file is a standard properties file similar to the Manager configuration file. The configuration file is not auto-reloadable, and therefore any changes made to the file do not take effect until the DIVAmigrate Service is restarted.

Note: The Windows Service Wrapper configuration must not be modified. The DIVA Service Options section also must not be modified unless instructed by Technical Support.

Use the following procedure to modify the DIVAmigrate configuration file:

Caution: Do not use Word, WordPad, or any other word processor or editing tool that adds extra characters to a file. Always use a plain text editor such as Notepad, or Notepad++.

1. Create a copy of the `migrate.conf.ini` file.

It is important to create a copy and keep the original file intact to refer back to in case the configuration you are working on either does not work, becomes corrupt, or has or creates errors.

2. Rename the copied file to `migrate.conf`.

3. Open the file with any plain text editor and populate the following parameters. These parameters are all mandatory unless otherwise noted in the description.

SERVICE_NAME

This parameter is the name for the Windows Service. The default value is DivaMigrate.

DIVAMANAGER_HOST

This parameter is the host name or IP address of the Core Manager. The default value is 127.0.0.1.

DIVAMANAGER_PORT

This parameter is the port number to connect to the Core Manager. The default value is 8000.

DIVA_MIGRATE_MANAGEMENT_PORT

This parameter is the management port number. The default value is 9191.

DIVAMANAGER_DBUSER

This parameter is the user name the Manager uses to connect to the Core Database. The default value is diva, and is case sensitive.

DIVAMANAGER_TNSNAME

This parameter is the TNS Name of the DIVA Core Schema within the Oracle Database. DIVAmigrate ignores this setting if the DIVAMNAGER_DBHOST and DIVAMANAGER_DBPORT settings are defined. There must be a corresponding entry in TNSNAMES.ORA found under the Oracle 11 Client installation. This is not a mandatory parameter.

DIVAMANAGER_DBHOST

This parameter specifies the host name or IP address of the computer containing the Core Database. If using a host name, this must be present in the hosts file on the computer where the Core Manager is installed. For example, you can use either 127.0.0.1 or localhost. This is not a mandatory parameter.

DIVAMANAGER_DBPORT

This parameter is the Oracle Listener Port you configured during the Core Database installation. The default value is port 1521. This is not a mandatory parameter.

DIVAMANAGER_DBSID

The Core Database instance SID (System Identifier) in the Oracle Database where the Core Manager connects. This value is typically lib5, which is the default value. Consult your location's System Configuration Plan for confirmation.

DIVAMANAGER_DBSERVICENAME

This parameter specifies one name for the database service to which this instance connects, and is listed in tnsnames.ora. Typically, lib5.world (the default), is used in most DIVA Core installations. Consult your delivery plan if you are unsure.

Note: Either this value must be set, or DIVAMANAGER_DBSID, when you use DIVAMANAGER_DBHOST and DIVAMANAGER_DBPORT for database connections. If you set both parameters, then SERVICENAME takes precedence over SID.

MAX_SIMULTANEOUS_REQUESTS

This parameter is the maximum number of simultaneous Manager requests processed by DIVAmigrate. The default value is 30. This is not a mandatory parameter.

DB_SCAN_PERIODICITY

This parameter (in seconds) determines how often DIVAmigrate looks for new Requests in the database. The default value is 60 seconds. This is not a mandatory parameter.

DIVA_RECONNECT_PERIODICITY

This parameter (in seconds) determines the time between reconnection attempts if connectivity with the Manager is lost. The default value is 30 seconds.

MAX_FAILED_REQUESTS_PAUSE

This parameter identifies the maximum number of sequential failure requests that can occur before DIVAmigrate pauses the Migration Request. DIVAmigrate pauses the Request if the configured number of requests fails sequentially. The default value is 10. This is not a mandatory parameter.

REQUEST_STATUS_CHECK_DELAY_SECS

This parameter (in seconds) determines how often DIVAmigrate scans manager for request status. The default value is 5 seconds. This is not a mandatory parameter.

REQUEST_MAX_INACTIVE_TIME

This parameter (in hours) determines the Migrations Plan's maximum inactivity time. If, after the service is restarted, it finds running jobs having the last access time greater than this value, the Migration Plan for those Requests are recreated. The default value is 24 hours. This is not a mandatory parameter.

MAX_SIMULTANEOUS_REQUEST

This parameter is the maximum number of simultaneous Requests or process request that can run at the same time. The default value is 15. This is not a mandatory parameter.

DIVAMANAGER_DB_SECURE_CONNECT

This parameter is for connecting securely to Oracle Database; it must be set to TRUE to connect securely. The value for the parameter DIVAMANAGER_DBPORT must be to the secure port number 1522 of the Oracle database. The default value is FALSE.

MAX_REQUESTS_TAPE_READ_PERMIT

This parameter restricts the number of Requests that are reading from tape. The default is 15.

MAX_REQUESTS_TAPE_WRITE_PERMIT

This parameter restricts the number of Requests that are writing to tape. The default is 15.

TAPE_READ_WRITE_LOCK_ACQUIRE_WAIT

This parameter identifies the acquired Read/Write permit timeout, before retrying. This parameter makes sure the Request is not waiting on acquiring a permit indefinitely. If the timeout occurs, the Request will check for any user actions on the Request (for example, pause, stop, cancel, or delete) before retrying to acquire the permit again. The default 30 seconds.

CLEAN_BUFFER_DELETE_REQUEST_CHUNK_SIZE

This parameter optimizes buffer cleaning when a user is performing a Migration Request cancel or stop, and if a cache buffer being used has a lot of object instances. This parameter will allow the Migration Request to send a chunk of delete requests rather than sending one at time. The default is 20 requests per chunk.

Configuring the Logging Settings

DIVAmigrate uses the same logging methods used for the Core Manager. However, DIVAmigrate logs are located in the %DIVA_HOME%\Program\log\migrate folder. DIVAmigrate logs are automatically archived and divided into separate files each time the current log file reaches its size limit. Set the following DIVAmigrate logging parameters in the migrate.conf file:

LOGGING_DIRECTORY

This parameter identifies the DIVAmigrate log file storage directory. The default is ../../log/migrate.

LOGGING_TRACE_LEVEL

You can modify this parameter to suit the required level of activity logging. The default value is INFO.

Tip: Only use the higher logging levels when instructed to do so by Technical Support to avoid large log files being created.

Valid options for each parameter are:

- DEBUG
- INFO
- WARN
- ERROR

LOGGING_MAXFILESIZE

This parameter identifies the maximum size of the log file before it is archived. When the current log file is archive, a new file is created. You must specify the file size using KB or MB to indicate Kilobytes or Megabytes respectively. The default value is 10MB. For example, LOGGING_MAXFILESIZE=10MB.

LOGGING_LIFETIME

All files older than the value of this parameter are removed. This includes trace, service, and .zip files. The value for this parameter is in hours, and the default value is 50.

Additional Functionality

This chapter describes DIVA Core additional functionality.

Topics:

- [Checksum Support Configuration](#)
- [Transcoder Installation and Configuration](#)
- [Disk Auto-Discovery](#)

Checksum Support Configuration

Overview

You configure the Checksum Support functions through the System Management App using the Engineer account. The following sections describe how to adjust the settings for each option.

Global Checksum Parameters

You must use the Engineer account in the System Management App to access and adjust the Global Checksum Parameters located under the Manager Setting area. Each of the global parameters affects all Checksum Support settings throughout the system. The following options are available:

Manager: Checksum feature is enabled

This setting enables (check box selected) or disables (check box deselected) the Checksum Support features throughout DIVA Core. The default setting is enabled (selected).

Manager: Default Checksum Type

There are several checksum algorithms supported by the system including MD2, MD5, SHA, SHA1, MDC2, and RIPEMD160. DIVA Core uses MD5 as the default checksum.

Each checksum type is associated with an ID Number. you use the menu list to change the default type and select the type of checksum desired.

The ID Number identifies the Checksum Type requested in the configuration as follows:

- MD2 is ID Number 1
- MD5 is ID Number 2
- SHA is ID Number 3
- SHA-1 is ID Number 4
- MDC2 is ID Number 5
- RIPEMD160 is ID Number 6

Manager: Number of retries following failed checksum

This parameter sets the number of times the system will retry the operation after a failed checksum. The default setting is one retry. Enter the number of retries allowable for your data and system in the Manager: Number of retries following failed checksum field. Technical Support recommends leaving this setting at the default value.

Manager: Select different drive per retry on failed checksum

This parameter distinguishes whether the retry (after a failed checksum) is attempted on the same drive (check box deselected), or if the system should attempt the operation using a different drive (check box selected). The default setting for this parameter uses the same drive (check box deselected).

Configuring Checksum Support for Servers

Adjust the Checksum Support configuration for Servers through the System Management App System page. In the Servers area, double-click the Servers requiring Checksum Support configuration. The Edit Servers Entry dialog box appears with several Checksum Support configuration options. These options are mainly associated with the Genuine Checksum Type.

The following list describes the options available:

External Checksum Source Server

You must use the External Checksum Source Server (Yes option) for the system to read the Checksum from the external Source Server providing the file. This initiates an immediate checksum calculation to compare the checksums and verify the initial transfer. Selecting the No option disables Genuine Checksum support from external Source Servers.

Checksum Type

You use the menu list to select the Checksum Type. All supported checksum types are listed. The Checksum Type and GC Mode (see the following description) must match the settings implemented at the Source Server.

The Genuine Checksum is only used for the first verification. Therefore, the checksum type selected is only used once and then discarded. Beyond the initial use of the selected checksum type (after this transfer), the default type is used.

GC Mode

You use the menu list to select the Genuine Checksum Mode. This notifies the Actor of the format of the files that contain the checksum data.

Verify Following Archive (VFA)

When Verify Following Archive (VFA) is turned on (check box selected), performing the initial transfer from the Source Server results in a read-back operation. Therefore, the data is read twice for verification. After the data is read twice, the two checksums are compared. If they are the same then verification is complete. If they are not identical then verification has failed.

Verify Following Archive is not compatible with Genuine Checksum (GC) or Complex Objects. There is no need to use VFA when GC is being used because the checksum is

already verified. The Genuine Checksum must be turned off to gain access to the VFA check box. If GC is turned on, the check box will be grayed out and not selectable.

Verify Following Restore (VFR)

When Verify Following Restore (VFR) is turned on (check box selected), performing the final transfer to the Destination Server results in a read-back operation. Therefore, the data is read twice for verification. After the data is read twice, the two checksums are compared. If they are the same then verification is complete. If they are not identical then verification has failed. The setting of GC has no bearing on the VFR setting.

Verify Following Restore is not compatible with Complex Objects or the -axf option. Verify Following Restore was designed to read back the restored content from a video server to confirm that it is not corrupt. Using the -axf option does not create a checksum verifiable restore. It creates an object export that is encompassed in an AXF wrapper. The VFR and -axf options are mutually exclusive and should not be part of the same workflow.

Configuring Checksum Support for Arrays and Disks

Checksum Support for Arrays and Disks is configured through the System Management App Disks page. Verify Write (VW) functionality can be turned on or off either on an array basis or disk by disk.

VW applies when you write to the final storage location in DIVA Core. When turned ON, the system will perform a read-back of what was just written and compare the checksums for verification.

The VW column in both the Arrays area and Disks area indicates whether the Verify Write function is on or off for the particular array and disk. The default setting is OFF.

If there is nothing defined in the VW column on the Disk area the system will use the setting defined in the Array VW column.

To override the setting defined in the Array VW column for a specific disk, you select the disk requiring configuration in the Disks area and click Edit located at the top of the area.

When the Edit Disks Entry dialog box appears, use the Verify Write menu list to select ON, OFF, or NONE (blank selection). If NONE is selected, Verify Write uses the setting identified in the array for this particular disk.

The selection made in the Edit Disk Entry dialog box is reflected in the Disks VW column.

Configuring Checksum Support for Tape Groups

Verify Write for Tape Groups is also configurable. The VW column displays in the System Management App Tape Groups page. This is the only place where configuration of Verify Write is available for the Tape Groups.

Similar to the configuration for disks, select the Tape Group requiring configuration. Click Edit and select ON or OFF using the Verify Write menu list. The selection is reflected in the Tape Groups VW column.

When DIVA Core writes a file to a particular Tape Group, the setting for that Tape Group is applied to the file. The default setting for Tape Groups is OFF.

Configuring Checksum Support for Actors

Verify Tape for Actors is also configurable. Similar to the configuration for disks and Tape Groups, select the Actor requiring configuration, click Edit, and then select Yes or No using the Verify Tape menu list.

This setting defines whether the Actor is automatically selected for the Verify Tape workflow. By default, all Actors are included, but you can exclude if necessary.

AXF and TEXT Genuine Checksum Modes

There are two additional Genuine Checksum modes as follows:

AXF Genuine Checksum Mode

This mode enables DIVA Core to archive all files and subfolders in a specified AXF file while comparing their checksum values against known values stored in the AXF file. This workflow is typically combined with a Restore request with `-axf` in the Request Options.

TEXT Genuine Checksum Mode

This mode enables DIVA Core to archive all files and subfolders in a specified folder while comparing their checksum values against known values stored in an external checksum file.

Configuring AXF Genuine Checksum Mode

There are specific requirements and limitations when using the AXF Genuine Checksum Mode as follows:

- The AXF file containing the files to be archived must contain checksum information for each file.
- The checksums must be the expected type as specified in the configuration.
- This workflow only works with AXF requests generated by DIVA Core.
- Verify Following Restore (VFR) is not compatible with the `-axf` option.

VFR was designed to read back the restored content from a video server to verify it has not been corrupted. Using the `-axf` option does not create a real restore; rather an object export in an AXF wrapper. These options are mutually exclusive and should not be part of the same workflow.

DIVA Core System Management App Settings

Use the following procedure to configure AXF Genuine Checksum Mode in the System Management App:

1. Create a new Server entry with the Source Server Type set to either DISK, FTP_STANDARD, or EXPEDAT as appropriate.
If you are required to specify an appropriate Root Path, this path along with the input files specified during the Archive request is used in determining the location of the checksum file.
For example, if the Source Server Type is DISK, you can set the Root Path to D:\root. If the Source Server Type is FTP_STANDARD, you can set the Root Path to /root.
2. Set the External Checksum Source Server to *YES*.
3. Set the Checksum Type to the expected checksum type (for example, MD5).
4. Set the GC Mode to *AXF*.
5. Click *SAVE*.
6. Notify the Manager of the configuration by selecting Tools > Notify Manager from the menu.

Configuring TEXT Genuine Checksum Mode

There are specific requirements and limitations when using the TEXT Genuine Checksum Mode as follows:

- A checksum file must be present in the folder specified by the Root File Path.
- Checksum files must end with a .md5 file extension.
- The checksum file name (excluding the extension) must be associated with the folder name that contains all files to be archived. This folder must exist.
For example, if the checksum file is D:\Data\Video\NewTitle.md5, then all files located in the D:\Data\Video\NewTitle folder will be archived.
- The checksum file must be present in the folder parent to the folder specified by the Root File Path.
- For a file to be archived with the Genuine Checksum value, the file must be referenced with a corresponding checksum within the checksum file.
- Absolute path names are supported on both Windows and Linux to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems (only).
- Linux paths, file names, and commands are case-sensitive.
- Only ASCII, non-UTF-8 encoded checksum files are supported.
- The format of the checksum file is that each line begins with an MD5 checksum, followed by 2 spaces, and then the file path to the referenced file.

System Management App Settings

Use the following procedure to configure TEXT Genuine Checksum Mode in the System Management App:

1. Create a new Server entry with Source Server Type set to either *DISK* or *FTP_STANDARD*.
2. Specify an appropriate Root Path. This path, along with the input files, specified during the Archive request is used in determining the location of the checksum file (see [Selecting the Root File Path](#) for further details).
For example, if the Source Server Type is *DISK*, you can set the Root Path to *D:\Data*. If the Source Server Type is *FTP_STANDARD*, you can set the Root Path to */Data*.
3. Set the External Checksum Source Server to *YES*.
4. Set the Checksum Type to *MD5*.
5. Set the GC Mode to *TEXT*.
6. Click *SAVE*.
7. Notify the Manager of the configuration by selecting *Tools > Notify Manager* from the menu.

Selecting the Root File Path

The Root File Path must point to the folder containing the checksum file. Therefore, the correct file and folder paths must be set in the Server and Archive request form so the checksum file can be located. For example, if the checksum file is located in *D:\Data\Video\NewTitle.md5* (or */Data/Video/NewTitle.md5* for *FTP* type), you set the appropriate file and folder paths as follows:

Server (Root Path)	Archive Request (File Path Root)	Archive Request (Files)
D:\	Data\Video\NewTitle	*
D:\Data	Video\NewTitle	*
D:\	Data\	Video\NewTitle*
D:\		Data\Video\NewTitle*

Server (Root Path)	Archive Request (File Path Root)	Archive Request (Files)
/	Data/Video/NewTitle	*
/Data	Video/NewTitle	*
/	Data/	Video/NewTitle/*
/		Data/Video/NewTitle/*

Transcoder Installation and Configuration

Transcoder Overview

The following instructions are directed toward servers running the Windows Server 2016 SP1 operating system. Linux-based Actors only support Telestream Vantage for transcoding operations.

Upgrading from Telestream Vantage 5.0 and Earlier

Upgrading from 5.0 or earlier releases of Vantage requires uninstalling and reinstalling the Vantage software. Refer to the Vantage 6.3 Installation Guide for details on the uninstall procedure.

Installing Telestream Vantage

Technical Support recommends that no anti-virus software is installed on the Vantage servers. Use the following procedure to install Vantage 6.3:

1. Download the Vantage 6.3 release from Telestream.
2. If you are uncertain of how to install the software, refer to the Quick Start Instructions in the downloaded file.
3. Install .NET 3.5 SP1, if not already installed, on the host computer that will be running the Vantage Database server.
4. Install QuickTime 7.6.9 if not already installed.
5. Install the Desktop Experience option. This is located in the Server Manager under Features.
6. Install the VantageDatabaseSetup_SQL2008_4.2.286.100451.exe, accepting the default settings.
7. Execute the Vantage_6.3_Setup.exe.
8. Select the Install Product(s) option.
9. Ensure the following options are selected:
 - Transcode/Transcode Pro
 - Web Applications
 - Workflow Portal Application
 - Vantage Domain Database
10. Enable any other options required for your installation.
11. Complete the installation.

Installing the Telestream License

Use the following procedure to install the Telestream license after the software is installed:

1. Launch the Vantage Workflow Designer.
2. If you are prompted to select a Domain, select the local computer.
3. If you are prompted for a Collection click Cancel (for now).
4. Click File, and then Add/Update License.

Vantage is now installed and configuration can continue for it to work with DIVA Core.

Technical Support recommends importing sample workflows in the Vantage Workflow Designer. You can view a demonstration at <http://www.telestream.net/vantage/demos.htm>.

Configuring DIVA Core and Transcoders

The following instructions identify the configuration of DIVA Core and transcoders to enable operation together. Starting with DIVA Core 7.3, it is no longer required to have Actor installed on the same computer as the transcode service.

A transcoder is no longer coupled to a single Actor. You select the transcoder after you select the Actor. Therefore, you no longer define a LOCAL transcode Actor as a Destination Server. A LOCAL Actor destination is dynamically and temporarily (only in memory, not stored in the database) created for the Actor that you chose as part of resource selection.

The transcoder server and cache locations are now embedded in the Working Directory on the Edit Transcoders Entry screen in the following format:

```
[actor:actor_name,actorPath:actor_transcoder_cache_path,transcoder:trancoder_ip_address],cifs://user_name\domain:password@\\transcoder_cache_ip_address\transcoder_cache
```

For example:

```
[actor:actor_001,actorPath:/tmp/vantagecache,transcoder:10.145.40.81],cifs://user:password@\\10.145.40.81\VantageCache
```

Parameters	Required / Optional	Description
actor	Optional	Specifies a list of one or more Actors that the Manager will select from to perform the transcoding. Multiple Actors are separated by a comma.
actorPath	Optional	Used for Linux-based Actors. Specifies a fixed and existing mount point to the CIFS transcoder cache folder (for example, /mnt/vantagecache). If this parameter is not specified, the Actor will automatically create its own mount point to the transcoder cache share folder.
transcoder	Optional	The IP address to the transcoder. If this parameter is not specified, 127.0.0.1 is assumed.

The following rules apply:

- The order of the actor, actorPath, and transcoder settings is important. The order of the parameters must be actor, followed by actorPath, and finally followed by transcoder.
- Multiple transcoders are not supported for Flip Factory. They are only supported for Vantage.
- Linux-based Actors only support Telestream Vantage for transcoding operations.
- If the transcoder parameter is not specified with the transcoder IP address, a local address of 127.0.0.1 is assumed.

For example:

```
[actor:actor_001_std,transcoder:127.0.0.1],d:\diva\local
```

- If the actor parameter is not specified with an Actor name, the transcoder is presumed to not be mapped to a specific Actor.
- The transcoder_cache folder is the location where both the Actor and Vantage use to perform the entire transcode operation. Because Vantage runs in the Windows platform, a CIFS formatted UNC path that is Windows compatible represents the transcoder_cache share folder. Vantage will use this path for transcoding.
- If the actorPath parameter is not specified, the Actor will use the same CIFS formatted UNC path.
- The original method of configuring a transcoder to a local Actor is still supported for legacy purposes
- The original method of configuring Local Servers tied to Actors is still supported so legacy configurations will continue to function.

Preparing a Fixed Mount Point for Linux-based Actors (optional)

Linux-based Actors must have access to the transcoder cache folder through a local mount point. You can either let the Actor dynamically create a mount point on its own (the path will be determined by the CIFS path of the transcoder cache), or you can create your own fixed mount point for the Actor to use.

If you let the Actor dynamically create a mount point automatically to a remote transcoder cache located at `\\hostname\vantagecache`, the Actor will create the mount point `{root_mount_point}/hostname/vantagecache`.

The `{root_mount_point}` is a mount point configurable in the System Management App. By default, the Actor will use `/mnt` if the root mount point is not configured. See [Advanced Actor Settings](#) for more information on where to configure this.

If you want to use a fixed mount point instead of having the Actor dynamically create one, use the following procedure to create a mount point to the remote transcoder cache:

1. Open a terminal window and execute the `id` command to confirm that you are logged in under the same user account that runs the Actor (typically `diva`) as follows:

```
[diva@Linux018 actor]$ id
```

The response will look similar to the following. Confirm the `uid` (User ID) and `gid` (Tape Group ID), and use these values in the mounting operation. In this example they are both 1000.

```
uid=1000 (diva) gid=1000 (diva)
tapegroups=1000 (diva), 10 (wheel), 30 (tape), 54321 (oinstall), 54322
(dba) context=unconfined_u:unconfined_r:unconfined_t:s0-
s0:c0.c1023
```

2. Execute the following command to create your local mount point directory (for example, `/mnt/vantagecache`). You may need to log in as the root user depending on where the directory is located.

```
mkdir /mnt/vantagecache
```

3. Execute the following command to create a local mount point (for example, `/mnt/vantagecache`) to the network share of the transcoder cache. Enter the appropriate share authentication information including the `user_id#`, `tapegroup_id#`, `remote_transcoder_cache_ip_address`, and `remote_transcoder_cache_path`.

```
mount -t cifs -o
username={user_name},password={password},uid={user_id#},gid={t
apegroup_id#} //transcoder_cache_ip_address/
remote_transcoder_cache_path /mnt/vantagecache
```

4. Set the `actorPath` parameter to `/mnt/vantagecache` when configuring the transcoder settings.

The following sections describe general transcoder configuration.

Configuring the Transcoder and Actor on a Single Computer

Use the following procedure to configure Vantage transcoders when the Actor is on the same computer as the transcode service:

1. Create a cache folder on the Actor computer. For a Vantage transcoder `M:\VantageCache` could be used.
2. Add the transcoder in the DIVA Core System Management App. with the following settings:
 - Transcoder Type: *vantage*
 - Working directory: *M:\VantageCache*
 - Leave the remaining options at the default settings.
3. Open the System Management App.
4. Navigate to the Transcoders area on the System page.
5. Ensure that the DIVA Core Transcoder configuration's Simul Transcodes value is less than or equal to the corresponding Vantage Session Limit value.
6. Open the Vantage Management Console.
7. Click Services in the left navigation tree.
8. Locate the transcoder you are configuring in the right area and then right-click the transcoder name.
9. Select Enter Maintenance Mode from the context menu.
10. Click Service Limits on the Setup page in the bottom area.
11. Confirm the Session Limit and the Target Resource Usage parameters are set correctly for your environment and adjust as necessary.

Configuring the Transcoder and Actor on Separate Computers

Use the following procedure to configure Vantage when the Actor is on a different computer than the Vantage Transcode service:

Caution: The cache folder must be located on a computer accessible by the Vantage SDK computer through a shared Windows path.

1. Create a cache folder on the remote computer. In the example `M:\VantageCache` is used.
2. In Windows, share this folder on the network and set the required access credentials.
3. Authorize the Vantage transcoder to access the shared Vantage Cache folder.
4. Open the Vantage Management Console on the Vantage SDK computer.
5. Navigate to the Settings & Options screen using the left navigation tree.

6. Click the Authorization tab.
7. Add a new entry with the Username, Password, and Folder. For example, \\10.145.50.81\VantageCache is the Windows UNC path for the shared Vantage Cache folder.
8. Open the System Management App.
9. Navigate to the Transcoders area on the System page.
10. Add the transcoder to the System Management App with the following settings:
 - Transcoder Type: vantage
 - Set the Working Directory as follows:
 - Use a CIFS UNC path pointing to the IP address of the Vantage Cache computer. Include the required authentication information for the shared Vantage Cache folder.
 - * Include the path to the shared Vantage Cache folder.
 - * If the Actor is Linux-based, you can have the Actor automatically create the mount point to the Vantage Cache, or you can set your own fixed mount point by setting the actorPath parameter. See the section [Preparing a Fixed Mount Point for Linux-based Actors \(optional\)](#) for more information.
 - If the Vantage Cache is located on a different computer than the Vantage SDK service (different IP address), you must tell the Actor the IP address where the transcoder service is located. Set the transcoder parameter to point to the address of the Vantage SDK service computer.
 - Leave the remaining options with the default settings. The following is an example Working Directory entry with a fixed mount point:


```
[actorPath:/mnt/vantagecache,transcoder:10.145.40.81],cifs://user\domain:pass@\\10.145.40.81\VantageCache
```

Configuring Telestream Vantage

The following sections describe only the configuration for the Vantage transcoder.

Creating the Output Path

Use the following procedure to create the output path in Vantage:

1. Open the Vantage Management Console and connect to the local computer.
2. In the left navigation tree, navigate to Workflow Design Items, Variables, Create New Variable.
3. Use the menu list to set the Select the variable type parameter to Path.
4. Click OK.
5. At the bottom of the screen, update the Name field to OutputPath.
6. Click the Save icon to save the variable.

Creating a Minimum Vantage Workflow

Use the following procedure to create the minimum Vantage workflow. First, create the workflow and link the Receive and Flip together as follows:

1. Open the Vantage Workflow Designer.
2. Create a New Collection.
In the example the Name is TESTMINWorkflow.

Note: No spaces or special characters are allowed in the Collection name.

3. Create a New Workflow and enter a name for it in the Enter a name field.
4. Select the Collection for the workflow from the Select a Collection for the new workflow list.
5. Optionally, enter a description in the Enter a description field if desired.
6. If desired, set the number of hours for the workflow to expire, and select the Expire after check box.
7. Click OK to save the workflow.
8. Click the Common icon, and then click Receive.
9. Click the Transcode icon, and then click Flip.
10. To link the Receive and Flip together, click the Receive yellow dot and drag it to the Flip yellow dot.

Next, configure the Flip options as follows:

1. Right-click Flip to configure the Flip options. For this example a media file is being configured using the following settings:
 - Encoder: Apple 3GP
 - Input media file nickname: Original
 - Output media file nickname: Mobile
2. Expand the Output Location section.
3. Select the Path option, and then enter, or browse, to select the output path (for example, E:\VantageStore).
4. Use the menu list to select the Collision Resolution. This identifies what the software will do if there is an existing file in the output path with the same file name. Initially set the Collision Resolution field to Overwrite.
5. Click Save to save the configuration.

Next, configure the Receive options as follows:

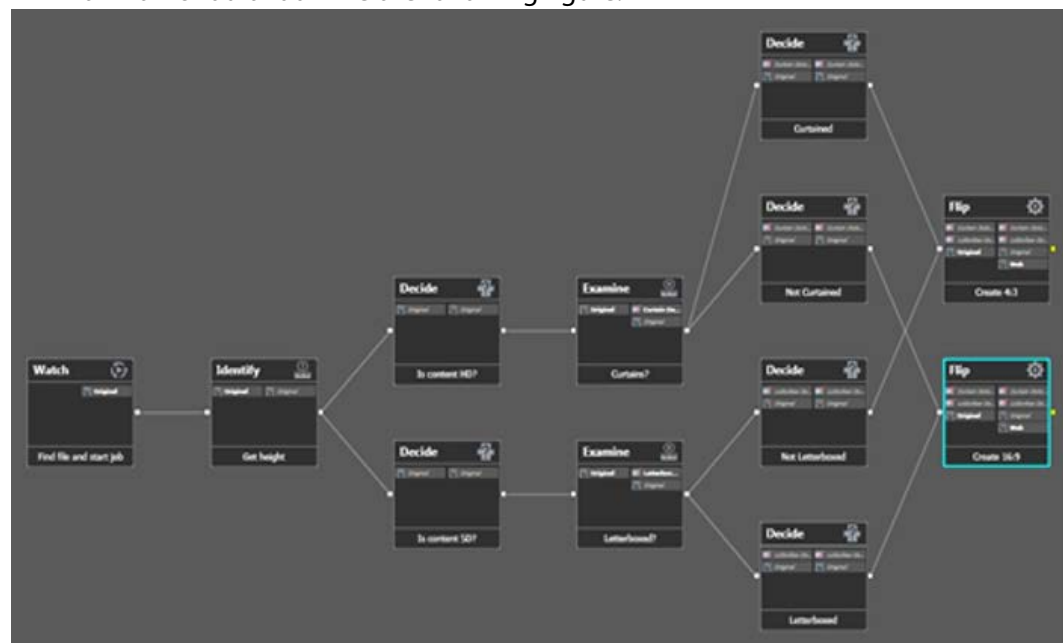
1. Right-click Receive to configure the Receive options.
2. Click the Media Files list and choose Vantage Proxy.
3. Click Save to save the changes.

When complete, click Release to enable DIVA Core to use the workflow.

Creating a Complex Vantage Workflow

This Vantage complex workflow example was created for documentation purposes; however it has not been tested with actual media files. Use the following procedure to create the complex Vantage workflow:

1. Open the Vantage Workflow Designer.
2. Navigate to File > create a New Collection. For this example the Collection created is named TESTComplex.
3. Navigate to File > Import Workflow.
4. Browse and select C:\Program Files (x86)\Telestream\Vantage\Samples\Analysis\Smart SD and HD Transcoding.xml.
5. Specify the Collection created in Step 2.
6. Technical Support recommends changing the Workflow Name to match the Collection. No spaces or special characters are allowed.
7. Delete the Watch and replace it with Receive.
8. Configure Receive and set MediaFiles to Original.
9. Link Receive with Identity.
10. Delete Deploy.
11. Configure both Flip Factories.
 - Change the Output Location to Path and then enter, or browse, to select the output path.
 - Change the Collision Resolution to Overwrite.
12. Click Release to enable DIVA Core to use the workflow. In this example, the workflow should look like the following figure.



Configuring Transcoders

Create a new Vantage transcoder as described in previous section.

Set the Working Directory to either a local folder, or a path on a remote system. Only a remote path for Vantage can be set. If setting a path to a remote system, a CIFS UNC path with the appropriate authentication credentials must be specified. The IP address specified in the UNC path must point to the remote computer running the Vantage SDK service.

Configuring Source and Destination Servers

Use the following procedure to configure a Source or Destination Server for use with transcoders:

1. Open the System Management App.
2. Navigate to the Servers area on the System page.
3. Create a LOCAL Server for the Actor using the following parameters:
 - Source Server Name: use the same name as the Actor name
 - IP Address: leave this field empty
 - Source Server Type: *LOCAL*
4. Configure the Destination Server to include the following transcode options along with any other required Connect Options:

```
-tr_names {TRANSCODER_NAME}  
-tr_restore_format {WORKFLOW_NAME}
```

Note: The auto format option is only valid for Telestream and BitScream.

For this example the Connect Options field is populated similar to the following:

```
-login diva -pass diva -tr_names vantage_001 -tr_restore_format  
TESTMINWorkflow
```

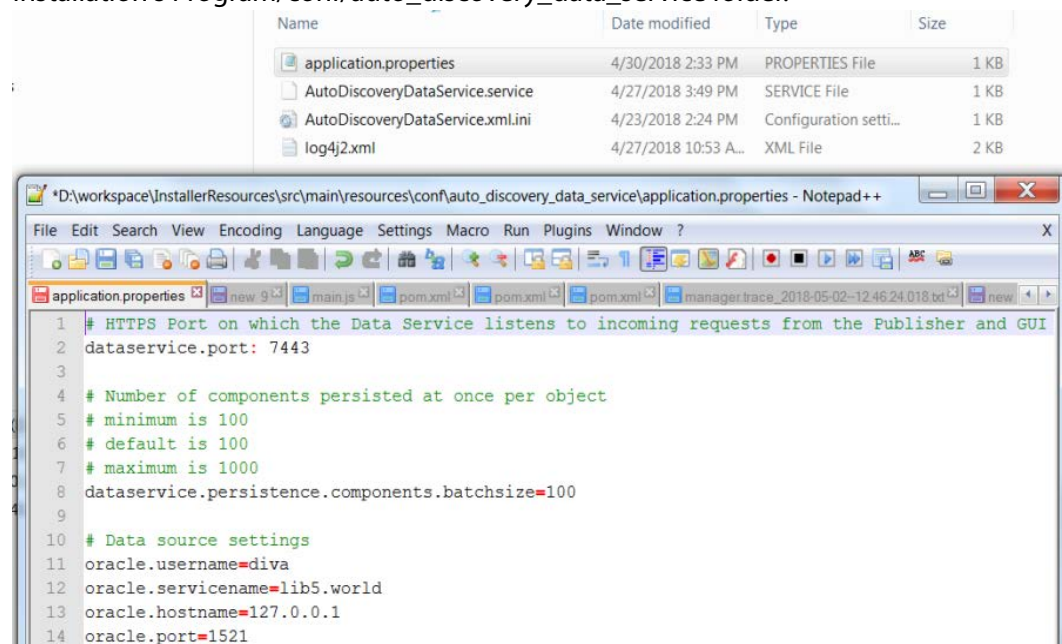
Disk Auto-Discovery

DIVA Core now supports the retrieval of non-complex objects and instance metadata from an OCI/S3 cloud account to a Core Database. This allows the user to update their on-premise DIVA Core system with content from the cloud. A new web-based GUI allows a user to start, resume, or stop a disk-discovery scan and to view the status of a scan. Refer to the DIVA Core Operations Guide for detailed information on how to use Disk Discovery.

Auto-Discovery Configuration

All Actors that can be used to discover metadata on disk must be enabled to do so by setting the Auto-Discovery flag for each Actor in the Actor Settings panel in the System Management App. Only Actors enabled for Auto-Discovery will be used by one of the Publisher's Request State Machines to discover instances and objects during the scan of a container. Technical Support recommends that multiple Actors are configured for Auto-Discovery to maximize performance.

Before to running the Data Service, the port and Core Database settings must be configured by editing the application.properties file located in the DIVA Core installation's Program/conf/auto_discovery_data_service folder.



Next, the Publisher's application.properties file is located in the DIVA Core installation's Program/conf/auto_discovery_publisher_service folder must be modified to update the port the Publisher listens to for HTTPS requests (GUI) and TCP messages (Actor). Users can also specify whether the disk discovery process should overwrite instances and/or objects. This process will delete any instances and (or) objects with the same name and Collection before writing the new metadata. Users can optionally

synchronize deleted instances by deleting all metadata related to instances that are no longer on the corresponding OCI disk.

application.properties	6/4/2018 8:13 AM	PROPERTIES File	1 KB
AutoDiscoveryPublisherService.service	4/27/2018 5:47 AM	SERVICE File	1 KB
AutoDiscoveryPublisherService.xml.ini	5/16/2018 5:43 PM	Configuration setti...	1 KB
log4j2.xml	5/14/2018 5:11 PM	XML File	2 KB


```

D:\workspace\InstallerResources\src\main\resources\conf\auto_discovery_publisher_service\application.properties - Notepad+
file Edit Search View Encoding Language Settings Macro Run Plugins Window ?
new 4 new 10.bt DiskDiscoverySQL.bt DIVA8.bt settings.xml tmp.bt NOSQL-METADB.bt Obj
1 # HTTPS Port on which the Publisher listens to incoming requests from the GUI
2 publisher.port: 8443
3
4 # Secure TCP Port on which Publisher listens for incoming messages from Actors.
5 actor.listen-port: 11443
6
7 # Determines whether the Publisher will overwrite existing instances.
8 statemachine.overwrite-instances=false
9
10 # Determines whether the Publisher will overwrite existing objects.
11 # Note: This will only overwrite objects with a single disk instance.
12 # The associated instance must be on the disk under scan.
13 statemachine.overwrite-objects=false
14
15 # Determines whether the Publisher will synchronize deleted instances.
16 # Note: This will delete all instances of a disk that are NOT reported
17 # by Auto-Discovery actors. Consequently, any objects with no instances
18 # will also be deleted if this parameter is set to true.
19 statemachine.synchronize-deleted-instances=false
20
21 # Request timeout
22 spring.mvc.async.request-timeout=600000
  
```

The log4j2.xml files located in both the Data Service and Publisher's configuration folders can be updated to one of the available log levels (ERROR, WARN, INFO, or DEBUG) by modifying the level attribute of the AsyncLogger element. It is info by default.

```

<Loggers>
  <AsyncLogger name="com.oracle.diva.autodiscovery.dataservice" level="info"
    additivity="false">
    <AppenderRef ref="FileAppender" />
  </AsyncLogger>

  <Root level="error">
    <AppenderRef ref="FileAppender" />
  </Root>
</Loggers>
  
```


Auto-Discovery Security

The Auto-Discovery GUI is accessible at the Auto-Discovery Publisher's address and port. To communicate securely with the Publisher and Data services, the browser hosting the GUI must import the AutoDiscoveryPubService.p12 and AutoDiscoveryDataService.p12 certificates from the usual %DIVA%/security/certificates folder as a Trusted Root Certificate Authority. The certificates must be generated by running the DIVASecurity Tool. However, before running the DIVASecurity Tool, the IPs in %DIVA%/security/conf/AutoDiscoveryHost.cnf file, and the browser running the GUI, must be updated to the IP addresses of the computers hosting the Publisher and Data services.

Insert or update the IPs with IPs of the services and browser running the GUI:

```

1 [req]
2 default_bits = 2048
3 prompt = no
4 default_md = sha256
5 x509_extensions = v3_req
6 distinguished_name = dn
7
8 [dn]
9 c = US
10 ST = NJ
11 L = Mount Laurel
12 O = Oracle
13 OU = Media
14 emailAddress = diva@oracle.com
15 CN=localhost
16
17 [v3_req]
18 subjectAltName = @alt_names
19
20 [alt_names]
21 DNS.1=localhost
22 # NOTE: Update the list of IPs below to the computers hosting the Publisher and Data services as well as the IP of the computer
23 # hosting the browser used to view the GUI.
24 # Then run the DIVASecurityTool to generate the AutoDiscoveryPubService.p12 and AutoDiscoveryDataService.p12 certificates.
25 # You must import these certificates into your browser to securely communicate with the Auto-Discovery services.
26 IP.1=10.39.227.78
27 IP.2=10.39.227.8
28

```

Run the DIVASecurity Tool:

```

C:\> Command Prompt
===== Security Administration Menu =====

1: Reset Key Store and Trust Store Password.
2: Generate new Keys & SSL Certificates for the services signed by DIUA_CA.
3: Generate Certificate Signing Request (CSR) for DIUA_CA.
4: Install External Certificate Authority.
5: quit.

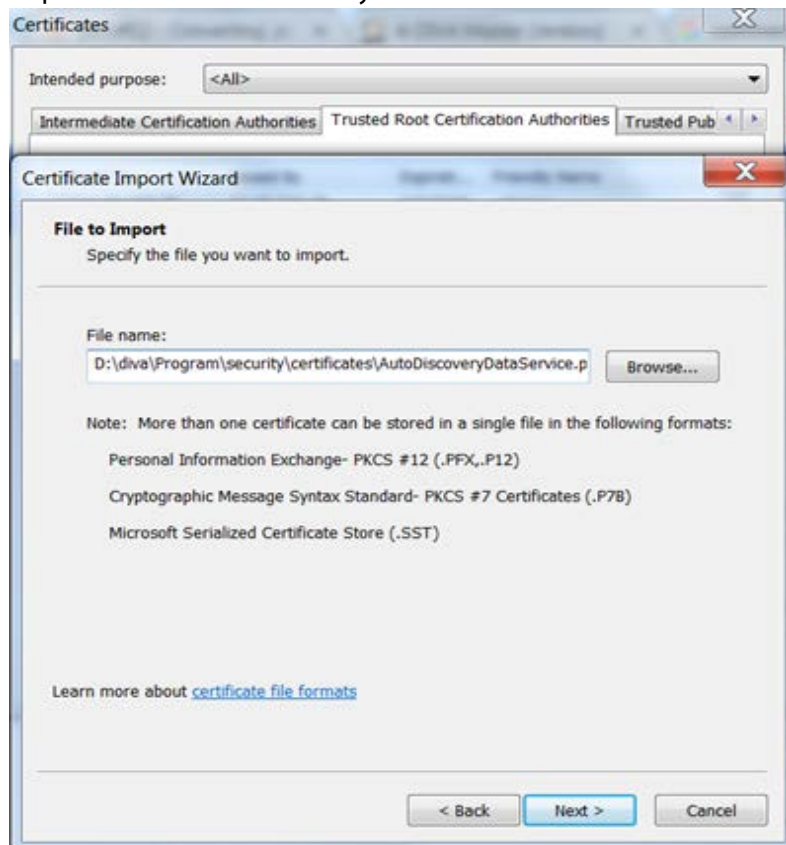
Please make a selection: 2
Please enter current password: xxxxxxxx
Generate ActorService Key
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'D:\security\keys\ActorServiceKey.pem'
-----
Sign ActorService CSR with DIUA_CA
Signature ok
subject=C = US, ST = NewJersey, L = MtLaurel, O = Oracle, OU = Media Storage, CN = ActorService
Getting CA Private Key
Generate RobotManagerService Key
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'D:\security\keys\RobotManagerServiceKey.pem'
-----
Sign RobotManagerService CSR with DIUA_CA
Signature ok
subject=C = US, ST = NewJersey, L = MtLaurel, O = Oracle, OU = Media Storage, CN = RobotManagerService
Getting CA Private Key
Generate SNMPService Key
Generating a 2048 bit RSA private key
.....+++
.....+++

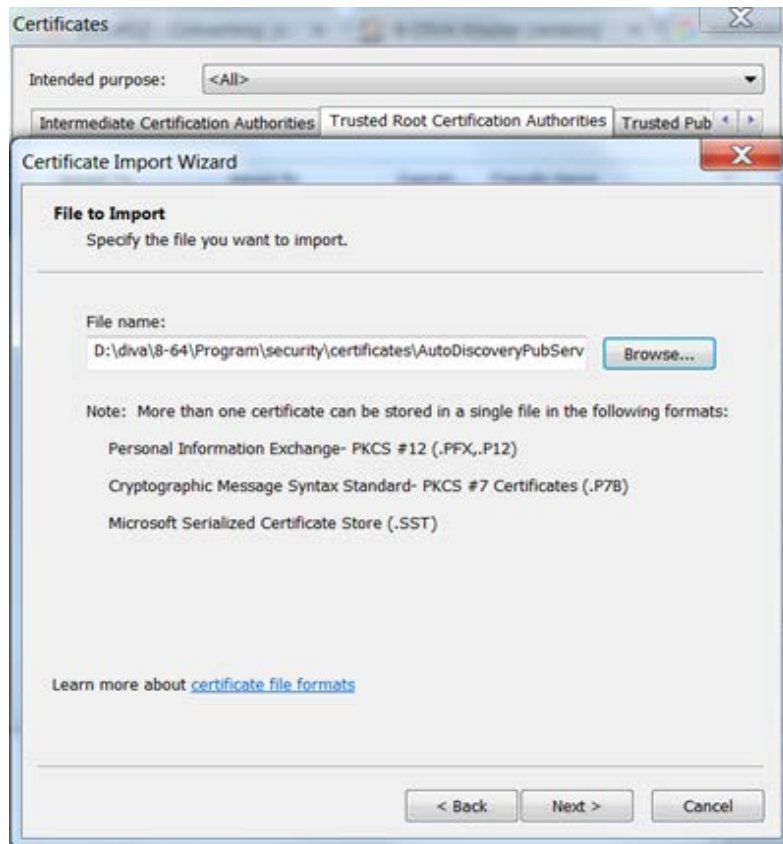
```

The tool generates the required AutoDiscoveryDataService.p12 and AutoDiscoveryPubService.p12 certificates:

```
Generating PKCS12 Certificate for AutoDiscoveryPubService. NOTE: THIS MUST BE IMPORTED INTO YOUR BROWSER AS A TRUSTED ROOT CERTIFICATE AUTHORITY.
Generating PKCS12 Certificate for AutoDiscoveryDataService. NOTE: THIS MUST BE IMPORTED INTO YOUR BROWSER AS A TRUSTED ROOT CERTIFICATE AUTHORITY.
Import ManagerService Certificate and keys to D:\security\certificates\DIUA.jks
Import AutoDiscoveryPubService Certificate and keys to D:\security\certificates\DIUA.jks
Import AutoDiscoveryDataService Certificate and keys to D:\security\certificates\DIUA.jks
Import ControlGUI Certificate and keys to D:\security\certificates\DIUA.jks
Import ConfigurationUtility Certificate and keys to D:\security\certificates\DIUA.jks
Import DBBackupService Certificate and keys to D:\security\certificates\DIUA.jks
Import MigrationService Certificate and keys to D:\security\certificates\DIUA.jks
Import DFMService Certificate and keys to D:\security\certificates\DIUA.jks
Import APIJava Certificate and keys to D:\security\certificates\DIUA.jks
Import ROTU Certificate and keys to D:\security\certificates\DIUA.jks
Import ManagerAdapterService Certificate and keys to D:\security\certificates\DIUANet.jks
Import DBSyncService Certificate and keys to D:\security\certificates\DIUANet.jks
Import ClientAdapterService Certificate and keys to D:\security\certificates\DIUANet.jks
Import DIUANetGUI Certificate and keys to D:\security\certificates\DIUANet.jks
Import DIUADatabaseServer Certificate and keys to D:\security\certificates\DIUADatabaseServer.jks
Certificate was added to keystore
Key refresh completed successfully
Press any key to continue . . .
```

Import Both certificates into your browser as Trusted Root Certification Authorities:





Note: The Data and Publisher services must be restarted to use these certificates for secure communication after generating the new certificates.

Auto-Discovery Install, Start, Stop and Uninstall


Run the menu script located in Program/AutoDiscovery folder as Administrator to install, start, stop, and uninstall the Data and Publisher Services.

```

===== Auto-Discovery Administration Menu =====
1: Install Data Service
2: Install Publisher Service
3: Start Data Service
4: Start Publisher Service
5: Stop Data Service
6: Stop Publisher Service
7: Uninstall Data Service
8: Uninstall Publisher Service
9: Install and Start All Services
10: Uninstall and Stop All Services
0: Quit

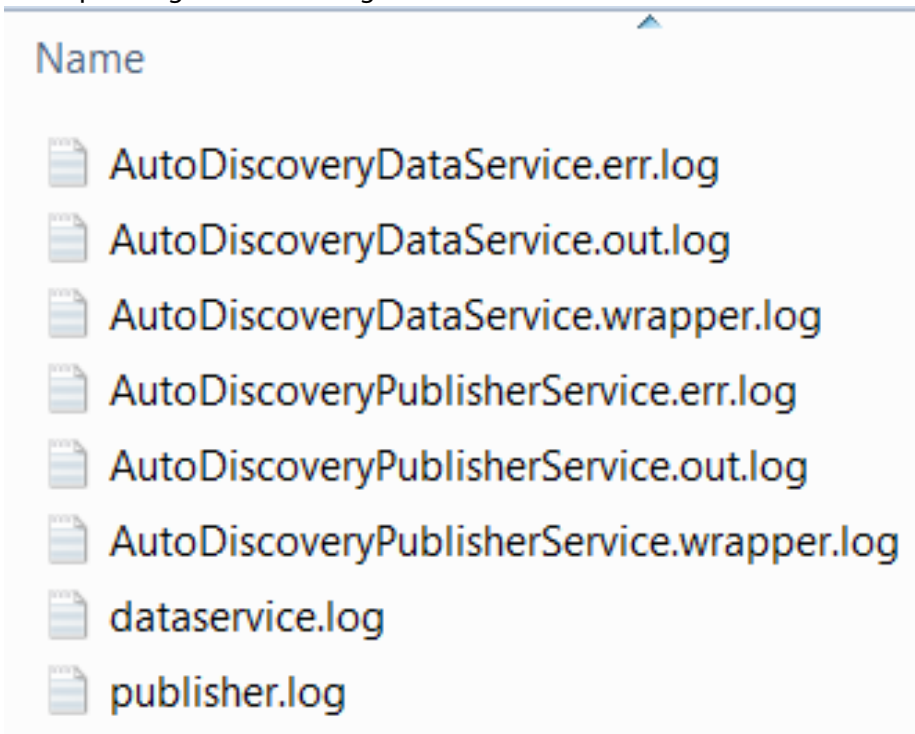
Please make a selection: 9
    
```

On Windows, the services can also be started and stopped from the Services window.



Auto-Discovery Logging

Both Service and Trace logs are generated and written to the Program/log/autodiscovery. The log file output is controlled by the log4j2.xml file in the corresponding service's configuration folder.



Frequently Asked Questions

Contact Technical Support for any additional questions not covered here.

Topics:

- [General DIVA Core Questions](#)
- [DIVA Core Database and Backup Service Questions](#)

General DIVA Core Questions

- What Happened to DIVA Command?
 DIVA Command has been replaced by the System Management App starting with the DIVA Core 8.3 release.
- What if the Customer Information Collection Tool does not work?
 Confirm that CygWin and the 7z archive programs are installed correctly. If the CygWin or the 7z programs are not installed, the Customer Information Collection Tool will stop running and display one of the following error messages:

```
Error: Cygwin environment could not be located at "...". Please check the configuration or reinstall Cygwin environment if necessary.
```

```
Error: 7Z archiver could not be located at "...". Please check the configuration or reinstall 7Z archiver if necessary.
```
- Should all operating systems be kept up to date with critical updates?
 Technical Support recommends applying all critical updates to all computers because some may include security updates. Windows operating system updates and patches are not tested by Telestream.
- Should all operating systems be kept up to date with optional updates?
 Optional operating system updates are not necessary in the DIVA Core environment and are not tested by Telestream. However the decision to apply optional updates is left to your System Administrator.
- Are there any operating system updates that should not be installed?
 Technical Support is not currently aware of any operating system updates that impact DIVA Core functionality or operations. However, operating system updates and patches are not tested by Telestream.
- Should the servers be restarted with any frequency?
 No, restarting the servers will cause downtime for the system and possibly cause data corruption if a process is executing when the server is restarted. Only restart a server when absolutely necessary and perform a normal system shutdown.

Caution: Do not just power off the computer unless absolutely necessary. Data and/or database corruption or loss could occur if normal shutdown procedures are not followed. See the DIVA Core Operations Guide for proper shutdown procedures.

- Should any services be restarted with any frequency?
 No, restarting the services will cause downtime for the system and possibly cause data corruption if a process is executing when the service is restarted. Only restart a service when absolutely necessary.
- Should any vendor applications be restarted with any frequency?
 No, only restart a vendor application when absolutely necessary.

- Should vendor applications always be updated to the latest version?
 No, only update vendor applications to benefit from new functionality or for bug fixes.
- What is the recommended frequency of database backups?
 The Core Database automatically backs up every fifteen minutes.
- Does Technical Support recommend any particular database backup application?
 A database backup service is provided in the DIVA Core package. You are welcome to use your own backup software as an additional security under the condition that you only backup the Core Database backup files (in H:\oraback) and not the database itself.
 Backing up the database directly is forbidden. For example, using Oracle RMAN or other non-Core Database backup applications. Backing up the database directly with another program may interfere with the DIVA Core Database Backup Service. This may render database restoration impossible using the embedded DIVA Core restore utility, and could possibly result in data losses for which our company will accept no responsibility.
- Where are the backup files located?
 The database backup files are located on the Main Manager computer in the H:\oraback folder. The files are synced to the Backup Manager and an Actor in the H:\oraback\mgr1 folder.
- Are there iterated versions of the database backup, and if so, how many are retained?
 The backup files are retained for the previous ten days. The retention period is configurable for the database backup files in the DIVA Core Backup Service configuration file. Contact Technical Support for assistance.
- Where are vendor-specific logs located?
 The vendor-specific log files are located in the %DIVA_HOME%\Program\log folder in Windows and in the /home/diva/DIVA/Program/log directory in Linux.
- How far back in time do the logs go?
 The log file retention period is configurable in the DIVA Core configuration file. Contact Technical Support for more information. The log files are retained as follows by default:
 - Manager, DIVA Connect: fifty hours
 - Actor, Robot Manager, Storage Policy Manager, Avid Transfer Manager Communicator, Avid Archive Manager Communicator: 10 days
 - Watch Folder Monitor: variable based on size
- What is the suggested log backup frequency?
 The log files do not require backup.

- Are there any special considerations regarding maintenance and backup of vendor servers and systems?

Technical Support only supports the DIVA Core software. You must contact the server supplier for any hardware issue. You must keep Technical Support in the loop for any issues on the DIVA Core solution (for example, loss of a RAID disk, failover to the backup manager, and so on).

- Are there any special considerations related to recovering from a server failure when the server is part of the vendor solution?

As previously mentioned, you must keep Technical Support in the loop if issues are encountered.

DIVA Core Database and Backup Service Questions

- How do I failover to a Backup System when the Main Manager System has failed?
See [Manager Failover Procedures](#) for the complete procedure.

- How do I recover when a Complex Object's Metadata file is corrupted in the Main Manager System?

The DIVA Core Backup Service backs up the Metadata Database file by file. After the file is backed up to the backup systems, any corruption to, or modifications of, the Metadata files are not propagated to the backup systems.

If a Complex Object Metadata file is corrupted, restore the Metadata file from one of the backup systems.

In the unlikely event of disk corruption due to hardware failure occurring before the Backup Service has backed up the Metadata files, the non-backed up Metadata files can only be restored from a tape or disk. The feature to restore Metadata files from tape or disk is not currently available in this DIVA Core release. Contact Technical Support for assistance.

- How do I recover a Complex Object's Metadata file when it is corrupted in the Backup Manager System?

Technical Support recommends always making backup copies to two separate backup systems to handle these scenarios. Restore the Metadata file from the Secondary Backup System or Main Manager System.

- When a Metadata file is manually deleted from Main Manager System, is it also deleted from all backup systems?

Manually deleted Metadata files are not propagated to any backup systems.

- How do I recover when a Complex Object's Metadata File is lost on the Main Manager System and all backup systems?

You can restore Metadata files from tape or disk. The feature to restore Metadata files from tape or disk is not currently available in this DIVA Core release. Contact Technical Support for assistance.

- How do I recover when the backup disk fails, or gets corrupted, on the Main Manager System?

Disk failures, or corruption, requires a failover to the Backup Manager. See [Manager Failover Procedures](#) for the complete procedure.

- How do I configure a full backup to start when the Backup Service starts?

The DIVA Core Backup Service automatically determines if a full backup is required when it starts. There is no configuration required.

- How do I locate a Complex Object's Metadata inside the Metadata Database?

Contact Technical Support for assistance.

- How do I turn off GUI Backup Service Notifications?
You can turn off notifications by deselecting the check box for the Database Backup Notification parameter under the Manager Setting panel in the System Management App.
- Can the Core Manager and Database be installed on separate servers?
No, they must be installed on the same server because the DIVA Core Backup Service does not support Manager and Oracle installations on separate servers in this DIVA Core release.
- Does the recovery window apply to both Oracle Secure Backups and Metadata Backups?
Yes, the recovery window setting applies to both backups.
- How do I estimate the size for the Metadata Database location?
See [Sizing the Metadata Database](#) for detailed information.
- Where do I configure the location of the Metadata Database?
The location of the Metadata Database is configured using the Complex Objects Metadata Location parameter in the Manager Setting panel in System Management App.
- What information is stored in the Metadata Database file?
All file details including file names, folder names, location, size, checksums, and so on.
- Is the information stored in the Metadata Database irreplaceable or mission critical?
Technical Support always recommends having at least two backup copies of the Metadata Database. Use the DIVA Core Backup Service to back up the Metadata Database. In a worst case scenario, use the Archive eXchange Format Explorer (AXF) to recover the object from tape if the Metadata Database file of a particular object is lost.
- Why is this information not being stored in the existing Database?
The amount of Metadata information is huge. Complex Objects are supported up to 1,000,000 files. Currently, the Database in use does not have any scalability features to support Complex Object workflows.
- What are the space requirements for the Metadata Database and data? Does it depend on the quantity of objects, the complexity of those objects, or something else?
See [Sizing the Metadata Database](#) for detailed information.
- What if a customer has, for example, 1,000,000 objects, each with 100,000 files?
The Metadata Database is very scalable and can handle this amount with no issues.
- What are the consequences of the Metadata Database becoming inoperable, corrupt, or missing? Will data loss, performance loss, or something else occur?
You will not be able to process Complex Object requests if the database becomes inoperable. You can restore from one of the backup copies if the database becomes corrupt, or is missing.

- What are the consequences of the Metadata Database running out of available storage space? Will data loss, performance loss, or something else occur?

In this case you will not be able to process any Complex Object requests. See [Sizing the Metadata Database](#) for detailed information.

- What tools exist for testing or verifying the integrity of the Metadata Database? Are the tools automatic, invoked manually, or can either method be used?

Currently there are no tools that exist to check the database integrity. Contact Technical Support if you need assistance.

- What tools exist for backing up the Metadata Database? Are the tools automatic, invoked manually, or can either method be used?

Always use the DIVA Core Backup Manager Service to back up the Metadata Database.

- What tools exist for recovering the Metadata Database if loss or corruption occurs? What is the procedure to execute recovery, and is any of the recovery automatic?

See [Metadata Database Failure Scenarios](#) for the complete procedure.

- Does the storage location of the live database affect performance or space, and is it critical?

Yes, it is both performance and space critical. See [Database Installation and Configuration](#) for installation and configuration procedures.

- Can the location of the Metadata Database backups be configured?

Yes, you can configure the backup location. See [Database Installation and Configuration](#) for DIVA Core Backup Service installation and configuration procedures.

Appendix A: Core Options and Licensing

The following table identifies Core options and licensing metrics.

Description	Licensing Metric
DIVA Core System	Per Server
DIVA Core Single	Per Server
DIVA Core Actor	Per Server
DIVA Core Avid Connector	Per Avid Archive Provider
DIVA Core Partial File Restore	Per System
DIVA Core Analytics	Per Server
DIVA View	Per Concurrent User
Managed Storage Capacity	Per 500 TB Block
Unlimited Storage Capacity	Per System

Appendix B:

Secure Deployment Checklist

1. Set strong passwords for Administrator (or root) and any other operating system accounts that have any DIVA Core administrator or service roles assigned to them, including:
 - DIVA, Oracle User IDs (if being used)
 - Any disk array administrative accounts
2. Do not use a local administrator operating system account. Assign roles as needed to other user accounts.
3. Set a strong password for Administrator and Operator for the System Management App. You must assign a password for these profiles in the System Management App before use.
4. Set a strong password for the Core Database login.
5. Install a firewall on every system and apply the default DIVA Core port rules. Restrict access to DIVA Core API (tcp/9000) to IPs that need access using firewall rules.
6. Install operating system and DIVA Core updates on a periodic basis since they include security updates.
7. Install Anti-virus and exclude the DIVA Core processes and storage (for performance reasons).
8. It is best practice to segregate FC disks and FC tape drives either physically or through FC Zoning so that disks and tape devices do not share the same HBA port. For Managed disks, only Core Actors should have access to disk and the tape drives. This security practice helps prevent loss-of-data accidents resulting from accidental overwriting of tape or disk.
9. Set up an appropriate set of backups of the DIVA Core configuration and database. Backups are part of security and provide a way of restoring data lost either accidentally, or through some type of breach. Your backup should include some policy while being transported to an off-site location. Backups need to be protected to the same degree as DIVA Core tape groups and disk.

Technical Support strongly recommends using an external CA for additional security.

Appendix C: Server Guide

This appendix describes Source and Destination Server configuration guidelines for each type of DIVA Core supported content server. See the DIVA Core Supported Environments Guide for detailed and up-to-date lists of supported content servers, formats, and related DIVA Core platforms.

Topics:

- [General Parameters](#)
- [Alto Disk Archive Integration](#)
- [Avid MSS \(Program Stream\) Servers](#)
- [Avid Airspace Servers](#)
- [Avid Transfer Manager DHM Interface](#)
- [Avid Transfer Manager DET Interface](#)
- [SeaChange BMS and BMC Servers](#)
- [SeaChange BML Servers](#)
- [SeaChange BMLe and BMLex Servers](#)
- [Leitch vR Series Servers](#)
- [Leitch Nexio Servers](#)
- [Grass Valley Profile Servers](#)
- [Grass Valley UIM Gateway](#)
- [Grass Valley K2 Servers](#)
- [Grass Valley M-Series iVDR Servers](#)
- [Sony MAV70 Servers](#)
- [Omneon Spectrum MediaDirector Servers \(QuickTime\)](#)
- [Omneon MediaGrid Content Storage System](#)
- [Quantel Power Portal Gateway](#)
- [Sony Hyper Agent Servers](#)
- [Standard FTP and SFTP Servers](#)
- [Local Source Servers](#)

- [Disk and CIFS Source Servers](#)
- [Metasources](#)
- [Expedat Servers](#)

General Parameters

This section introduces general items that may apply to any, or most, servers including features, configuration attributes, and connection options.

Files Path Root Parameter

The Files Path Root (FPR) parameter is for Archive and Restore requests. This parameter specifies the root folder for data transfers and applies to any type of Server.

An absolute or relative path can be entered in the Files Path Root field. This parameter is limited to 260 characters.

Each content server section of this appendix specifies the expected format of the Files Path Root and related File Names parameters for Archive requests.

For Partial File Restore requests, the file names on the destination are those specified when archiving. If no Files Path Root is entered, DIVA Core uses the one specified during archiving.

Root Path Parameter

The Root Path is a Server attribute that can use as a default path for FTP-like Servers, or as a disk mount point for disk and local Source Servers. This applies to any type of Server. The path is appended before any Files Path Root specified in requests, unless the path specified in a request is an absolute path.

This approach provides better Server abstraction. You specify the server directories used by DIVA Core at the configuration level, not at the request level. They can be changed at any time without requiring a change to DIVA Core clients.

The Root Path value is always an absolute path defined by the operating system. An Omneon Path is the player name and always considered an absolute path.

Absolute path names are supported on both Windows and Linux to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems (only).

If you leave the Root Path field empty, DIVA Core ignores the parameter. However, if you do specify a Root Path its value is combined with the Files Path Root you specified in a request to give the final Server path. This process is performed according to the following rules:

- Relative paths are added to the absolute path, absolute paths override preceding absolute paths (standard Path Arithmetic).
- If the Root Path and Files Path Root have different operating system types, the second path (Files Path Root) is converted to the operating system type specified by the first path (Root Path) by replacing \ with / (and vice versa). The converted path is then considered the relative path.

- If the Root Path ends with a > character, the Files Path Root is always considered to be a relative path, and the > character is omitted during concatenation.

Server ROOT_PATH	Object: Original_FPR recorded in database & metadata	Request Type	Files Path Root (FPR)	Resulting rule applied to create actual path for the transfer	Resulting path considered for the transfer	Resulting original Files Path Root (FPR) recorded in database and metadata
Null		Archive	Null	ROOT_PATH +FPR	Null	Null
Null		Archive	Set	ROOT_PATH +FPR	FPR	FPR
Set		Archive	Null	ROOT_PATH +FPR	ROOT_PATH	Null
Set		Archive	Set	ROOT_PATH +FPR	ROOT_PATH+FPR	FPR
Null		Archive with tr_arch format	Null	ROOT_PATH +FPR	Null	Null
Null		Archive with tr_arch format	Set	ROOT_PATH +FPR	FPR	Null
Set		Archive with tr_arch format	Null	ROOT_PATH +FPR	ROOT_PATH	Null
Set		Archive with tr_arch format	Set	ROOT_PATH +FPR	ROOT_PATH+FPR	Null
Null	Null	Restore	Null	(ROOT_PATH +FPR) Original_FPR	Null	
Null	Null	Restore	Set	(ROOT_PATH +FPR) Original_FPR	FPR	

Server ROOT_PATH	Object: Original_FPR recorded in database & metadata	Request Type	Files Path Root (FPR)	Resulting rule applied to create actual path for the transfer	Resulting path considered for the transfer	Resulting original Files Path Root (FPR) recorded in database and metadata
Set	Null	Restore	Null	(ROOT_PATH+ FPR) Original_FPR	ROOT_PATH	
Set	Null		Set		ROOT_PATH+FPR	
Null	Set		Null		Original FPR	
Null	Set		Set		FPR	
Set	Set		Null		ROOT_PATH	
Set	Set		Set		ROOT_PATH+FPR	
	Null	Transcode Archive				Null
	Set	Transcode Archive				Null

UNIX Style Paths

The following table describes UNIX style paths for the Root Path, File Path Root, and the actual path to the files.

Root Path (Server)	File Path Root (Request)	Actual Path to Files
/diva/upload	tmp	/diva/upload/tmp
/diva/upload	/tmp	/tmp
/diva/upload		/diva/upload
/diva/upload	C:\tmp	/diva/upload/C:\tmp (!!!)
/diva/upload>	/tmp	/diva/upload/tmp
/diva/upload>	\tmp	/diva/upload/tmp
/diva/upload>		/diva/upload

Windows Style Paths

The following table describes Windows style paths for the Root Path, File Path Root, and the actual path to the files.

Root Path (Server)	File Path Root (Request)	Actual Path to Files
D:\diva\upload	tmp	D:\diva\upload\tmp
D:\diva\upload	C:\tmp	C:\tmp
D:\diva\upload		D:\diva\upload
D:\diva\upload>	/tmp	D:\diva\upload\tmp
D:\diva\upload>	C:\tmp	D:\diva\upload\tmp
D:\diva\upload>	C:/tmp	D:\diva\upload\C:\tmp
D:\diva\upload>		D:\diva\upload

Metasource Parameter

The Metasource parameter is a specific type of Server to manage several Servers sharing the same online storage as one (or multiple Watch Folder Monitors) with failover and load-balancing features. This applies to any type of Server. See [Metasources](#) for more information on the Metasource Server types.

Connect Options Parameter

Connect Options are a Server parameter used to specify the communication protocol with the Server or to modify the protocol's defaults.

Some options exclusively apply to a specific Server type, and are documented as part of that specific Server type later in this appendix. Others options are for general use and are documented in this section.

Some Connect Options (explicitly or implicitly) specified for the Server may be superseded by those specified in requests. This section also specifies, for each Connect Option, whether it can be superseded at the request level.

Quality of Service (qos=)

This option specifies the transfer mode used when transferring from this specific Server when the archive initiator sets the QualityOfService parameter in Archive or Restore parameters to DEFAULT.

This parameter applies to any type of Server, and cannot be superseded by the request option.

If the archive initiator sets the QualityOfService to something other than DEFAULT, DIVA Core ignores the qos= Connect Option.

The format for the parameter is qos=[DIRECT_AND_CACHE|CACHE_AND_DIRECT].

Note: This option must be the first one in place in the Server Connect Options field, and must always be specified in lowercase.

The valid values for Quality of Service are as follows:

DIRECT_AND_CACHE

Direct transfers from (or to) a Server to (or from) DIVA Core are preferred, but cache transfers will occur if processing the request in direct mode is not possible.

CACHE_AND_DIRECT

Cache transfers from (or to) a Server to (or from) DIVA Core are preferred, but direct transfers will occur if processing the request in cache mode is not possible.

The following table describes sample Quality of Service connections:

QOS Connect Option	QOS Set by the Archive Initiator	Actual Transfer Mode Applied by the Core Manager
DIRECT_AND_CACHE	DEFAULT	DIRECT_AND_CACHE
DIRECT_AND_CACHE	DIRECT_ONLY	DIRECT_ONLY
DIRECT_AND_CACHE	CACHE_ONLY	CACHE_ONLY
CACHE_AND_DIRECT	DEFAULT	CACHE_AND_DIRECT
CACHE_AND_DIRECT	DIRECT_ONLY	DIRECT_ONLY
CACHE_AND_DIRECT	CACHE_ONLY	CACHE_ONLY
	DEFAULT	DEFAULT (that is, DIRECT_AND_CACHE)
	DIRECT_ONLY	DIRECT_ONLY
	CACHE_ONLY	CACHE_ONLY

Server FTP User Log In (-login)

This option is generally used to specify a user name to connect to a Server when the transfer protocol is FTP or FTP-like, and is in the format -login {user_name}.

This option applies when specified in Server type description, and can be superseded by the request option.

Possible values applicable to a specific Server type are detailed in the related section later in this appendix.

Server Swift (-oracle_storage_class)

This option is generally used to specify the class of storage to connect to a SWIFT Server and is in the format `oracle_storage_class={ARCHIVE|STANDARD}`.

Server CIFS User Log In (-user)

This option is generally used to specify a user name to connect to a CIFS Server, and is in the format `-user {user_name@domain}`.

This option applies when specified in Server type description, and can be superseded by the request option.

Possible values applicable to a specific Server type are detailed in the related section later in this appendix.

Server Password (-pass)

This option is generally used in combination with the `-login` option, and is in the format `-pass [password]`.

This option applies when specified in Server type description, and can be superseded by the request option.

Possible values applicable to a specific Server type are detailed in the related section later in this appendix.

Server Connection Port (-port)

This option is used when a port parameter is required to connect to a Server, and specifies the port number in the format `-port [port_number]`.

This is an integer value that applies when specified in Server type description, and can be superseded by the request option.

Possible values applicable to a specific Server type are detailed in the related section later in this appendix.

Deleting Source Server Content after Archiving (-allow_delete_on_source)

This parameter specifies if an Archive request can use the Delete on Source QOS option, and is in the format `-allow_delete_on_source`.

The Archive request optional parameter `delete_on_source` instructs DIVA Core to delete the original asset on the Source Server after the archive of the object is successfully completed.

If this option is specified in an Archive request and the Source Server Type is not LOCAL, DISK or CIFS, DIVA Core automatically terminates the request.

This parameter applies to the FTP_STANDARD Source Server Type. you can change this behavior so that requests will not fail when delete_on_source is specified in an Archive request.

If the -allow_delete_on_source option is specified, and the delete_on_source parameter is specified in an Archive request, DIVA Core will attempt to delete the asset from the Source Server after the archive has been completed successfully.

This option cannot be superseded by the request option.

Archiving and Restoring Filename and Path Renaming Rules (-arch_renaming, -rest_renaming, -arch_path_renaming, -rest_path_renaming)

This feature is available for Archive and Restore requests. There are no pre-defined set of values for these options. Option values are based on regular expressions. Possible values for these options are infinite and fully customizable.

Renaming rules are associated with Server. Configure filename or path renaming during archive or restore using the System Management App. The configuration can be superseded by the request option.

Use these parameters when a workflow implementation requires automatic filename or path renaming during object archiving, when the object is (partially) restored, or when a transcoded object is re-archived or restored.

Rename files at archive time (-arch_renaming) or at restore time (-rest_renaming). Rename relative path at archive time (-arch_path_renaming) or at restore time (-rest_path_renaming). The format for these parameters are as follows:

```
-arch_renaming [renaming_rule]+
-rest_renaming [renaming_rule]+

-arch_path_renaming [renaming_rule]+
-rest_path_renaming [renaming_rule]+

renaming_rule =
[activation_format:expression_patterns:output_format]
```

The -arch_renaming option enables renaming files during the archive process. This option is typically used for the following example cases:

- A file extension must be added to archived files.
- When associated to a transcoder cache (Local Server), archive renaming rules can be set to rename the files of a transcoded clip. This is useful when files created by the transcoder do not have the expected names.

The -rest_renaming option enables renaming of files during the restore process. This is typically used when the Server requires strict naming of files, and the files being transferred do not follow these rules.

This option is available for Restore, Partial File Restore (this is an alternate way to rename partially restored files), and N-Restore. If multiple renaming rules are defined, DIVA Core will process the rule for each Server independently.

The `-arch_path_renaming` and `-rest_path_renaming` options enable renaming relative paths for files at archive and restore time. The relative path to be renamed is not the Path Root, it is the relative path between the Path Root and the files.

At least one `renaming_rule` must be specified for the option. All renaming rules are located in the System Management App except the Service Name and Port parameters. DIVA Core goes through each `renaming_rule` for each file on the list to be transferred starting with the first one:

- The rule is applied if the file name matches this rule's `activation_format`.
- The condition is satisfied if the beginning of a file name matches the evaluation condition of the first rule.

For example, a condition such as `.*\track` will be satisfied by all of the following file names - `audio.track1`, `audio.track2`, `video.track`.

- As soon as a rule is applied for a given file, other rules from the list are no longer considered.

If none of the rules can be applied, the file is not renamed. An `activation_format` is a regular expression (regexp) to check whether the renaming rule must apply. This is useful when renaming paths because the relative path of each file is checked using the activation format. For example, DIVA could rename the path of some files depending on file extensions.

The `expression_patterns` parse the file name. It is a regular expression, which will include up to nine special symbols to identify different parts of the file name: `\1 \2 \3 \4 \5 \6 \7 \8 \9`.

The `output_format` is an expression that qualifies the format of a renamed file, based on atomic items (`\1` through `\9`) previously identified when applying `expression_patterns` to the original file name. Two additional specific symbols can be used:

- `\o` indicates the object name
- `\c` indicates the object Collection

Note: Describing how regular expression pattern matching works is beyond the scope of this document. There are many web sites on this subject such as <http://www.regular-expressions.info/>.

The following examples describe different possible scenarios and their associated outcomes using these parameters.

Example One

To add the `.gxf` extension to all files archived from GVG Profile (by default, these files do not have an extension). If a file does have an extension, the `.gxf` extension will not be added. Use the following statement to achieve this:

```
-arch_renaming <.*\..*:(.*)\.(.*):\1.\2><.*:(.*):\1.gxf>
```


With this option all the files under media will be moved to media.dir.

To help regular expression development, regular expression exercisers are available online at <http://regexone.com/> and <http://www.lornajane.net/posts/2011/simple-regular-expressions-by-example>.

To use this feature, you must know the basic regular expression syntax. You can find Regular Expression introductory information online at <http://www.hathitrust.org/>, <http://books.google.com/>, and <http://www.gutenberg.org/>.

Using a Temporary Filename when Restoring to a Destination

By default, DIVA Core restores objects to a destination using the destination filename. There are situations where it could be desired to restore to a temporary file first, and then rename the temporary file at the end of the transfer. To achieve this behavior, add `-restore_to_temp_files` to the Unmanaged Storage (source/destination) Options field.

This option is supported by LOCAL or SMB based disk Unmanaged Storage, and also FTP-based.

Just before renaming a file to its final name, DIVA Core will check whether this file already exists. If it does exist, DIVA will delete it so that the renaming operation doesn't fail.

Skipping Files During Restore (-rest_ignoring)

This option is available for Restore, Partial File Restore, and N-Restore requests. It ignores some files during restore based on one or more regular expression rules. The possibilities offered by regular expressions are versatile and enable many different types of filtering.

Files matching one of the regular expressions are ignored by the Server. The rule supports Unicode characters to offer maximum flexibility. Use the following statement to ignore files during restore:

```
-rest_ignoring {<rule>} [<rule>|<rule>|<rule>]
```

Continue to add <rules> as necessary in the previous statement.

There are no predefined set of values for these options. Possible values for this option are infinite and fully customizable.

The files being ignored are still read from the disk or tape instance. If the set of rules is designed to ignore all the files of an object, then no file is restored and the request will be complete.

During an N-Restore, if multiple renaming rules are defined, DIVA Core will process the rule for each Server independently.

Example

A typical use case is restoring a SeaChange clip to a destination that does not support SeaChange special files (private data and video index files). The following statement prevents a Server from restoring files with `.pd` or `.vix` extension:

```
-rest_ignoring <.*\.pd><.*\.vix>
```

The results if the previous statement are as follows:

DIVA Core Object	Destination Server
Clipname.pd	
Clipname.vix	
Clipname	Clipname

Ignoring File Relative Paths (-ignore_relative_path)

If this option is specified on a Source or Destination Server, DIVA will restore all the files of an object directly to the Files Path Root, ignoring any relative paths found in the list of files. This option is useful when absolute paths were specified in the list of files when objects were archived.

Archiving Files in a Specific Order (-file_order)

Use this option archiving or restoring files that are MSS files (Omneon QuickTime files), but the archiving Source Server is not an AVID (Pinnacle) MSS Server (an Omneon server).

This option is not limited to specific Server types, but it is only meaningful for LOCAL, DISK, CIFS, and FTP_STANDARD Servers. This option can be superseded by the request option.

Specify the file sequence during archiving or restoring using the following statement:

```
-file_order {MSS|OMNEON|DIFWAV|SEACHANGE DIRS_FIRST|FILES_FIRST}
```

The following list identifies the archive sequence for specific formats:

MSS

The sequence is header, ft, info, and then std.

OMNEON

The sequence is clip.mov, and then essence files (.wav, .aiff, .m2v, .mpeg, .diff, and so on).

DIFWAV

The sequence is clip.dif, and then .wav files.

SEACHANGE

The sequence is clip.pd, clip.vix, and then clip.

DIRS_FIRST

The sequence places directories first and is as follows:

```

Folder test_1;
Folder test_1\test_2;
File test_1\test_2\1.txt;
File test_1\test_2\_A2.txt;
File test_1\test_2\test.txt;
File test_1\test_2\test1.txt;
File test_1\test_2\test2.txt;
File test_1\1.txt;
File test_1\_A2.txt;
File test_1\test.txt;
File test_1\test1.txt;
File test_1\test2.txt;
File 1.txt;
File _A2.txt;
File test.txt;
File test1.txt;
File test2.txt;

```

FILES_FIRST

The sequence places files first and is as follows:

```

File 1.txt;
File _A2.txt;
File test.txt;
File test1.txt;
File test2.txt;
Folder test_1;
File test_1\1.txt;
File test_1\_A2.txt;
File test_1\test.txt;
File test_1\test1.txt;
File test_1\test2.txt;
Folder test_1\test_2;
File test_1\test_2\1.txt;
File test_1\test_2\_A2.txt;
File test_1\test_2\test.txt;
File test_1\test_2\test1.txt;
File test_1\test_2\test2.txt;

```

This ensures that files are archived in the correct sequence so that they are restored in the correct sequence when restoring them to a real Pinnacle MSS Server (a real Omneon server).

DPX Partial File Restore does not examine the file name or the DPX header information to determine which file is assigned to which frame. The assignment is based purely on the sequence in which the .dpx files appear within the archive. By default this sequence is based on ordering established by the Source Server, and is typically alphanumeric. For example, NTFS DISK Servers order files and folders are not case-sensitive as a general rule (but not where diacritical marks, such as ` , ^, and so on are applied). By default, when DIVA Core encounters a subfolder it recursively processes all of the children of that folder (including subfolders) before continuing with other files. If a

folder appears in the alphanumeric folder listing, it is archived recursively in the order it appears.

However, this can create some issues. A user may want all of the subdirectories of a given directory processed first followed by the files in the directory, or the user might want all files processed first followed by subdirectories. In DIVA Core 7.0 and later, the Actor allows the archive options `-file_order DIRS_FIRST` or `-file_order FILES_FIRST` to address these issues as previously described.

Example

An archive contains SeaChange SAF files. These files must be transcoded, and then restored to a Pinnacle MSS Server. In this case, the LOCAL Source Server used by the transcoding process is defined with the `-file_order MSS` option (among others).

This ensures the files coming out of the transcoder are archived and restored in the correct sequence. That is, header, ft, info and then std.

Specifying the Transcode Format (`-tr_archive_format`, `-tr_restore_format`)

Each factory in a transcoder determines the format of the output file. These options allow you to define the factory and output format.

They apply to any Server type, and have no fixed list of values. This option cannot be superseded by the request option unless used in a TranscodeArchived request.

These options specify the transcode operation to apply to essence files during archive (`-tr_archive_format`) or restore (`-tr_restore_format`).

```
-tr_archive_format {factory_name}  
-tr_restore_format {factory_name}
```

The `{factory_name}` is the name of a Flip Factory factory, or the name of a BitScream output format.

Specifying a Transcoder Name (`-tr_names`)

Use this option to specify the transcoder to use for transcode operations. It applies to any Server type and cannot be superseded by the request option, unless used in a TranscodeArchived request.

Either the `-tr_archive_format` or the `-tr_restore_format` option must always be used with `-tr_names`. When transcoding is applied, one of the transcoders defined by `-tr_names` is selected by DIVA Core according to the transcoders defined in the DIVA Core configuration based on the availability, configured queue size, and configured performance of the transcoder.

The format for this option is as follows:

```
-tr_names {transcoder_name} [transcoder_name]
```

The `{transcoder_name}` is the name of a DIVA Core Transcoder defined in the Transcoders area of the Systems page of the System Management App.

If this option is not present, DIVA Core will select one of the transcoders defined in the DIVA Core Configuration based on the availability, configured queue size, and configured performance of the transcoder.

Restoring Metadata (-rest_metadata, -rm)

This option specifies that a metadata file must be generated and restored on every Restore request. This option applies to any Server type. Because video servers may reject the metadata file, this option actually applies mainly to LOCAL, DISK and FTP_STANDARD types.

Either form of the option can be used as follows:

```
-rest_metadata  
-rm
```

When an object is restored, it is first restored normally. After the regular restore has completed, a metadata file is generated and restored on the specified destination in the specified (or implicit) FilePathRoot of the related Restore request.

The metadata file format is compliant with the DIVA Core File Set Watch Folder Metadata File specification. The metadata file name is virtualobjectname.mdf.

Restricting the Number of Actors to Retry (-num_actors_to_retry)

Use this option to limit the number of Actors that an Archive, Restore, or Partial File Restore request will be retried on. By default, this option is not specified and there is no limit. Therefore, all Actors will be tried in case the request constantly fails.

This option applies to any Server type and cannot be superseded by the request option.

This option uses the following statement:

```
-num_actors_to_retry {number}
```

The {number} is the number of retries to attempt and can include zero.

Example

The option -num_actors_to_retry 3 means that the Core Manager will perform no more than four operations (total) with different Actors, even if there are more than four Actors configured. That is, the initial request plus three retries for a total of four attempts.

MSS Server in MXF Mode (-mxf)

This option specifically applies only to MSS Server types, otherwise DIVA Core ignores it. You use this option to indicate when a MSS Server is configured to import and export MXF wrapped clips.

There are no additional parameters for this option and you include it in the following format:

-mxf

FTP Socket Window Size (-socket_window_size)

This option specifies the total buffer space per data socket reserved for send and receive. This option applies to some Server types using FTP protocol, such as FTP_STANDARD, OMNEON, PDR, MSS, and so on.

This parameter has a direct effect on transfer performance. Its value depends on the operating system and is usually set between 2048 and 65536 bytes. When this option is not set DIVA Core uses the default value set at the operating system level. Technical Support recommends increasing this value to 32768 or more on fast networks. Some performance tests must be run to identify the best setting.

The TCP Window Scale option increases the TCP receive window size above its maximum 65536 bytes value. This option is recommended when dealing with Long-Fat Networks, or LFN.

Use the following statement for this option:

```
-socket_window_size {number}  
-socket_bufsize {number}
```

The {number} is the buffer size in bytes.

Note: The -socket_bufsize syntax deprecated but still available. Technical Support recommends not using it in DIVA Core releases later than 6.2.2 because it may conflict with the -socket_block_size parameter.

FTP Socket Block Size (-socket_block_size)

This option defines how much data (in kilobytes) the Actor tries to send and receive in a single system call during FTP transfers. For example, if the internal buffer size of the Actor is set to 2 Mb and -socket_block_size is set to 64 KB, 32 system calls are required to write a single buffer to a data socket.

This option applies to some Server types using FTP, such as FTP_STANDARD, OMNEON, PDR, MSS, and so on.

Use the following statement for this option:

```
-socket_block_size {number}
```

The {number} is the buffer size in kilobytes, ranging from 32 to 2048 kilobytes.

FTP Passive Mode Transfers (-pasv)

This option specifies that the FTP data connection must be opened in passive mode (as opposed to active mode) for the associated Server. This may be necessary if a firewall is between the Actor and the Server.

This option applies to some Server types using FTP, such as FTP_STANDARD, OMNEON, PDR, MSS, and so on.

Use one of the following statements for this option (not case-sensitive):

```
-pasv
-PASV
```

Restoring in AXF Mode (-axf)

The -axf parameter is optional for Restore requests and instructs DIVA Core to restore the original asset into an AXF File. Instead of purely restoring the content of an object to the destination, DIVA Core restores the content into a new AXF File.

Combined with the -rm or -rmxl parameters, you can use this option to export an object with metadata information and then drop it into a WFM Watch Folder.

This option applies to FTP_STANDARD, SFTP, LOCAL, DISK, and EXPEDAT Server types.

You use the following statement to restore an asset in AXF mode:

```
-axf
```

Specifying Connection Timeouts (-list_timeout, -transfer_timeout, -control_timeout)

These options specify the maximum timeout values allowed for different FTP connection operations, and override the default timeout settings. The timeout value for directory and file listings (-list_timeout), file transfers (-transfer_timeout), and control port connections (-control_timeout) can be set.

If an operation exceeds the set timeout value the operation is terminated.

The default value is used if a timeout parameter is not specified, or if the timeout value is set to zero.

Use the following statement for each of these options:

```
-list_timeout {number}
-transfer_timeout {number}
-control_timeout {number}
```

The {number} is the maximum allowed timeout in seconds.

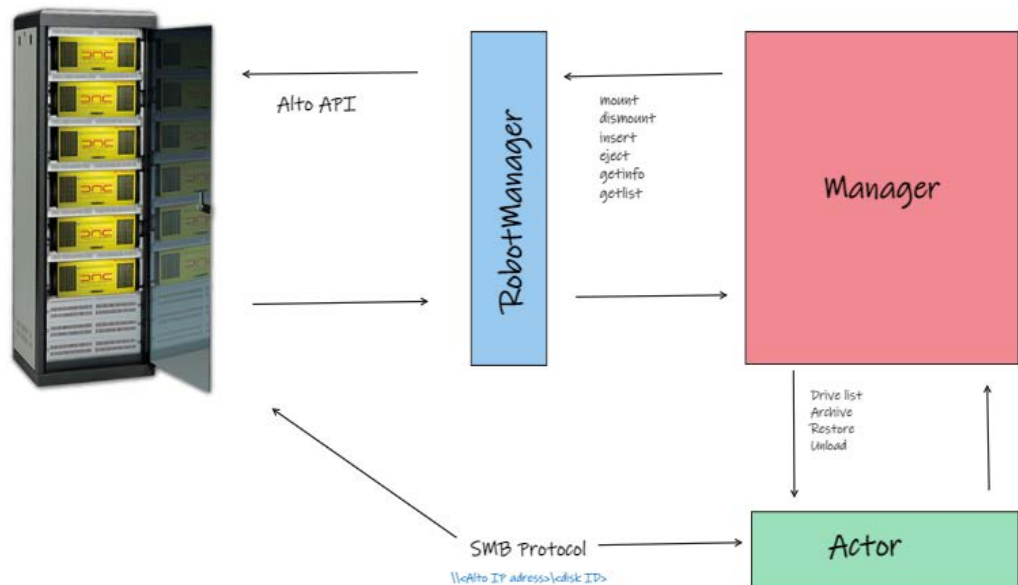
The default timeout values for each FTP connect operation are as follows:

Statement	Default Timeout
-list_timeout	120 seconds
-transfer_timeout	180 seconds
-control_timeout	120 second

Alto Disk Archive Integration

Alto Disk Archive is a type of library of disks. Instead of mounting/dismounting tape, this library is designed to mount/dismount the filesystem of unpluggable disks and map them to a SMB share.

DIVA Core supports Alto like a tape library. A disk of the library is seen as a tape in DIVA Core with its UUID as label.



Configuration

The following subsections describe Alto configuration.

Robot Manager

Set the following parameters in the Robot Manager configurations file:

- Set `RM_MODULE` to `ALTO_Robot.dll` so that RobotManager loads `ALTO_Robot.dll`.

These parameters use the typical settings:

- Set `SERVICE_NAME` to the name of RobotManager.
- If necessary, modify `RM_PORT` so that RobotManager listens on different TCP port.
- `RM_ACS` can be used to setup multiple Alto servers in the same configuration.
- `RM_ALTO_SERVER_HOSTNAME` must be set to the IP address of the Alto server.
- `RM_ALTO_SERVER_PORT` is the TCP port of the Alto API. The default value is 6480.
- `RM_ALTO_GROUP` can be set to only use a specific group of disks with DIVA Core. If this parameter is not set, or empty, RobotManager will use the group named "default".

- RM_ALTO_GROUP_CREATE_IF_NOT_EXISTS can be set to 1 and then RobotManager will create the group at startup. This parameter is set to 0 by default.
- RM_ALTO_NUMBER_OF_DRIVES represents the number of virtual tape drives. This is the virtual entity for the Alto integration, but is required by DIVA Core because tapes are mounted into drives. This number also defines the number of disks that DIVA Core can mount at the same time.

Core Actor Configuration

There is a new parameter in the configuration of Core Actor:

AltoVirtualDrives: This parameter contains a pattern string so that Core Actor can build the list of ALTO Virtual Tape Drives during scandrive for each ACS.

Syntax:

```
<pattern1>[, <pattern2>] [...]
```

Pattern syntax:

```
<alto-hostname>:<alto-port>-ACS<acs number>-<number of virtual drives>
```

Example:

```
192.168.59.102:6480-ACS1-2,192.168.59.102:6480-ACS0-10
```

Service Configuration

RobotManager and Core Actor services must be set to log on under the same account (that is, Administrator).

System Management App Support

Alto Virtual Drives value can be set for a Core Actor from the System Management App Actors Configuration page. It is on the Advanced Settings tab, in the Disk Operations Section.

The screenshot shows the 'EDIT ACTOR' configuration page in the System Management App. The breadcrumb navigation is 'SYSTEM SETTINGS > ACTORS > EDIT ACTOR'. The 'ADVANCED SETTINGS' tab is selected, showing the 'Disk Operations' section. The configuration fields are as follows:

- Profile Read Block Size (B):** 1000
- Profile Write Block Size (B):** 1500
- Profile Write Block Size (B):** 32768
- QT Self-Contained Threshold (MB):** 50
- Disk FTP Block Size (KB):** 32
- Disk FTP Socket Window Size (B):** 65536
- Actor To Actor Connect Timeout (s):** 1200
- Actor To Actor Transfer Timeout (s):** 30
- Linux SMB Mount Point:** /mnt
- Alto Virtual Drives:** (Empty text field)
- Auto Discovery:**
- Do Not Check Object Name:**
- Do Not Check Category:**
- Directory Server:**
- Disable Disk Preallocation:**
- Disk FTP Passive Mode:**
- Quantel Rename Clips:**

At the bottom of the page, there are three buttons: 'Save' (green), 'Reset', and 'Cancel'.

Avid MSS (Program Stream) Servers

Avid (previously Pinnacle) MSS Video Servers can be installed in one of the following configurations:

Independent Storage

The video server (itself) includes its own fault tolerant disk storage.

Shared Storage

The video servers are connected to a SAN where the fault tolerant disk storage is based.

In both cases, external connectivity is provided by one (or several) Connect+ gateways supporting the FTP protocol over a Gigabit Ethernet Network. A clip on the MSS storage is always comprised of three files as listed below (or four if the optional information file is present). They are always archived and restored in the following sequence:

header

This is the first file and the clip's header.

ft

This is the second file and the frame table.

std

This is the third file and the video and audio essence.

info

When present, this is the fourth file. It is an optional information file.

All files are located in a folder that matches the name of the clip (that is, if the clip name is FOO, the files are located in a folder also named FOO).

Newer MediaStream servers can export and import clips with a MXF wrapper. When configured for MXF, the server generates a single file (std) which is the MXF file. DIVA Core only archives one file (std) in MXF Mode. The file is automatically renamed to {clipname}.mxf. This mode is not supported by independent storage servers.

See [Appendix A: Core Options and Licensing](#) for DIVA Core options and licensing information.

MSS with Independent Storage

One record is created for each MSS that DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	MSS IP address	10.80.114.21
Source Server Type	MSS	MSS
Connect Options for Systems with One Gateway	-login {gw_host_name} -pass .video_fs	-login fcgate1 -pass .video_fs
Connect Options for Systems with Two Gateways	-login {gw1_host_name}[, gw2_host_name] -pass .video_fs	-login fcgate1,fcgate2 -pass .video_fs

In a system with two gateways, fcgate1 and fcgate2, DIVA Core manages failover between the two when a connect option such as -login fcgate1, fcgate2 is specified. If the initial FTP connection fails with fcgate1, it will be retried on fcgate2.

Note: This feature has been deprecated and is now implemented using the METASOURCE Source Server Type.

MSS with Shared Storage

One record is created for each gateway connected to the storage network that DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	IP Address of the gateway through which DIVA Core accesses the shared storage.	10.80.114.28
Source Server Type	MSS	MSS
Connect Options	-login video_fs -pass .video_fs	-login video_fs -pass .video_fs

MSS with Shared Storage in MXF Mode

One record is created for each gateway connected to the storage network DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	IP Address of the gateway through which DIVA Core accesses the shared storage.	10.80.114.28
Source Server Type	MSS	MSS
Connect Options	-login video_fs (or -login mxf_fs) -pass .video_fs (or -pass .mxf_fs) -mxf	-login video_fs -pass .video_fs

Using MSS with DIVA_archiveVirtualObject

The following table describes typical Server example parameters.

Parameter	Value	Example
FilePathRoot	The name of the clip.	CITIZENKANE
FileNames	*	*

Avid Airspace Servers

Avid Airspace (previously known as Pluto) is a video server with independent storage. Each clip deals with a single essence file located on the storage root. Airspace uses standard FTP protocol to transfer files to and from the video server internal storage over a Gigabit Ethernet Network.

One record is created for each video server DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	IP address of the video server.	10.80.114.28
Source Server Type	FTP_STANDARD	FTP_STANDARD
Connect Options	-login {FTP_user_name} -pass {FTP_password} -port {FTP_port_number}	-login ftpuser -pass Pa\$\$word -port 6530

The following table describes an Avid Airspace Server use example:

DIVA_archiveVirtualObject Parameter	Value	Example
FilesPathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	TRAFFIC

Avid Transfer Manager DHM Interface

The Avid Transfer Manager is the Avid Unity Outer Gateway, which you can address using two different interfaces. One interface is called the DHM (Data Handler Module) and the other called DET (Dynamically Extensible Transfer). Each interface has a specific purpose.

For this Source Server Type the DHM interface is used for transfer of video and audio content to and from external devices (for example, an archive system).

See the DIVA Core Avid Connectivity User Guide for detailed information.

One record is created for each video server DIVA Core has to move data to and from.

Attributes	Value	Example
IP Address	IP address of the Avid Transfer Manager	10.80.114.28
Source Type	AVID_DHM	AVID_DHM
Connect Options	-port {FTP_port_number} -login {FTP_user_name} -pass {FTP_password}	-port 6021 -login diva -pass diva

The Connect Option values indicated in the previous table are as follows:

-port

This is the TM Communicator FTP service port number.

-login

This is the TM Communicator FTP service user log in.

-pass

This is the TM Communicator FTP service user password associated with the log in.

Archive requests are initiated from Avid Edit Stations using Send to Playback. The TM Communicator supports setting custom titles for ingested (restored) clips. If the -title option is specified with a title name in a DIVA Core Restore or Partial File Restore request, this option's value is used as the clip title, otherwise the original clip name is used. The original clip name is stored in the Video ID field of the Avid metadata.

The following rules apply to custom title settings:

- Custom titles can consist of one or more words separated by spaces and (or) tabulation characters.
- Technical Support strongly recommends single word titles, and absolutely requires that multiple word titles are enclosed in double quotes to ensure proper processing.
- New line (\x0A) and carriage return (\x0D) characters are not allowed in titles.

- Single quote, ampersand, dash, slash, asterisk, and other special characters are supported.
- Double quote characters must be escaped with a backslash to be included in the title.
- Titles composed of one or more spaces enclosed in double quotes are not considered empty.

The following table describes a Server use example:

Restore Option Values	Ingested Clip Title
-title Clip	Clip
-title "Clip"	Clip
-title "My clip"	My clip
-title "My \"special\" clip"	My "special" clip

Avid Transfer Manager DET Interface

Avid Transfer Manager is the Avid Unity Outer Gateway. It can be addressed through two different interfaces called the DHM (Data Handler Module) and DET (Dynamically Extensible Transfer). Each interface has a specific purpose.

For this Source Server type, the DET interface is used for transfer of Metadata and Media Files to Unity Workgroups (or an archive system, seen as an external workgroup / Unity storage extender).

See the DIVA Core Avid Connectivity User's Guide for detailed information.

One record is created for each video server DIVA Core has to move data to and from.

Attributes	Value	Example
IP Address	IP address of the Avid Transfer Manager	10.80.114.28
Source Type	AVID_DET	AVID_DET
Connect Options	-port {FTP_port_number} -login {FTP_user_name} -pass {FTP_password}	-port 6021 -login det -pass diva

The Connect Option values indicated in the previous table are as follows:

-port

This is the TM Communicator FTP service port number.

-login

This is the TM Communicator FTP service user log in.

-pass

This is the TM Communicator FTP service user password associated with the log in.

Archive requests are initiated from Avid Edit Stations using Send to Workgroup.

SeaChange BMS and BMC Servers

A SeaChange BMS (Broadcast Media Server) is a standalone video server equipped with a fast-Ethernet Interface and its own storage.

A SeaChange BMC (Broadcast Media Cluster) is a cluster of video servers providing unified storage based on SeaChange RAID2 technology. Each server of the BMC can deliver files stored on RAID2 to DIVA Core using the FTP protocol. The file transfer format is SAF (SeaChange Archive Format) only.

Note: The SeaChange FTP servers do not support directories. All files must be listed under the FTP root directory.

By default, a SeaChange BMC node offers Automatic Load Balancing management for data transfer across all nodes of the cluster.

If you want to use this feature, you must only declare the last node of the BMC in the DIVA Core configuration. In this case, DIVA Core will always connect to the same node of the cluster. This node will transparently redirect transfers to other nodes as required.

This feature can be disabled by using a special IP address setting in the DIVA Core configuration (see the following table). In this case, all nodes of the BMC must be declared in the DIVA Core configuration.

You can also add a Metasource that encompasses all nodes of the cluster to enable load balancing and failover from within DIVA Core.

Attribute	Value	Example
IP Address	IP address of the BMS or BMC node. You can disable the SeaChange Automatic Load Balancing by placing a \$ in front of the IP address of all BMC nodes. The syntax for this is \$IP_Address.	10.80.114.26 \$10.80.114.26
Source Type	SEACHANGE_BMC	SEACHANGE_BMC
DIVAActor_SEACHANGECHECKDELAY	Identifies the delay before checking if a video was not deleted by SeaChange just after a restore service. The default value is 1000.	DIVAActor_SEACHANGECHECKDELAY=1000

SeaChange uses a flat file system. You must specify the parameters as shown in the following table when archiving a clip.

DIVA_archiveVirtualObject Parameter	Value	Example
FilePathRoot	Leave this field empty	
FileNames	Enter the name of the clip in this field.	POKEMON

SeaChange BML Servers

The SeaChange BML (Broadcast Media Library) is a large storage system for SAF (SeaChange Archive Format) files and is based on the RAID2 technology of the SeaChange BMC platform.

A SeaChange BMC (Broadcast Media Cluster) is a cluster of video servers providing unified storage based on SeaChange RAID2 technology. Each server of the BMC can deliver files stored on RAID2 to DIVA Core using the FTP protocol.

DIVA Core uses the FTP protocol to communicate with either a BMS or BMC. You can only overwrite the files when the Actor service is stopped. The file transfer format is SAF (SeaChange Archive Format) only.

Note: The SeaChange FTP servers do not support directories. All files must be listed under the FTP root directory.

The Automatic Load Balancing feature as described for BMC also exists for BML and operates in a similar fashion.

Attribute	Value	Example
IP Address	IP address of the BML Node. You can disable the SeaChange Automatic Load Balancing by placing a \$ in front of the IP address of all BMC nodes. The syntax for this is \$IP_Address.	10.80.114.26 \$10.80.114.26
Source Server Type	SEACHANGE_BML	SEACHANGE_BML
DIVAActor_SEACHANGECHECKDELAY	Identifies the delay before checking if a video was not deleted by SeaChange just after a restore service. The default value is 1000.	DIVAActor_SEACHANG ECHECKDELAY=1000
DIRECTORY_SERVER_ENABLED	Identifies whether the BML directory server is enabled or disabled.	Valid values are 1 (enabled) and 0 (disabled). The default value is 1 (enabled).

SeaChange BML clip storage is flat. You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtualObject Parameter	Value	Example
FilePathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	OFFICESPACE

SeaChange BMLe and BMLex Servers

The SeaChange BMLe is the storage subsystem of the latest SeaChange MediaClient architecture. SeaChange BMLe is superseded by the BMLex series.

Both the BMLe and BMLex servers are based on the BML architecture. However Infiniband is used for the cluster interconnect rather than the earlier IOP interfaces. Each node of the cluster is equipped with four FSI ports to provide high speed transfers to and from the BMLe and BMLex.

DIVA Core uses CIFS or FTP protocols to communicate with BMLe and BMLex.

File transfer format is the native format of the files stored on the BMLe and BMLex. Each asset consists of:

MPEG2 Files

MPEG essence, private data (.pd) and video index (.vix) files.

MXF Files

MXF file (.mxf), private data (.pd) and video index (.vix) if the MXF essence is MPEG2.

When the clip consists of three files (that is, the essence file, .vix, and .pd), the files are always archived and restored by DIVA Core in the following sequence:

.pd

This is the private data file and the first file archived or restored.

.vix

This is the index file and the second file archived or restored.

Essence File

There is no extension on this file and it is the last one archived or restored.

DIVA Core can restore SAF (SeaChange Archive Format) files from the archive to the BMLe or BMLex. When a SAF clip is restored to a BMLe or BMLex, the SAF file is automatically unwrapped by DIVA Core and the three files are restored to BMLe or BMLex (that is, the essence file, .pd file, and .vix file). This Server can also restore SAF files from an archived SAF Object to BMLe.

This feature is transparent to you because DIVA Core automatically detects SAF and unwraps it in real time. When a SAF clip is restored to the BMLe, the SAF file is unwrapped by DIVA Core and the name of each file is extracted from the SAF file header. The content is restored to BMLe in the separate files previously described.

BMLe and BMLex generated files support SAF releases SAF 0.1, SAF 1.0, and SAF. SAF may contain two consecutive private data files including a 12 byte .pd file, and a 28

byte .pd file. In this case, DIVA Core will only restore the 28 byte file while ignoring the 12 byte file.

You must declare one Server for each FSI of each node:

Attribute	Value	Example
IP Address	IP address FSI	10.80.114.26
Source Server Type	SEACHANGE_BML	SEACHANGE_BML
Connect Options	-ftp or -cifs -login {FTP_user_name} -user {cifs_user_name@domain} -pass {password} -nometadata	-cifs -user me@ourdomain.com -pass Pa\$\$word
DIVAActor_SEACHANGECHECKDELAY	Identifies the delay before checking if a video was not deleted by SeaChange just after a restore service. The default value is 1000.	DIVAActor_SEACHANG ECHECKDELAY=1000

-ftp or -cifs

One of these two options must be specified. Otherwise, Streaming API protocol is assumed, which is not supported by DIVA Core for BMLe and BMLex. This option cannot be superseded by the request option.

-ftp

FTP protocol is used for data transfer to and from BMLe and BMLex.

-cifs

CIFS protocol is used for data transfer to and from the BMLe and BMLex FSI cards. The implicit CIFS path to BMLe is \\fsi_ip_address\vstrm.

-nometadata

This option prevents DIVA Core from archiving the .vix and .pd files when the clip being transferred includes essence, .vix, and .pd files. This option cannot be superseded by the request option.

You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtualObject Parameter	Value	Example
FilePathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	HANNITY

Leitch vR Series Servers

The Leitch vR series video server is connected to external storage that is usually shared with other video servers of the same brand. Clips are stored on Leitch storage as flat files, one file per clip, without any folder structure.

To move clips in and out of the shared storage, Leitch provides a dedicated gateway called the Archive Streamer. The Archive Streamer offers standard FTP protocol over a Gigabit Ethernet network.

Note: The Leitch vR Source Type is deprecated. It was initially created to follow the first Archive Streamer releases that did not correctly report the size of the file to be transferred.

One record must be created for each Archive Streamer DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	IP address of Leitch Archive Streamer	10.80.114.21
Source Type	FTP_STANDARD	FTP_STANDARD
Connect Options	-login {FTP_user_name} -pass {FTP_password} -port {FTP_port}	-login ftpuser -pass Pa\$\$word -port 6021

You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtualObject Parameter	Value	Example
FilesPathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	FRIENDS

Leitch Nexio Servers

The Leitch Nexio video server is connected to external storage that is usually shared with other video servers of the same brand. Clips are stored on Leitch storage as flat files, one file per clip, without any folder structure.

To move clips in and out of the shared storage is possible directly from the video server using the standard FTP protocol over a Gigabit Ethernet network.

Note: The Leitch Nexio Source Type is deprecated.

You must create one record for each video server DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	IP address of Leitch Nexio video server.	10.80.114.21
Source Server Type	FTP_STANDARD	FTP_STANDARD
Connect Options	-login {FTP_user_name} -pass {FTP_password} -port {FTP_port}	-login ftpuser -pass Pa\$\$word -port 6021

You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtualObject Parameter	Value	Example
FilesPathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	ENEMIES

Grass Valley Profile Servers

Grass Valley Profile video servers are provided in one of two ways; with independent storage, where the video server includes its own fault tolerant disk storage, or as part of a MAN, where video servers are connected to a SAN where the fault tolerant disk storage resides.

Irrespective of the storage mechanism, the Core Actor always connects to a specific Profile server. The exchange format is GXF only.

Profile Storage consists of one master disk (for example, EXT: or INT1:), and one level of folders where one clip is seen as one file. One folder called default always exists.

The network infrastructure between GVG Profiles and Core Actors is an IP/FC network.

You must create one record for each video server DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	IP address of the video server.	10.80.114.21
Source Server Type	PDR	PDR
Name	Logical name for the video server.	GVG-Profile-1

The Actor configuration parameters are located in the Actor area of the DIVA Core System Management App. The two parameters in the following table directly influence transfer performance. Technical Support recommends trying several value combinations on the target platform.

In addition to these two parameters, the MTU size setting for the HBA used for IP/FC traffic to the Profile servers may also have an influence on transfer performance.

Grass Valley does not provide any recommendation for MTU size. However, Technical Support recommends setting the MTU size on the Actor HBA to the same value as the MTU size of the Profile HBA. This is only a recommended setting and not an absolute rule.

Attribute	Description	Recommended Values
DIVAActor_ProfileReadingBS	The FTP block size (in bytes) used for transfers on Profile video servers in reading.	1500 16374 32768 (default)
DIVAActor_ProfileWritingBS	FTP block size (in bytes) used for transfers on profile video servers in writing.	16374 32768 (default)

You must specify the parameters as described in the following table when archiving a clip:

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	/explodedFile/disk:/folder	/explodedFile/INT1:/default
FileNames	Enter the name of the clip in this field.	MyClip

Grass Valley UIM Gateway

UIM is a gateway to standalone or MAN Grass Valley Profile servers. It provides TCP/IP over Gigabit Ethernet connections to external systems (such as DIVA Core). For legacy purposes, the connection can also be IP/FC for regular profiles.

UIM also provides real-time format conversion (to MXF). The UIM exchange format is GXF (by default), or alternately MXF.

You must create one record for each UIM DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	IP address for the UIM.	10.80.114.21
Source Server Type	PDR	PDR
Connect Options	-login {movie mxfmovie} -format {?D10AES3} -extension {file_extension}	-login mxfmovie -format ?D10AES3 -extension .mxf

-login

Specifies the FTP user for logging onto the UIM to achieve transfers in the desired format. The two available logins are movie (for GXF exchange format), and mxfmovie (for MXF exchange format). The movie user is assumed if -login is not specified.

-format

The UIM supported options for some file formats. This depends on -login option. The only available option is ?D10AES3. The ?D10AES3 option is an e-VTR compliant file format used with the -login mxfmovie option. If this option is not specified, MXF files will be processed in Grass Valley OP1a format. This option is not specified by default.

This option can be superseded by the request option.

-extension

This option adds the specified extension to the original clip name in the archive. For example, if the original clip is clip1 and the -extension .mxf option is specified, then the archived file will be clip1.mxf.

You must suppress the specified extension before restoring to the destination if it already exists. For example, if the archived file is clip1.mxf and -extension .mxf option is specified, the restored file on the destination will be clip1.

This option is deprecated and replaced by the -arch_renaming and the -rest_renaming options. This option can be superseded by the request option.

UIM are gateways to the Profile server. Use this the same way for UIM and Profile servers regardless of the transfer format (GXF or MXF).

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	/explodedFile/disk:/folder	/explodedFile/INT1:/default
FileNames	Enter the name of the clip in this field.	MyClip

Grass Valley K2 Servers

From DIVA Core's standpoint, K2 servers are similar to Profiles and UIM combined. K2 servers offer Gigabit Ethernet connections to external systems, and the exchange format is GXF (default), and alternately MXF.

You must create one record for each K2 server DIVA Core must move data to and from.

Attribute	Value	Example
IP Address	IP address of the K2 server.	10.80.114.21
Source Server Type	PDR	PDR
Connect Options	-k2 -login {movie mxfmovie} -format {?D10AES3} -extension {file_extension}	-k2 -login mxfmovie -format ?D10AES3 -extension .mxf

-k2

This specifies the interface with the K2 servers. When this option is set, DIVA Core will retrieve the size of the file to be transferred before the actual archive transfer (K2 FTP does support the SIZE command). Correct transfer progress is reported by DIVA Core.

When this option is not set, DIVA Core will assume that servers are Profile, and will not retrieve the file size before archive transfers. Progress will then remain at 0% before suddenly jumping to 100% when the transfer is complete.

This option has no impact on transferred content, and can be superseded by the request option.

-login

This option specifies the FTP user for logging onto the K2 Server to achieve transfers in the desired format. The two available logins are movie (for GXF exchange format), and mxfmovie (for MXF exchange format). The movie user is assumed if -login is not specified.

-format

The K2 supported options for some file formats. This depends on -login option. The only available option is ?D10AES3. The ?D10AES3 option is an e-VTR compliant file format used with the -login mxfmovie option. If this option is not specified, MXF files will be processed in Grass Valley OP1a format. This option is not specified by default, and can be superseded by the request option.

-extension

This option adds the specified extension to the original clip name in the archive. For example, if the original clip is clip1 and the -extension .mxf option is specified, then the archived file will be clip1.mxf.

If the specified extension already exists it must be suppressed before restoring to the destination. For example, if the archived file is clip1.mxf and -extension .mxf option is specified, the restored file on the destination will be clip1.

This option is deprecated and replaced by the -arch_renaming and the-rest_renaming options. This option can be superseded by the request option.

You use this the same way for K2 and Profile servers regardless of the transfer format (GXF or MXF).

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	/explodedFile/disk:/folder	/explodedFile/INT1:/default
FileNames	Enter the name of the clip in this field.	MyClip

Grass Valley M-Series iVDR Servers

Grass Valley iVDR is an analog and digital VTR that includes a Gigabit connection for material exchange of GXF files. The iVDR exchange protocol is similar to the exchange protocol for Profile servers.

One record must be created for each video server DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	IP address of the iVDR.	10.80.114.21
Source Server Type	PDR	PDR
Name	Logical name for the iVDR.	GVG-iVDR

You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	/explodedFile/disk:/folder	/explodedFile/INT1:/default
FileNames	Enter the name of the clip in this field.	MyClip

Sony MAV70 Servers

The Sony MAV70 video server has its own independent storage. MAV70 storage organization is flat and all files reside in the storage root. A Linux computer in front of each MAV70 provides a standard FTP connection for moving data to and from the video server over a Gigabit Ethernet Network.

One record must be created for each MAV70 server DIVA Core has to move data to and from.

Attributes	Value	Example
IP Address	IP address of the MAV70 server.	10.80.114.21
Source Server Type	FTP_STANDARD	FTP_STANDARD
Connect Options	-login {user_name} -pass {password}	-login wing -pass mpegworld

You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	MyClipName

Omneon Spectrum MediaDirector Servers (QuickTime)

The Omneon MediaDirector is the heart of the Omneon Spectrum architecture. It is connected to MediaPorts or MultiPorts which handle isochronous ingest and playback, and to external storage that is usually shared with other Omneon MediaDirectors.

You can use either MediaStore or MediaGrid for external storage. This section describes connecting MediaDirector to MediaStore storage for MediaGrid support in DIVA Core.

Note: MediaGrid is not supported in the Linux environment.

DIVA Core interfaces with an Omneon MediaDirector to move clips in and out of the shared storage, using standard FTP protocol, over a Gigabit Ethernet Network.

When Omneon Spectrum Servers are configured to ingest material in QuickTime format, essence files are stored in a specific folder structure. The Core Actors use unique FTP site commands for smart and transparent access to essence files (in particular, the automatic discovery of a folders structure and collision-avoidance at restore time).

One record must be created for each MediaDirector DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	IP address of Omneon Director.	10.80.114.21
Source Server Type	OMNEON	OMNEON
Root Path	Either leave this field empty or enter an absolute clip directory.	/default/clip.dir
Connect Options	-streaming_mode -sm -tempdir_mode	-streaming_mode -sm -tempdir_mode

-streaming_mode or -sm

This option is QuickTime specific and has no effect on the MXF content. If this option is set, DIVA Core will restore the QuickTime reference file in the following sequence:

1. Audio Tracks
2. QuickTime File
3. Video track

The restore workflow is specific when this option is set. DIVA Core uses the temporary folder to cache the QuickTime file.

-tempdir_mode

This option performs a Partial File Restore of MXF files, and is applicable only to Omneon servers. The MXF Partial File Restore request will terminate if this option is not included in the request.

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	Enter the absolute clip directory in this field, or leave this field empty to use the configured Root Path.	/default/clip.dir
FileNames	Enter the name of the clip in this field.	MyClip

Omneon MediaGrid Content Storage System

MediaGrid is the Content Storage System from Omneon to which Omneon Spectrum servers can be connected.

Note: MediaGrid is not supported in the Linux environment.

The MediaGrid system consists of two major components; ContentServers that store and provide access to media, and ContentDirectors that act as overall file system controllers. ContentDirectors manage the distribution of data throughout the system.

Like any other client system, DIVA Core gets access to the media through a MediaGrid ContentDirector. DIVA Core interfaces with MediaGrid using the CIFS protocol exclusively over a Gigabit Ethernet Network.

The MediaGrid ContentDirector manages data access while the data transfer occurs directly to/from the ContentServers. The Omneon FSD (File System Driver), installed on MediaGrid clients hides this complexity to client systems.

Note: The Omneon FSD must be installed on each Actor exchanging assets with MediaGrid.

The latest release of Omneon FSD for Windows is available for download at <http://support.omneon.com/Updates/Omneon/Current/MediaGrid/WinFSD>. The password for the site (if required) is "alloyparka".

When material is wrapped in QuickTime format, the essence files are stored using a specific folder structure. DIVA Core also uses unique FTP site commands for smart and transparent access to the essence files (in particular, automatic discovery of folders structure and collision-avoidance at restore time).

When the Actor is running as a Windows service, MediaGrid shares are accessed through a UNC path because drive letters mapped to network drives are not accessible by Windows services. In this case ensure the following:

- Omneon MediaGrid folders being accessed by DIVA Core are properly shared for a given Windows user.
- The Core Actor service is configured to run under this user account.
- The user has local administrative rights on the Core Actor.

One record must be created for each ContentDirector DIVA Core has to move data to and from.

Attribute	Value	Example
IP Address	Leave this field empty.	
Source Server Type	MEDIAGRID	MEDIAGRID
Root Path	\\ContentDirector\filesystem\clip.dir	\\10.30.0.200\cldev4\clip.dir \\mycontentdir\fs5\clip.dir

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	Leave this field empty.	
FileNames	Enter the name of the clip in this field.	MyClip

In cases where the asset is wrapped as QuickTime, DIVA Core searches for files matching the format clipname.mov or clipname.MOV. DIVA Core automatically retrieves and processes all of the potentially referenced files.

In cases where the material is wrapped as MXF, DIVA Core will search for a file matching the format clipname.mxf or clipname.MXF. There is only one file per clip.

Quantel Power Portal Gateway

The Quantel Power Portal was previously called the *ISA Gateway*. An ISA system consists of the following components:

ISA Manager

The ISA Manager contains the Clip Database. Clips are identified using a unique FID (File Identifier) in the ISA System.

Q or sQ Servers

One or more Q or sQ servers. These servers contain video cards and disk arrays. Each disk array is associated to a POOL ID, and a single sQ Server can have several POOL IDs. For example, sQ Server ID 1 contains POOL ID 1 and POOL ID 2, sQ Server ID 2 contains POOL ID 3, and sQ Server 3 contains POOL ID 4.

ISA Gateway (Power Portal)

This gateway is a FTP server that imports and exports clips.

One record must be created for each Power Portal (ISA Gateway).

Attribute	Value	Example
IP Address	IP address of the video server.	10.80.114.21
Source Server Type	QUANTEL_ISA	QUANTEL_ISA
Connect Options	-login {FTP_user_name} -pass {FTP_password}	-login ftpuser -pass Pa\$\$word

The Actor configuration parameters are located in the Actor area of the DIVA Core System Management App.

Parameter	Description	Suggested Values
DIVAActor_QUANTELRENAMECLIPS	Enables and disables the file renaming feature.	0 indicates the renaming feature is disabled. 1 indicates the renaming feature is enabled.

DIVAActor_QUANTELRENAMECLIPS applies to Restore requests only. If this parameter is set to 1, and the Object Name format is clipName,UID (this is Omnibus naming), then object related files are renamed using clipName as the Name Root.

For example, if the object Superman,01AB45 is composed of files 8152.D10 and 8152.WAV, and is restored to a QUANTEL_ISA destination, the following is true:

- If DIVAActor_QUANTELRENAMECLIPS is set to 0 (disabled), DIVA Core transfers files called 8152.D10 and 8152.WAV to Power Portal.

- If DIVAACTOR_QUANTELRENAMECLIPS is set to 1 (enabled), DIVA Core transfers files called Superman.D10 and Superman.WAV to Power Portal.

Quantel storage is a flat structure. You must specify the parameters as follows when archiving a clip:

DIVA_archiveVirtualObject Parameter	Value	Example
FilePathRoot	Leave this field empty.	
FileNames	FID1.ext1[,FID1.ext2,] and so on.	clip.mxf,clip1.tar

Files coming from Power Portal can be different file types including: D10+WAV (file names similar to 8152.D10 and 8152.WAV), MXF (TestClip.mxf), and TAR (FramesDifference.tar).

If a file is restored twice to Power Portal, the first file is not overwritten. The second restore creates a new file that is identified by a new FID. The Core Actor captures the new FID after the transfer and forwards it to the Core Manager.

DIVA_GetRequestInfo must be called to obtain the new FID using the DIVA Core API. If the request is completed, the new FID is in the request's ADDITIONAL_INFO field within ClipID tags. The ClipID tag is encapsulated in the ADDITIONAL_INFO tag.

```
<ADDITIONAL_INFO>
  <ClipID>8546</ClipID>
</ADDITIONAL_INFO>
```

Automation is also free to specify a POOL ID in the FilePathRoot Restore request parameter. If no POOL ID is specified, Power Portal will automatically assign one at restore time.

Sony Hyper Agent Servers

Hyper Agent is the name given to Newsbase's FTP server from Sony. The implementation of this FTP server is specific because the LIST command returns a proprietary formatted list of files. This list contains duration, and start and end time codes, but not the size of the file in bytes. The size of each clip is calculated by the Actor using three values; duration, frame rate and bitrate. The resultant size is not accurate, but it is enough for the Manager to allocate a tape for all Archive requests. The progress bar is not affected by the approximated size.

Duration, frame rate and bitrate are retrieved using the following two commands, which are set by the Actor at the beginning of each Archive request:

SITE FSIZ {Clip ID}

This SITE command returns the duration of the specified clip.

SITE GCNF

This SITE command returns the current system configuration of the server. This system configuration must remain the same to ensure that all of the clips on the server are the same.

The following example is a log entry of communications between Actor and the Hyper Agent FTP:

Note: In the following log example the word configuration is misspelled; this is a bug in the FTP server and appears in logs as shown in the example.

```
SITE FSIZ 1444247
200 150 (the duration is 150 frames)
SITE GCNF
213-System configureation
PAL (the frame rate is 25 frames per second)
20
30.0 (the Bitrate is 30 Mbps)
D10
SD_IFRAMEONLY
213 End of system configuration
```

You must create one record for each ClipBox DIVA Core must moved data to and from.

Attribute	Value	Example
IP Address	IP address of the Newsbase server.	10.80.114.21
Source Server Type	SONY_HYPER_AGENT	SONY_HYPER_AGENT
Connect Options	-login {user_name} -pass {password}	-login sony -pass sony

DIVA_archiveVirtualObject Parameter	Value	Example
FilePathRoot	Leave this field empty.	
FileNames	Enter the Clip ID in this field.	1444247

Standard FTP and SFTP Servers

DIVA Core running in a Windows environment can interface with any standard FTP server (Linux or Windows), and SFTP servers (known as SSH FTP or Secure FTP). The Windows-based FileZilla and IIS FTP servers are not supported in Linux because these servers are incapable of handling large numbers of files.

Video servers supporting a fully RFC-959 compliant FTP server are considered standard FTP servers. The only restriction that applies is that Linux-style directory listings are required. You set this parameter in the Home Directory section of the FTP Site Properties for Microsoft IIS FTP servers.

You must create one record for each video server DIVA Core must transfer data to and from.

Attribute	Value	Example
IP Address	IP address of the FTP server.	10.80.114.21
Source Server Type	FTP_STANDARD or SFTP	FTP_STANDARD
Connect Options	-login {user_name} -pass {password} -port {port_number}	-login moon -pass mars -port 27

-login

This is the FTP or SFTP user name. The default value is anonymous.

-pass

This is the FTP or SFTP user's associated password. The default value is anonymous.

-port

This is the port number the FTP or SFTP server is listening on for connections. The default value for FTP_STANDARD is 21, and for SFTP is 22.

You can specify parameters three different ways for Archive requests as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	Full path to files	/my_videos/movies
	Partial path to files	/my_videos
	No path entry	
FileNames	Names of files	maniolia, matrix
	Partial path and names of files	movies/maniolia, movies/matrix
	Full path and names of files	/my_videos/movies/maniolia, /my_videos/movies/matrix

DISK_FTP_PASSIVE_MODE

By default, data connections are created in active mode. In active mode, the DivaFtp client connects from a random, unprivileged port that is higher than port 1023. Then, it starts listening on the port and sends a PORT command to the FTP server. Valid values for this parameter are 0 (disabled) and 1 (enabled).

When DISK_FTP_PASSIVE_MODE is set to 1 (enabled), data connections are created in passive mode. In passive mode, DivaFTP sends a PASV command and the server (not the client) creates the socket.

DISK_FTP_BLOCK_SIZE

The DISK_FTP_BLOCK_SIZE parameter defines how much data Actor tries to send and receive with a single system call during FTP transfers. For example, if the internal buffer size of Actor is set to 2 MB and DISK_FTP_BLOCK_SIZE is set to 32768 bytes, 64 system calls are required to write a single buffer to a data socket. The default value is 32768 bytes.

DISK_FTP_SOCKET_WINDOW_SIZE

The DISK_FTP_SOCKET_WINDOW_SIZE parameter adjusts the normal buffer sizes allocated for output and input buffers. DISK_FTP_SOCKET_WINDOW_SIZE is internally used to set SO_SNDBUF and SO_RCVBUF for FTP managed disk types. The default value is 65536 bytes.

Local Source Servers

A local Source Server represents a disk partition for a specific Actor (internal disks, NAS or SAN disks), and is tied to a specific Actor (versus a disk Source Server not tied to any particular Actor).

One record must be created for each local Source Server DIVA Core must transfer data to and from.

Attribute	Value	Example
Name	Enter the same name as the Actor this Source Server is bound to.	actor1
IP Address	Enter the same IP address as the Actor this Source Server is bound to.	10.80.114.21
Source Server Type	LOCAL	LOCAL

Parameters can be specified in three different ways for Archive requests as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	Full path to files Partial path to files No path entry	/my_videos/movies /my_videos
FileNames	Names of files Partial path and filenames Full path and names of files	maniolia, matrix movies/maniolia, movies/matrix /my_videos/movies/maniolia /my_videos/movies/matrix

If NT drive letters (for example E:) are used, Technical Support **highly recommends** leaving them in the FilePathRoot section (that is, use scheme 1 or 2 in the previous table). Including them in the FileNames section prevents the request from replacing them with another path at restore time. Therefore, these objects cannot be restored on a different platform (for example a Linux-based FTP server) where drive letters are not considered valid paths.

Disk and CIFS Source Servers

A DISK or CIFS Source Server represents a disk partition assumed to be visible from all Network Actors. The only difference between DISK and CIFS is the way blocks of data are read and written:

- DISK instructs Actors to use (Windows) Direct I/O.
- CIFS instructs Actors to use (Windows) Buffered I/O.
- Both DISK and CIFS Source Servers support UNC paths.

One record must be created for each DISK or CIFS Source Server DIVA Core has to move data from the Source Server to the Destination Server. A generic Source Server can also be created to represent any type of DISK or CIFS Source Server.

Attribute	Value	Example
Name	Enter a nickname for the Source Server.	generic-disk
IP Address	Enter the IP address for the Source Server.	10.80.114.21
Source Server Type	DISK or CIFS	DISK

Parameters can be specified three different ways for Archive requests as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilePathRoot	Full path to files	/my_videos/movies
	Partial path to files	/my_videos
	No path entry	
FileNames	Names of files	maniolia, matrix
	Partial path and filenames	movies/maniolia, movies/matrix
	Full path and names of files	/my_videos/movies/maniolia, /my_videos/movies/matrix

If NT drive letters (for example E:) are used, Technical Support **highly recommends** leaving them in the FilePathRoot section (that is, use scheme 1 or 2 in the previous table). Including them in the FileNames section prevents the request from replacing them with another path at restore time. Therefore, you cannot restore these objects on a different platform (for example a Linux-based FTP server) where drive letters are not considered valid paths. Technical Support only supports Linux-based FTP servers when operating in a Linux environment. The Windows-based FileZilla and IIS FTP servers are not supported in Linux because these servers are incapable of handling large numbers of files.

Metasources

A Metasource is a collection of several (single) Source Servers of the same type. It is assumed that all Source Servers of the Metasource are sharing the same online storage. Each Source Server of the Metasource should be of the same regular type (that is, any type except METASOURCE), aka Metasource Base Type. A Metasource provides load-balancing and failover mechanisms across all single Source Servers of the Metasource.

One record must be created for each Metasource DIVA Core has to transfer data to and from.

Attribute	Value	Example	Comments
Name	Name for video server's shared storage.	gvg-man-production	
IP Address	server1 [,server2,server3] and so on	10.158.1.10,10.2.5.60,97.64.52.3	server1, server2,server3 must also be defined in the configuration as regular Source Servers of the same type (all types except METASOURCE, LOCAL, and DISK are permitted, for example, OMNEON, PDR, and so on).
Source Server Type	METASOURCE	METASOURCE	
Network	Must be the same for Metasource and all single Source Servers.		Manager will not start if there is no match.
Site	Either one or the other of the sites from Metasource single Source Servers.		Site specified for Metasource is considered by Manager for resource selection.

Attribute	Value	Example	Comments
Root Path	You can specify a Root Path at the Metasource level.		If the Metasource Root Path is null, the Root Path from the selected single Source Server is considered.
Max Accesses Max Write Acc. Max Read Acc. Max Throughput	Actual value for Metasource does not matter.		The value from the selected single Source Server is considered. You cannot leave these fields empty. Technical Support suggests setting traffic regulation parameters to the sum of all single Source Server's respective parameters. Technical Support also recommends that you do not make any changes to this parameter while there are active requests being processed because it can lead to request termination.
Connect Options	-failover_time={time_in_milliseconds} -retry_actor={number_of_retries}	-failover_time=300 -retry_actor=3	

-failover_time={time_in_milliseconds}

When a single Source Server is selected to process a request and it fails, the single Source Server is temporarily not considered part of the Metasource for 600 milliseconds. This default value can be changed using this option. This option cannot be superseded by the request option.

-retry_actor={number_of_retries}

Use this option to specify the number of Metasource single Source Servers to be tried for each Actor that can be part of the request processing. The default, when this option is not specified, is 2.

For example, if the Metasource is defined as sd1, sd2, sd3, the set of possible Actors is a1, a2, and -retry_actor is set to 2, DIVA Core will try a maximum of four combinations; most likely a0-sd1, a0-sd2, a1-sd3, a1-sd1.

This option cannot be superseded by the request option.

Other single Source Server connection options can also be specified for the Metasource. The following table indicates the effects for each possible option when specified at the Metasource level:

Connect Option	Considered?	Comments
qos=	No	The qos value should be the same for all Metasource single Source Servers, otherwise Manager will not start.
-login	No	Value from selected single Source Server is considered. Applicable to FTP Servers.
-user	No	Value from selected single Source Server is considered. Applicable to CIFS Servers.
-pass	No	Value from selected single Source Server is considered.
-port	No	Value from selected single Source Server is considered.
-allow_delete_on_source	No	Implicitly assumed to be true if all single Source Servers (implicitly or explicitly) allow deleting on Source Server. Otherwise, assumed to be false.
-arch_renaming	No	Value from selected single Source Server is considered.
-rest_renaming	No	Value from selected single Source Server is considered.
-file_order	No	Value from selected single Source Server is considered.

Connect Option	Considered?	Comments
-tr_archive_format	Yes	Values specified for single Source Servers do not matter.
-tr_restore_format	Yes	Values specified for single Source Servers do not matter.
-tr_names	Yes	Values specified for single Source Servers do not matter.
-rest_metadata	No	Value from selected single Source Server is considered.
-num_actors_retry=	Yes	Values specified for single Source Servers do not matter.
-ftp	No	Value from selected single Source Server is considered.
-cifs	No	Value from selected single Source Server is considered.
-nometadata	No	Value from selected single Source Server is considered.
-format	No	Value from selected single Source Server is considered.
-extension	No	Value from selected single Source Server is considered.
-k2	No	Value from selected single Source Server is considered.

A Metasource is used the same as any Source Server of Metasource Base Type.

There are instances where it is required to delete content, and possibly the parent folder, on a server. To satisfy all possible scenarios there are two options available:

- -r deletes recursively
- -delete_fpr includes deletion of the parent folder

The two options, -r and -delete_fpr, work either separately or together, as described in the following workflow examples:

FilesPathRoot	Files	Options	Result
C:\sourceserver\root	*	-r	DIVA Core deletes the content of C:\sourceserver\root recursively.
C:\sourceserver\root	*	-r -delete_fpr	DIVA Core deletes the content of C:\sourceserver\root recursively, and then deletes root.
C:\sourceserver\root	*		DIVA Core deletes the content of C:\sourceserver\root (first level only).
C:\sourceserver\root	*	-delete_fpr	DIVA Core deletes the content of C:\sourceserver\root (first level only), and then eventually deletes root if it is empty.
C:\sourceserver\root	obj*	-r	DIVA Core deletes the content of C:\sourceserver\root\obj recursively, and then deletes C:\sourceserver\root\obj.
C:\sourceserver\root	obj*	-r -delete_fpr	DIVA Core deletes the content of C:\sourceserver\root\obj recursively, then deletes C:\sourceserver\root\obj, and finally deletes C:\sourceserver\root if it is empty.

FilePathRoot	Files	Options	Result
C:\sourceserver\root	obj1\ obj2\ *	-r	DIVA Core deletes the content of C:\sourceserver\root\obj1 recursively, then deletes C:\sourceserver\root\obj1, and then deletes the content of C:\sourceserver\root\obj2 recursively, and finally deletes C:\sourceserver\root\obj2.
C:\sourceserver\root	obj1\ obj2\ *	-r -delete_fpr	DIVA Core deletes the content of C:\sourceserver\root\obj1 recursively, then deletes C:\sourceserver\root\obj1, then deletes the content of C:\sourceserver\root\obj2 recursively, then deletes C:\sourceserver\root\obj2, and finally deletes C:\sourceserver\root if it is empty.
C:\sourceserver\root	obj1\ obj2\subfolder\clip.mov	-r -delete_fpr	DIVA Core deletes the content of C:\sourceserver\root\obj1 recursively, then deletes C:\sourceserver\root\obj1, then deletes the content of C:\sourceserver\root\obj2\subfolder\clip.mov, then deletes C:\sourceserver\root\obj2\subfolder if it is empty, and then deletes C:\sourceserver\root\obj2 if it is empty, and finally deletes C:\sourceserver\root if it is empty.

Expedat Servers

DIVA Core can interface with DataExpedition Expedat servers (up to release 1.17), also known as servedat. This solution uses MTP, which is a high performance file transfer protocol. This WAN acceleration software can use 100 percent of the bandwidth of any long distance or high latency networks.

See the DataExpedition Expedat Server Installation Manual for detailed information on installation and configuration.

This Server works similar to the FTP_STANDARD Server in terms of the FilesPathRoot and list of files.

When Expedat Server is configured with folders having the RestrictHome setting enabled, the RootPath for the Data Expedition Server entry must not reference an absolute path. The RootPath may be interpreted as a path that is not accessible from the Expedat home directory. For example, the Root Path / is interpreted as C:\. However, if the Expedat home directory is D:\folder, Expedat will attempt to access the path D:\folder on C:\, which is not valid. If the home directory is C:\folder, using the Root Path / is acceptable.

Instead of using an absolute path, relative path addressing must be used to resolve this situation. You accomplish relative path addressing by leaving the Root Path field empty in the System Management App, or specifying the relative path in the FilesPathRoot field of the GUI Manager or API request for the archive or restore operation.

To set up a default home location so that an API request can always use "" files path, the Expedat cv_password.txt file must contain a log in account assigned to a folder with the RestrictHome option set.

For example:

```
diva:diva:::S:\WFM:RestrictHome
diva1:diva:::S:\WFM1:RestrictHome
diva2:diva:::S\some_other_folder:RestrictHome
```

The separate user log in and password accounts allow for the creation of more than one EXPEDAT Server entry with different home locations. The API request can then reference the EXPEDAT Server pointing to the desired home location.

When WFM is used to monitor an FTP folder in a Linux environment, it must be configured similar to the following example:

User: diva

User home directory: /ifs

Folder to be Monitored: /ifs/folder1

Correct WFM Configuration: ftp://diva:password@host_ip/folder1

Incorrect WFM Configuration: ftp://diva:password@host_ip/ifs/folder1

One record must be created for each Expedat server DIVA Core must transfer data to and from.

Attribute	Value	Example
IP Address	IP address of the Expedat server.	10.80.114.21
Source Server Type	EXPEDAT	EXPEDAT
Connect Options	-login {user_name} -pass {password} -port {port_number} -license {license_code} -encryption -seq_buffer_size {size_in_megabytes} -exp_maxrate {size_in_kilobytes} -exp_mindatagram {size_in_bytes}	-login moon -pass mars -port 8080 -license 46FE464A98 -encryption -seq_buffer_size 16 -exp_maxrate 1024 -exp_mindatagram 2848

-login and -pass

These options are mandatory if the server is configured with authentication parameters.

-port

This option should always be present because there is no default value.

-license

This is a mandatory parameter to use the DIVA Core Expedat Client. Without the license code the EXPEDAT Server is unusable. You can only configure one Expedat license key per Network.

-encryption

This option works with the Expedat Server, is optional, and enables Expedat content encryption during transfers.

-seq_buffer_size {size_in_megabytes}

This option defines the size of the DataExpedition internal buffer for each transfer. The default value is 16 MB and should be sufficient for most transfers. A large buffer allows DataExpedition to continue moving data during times when the sender or receiver may not be able to process it. However, a small buffer consumes less memory.

-exp_maxrate {size_in_kilobytes}

This option sets an approximate limit on the number of kilobytes per second, per transfer. The default is unlimited, but can be used as an alternate method of controlling bandwidth.

-exp_mindatagram {size_in_bytes}

This transfer protocol is over UDP. This option defines a minimum size for each network datagram payload that DataExpedition sends. The purpose is to prevent DataExpedition from sending too small of a packet over the network. You may want to set this value between 2848 and 8544 when using a very fast network path (Gigabit or higher) and every device along the path supports Jumbo Frames (MTU 9000). Using large datagrams can greatly reduce CPU overhead. However, using this setting without Jumbo Frames being fully supported can cause severe performance issues or loss of connectivity.

Appendix D: Dynamic Configuration Changes

This appendix lists the currently supported changes to DIVA Core configuration that become effective while the Manager is running, and those that require a software component or the Core Manager to be restarted.

Topics:

- [Updates in the Manager Configuration](#)
- [Updates in the System Management App System Page](#)
- [Updates in the System Management App Robots Page](#)
- [Updates in the System Management App Disks Page](#)
- [Updates in the System Management App Drives Page](#)
- [Updates in the System Management App Tapes Page](#)
- [Updates in the System Management App Sets, Tape Groups & Media Mapping Page](#)
- [Updates in the System Management App Analytics App Page](#)
- [Updates in the System Management App Storage Plans Page](#)
- [Updates in the System Management App Slots Page](#)
- [Event Fields](#)
- [Metrics Definitions](#)
- [Configuration Parameter Defaults and Values](#)

Updates in the Manager Configuration

If a parameter in the Core Manager configuration file is changed, the following list identifies what is currently required for the change to take effect.

The manager restart command must be used for the following parameter changes to take effect:

- SERVICE_NAME (also effective after reinstall)
- DIVAMANAGER_NAME
- DIVAMANAGER_PORT
- DIVAMANAGER_TNSNAME
- DIVAMANAGER_DBHOST
- DIVAMANAGER_DBPORT
- DIVAMANAGER_DBSID
- DIVAMANAGER_DBUSER
- DIVAMANAGER_MAX_CONNECTIONS
- DIVAMANAGER_TYPICAL_VIRTUALOBJECT_SIZE
- DIVAMANAGER_CAPACITY_LOW_WATER_MARK
- DIVAMANAGER_STOP_IMMEDIATELY_FOR_REPACK
- DIVAMANAGER_TIME_TO_WAIT_FOR_GRACEFUL_SHUTDOWN
- DIVAMANAGER_DISMOUNT_AFTER
- DIVAMANAGER_UPDATE_PRIORITIES_PERIOD
- DIVAMANAGER_PING_INTERVAL
- DIVAMANAGER_ETC_FEATURE
- DIVAMANAGER_ETC_CONFIDENCE_LEVEL

The manager reload command must be used for the following parameter changes to take effect:

- DIVAMANAGER_TO_LOWER
- DIVAMANAGER_MAX_SIMULTANEOUS_REQUESTS
- DIVAMANAGER_MAX_INACTIVE_REQUESTS
- DIVAMANAGER_MAX_SPAN_SEGMENTS
- DIVAMANAGER_MAX_VIRTUALOBJECTS_FOR_REPACK
- DIVAMANAGER_MAX_DELAY_BETWEEN_SCHEDULER
- DIVAMANAGER_SCHEDULER_AFTER_INACTIVITY
- DIVAMANAGER_EXPORT_ROOT_DIR
- DIVAMANAGER_MAX_RESTORE_SERVERS
- DIVAMANAGER_MAX_EXPORT_TAPES

- DIVAMANAGER_MAX_EXPORT_ELEMENTS
- DIVAMANAGER_MAX_FILES_IN_ARCHIVE
- DIVAMANAGER_MAX_FILES_IN_PARTIAL_RESTORE
- USE_IMPROVED_BEST_WORST_FIT_ALGORITHM
- DIVAMANAGER_SITE_SUPPORT_ENABLED
- DIVAMANAGER_CACHE_QOS_USE_DISK
- DIVAMANAGER_PRIORITY_TIER
- DIVAMANAGER_OVERWRITE_POLICY
- DIVAMANAGER_OVERWRITE_OVERRIDE
- ATTEMPT_ACCESS_TO_OFFLINE_DISK
- CHANGE_DISK_STATE_ON_ERROR
- MANAGER_ACTOR_DISK_RETRY_NUMBER
- DISK_STATUS_POLLING_RATE
- DISK_BUFFER_SPACE
- DISK_CONNECTION_STATE_RESET_DELAY
- DIVAMANAGER_MAX_EXCLUDED_INSTANCES
- DIVAMANAGER_REQUEST_SCHEDULING_QUEUE_SIZE
- DIVAMANAGER_API_TASK_QUEUE_SIZE
- DIVAMANAGER_MAX_CONCURRENT_REQUESTS
- DIVAMANAGER_MIN_DB_CONNECTION_LIMIT
- DIVAMANAGER_MAX_DB_CONNECTION_LIMIT
- DIVAMANAGER_INITIAL_DB_CONNECTION_LIMIT
- DIVAMANAGER_INACTIVITY_TIMEOUT
- DIVAMANAGER_SIZE_OF_STATEMENT_CACHE
- DIVAMANAGER_DEFAULT_ROW_PREFETCH
- DIVAMANAGER_FAILOVER_ENABLED
- DIVAMANAGER_NUM_RS_SOLUTIONS_TO_EVALUATE
- DIVAMANAGER_DBSERVICENAME
- ABORT_ARCHIVES_ON_EMPTY_FILES (reloadable in service mode)
- TAPE_FULL_ON_SPAN_REJECTED (reloadable in service mode)

Updates in the System Management App System Page

The following sections describe updates made in the various areas on the System page.

Networks Area

If one of the following parameters or actions in the Networks area of the Systems page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Network Name

Sites Area

If one of the following parameters or actions in the Sites area of the Systems page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Site Name
- Is Main Site
- Comments

Servers Area

If one of the following parameters or actions in the Servers area of the Systems page is changed, Manager must be notified for the changes to take effect.

Some changes only take effect after notifying the Manager, and currently executing requests are complete.

- Add
- Delete (Notify Manager and after requests complete)
- Source Server Name (Notify Manager and after requests complete)
- IP Address (Notify Manager and after requests complete)
- Source Server Type (Notify Manager and after requests complete)
- Network (Notify Manager and after requests complete)
- Site (Notify Manager and after requests complete)
- Connect Options (Notify Manager and after requests complete)
- Root Path (Notify Manager and after requests complete)

- Max Throughput (Notify Manager and after requests complete)
- Max Accesses (Notify Manager and after requests complete). You must not make changes to this parameter while there are active request because it could lead to the request being terminated.
- Max Read Accesses (Notify Manager and after requests complete). You must not make changes to this parameter while there are active request because it could lead to the request being terminated.
- Max Write Accesses (Notify Manager and after requests complete). You must not make changes to this parameter while there are active request because it could lead to the request being terminated.

Actors Area

If one of the following parameters or actions in the Actors area of the Systems page is changed, Manager must be notified for the changes to take effect.

Before the change becomes effective on several of the parameters or actions, the Actor must be disconnected. Also, some changes only take effect after notifying the Manager, and currently executing requests are complete.

- Add
- Delete (must disconnect Actor first and Notify Manager)
- Actor Name (must disconnect Actor first and Notify Manager)
- IP Address (must disconnect Actor first and Notify Manager)
- Port (must disconnect Actor first and Notify Manager)
- Network (Notify Manager and after requests complete)
- Site (Notify Manager and after requests complete)
- Max Drive Operations (Notify Manager and after requests complete)
- Max Server Operations (Notify Manager and after requests complete)
- Max Disk Operations (Notify Manager and after requests complete)
- Direct Restore (Notify Manager and after requests complete)
- Cache Restore (Notify Manager and after requests complete)
- Copy to Tape Group (Notify Manager and after requests complete)
- Associative Copy (Notify Manager and after requests complete)
- Repack (Notify Manager and after requests complete)
- Delete (Notify Manager and after requests complete)
- Direct Archive (Notify Manager and after requests complete)
- Cache Archive (Notify Manager and after requests complete)

Transcoders Area

If one of the following parameters or actions in the Transcoders area of the Systems page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Transcoder Name
- Transcoder Type
- Transcoder Port
- Working Directory
- Executable Path
- Performance

Updates in the System Management App Robots Page

The following sections describe updates made in the various areas on the Robots page.

Robot Managers Area

If one of the following parameters or actions in the Robot Managers area of the Robots page is changed, Manager must be notified for the changes to take effect.

Before the change becomes effective on several of the parameters or actions, the Robot Manager must be disconnected.

- Add
- Delete
- Robot Manager Name
- Address (must disconnect Robot Manager first and Notify Manager)
- Port (must disconnect Robot Manager first and Notify Manager)
- Site

Media Compatibility Area

If an entry is deleted in the Media Compatibility area of the Robots page, Manager must be notified for the changes to take effect.

Robot Managers-ACS Area

If an entry is deleted in the Robot Managers-ACS area of the Robots page, Manager must be notified for the changes to take effect.

Updates in the System Management App Disks Page

The following sections describe updates made in the various areas on the Disks page.

Arrays Area

If one of the following parameters or actions in the Arrays frame of the Disks page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Array Name
- Description

Disks Area

If one of the following parameters or actions in the Disks area of the Disks page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Disk Name
- Array
- Site
- Status
- Min Free Space

Actor-Disk Connections Area

If one of the following parameters or actions in the Actor-Disk Connections area of the Disks page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Disk
- Actor
- Interface
- Mount Point
- Max Throughput
- Access

- Used For

Object Storage Accounts Area

If one of the following parameters or actions in the Object Storage Accounts area of the Disks page is changed, Manager must be notified for the changes to take effect.

- Add
- Delete
- Account Name
- Login
- Password
- URL
- Proxy
- Service Name
- Identity Domain
- Threads Per Transfer
- Type
- Vendor

Updates in the System Management App Drives Page

The following sections describe updates made in the various areas on the Drives page.

Drives Area

If one of the following parameters or actions in the Drives area of the Drives page is changed, the noted action must be performed for the changes to take effect.

- Delete (Notify Manager)
- Serial Number (Notify Manager)
- Status (Notify Manager)
- Enabled Operations (Notify Manager)
- Used (manager restart)
- Installation Date (no action required, effective immediately)
- Last Upgrade Date (no action required, effective immediately)
- Last Cleaning Date (no action required, effective immediately)

Managed Storage Area

If one of the following parameters or actions in the Managed Storage area of the Drives page is changed, Manager must be notified for the changes to take effect.

- Delete
- Name
- Serial Number
- Status

Drive Properties Area

If one of the following parameters or actions in the Drive Properties area of the Drives page is changed, Manager must be notified for the changes to take effect.

- Add (through syncDB)
- Delete

Actor-Drives Area

If one of the following parameters or actions in the Actor-Drives area of the Drives page is changed, Manager must be notified for the changes to take effect.

- Add

- Delete
- Actor
- Drive

Updates in the System Management App Tapes Page

If one of the following parameters or actions in the Tapes page is changed, the noted action must be preformed for the changes to take effect.

- Tape Properties (Notify Manager)
- Empty Ejected Tapes (no action required, effective immediately)
- Inserted Protected Tapes (no action required, effective immediately)
- Tape States (no action required, effective immediately)

Updates in the System Management App Sets, Tape Groups & Media Mapping Page

Changes made in this page are effective as soon as they are applied. No manual update is necessary.

Updates in the System Management App Analytics App Page

If one of the following parameters or actions in the Analytics App page is changed, the noted action must be performed for the changes to take effect.

- Configuration (Notify Manager)
- Event Definitions (currently cannot be altered)
- Metric Definitions (no action required, effective immediately)

Updates in the System Management App Storage Plans Page

Changes made in this page are effective immediately. It is **highly recommended** that the Storage Policy Manager Service be stopped before altering any setting in this page.

Updates in the System Management App Slots Page

Changes made in this page are effective immediately. It is highly recommended that the Storage Policy Manager Service be stopped before altering any setting in this page.

Event Fields

The following three tables identify event fields and the types of events associated with them. There are three tables only due to the amount of entries. Locate the desired field on the top row of the table, and then follow down the column to identify which events are valid for the selected field.

	Event Type	Tape Type	Tape Barcode	Drive Type	Drive Name	Disk Name	Drive Serial Number	Library Serial Number	SD Name	Actor Name
TAPE_INSERT	Yes	Yes	Yes					Yes		
TAPE_INSERT_ERR	Yes							Yes		
TAPE_MOUNT	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_MOUNT_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_POSITION	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_POSITION_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_READ	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_READ_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_WRITE	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_WRITE_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
DISK_READ ¹	Yes					Yes				Yes
DISK_READ_ERR ¹	Yes					Yes				Yes
DISK_WRITE ¹	Yes					Yes				Yes
DISK_WRITE_ERR ¹	Yes					Yes				Yes
SD_READ	Yes								Yes	Yes

	Event Type	Tape Type	Tape Barcode	Drive Type	Drive Name	Disk Name	Drive Serial Number	Library Serial Number	SD Name	Actor Name
SD_READ_ERR	Yes								Yes	Yes
SD_WRITE	Yes								Yes	Yes
SD_WRITE_ERR	Yes								Yes	Yes
TAPE_UNLOAD	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_UNLOAD_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_DISMOUNT	Yes	Yes	Yes	Yes	Yes		Yes	Yes		
TAPE_DISMOUNT_ERR	Yes	Yes	Yes	Yes	Yes		Yes	Yes		
TAPE_EJECT	Yes	Yes	Yes					Yes		
TAPE_EJECT_ERR	Yes	Yes	Yes					Yes		
END_OF_TAPE	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
TAPE_REPACK	Yes							Yes		
ARCHIVE_REQUEST	Yes								Yes	
COPY_REQUEST	Yes									
COPY_AS_REQUEST (to new)	Yes									
CREATE_INSTANCE	Yes									
RESTORE and PARTIAL_RESTORE	Yes								Yes	
DELETE_VIRTUAL_OBJECT	Yes									
DELETE_INSTANCE	Yes									
TRANSCODE_END	Yes									Yes

	Event Type	Tape Type	Tape Barcode	Drive Type	Drive Name	Disk Name	Drive Serial Number	Library Serial Number	SD Name	Actor Name
TRANSCODE_ERROR	Yes									Yes
STOPPED_ON_CANCEL	Yes									
CHECKSUM_ERROR_TAPE	Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes
CHECKSUM_ERROR_DISK	Yes					Yes				Yes
CHECKSUM_ERROR_SD	Yes								Yes	Yes
TAPE_IMPORT	Yes		Yes							
TAPE_EXPORT	Yes		Yes							

1 The transcoder work directory is not a DIVA Core disk. No DISK READ or DISK WRITE events are created when accessing this directory.

The presence of Optional in the following table indicates that it is optional. New Instance IDs are only generated after the final write to the Destination Server media. Instance ID is not available in the following cases:

- Temporary instances created in cache disk by an Archive request
- SD READ or SD WRITE during the transcode phase of an archive when transferring to or from the transcoder work directory
- Cache DISK READ or DISK WRITE when performing a tape to tape Copy request
- Tape positioning before a tape write (Archive request)
- End Of Tape (EOT exception) encountered during an Archive request

	Object Name ¹	Object Collection ¹	Object Instance ¹	Media (Tape Group or array)	Request ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
TAPE_INSERT						Yes	Yes		
TAPE_INSERT_ERROR				Yes		Yes			
TAPE_MOUNT				Yes		Yes	Yes		

	Object Name ¹	Object Collection ¹	Object Instance ¹	Media (Tape Group or array)	Request ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
TAPE_MOUNT_ERR				Yes		Yes			
TAPE_POSITION	Yes	Yes	Optional	Yes	Yes	Yes	Yes		
TAPE_POSITION_ERR	Yes	Yes	Optional	Yes	Yes	Yes			
TAPE_READ	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
TAPE_READ_ERR	Yes	Yes	Yes	Yes	Yes	Yes		Yes	
TAPE_WRITE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
TAPE_WRITE_ERR	Yes	Yes		Yes	Yes	Yes		Yes	
DISK_READ ²	Yes	Yes	Optional	Yes	Yes	Yes	Yes	Yes	Yes
DISK_READ_ERR ²	Yes	Yes	Optional	Yes	Yes	Yes		Yes	
DISK_WRITE ²	Yes	Yes	Optional	Yes	Yes	Yes	Yes	Yes	Yes
DISK_WRITE_ERR ²	Yes	Yes		Yes	Yes	Yes		Yes	
SD_READ	Yes	Yes	Optional		Yes	Yes	Yes	Yes	Yes
SD_READ_ERR	Yes	Yes	Optional		Yes	Yes		Yes	
SD_WRITE	Yes	Yes	Optional		Yes	Yes	Yes	Yes	Yes

	Object Name ¹	Object Collection ¹	Object Instance ¹	Media (Tape Group or array)	Request ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
SD_WRITE_ERR	Yes	Yes			Yes	Yes		Yes	
TAPE_UNLOAD				Yes		Yes	Yes		
TAPE_UNLOAD_ERR				Yes		Yes			
TAPE_DISMOUNT				Yes		Yes	Yes		
TAPE_DISMOUNT_ERR				Yes		Yes			
TAPE_EJECT						Yes	Yes		
TAPE_EJECT_ERR						Yes			
END_OF_TAPE	Yes	Yes	Optional	Yes	Yes	Yes			
TAPE_REPACK					Yes	Yes			
ARCHIVE_REQUEST	Yes	Yes		Yes	Yes	Yes	Yes	Yes	
COPY_REQUEST	Yes	Yes		Yes	Yes	Yes	Yes	Yes	
COPY_AS_REQUEST (to new)	Yes	Yes		Yes	Yes	Yes	Yes	Yes	
CREATE_INSTANCE	Yes		Yes	Yes	Yes	Yes		Yes	

	Object Name ¹	Object Collection ¹	Object Instance ¹	Media (Tape Group or array)	Request ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
RESTORE and PARTIAL_RESTORE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
DELETE_VIRTUALOBJECT	Yes	Yes			Yes	Yes			
DELETE_INSTANCE	Yes	Yes	Yes	Yes	Yes	Yes		Yes	
TRANSCOPE_END	Yes	Yes	Yes		Yes	Yes	Yes	Yes	Yes
TRANSCOPE_ERR	Yes	Yes	Yes		Yes	Yes			
STOPPED_ON_CANCEL	Yes	Yes			Yes	Yes			
CHECKSUM_ERROR_TAPE	Yes	Yes	Optional	Yes	Yes	Yes			
CHECKSUM_ERROR_DISK	Yes	Yes	Optional	Yes	Yes	Yes			
CHECKSUM_ERROR_SD	Yes	Yes	Optional		Yes	Yes			
TAPE_IMPORT				Yes		Yes			
TAPE_EXPORT				Yes	Yes	Yes			

1. Information is not provided for Repack requests.
2. The transcoder work directory is not a DIVA Core disk. No DISK READ or DISK WRITE events are created when accessing this directory.

	Transfer Error Rate	Error Code	Error Message	Transcoder or Analyzer Name	Number of Archive Operations	Data Size
TAPE_INSERT						
TAPE_INSERT_ERR		Yes	Yes			
TAPE_MOUNT						
TAPE_MOUNT_ERR		Yes	Yes			
TAPE_POSITION						
TAPE_POSITION_ERR		Yes	Yes			
TAPE_READ	Yes					
TAPE_READ_ERR		Yes	Yes			
TAPE_WRITE	Yes					
TAPE_WRITE_ERR		Yes	Yes			
DISK_READ ¹						
DISK_READ_ERR ¹		Yes	Yes			
DISK_WRITE ¹						
DISK_WRITE_ERR ¹		Yes	Yes			
SD_READ						
SD_READ_ERR		Yes	Yes			
SD_WRITE						
SD_WRITE_ERR		Yes	Yes			
TAPE_UNLOAD						
TAPE_UNLOAD_ERR		Yes	Yes			
TAPE_DISMOUNT						
TAPE_DISMOUNT_ERR		Yes	Yes			
TAPE_EJECT						
TAPE_EJECT_ERR		Yes	Yes			

	Transfer Error Rate	Error Code	Error Message	Transcoder or Analyzer Name	Number of Archive Operations	Data Size
END_OF_TAPE						
TAPE_REPACK						
ARCHIVE_REQUEST					Yes	
COPY_REQUEST					Yes	
COPY_AS_REQUEST (to new)					Yes	
CREATE_INSTANCE						
RESTORE and PARTIAL_RESTORE					Yes	
DELETE_VIRTUALOBJECT						
DELETE_INSTANCE						
TRANSCODE_END				Yes		
TRANSCODE_ERR		Yes	Yes	Yes		
STOPPED_ON_CANCEL						
CHECKSUM_ERROR_TAPE						
CHECKSUM_ERROR_DISK						
CHECKSUM_ERROR_DISK						
TAPE_IMPORT						Yes
TAPE_EXPORT						Yes

1. The transcoder work directory is not a DIVA Core disk. No DISK READ or DISK WRITE events are created when accessing this directory.

Metrics Definitions

The following table identifies DIVA Core metrics definitions. By default, all definitions are enabled.

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_READ_WRITE	Actor: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Actor Name	Lifetime
ACTOR_READ_WRITE_ABORTED_NUMBER	Actor: number of terminated READ and terminated WRITE operations with drives	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Actor Name	Lifetime
ACTOR_READ_WRITE_ABORTED_NUMBER_DAY	Actor: number of terminated READ and terminated WRITE operations with drives	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Actor Name	Day
ACTOR_READ_WRITE_ABORTED_NUMBER_SD	Actor: number of terminated READ and terminated WRITE operations with SD	SD_READ_ERR SD_WRITE_ERR	Count	Null	Event ID	Actor Name	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_READ_WRITE_ABORTED_NUMBER_SD_DAY	Actor: number of terminated READ and terminated WRITE operations with SD	SD_READ_ERR SD_WRITE_ERR	Count	Null	Event ID	Actor Name	Day
ACTOR_READ_WRITE_DAY	Actor: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Actor Name	Day
ACTOR_READ_WRITE_MONTH	Actor: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Actor Name	Month
ACTOR_READ_WRITE_NUMBER	Actor: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Actor Name	Lifetime
ACTOR_READ_WRITE_NUMBER_DAY	Actor: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Actor Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_READ_WRITE_NUMBER_MONTH	Actor: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Actor Name	Month
ACTOR_TIME_ALL_OPERATION	Actor: time in all operations	DISK_READ DISK_READ_ERR DISK_WRITE DISK_WRITE_ERR SD_READ SD_READ_ERR SD_WRITE SD_WRITE_ERR TAPE_END_OF_TAPE TAPE_MOUNT_ERR TAPE_POSITION TAPE_POSITION_ERR TAPE_READ TAPE_READ_ERR TAPE_UNLOAD TAPE_UNLOAD_ERR TAPE_WRITE TAPE_WRITE_ERR	Sum	Null	Duration	Actor Name	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_TIME_ALL_OPERATION_DAY	Actor: time in all operations	DISK_READ DISK_READ_ERR DISK_WRITE DISK_WRITE_ERR SD_READ SD_READ_ERR SD_WRITE SD_WRITE_ERR TAPE_END_OF_TAPE TAPE_MOUNT_ERR TAPE_POSITION TAPE_POSITION_ERR TAPE_READ TAPE_READ_ERR TAPE_UNLOAD TAPE_UNLOAD_ERR TAPE_WRITE TAPE_WRITE_ERR	Sum	Null	Duration	Actor Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_TIME_ALL_OPERATION_MONTH	Actor: time in all operations	DISK_READ DISK_READ_ERR DISK_WRITE DISK_WRITE_ERR SD_READ SD_READ_ERR SD_WRITE SD_WRITE_ERR TAPE_END_OF_TAPE TAPE_MOUNT_ERR TAPE_POSITION TAPE_POSITION_ERR TAPE_READ TAPE_READ_ERR TAPE_UNLOAD TAPE_UNLOAD_ERR TAPE_WRITE TAPE_WRITE_ERR	Sum	Null	Duration	Actor Name	Month
ACTOR_TIME_READ	Actor: time in READ operations	DISK_READ SD_READ TAPE_READ	Sum	Null	Duration	Actor Name	Lifetime
ACTOR_TIME_READ_DAY	Actor: time in READ operations	DISK_READ SD_READ TAPE_READ	Sum	Null	Duration	Actor Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
ACTOR_TIME_READ_MONTH	Actor: time in READ operations	DISK_READ SD_READ TAPE_READ	Sum	Null	Duration	Actor Name	Month
ACTOR_TIME_WRITE	Actor: time in WRITE operations	DISK_WRITE SD_WRITE TAPE_WRITE	Sum	Null	Duration	Actor Name	Lifetime
ACTOR_TIME_WRITE_DAY	Actor: time in WRITE operations	DISK_WRITE SD_WRITE TAPE_WRITE	Sum	Null	Duration	Actor Name	Day
ACTOR_TIME_WRITE_MONTH	Actor: time in WRITE operations	DISK_WRITE SD_WRITE TAPE_WRITE	Sum	Null	Duration	Actor Name	Month
DISK_AVERAGE_TRANSFER_RATE_READ	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Lifetime
DISK_AVERAGE_TRANSFER_RATE_READ_DAY	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Day
DISK_AVERAGE_TRANSFER_RATE_READ_MONTH	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Month
DISK_AVERAGE_TRANSFER_RATE_WRITE	DISK: average transfer rate of WRITE	DISK_WRITE	Average	Null	Transfer Rate	Disk Name	Lifetime
DISK_AVERAGE_TRANSFER_RATE_WRITE_DAY	DISK: average transfer rate of WRITE	DISK_WRITE	Average	Null	Transfer Rate	Disk Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DISK_AVG_TRANSFER_RATE_WRITE_MONTH	DISK: average transfer rate of WRITE	DISK_WRITE	Average	Null	Transfer Rate	Disk Name	Month
DISK_CHECKSUM_FAILURE_COUNT_DAY	DISK: Checksum failure operations count	CHECKSUM_ERROR_DISK	Count	Null	Event ID	Disk Name	Day
DISK_CHECKSUM_FAILURE_COUNT_MONTH	DISK: Checksum Failure Operations Count	CHECKSUM_ERROR_DISK	Count	Null	Event ID	Disk Name	Month
DISK_NUMBER_READ	Disk: Total number of READ operations	DISK_READ DISK_READ_ERR	Count	Null	Event ID	Disk Name	Lifetime
DISK_NUMBER_READ_ABORTED	Disk: Total number of terminated READ operations	DISK_READ_ERR	Count	Null	Event ID	Disk Name	Lifetime
DISK_NUMBER_READ_ABORTED_DAY	Disk: Total number of terminated READ operations	DISK_READ_ERR	Count	Null	Event ID	Disk Name	Day
DISK_NUMBER_READ_ABORTED_MONTH	Disk: Total number of terminated READ operations	DISK_READ_ERR	Count	Null	Event ID	Disk Name	Month
DISK_NUMBER_READ_DAY	Disk: Total number of READ operations	DISK_READ DISK_READ_ERR	Count	Null	Event ID	Disk Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DISK_NUMBER_READ_MONTH	Disk: Total number of READ operations	DISK_READ DISK_READ_ERR	Count	Null	Event ID	Disk Name	Month
DISK_NUMBER_WRITE	Disk: Total number of WRITE operations	DISK_WRITE DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Lifetime
DISK_NUMBER_WRITE_ABORTED	Disk: Total number of terminated WRITE operations	DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Lifetime
DISK_NUMBER_WRITE_ABORTED_DAY	Disk: Total number of terminated WRITE operations	DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Day
DISK_NUMBER_WRITE_ABORTED_MONTH	Disk: Total number of terminated WRITE operations	DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Month
DISK_NUMBER_WRITE_DAY	Disk: Total number of WRITE operations	DISK_WRITE DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Day
DISK_NUMBER_WRITE_MONTH	Disk: Total number of WRITE operations	DISK_WRITE DISK_WRITE_ERR	Count	Null	Event ID	Disk Name	Month
DISK_READ	DISK: total amount of data READ	DISK_READ	Sum	Null	Transfer Size	Disk Name	Lifetime
DISK_READ_DAY	DISK: total amount of data READ	DISK_READ	Sum	Null	Transfer Size	Disk Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DISK_READ_MONTH	DISK: total amount of data READ	DISK_READ	Sum	Null	Transfer Size	Disk Name	Month
DISK_TIME_ALL_OPERATION	DISK: total time of ALL operations	DISK_READ DISK_WRITE	Sum	Null	Duration	Disk Name	Lifetime
DISK_TIME_ALL_OPERATION_DAY	DISK: total time of ALL operations	DISK_READ DISK_WRITE	Sum	Null	Duration	Disk Name	Day
DISK_TIME_ALL_OPERATION_MONTH	DISK: total time of ALL operations	DISK_READ DISK_WRITE	Sum	Null	Duration	Disk Name	Month
DISK_TIME_READ	DISK: total time of READ operations	DISK_READ	Sum	Null	Duration	Disk Name	Lifetime
DISK_TIME_READ_DAY	DISK: total time of READ operations	DISK_READ	Sum	Null	Duration	Disk Name	Day
DISK_TIME_READ_MONTH	DISK: total time of READ operations	DISK_READ	Sum	Null	Duration	Disk Name	Month
DISK_TIME_WRITE	DISK: total time of WRITE operations	DISK_WRITE	Sum	Null	Duration	Disk Name	Lifetime
DISK_TIME_WRITE_DAY	DISK: total time of WRITE operations	DISK_WRITE	Sum	Null	Duration	Disk Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DISK_TIME_WRITE_MONTH	DISK: total time of WRITE operations	DISK_WRITE	Sum	Null	Duration	Disk Name	Month
DISK_WRITE	DISK: total amount of data WRITE	DISK_WRITE	Sum	Null	Transfer Size	Disk Name	Lifetime
DISK_WRITE_DAY	DISK: total amount of data WRITE	DISK_WRITE	Sum	Null	Transfer Size	Disk Name	Day
DISK_WRITE_MONTH	DISK: total amount of data WRITE	DISK_WRITE	Sum	Null	Transfer Size	Disk Name	Month
DIVA_SYSTEM_ACTIVE_ARCHIVE_NUMBER	DIVA Core System: number of active Archive requests	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Lifetime
DIVA_SYSTEM_ACTIVE_ARCHIVE_NUMBER_DAY	DIVA Core System: number of active Archive requests	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Day
DIVA_SYSTEM_ACTIVE_ARCHIVE_NUMBER_MONTH	DIVA Core System: number of active Archive requests	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Month
DIVA_SYSTEM_ACTIVE_COPY_AS_NEW_NUMBER	DIVA Core System: number of active Copy As New object requests	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_ACTIVE_COPY_AS_NUMBER_DAY	DIVA Core System: number of active Copy As New object requests	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Day
DIVA_SYSTEM_ACTIVE_COPY_AS_NUMBER_MONTH	DIVA Core System: number of active Copy As New object requests	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Month
DIVA_SYSTEM_ACTIVE_COPY_NUMBER	DIVA Core System: number of active Copy requests	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Lifetime
DIVA_SYSTEM_ACTIVE_COPY_NUMBER_DAY	DIVA Core System: number of active Copy requests	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Day
DIVA_SYSTEM_ACTIVE_COPY_NUMBER_MONTH	DIVA Core System: number of active Copy requests	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA Core System	Month
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER	DIVA Core System: number of active Restore requests	RESTORE	Maximum	Null	Number of Operations	Local DIVA Core System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER_DAY	DIVA Core System: number of active Restore requests	RESTORE	Maximum	Null	Number of Operations	Local DIVA Core System	Day
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER_MONTH	DIVA Core System: number of active Restore requests	RESTORE	Maximum	Null	Number of Operations	Local DIVA Core System	Month
DIVA_SYSTEM_AVG_READ_WRITE	DIVA Core System: average amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Weighted Average	Duration	Transfer Size	Local DIVA Core System	Lifetime
DIVA_SYSTEM_AVG_READ_WRITE_DAY	DIVA Core System: average amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Weighted Average	Duration	Transfer Size	Local DIVA Core System	Day
DIVA_SYSTEM_AVG_READ_WRITE_MONTH	DIVA Core System: average amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Weighted Average	Null	Transfer Size	Local DIVA Core System	Month
DIVA_SYSTEM_NUMBER_VIRTUAL_OBJECT_ARCHIVE	DIVA Core System: number of objects archived	ARCHIVE_REQUEST	Count	Null	Transfer Size	Local DIVA Core System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ ARCHIVE _DAY	DIVA Core System: number of objects archived	ARCHIVE_RE QUEST	Count	Null	Transfer Size	Local DIVA Core System	Day
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ ARCHIVE _MONTH	DIVA Core System: number of objects archived	ARCHIVE_RE QUEST	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ CREATED	DIVA Core System: number of objects created	ARCHIVE_RE QUEST COPY_AS_R EQUEST	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ CREATED _DAY	DIVA Core System: number of objects created	ARCHIVE_RE QUEST COPY_AS_R EQUEST	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ CREATED _MONTH	DIVA Core System: number of objects created	ARCHIVE_RE QUEST COPY_AS_R EQUEST	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ DELETED	DIVA Core System: number of objects deleted	DELETE_VIR TUALOBJECT	Count	Null	Event ID	Local DIVA Core System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ DELETED _DAY	DIVA Core System: number of objects deleted	DELETE_VIRTUALOBJECT	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ DELETED _MONTH	DIVA Core System: number of objects deleted	DELETE_VIRTUALOBJECT	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_COPY	DIVA Core System: number of object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_COPY_ DAY	DIVA Core System: number of object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_COPY_ MONTH	DIVA Core System: number of object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA Core System	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_CREAT ED	DIVA Core System: number of object instances created	CREATE_INS TANCE	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_CREAT ED_DAY	DIVA Core System: number of object instances created	CREATE_INS TANCE	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_CREAT ED_MON TH	DIVA Core System: number of object instances created	CREATE_INS TANCE	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SY STEM_N UMBERVI RTUAL_ OBJECT_ INSTANC E_DELET ED	DIVA Core System: number of object instances deleted	DELETE_INS TANCE	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SY STEM_N UMBER_ VIRTUAL OBJECT_ INSTANC E_DELET ED_DAY	DIVA Core System: number of object instances deleted	DELETE_INS TANCE	Count	Null	Event ID	Local DIVA Core System	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_NUMBER_VIRTUAL_OBJECT_INSTANCE_DELETED_MONTH	DIVA Core System: number of object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SYSTEM_NUMBER_VIRTUAL_OBJECT_RESTORE	DIVA Core System: number of objects restored	RESTORE	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL_OBJECT_RESTORE_DAY	DIVA Core System: number of objects restored	RESTORE	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SYSTEM_NUMBER_VIRTUAL_OBJECT_RESTORE_MONTH	DIVA Core System: number of objects restored	RESTORE	Count	Null	Event ID	Local DIVA Core System	Month
DIVA_SYSTEM_READ_WRITE	DIVA Core System: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Local DIVA Core System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER	DIVA Core System: number of terminated READ and terminated WRITE operations	DISK_READ_ERR DISK_WRITE_ERR SD_READ_ERR SD_WRITE_ERR TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER_DAY	DIVA Core System: number of terminated READ and terminated WRITE operations	DISK_READ_ERR DISK_WRITE_ERR SD_READ_ERR SD_WRITE_ERR TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER_MONTH	DIVA Core System: number of terminated READ and terminated WRITE operations	DISK_READ_ERR DISK_WRITE_ERR SD_READ_ERR SD_WRITE_ERR TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Local DIVA Core System	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SY STEM_READ_WRITE_DAY	DIVA Core System: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Local DIVA Core System	Day
DIVA_SY STEM_READ_WRITE_MONTH	DIVA Core System: amount of data READ and written	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Local DIVA Core System	Month
DIVA_SY STEM_READ_WRITE_NUMBER	DIVA Core System: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Local DIVA Core System	Lifetime
DIVA_SY STEM_READ_WRITE_NUMBER_DAY	DIVA Core System: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Local DIVA Core System	Day
DIVA_SY STEM_READ_WRITE_NUMBER_MONTH	DIVA Core System: number of READ and WRITE operations	DISK_READ DISK_WRITE SD_READ SD_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Local DIVA Core System	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
MEDIA_ARCHIVE_VIRTUAL_OBJECT_DATASIZE_DAY	Media: data size of all objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	Media Name	Day
MEDIA_ARCHIVE_VIRTUAL_OBJECT_DATASIZE_MONTH	Media: data size of all objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	Media Name	Month
MEDIA_VIRTUAL_OBJECT_INSTANCE_CREATE	Media: number of object instances created	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Lifetime
MEDIA_VIRTUAL_OBJECT_INSTANCE_CREATE_DAY	Media: number of object instances created	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Day
MEDIA_VIRTUAL_OBJECT_INSTANCE_CREATE_MONTH	Media: number of object instances created and deleted	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Month
MEDIA_VIRTUAL_OBJECT_INSTANCE_DELETE	Media: number of object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Media Name	Lifetime
MEDIA_VIRTUAL_OBJECT_INSTANCE_DELETE_DAY	Media: number of object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Media Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
MEDIA_VIRTUALOBJECT_INSTANCE_DELETE_MONTH	Media: number of object instances created and deleted	DELETE_INSTANCE	Count	Null	Event ID	Media Name	Month
MEDIA_READ_WRITE	Media: amount of data READ and written	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Media Name	Lifetime
MEDIA_READ_WRITE_DAY	Media: amount of data READ and written	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Media Name	Day
MEDIA_READ_WRITE_MONTH	Media: amount of data READ and written	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Media Name	Month
MEDIA_READ_WRITE_NUMBER	Media: number of READ and WRITE operations	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Media Name	Lifetime
MEDIA_READ_WRITE_NUMBER_DAY	Media: number of READ and WRITE operations	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Media Name	Day
MEDIA_READ_WRITE_NUMBER_MONTH	Media: number of READ and WRITE operations	DISK_READ DISK_WRITE TAPE_READ TAPE_WRITE	Count	Null	Event ID	Media Name	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
MEDIA_RESTORE_VIRTUAL_OBJECT_DATASIZE_DAY	Media: data size of all objects restored	RESTORE	Sum	Null	Transfer Size	Media Name	Day
MEDIA_RESTORE_VIRTUAL_OBJECT_DATASIZE_MONTH	Media: data size of all objects restored	RESTORE	Sum	Null	Transfer Size	Media Name	Month
MEDIA_TAPE_EXPORT_NUMBER_DAY	Media: Number of tapes EXPORTED	TAPE_EXPORT	Count	Null	Event ID	Media Name	Day
MEDIA_TAPE_EXPORT_NUMBER_MONTH	Media: Number of tapes EXPORTED	TAPE_EXPORT	Count	Null	Event ID	Media Name	Month
MEDIA_TAPE_IMPORT_NUMBER_DAY	Media: Number of tapes IMPORTED	TAPE_IMPORT	Count	Null	Event ID	Media Name	Day
MEDIA_TAPE_IMPORT_NUMBER_MONTH	Media: Number of tapes IMPORTED	TAPE_IMPORT	Count	Null	Event ID	Media Name	Month
SD_ARCHIVE_VIRTUALOBJECT_DATASIZE_DAY	SD: data size of all objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	SD Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
SD_ARCHIVE_VIRTUALOBJECT_DATASIZE_MONTH	SD: data size of all objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	SD Name	Month
SD_CHECKSUM_FAILURE_COUNT_DAY	SD: checksum failure operations count	CHECKSUM_ERROR_SD	Count	Null	Event ID	SD Name	Day
SD_READ	SD: amount of data READ	SD_READ	Sum	Null	Transfer ID	SD Name	Lifetime
SD_READ_DAY	SD: amount of data READ	SD_READ	Sum	Null	Transfer ID	SD Name	Day
SD_READ_MONTH	SD: amount of data READ	SD_READ	Sum	Null	Transfer Id	SD Name	Month
SD_READ_NUMBER	SD: number of READ operations	SD_READ	Count	Null	Event ID	SD Name	Lifetime
SD_READ_NUMBER_DAY	SD: number of READ operations	SD_READ	Count	Null	Event ID	SD Name	Day
SD_READ_NUMBER_MONTH	SD: number of READ operations	SD_READ	Count	Null	Event ID	SD Name	Month
SD_RESTORE_VIRTUALOBJECT_DATASIZE_DAY	SD: data size of all objects restored	RESTORE	Sum	Null	Transfer Size	SD Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
SD_RESTORE_VIRTUALOBJECT_DATASIZE_MONTH	SD: data size of all objects restored	RESTORE	Sum	Null	Transfer Size	SD Name	Month
SD_TIME	SD: time in operation	SD_READ SD_WRITE	Sum	Null	Duration	SD Name	Lifetime
SD_TIME_DAY	SD: time in operation	SD_READ SD_WRITE	Sum	Null	Duration	SD Name	Day
SD_TIME_MONTH	SD: time in operation	SD_READ SD_WRITE	Sum	Null	Duration	SD Name	Month
SD_WRITE	SD: amount of data written	SD_WRITE	Sum	Null	Transfer Size	SD Name	Lifetime
SD_WRITE_DAY	SD: amount of data written	SD_WRITE	Sum	Null	Transfer Size	SD Name	Day
SD_WRITE_MONTH	SD: amount of data written	SD_WRITE	Sum	Null	Transfer Size	SD Name	Month
SD_WRITE_NUMBER	SD: number of WRITE operations	SD_WRITE	Count	Null	Event ID	SD Name	Lifetime
SD_WRITE_NUMBER_DAY	SD: number of WRITE operations	SD_WRITE	Count	Null	Event ID	SD Name	Day
SD_WRITE_NUMBER_MONTH	SD: number of WRITE operations	SD_WRITE	Count	Null	Event ID	SD Name	Month
TAPE_CHECKSUM_FAILURE_COUNT_DAY	Tape: checksum failure operations count	CHECKSUM_ERROR_TAPE TAPE_DISMOUNT_ERR TAPE_MOUNT_ERR	Count	Null	Event ID	Tape Barcode	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_DRIVE_DATA_RATE	Tape Drive: data rate	TAPE_READ TAPE_WRITE	Average	Null	Transfer Rate	Drive Serial Number	Day
TAPE_DRIVE_DATA_RATE_MONTH	Tape Drive: data rate	TAPE_READ TAPE_WRITE	Average	Null	Transfer Rate	Drive Serial Number	Month
TAPE_DRIVE_ERROR_RATE	Tape Drive: internal error rate	TAPE_READ TAPE_WRITE	Average	Null	Error Rate	Drive Serial Number	Day
TAPE_DRIVE_ERROR_RATE_MONTH	Tape Drive: internal error rate	TAPE_READ TAPE_WRITE	Average	Null	Error Rate	Drive Serial Number	Month
TAPE_DRIVE_LAST_OPERATION_DATE	Tape Drive: date of last MOUNT, DISMOUNT, READ or WRITE	TAPE_DISMOUNT TAPE_MOUNT TAPE_READ TAPE_WRITE	Maximum	Null	Event Time	Drive Serial Number	Lifetime
TAPE_DRIVE_NUMBER_MOUNTS	Tape Drive: number of mounts	TAPE_MOUNT	Count	Null	Event ID	Drive Serial Number	Lifetime
TAPE_DRIVE_NUMBER_MOUNT_ABORTED	Tape Drive: number of terminated MOUNT and DISMOUNT operations (<i>together</i>)	TAPE_DISMOUNT_ERR TAPE_MOUNT_ERR	Count	Null	Event ID	Drive Serial Number	Lifetime
TAPE_DRIVE_NUMBER_READ_WRITE_ABORTED	Tape Drive: number of terminated READ and WRITE operations (<i>together</i>)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Drive Serial Number	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_DRIVE_NUMBER_READ_WRITE_ABORTED_DAY	Tape Drive: number of terminated READ and WRITE operations (together)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Drive Serial Number	Day
TAPE_DRIVE_NUMBER_READ_WRITE_ABORTED_MONTH	Tape Drive: number of terminated READ and WRITE operations (together)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Drive Serial Number	Month
TAPE_DRIVE_OPERATION_TOTAL_TIME	Tape Drive: total time of drive operations	TAPE_READ TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Lifetime
TAPE_DRIVE_OPERATION_TOTAL_TIME_DAY	Tape Drive: total time of drive operations	TAPE_READ TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Day
TAPE_DRIVE_READ_WRITE	Tape Drive: amount of data READ and written (together)	TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Drive Serial Number	Lifetime
TAPE_DRIVE_READ_WRITE_DAY	Tape Drive: amount of data READ and written (together)	TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Drive Serial Number	Day
TAPE_DRIVE_READ_WRITE_MONTH	Tape Drive: amount of data READ and written (together)	TAPE_READ TAPE_WRITE	Sum	Null	Transfer Size	Drive Serial Number	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_DRIVE_READ_WRITE_NUMBER	Tape Drive: number of READ and WRITE operations (together)	TAPE_READ TAPE_WRITE	Count	Null	Event ID	Drive Serial Number	Lifetime
TAPE_DRIVE_READ_WRITE_NUMBER_DAY	Tape Drive: number of READ and WRITE operations (together)	TAPE_READ TAPE_WRITE	Count	Null	Event ID	Drive Serial Number	Day
TAPE_DRIVE_READ_WRITE_NUMBER_MONTH	Tape Drive: number of READ and WRITE operations (together)	TAPE_READ TAPE_WRITE	Count	Null	Event ID	Drive Serial Number	Month
TAPE_DRIVE_TIME_ALL_OPERATION	Tape Drive: time in all operations	TAPE_DISMOUNT TAPE_EJECT TAPE_INSERT TAPE_MOUNT TAPE_POSITION TAPE_READ TAPE_UNLOAD TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_DRIVE_TIME_ALL_OPERATION_DAY	Tape Drive: time in all operations	TAPE_DISMOUNT TAPE_EJECT TAPE_INSERT TAPE_MOUNT TAPE_POSITION TAPE_READ TAPE_UNLOAD TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Day
TAPE_DRIVE_TIME_ALL_OPERATION_MONTH	Tape Drive: time in all operations	TAPE_DISMOUNT TAPE_EJECT TAPE_INSERT TAPE_MOUNT TAPE_POSITION TAPE_READ TAPE_UNLOAD TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Month
TAPE_DRIVE_TIME_READ	Tape Drive: time in READ operation	TAPE_READ	Sum	Null	Duration	Drive Serial Number	Lifetime
TAPE_DRIVE_TIME_READ_DAY	Tape Drive: time in READ operation	TAPE_READ	Sum	Null	Duration	Drive Serial Number	Day
TAPE_DRIVE_TIME_READ_MONTH	Tape Drive: time in READ operation	TAPE_READ	Sum	Null	Duration	Drive Serial Number	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_DRIVE_TIME_WRITE	Tape Drive: time in WRITE operation	TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Lifetime
TAPE_DRIVE_TIME_WRITE_DAY	Tape Drive: time in WRITE operation	TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Day
TAPE_DRIVE_TIME_WRITE_MONTH	Tape Drive: time in WRITE operation	TAPE_WRITE	Sum	Null	Duration	Drive Serial Number	Month
TAPE_EXTERNALIZATION_NUMBER	Tape: number of externalizations	TAPE_EJECT	Count	Null	Event ID	Tape Barcode	Lifetime
TAPE_LAST_DISMOUNT	Tape: date of last DISMOUNT	TAPE_DISMOUNT	Maximum	Null	Event Time	Tape Barcode	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LAST_EVENT_ID	Tape: Analytics App Event ID of the last tape or drive operation	TAPE_DISMOUNT TAPE_DISMOUNT_ERR TAPE_MOUNT TAPE_MOUNT_ERR TAPE_POSITION TAPE_POSITION_ERR TAPE_READ TAPE_READ_ERR TAPE_UNLOAD TAPE_UNLOAD_ERR TAPE_WRITE TAPE_WRITE_ERR	Maximum	Null	Event ID	Tape Barcode	Lifetime
TAPE_LAST_MOUNT_DATE	Tape: date of last MOUNT	TAPE_MOUNT	Maximum	Null	Event Time	Tape Barcode	Lifetime
TAPE_LAST_READ	Tape: date of last READ	TAPE_READ	Maximum	Null	Event Time	Tape Barcode	Lifetime
TAPE_LAST_WRITE	Tape: date of last WRITE	TAPE_WRITE	Maximum	Null	Event Time	Tape Barcode	Lifetime
TAPE_LIBRARY_NUMBER_DISMOUNT_ABORTED	Tape Library: total number of terminated DISMOUNT operations	TAPE_DISMOUNT_ERR	Count	Null	Event ID	Library Serial Number	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_NUMBER_DISMOUNT_ABORTED_DAY	Tape Library: total number of terminated DISMOUNT operations	TAPE_DISMOUNT_ERR	Count	Null	Event ID	Library Serial Number	Day
TAPE_LIBRARY_NUMBER_DISMOUNT_ABORTED_MONTH	Tape Library: total number of terminated DISMOUNT operations	TAPE_DISMOUNT_ERR	Count	Null	Event ID	Library Serial Number	Month
TAPE_LIBRARY_NUMBER_MOUNT	Tape Library: total number of MOUNT operations	TAPE_MOUNT	Count	Null	Event ID	Library Serial Number	Lifetime
TAPE_LIBRARY_NUMBER_MOUNT_ABORTED	Tape Library: total number of terminated MOUNT operations	TAPE_MOUNT_ERR	Count	Null	Event ID	Library Serial Number	Lifetime
TAPE_LIBRARY_NUMBER_MOUNT_ABORTED_DAY	Tape Library: total number of terminated MOUNT operations	TAPE_MOUNT_ERR	Count	Null	Event ID	Library Serial Number	Day
TAPE_LIBRARY_NUMBER_MOUNT_ABORTED_MONTH	Tape Library: total number of terminated MOUNT operations	TAPE_MOUNT_ERR	Count	Null	Event ID	Library Serial Number	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_NUMBER_MOUNT_DAY	Tape Library: total number of MOUNT operations	TAPE_MOUNT	Count	Null	Event ID	Library Serial Number	Day
TAPE_LIBRARY_NUMBER_MOUNT_MONTH	Tape Library: total number of MOUNT operations	TAPE_MOUNT	Count	Null	Event ID	Library Serial Number	Month
TAPE_LIBRARY_NUMBER_READ	Tape Library: total number of READ operations	TAPE_READ TAPE_READ_ERR	Count	Null	Event ID	Library Serial Number	Lifetime
TAPE_LIBRARY_NUMBER_READ_DAY	Tape Library: total number of READ operations	TAPE_READ TAPE_READ_ERR	Count	Null	Event ID	Library Serial Number	Day
TAPE_LIBRARY_NUMBER_READ_MONTH	Tape Library: total number of READ operations	TAPE_READ TAPE_READ_ERR	Count	Null	Event ID	Library Serial Number	Month
TAPE_LIBRARY_NUMBER_WRITE	Tape Library: total number of WRITE operations	TAPE_WRITE TAPE_WRITE_ERR	Count	Null	Event ID	Library Serial Number	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_NUMBER_WRITE_DAY	Tape Library: total number of WRITE operations	TAPE_WRITE TAPE_WRITE_ERR	Count	Null	Event ID	Library Serial Number	Day
TAPE_LIBRARY_NUMBER_WRITE_MONTH	Tape Library: total number of WRITE operations	TAPE_WRITE TAPE_WRITE_ERR	Count	Null	Event ID	Library Serial Number	Month
TAPE_LIBRARY_READ	Tape Library: total amount of data READ	TAPE_READ	Sum	Null	Transfer Size	Library Serial Number	Lifetime
TAPE_LIBRARY_READ_DAY	Tape Library: total amount of data READ	TAPE_READ	Sum	Null	Transfer Size	Library Serial Number	Day
TAPE_LIBRARY_READ_MONTH	Tape Library: total amount of data READ	TAPE_READ	Sum	Null	Transfer Size	Library Serial Number	Month
TAPE_LIBRARY_WRITE	Tape Library: total amount of data WRITE	TAPE_WRITE	Sum	Null	Transfer Size	Library Serial Number	Lifetime
TAPE_LIBRARY_WRITE_DAY	Tape Library: total amount of data WRITE	TAPE_WRITE	Sum	Null	Transfer Size	Library Serial Number	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_WRITE_MONTH	Tape Library: total amount of data WRITE	TAPE_WRITE	Sum	Null	Transfer Size	Library Serial Number	Month
TAPE_MOUNT_DISMOUNT_NUMBER	Tape: number of MOUNT and DISMOUNT operations (<i>together</i>)	TAPE_DISMOUNT TAPE_MOUNT	Count	Null	Event ID	Tape Barcode	Lifetime
TAPE_MOUNT_NUMBER	Tape: number of MOUNT operations	TAPE_MOUNT	Count	Null	Event Id	Tape Barcode	Lifetime
TAPE_READ_WRITE_ABORTED_NUMBER	Tape: number of terminated READ and WRITE operations (<i>together</i>)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Tape Barcode	Lifetime
TAPE_READ_WRITE_ABORTED_NUMBER_DAY	Tape: number of terminated READ and WRITE operations (<i>together</i>)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Tape Barcode	Day
TAPE_READ_WRITE_NUMBER	Tape: number of READ and WRITE operations (<i>together</i>)	TAPE_READ TAPE_WRITE	Count	Null	Event Id	Tape Barcode	Lifetime
TAPE_READ_WRITE_NUMBER_DAY	Tape: number of READ and WRITE operations	TAPE_READ TAPE_WRITE	Count	Null	Event ID	Tape Barcode	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_REPACK_NUMBER	Tape: number of REPACK, REUSE, and REFORMAT operations (together)	TAPE_REPACK	Count	Null	Event ID	Local DIVA Core System	Lifetime
TRANSCODE_ABORTED_NUMBER	Transcoder: number terminated TRANSCODE operations	TRANSCODE_ERR	Count	Null	Event ID	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_ABORTED_NUMBER_DAY	Transcoder: number terminated TRANSCODE operations	TRANSCODE_ERR	Count	Null	Event ID	Transcoder Name or Analyzer Name	Day
TRANSCODE_AVG_DATA	Transcoder: average amount of data TRANSCODED	TRANSCODE_END	Weighted Average	Duration	Transfer Size	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_AVG_DATA_DAY	Transcoder: average amount of data TRANSCODED	TRANSCODE_END	Weighted Average	Duration	Transfer Size	Transcoder Name or Analyzer Name	Day
TRANSCODE_AVG_THROUGHPUT	Transcoder: average transcoding throughput	TRANSCODE_END	Average	Null	Transfer Rate	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_AVG_THROUGHPUT_DAY	Transcoder: average transcoding throughput	TRANSCODE_END	Average	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TRANSCODE_DATA	Transcoder: amount of data TRANSCODED	TRANSCODE_END	Sum	Null	Transfer Size	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_DATA_DAY	Transcoder: amount of data TRANSCODED	TRANSCODE_END	Sum	Null	Transfer Size	Transcoder Name or Analyzer Name	Day
TRANSCODE_DATA_MONTH	Transcoder: amount of data TRANSCODED	TRANSCODE_END	Sum	Null	Transfer Size	Transcoder Name or Analyzer Name	Month
TRANSCODE_MAXIMUM_THROUGHPUT	Transcoder: maximum transcoding throughput	TRANSCODE_END	Maximum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_MAXIMUM_THROUGHPUT_DAY	Transcoder: maximum transcoding throughput	TRANSCODE_END	Maximum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day
TRANSCODE_MINIMUM_THROUGHPUT	Transcoder: minimum transcoding throughput	TRANSCODE_END	Minimum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_MINIMUM_THROUGHPUT_DAY	Transcoder: minimum transcoding throughput	TRANSCODE_END	Minimum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day
TRANSCODE_NUMBER	Transcoder: number TRANSCODE operations	TRANSCODE_END	Count	Null	Event ID	Transcoder Name or Analyzer Name	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TRANSCODE_NUMBER_DAY	Transcoder: number TRANSCODE operations	TRANSCODE_END	Count	Null	Event Id	Transcoder Name or Analyzer Name	Day
TRANSCODE_NUMBER_MONTH	Transcoder: number TRANSCODE operations	TRANSCODE_END	Count	Null	Event ID	Transcoder Name or Analyzer Name	Month
TRANSCODE_TIME	Transcoder: time in transcoding operations	TRANSCODE_END	Sum	Null	Duration	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_TIME_DAY	Transcoder: time in transcoding operations	TRANSCODE_END	Sum	Null	Duration	Transcoder Name or Analyzer Name	Day
TRANSCODE_TIME_MONTH	Transcoder: time in transcoding operations	TRANSCODE_END	Sum	Null	Duration	Transcoder Name or Analyzer Name	Month

Configuration Parameter Defaults and Values

Parameter	Default	Values
Manager: Enable/Disable Analytics App Data Collection	1	0 or 1
Manager: Size of the event batch download (number of events)	100	Integer
Manager: Max timeout in the event there are not events to fill the above batch (seconds)	15	Integer
Conf Utility GUI: Enable/Disable Analytics App Configuration	0	0 or 1
DB: Maximum possible history of Events in Months	12	Integer
DB: Maximum possible number of Metrics in DB	1000000	Integer

Appendix E: ADIC SDLC Installation and Configuration

This appendix describes installation and configuration of the SDLC Server and SDLC Client.

Topics:

- [SDLC Server](#)
- [SDLC Client](#)
- [Using dasadmin Commands](#)
- [Troubleshooting](#)

SDLC Server

The following sections describe prerequisites and configuration of the SDLC Server.

Prerequisites

The SDLC Server process is called supervisor. The SDLC GUI is also available as an applet in your web browser address bar. You access the GUI by entering the IP address of the computer on which SDLC Server is running in the browser address bar.

Avoid stopping the SDLC Manager (that is, the NobleNet PortMapper for TCP Windows service) while SDLC Clients are currently connected (for example, the SDLC GUI connection). If the service is stopped, the SDLC Server will vary to a transient state making it temporarily impossible to restart.

Configuration

First define a physical resources partition (in the SDLC GUI Managed Storage tab, then the Wizard tab) to make the SDLC usable. After the physical resources are defined, then define a logical library with its slots and drives. When the wizard completes, the partition is automatically bound to an ADIC Client. The Core ADIC Robot Manager uses the client to obtain status information about library items, and to send mount and dismount commands.

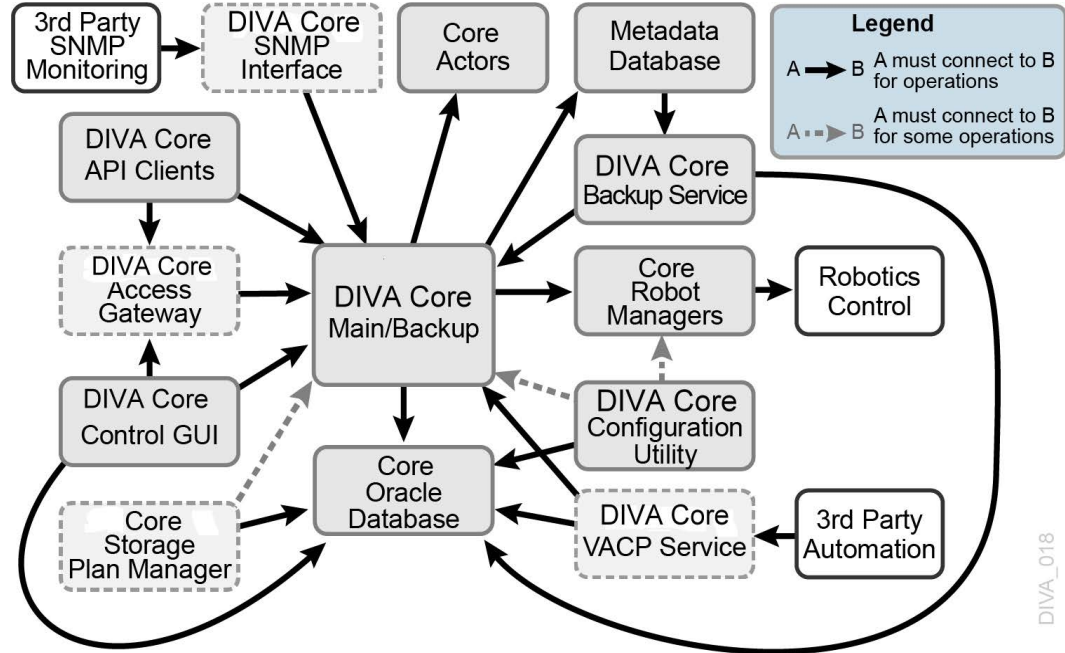
Confirm that the drives being used by DIVA Core are bound to the client dedicated to the Core ADIC Robot Manager. In the following figure, two drives are bound to Client2.

Use the following procedure to bind a drive to a specific client:

1. Open the SDLC GUI.
2. Click the Clients tab.
3. Select the client from the Clients navigation tree on the left.
4. Right-click the desired drive in the right Drives area.

5. Select Up from the menu.

After a drive is bound to a client, the name of the client is appended to the name of the drive.



6. Confirm that for the given client (Client2 in the previous example), the Client Hostname field is configured with the host name or IP address of the client computer; that is, the client that will use the Client2 connection when communicating with the SDLC Server.

You can specify any in this field to accept any incoming connections from any client computer that provides Client2 as the client name when connecting through the SDLC API. You must use the any keyword to use the dasadmin tool from a supervising computer.

7. Confirm the NobleNet PortMapper for TCP Windows service is started. You must start the service if it is not running.

SDLC Client

The SDLC Client must be installed on the computer where the Core ADIC Robot Manager is installed.

Installation

Install the SDLC Client from the SDLC distribution. You are prompted for the name of the client being used by the ADIC Robot Manager to connect to SDLC Server during installation. You must use the client you created in the SDLC Server. This is Client2 in the example.

Note: The client name is case-sensitive.

Confirm the NobleNet PortMapper for TCP Windows service is started. You must start the service if it is not running.

Configuration

Two Windows environmental variables must be defined on a Windows system as follows:

Environment Variable	Definition	Example
DAS_SERVER	Host name or IP address of the computer where the SDLC Server has been installed.	10.201.10.100
DAS_CLIENT	Name of the client that the Core Robot Manager uses to connect to SDLC.	Client2

Use the following procedure to test the SDLC Client connection to the SDLC Server:

1. Open a Windows command line window.
2. Change to the C:\Program Files\ADIC\SDLC\Bin folder.
3. Execute dasadmin qversion.

The output will be similar to the following, and then the command prompt is displayed.

```
ACI-Version: 3.10E
DAS-Version: 3.10
```

Using dasadmin Commands

The following is a list of commands used when executing the dasadmin application. You must always execute dasadmin from the C:\Program Files\ADIC\SDLC\Bin folder.

Getting Help

```
dasadmin -h
```

Mounting a Tape

```
dasadmin mount {tape_id} [drive]
```

The `tape_id` is required. If `drive` is not specified, the first free drive is chosen automatically.

Dismounting a Tape

```
dasadmin dism {tape_id}
```

Alternatively you can execute `dasadmin dism {drive_name}`. The `drive_name` is the name of the drive to dismount.

Note: The tape must first be ejected with a SCSI unload before dismounting.

Ejecting a Tape

```
dasadmin eject2 {tape_name} {eject_or_insert_slot_name}
```

Note: Depending on the server configuration, the eject and insert area (that is, slots from the CAP) can have different names.

Inserting a Tape

```
dasadmin insert2 {-n|-c} {eject_or_insert_slot_name}
```

You use the `-n` to specify data tapes and the `-c` to specify cleaning tapes.

Note: Depending on the server configuration, the eject and insert area (that is, slots from the CAP) can have different names.

Querying Drives

```
dasadmin ld
```

Retrieving the Tapes List

```
dasadmin qvolsrange """" {number_of_tapes_to_list}
```

Parking the Robot Arm

```
dasadmin robhome {robot_number}
```

Synchronizing the SDLC Database and Library

```
dasadmin inventory
```

Retrieving Tape Information

```
dasadmin view {tape_id}
```

dasadmin Release Information

```
dasadmin qversion
```

Library Configuration Information

```
dasadmin eif_conf
```

Note: This command is not supported in SDLC 2.1 and later.

dasadmin References

See the `sdlc_doc.pdf` file on the SDLC Installation CD.

Troubleshooting

The dasadmin qversion command may not respond as previously stated. The following list identifies the most common cases and remedies:

RPC failure error dialog box appears

A dialog box appears on the screen with the title ACI0004 Function clnttcp_create Failed, and the following error displays in the command window:

```
version failed: An RPC failure occurred.  
ACI-Version: 3.10E  
DAS error = 1
```

To resolve this issue, confirm on the server that a connection to this client can be established from the computer where dasadmin was launched.

Invalid host name or IP Address error in command window

The following error appears in the command window:

```
version failed: Invalid hostname or IP Address  
ACI-Version: 3.10E  
DAS error = 14
```

To resolve this issue, confirm on the server that a connection to this client can be established from the computer where dasadmin was launched. The client host name is probably set to localhost.

Invalid pointer to IDAS interface error in command window

The following error appears in the command window:

```
version failed: Invalid pointer to IDAS interface  
ACI-Version: 3.10E  
DAS error = 28
```

To resolve this issue confirm the DAS_CLIENT environment variable is set properly.

The command never ends (endless loop)

If the command results in an endless loop and never stops, confirm the following:

- Confirm the SDLC Server is started.
- Confirm the DAS_SERVER environment variable is set properly.

Confirm the NobleNet PortMapper for TCP Windows service is running.

Appendix F: Backup Service and DBAgent Configuration

Topics:

- [Sample BKS Configuration File](#)
- [Sample DBAgent Configuration File](#)

Sample BKS Configuration File

The following is a sample Backup Service configuration file with field descriptions and is located in \$DIVA_HOME\Program\conf\db_agent\appsettings.json:

```
{
  <=== CONFIGURES THE LOGGING LEVEL ===>
  "Logging": {
    "LogLevel": {
      "Default": "Debug", <=== CONTROLS THE GENERAL LOGGING LEVEL
      "Microsoft": "Warning",
      "Microsoft.Hosting.Lifetime": "Information",
      "System.Net.Http.HttpClient": "None"
    },
    "Configuration": {
      "RetentionDays": 7 <=== NUMBER OF DAYS TO KEEP SERVICE LOGS
    }
  },
  "Version": "0.0.0.0",
  "AllowedHosts": "*",
  <=== CONFIGURES THE DATABASES MANAGED BY THE BACKUP SERVICE ===>
  <=== NOTE: THIS IS CONFIGURABLE VIA THE API ===>
  "DatabaseSettings": {
    "Databases": [
      {
        "ProfileName": "MetadataDatabase", <=== PROFILE NAME FOR A
        SPECIFIC DB
        "DatabaseName": "Core", <=== NAME OF THE DATABASE
        "DatabaseType": "MongoDB", <=== TYPE OF DB (MongoDB,
        Oracle, Postgres, ElasticSearch)
        "DatabaseVersion": "5.0", <=== DB VERSION
        "ConnectionString": "mongodb://127.0.0.1:27017/
        ?replicaSet=rs0", <=== CONNECTION STRING
        "RootDirectory": "", <=== ROOT DIRECTORY FOR
        THE DB BINARIES
        "DataDirectory": "", <=== DATA DIRECTORY FOR
        THE DB DATA FILES
        "User": "MongoAdmin", <=== ADMIN DATABASE USER
        "Password": "<A password>" <=== ADMIN PASSWORD, WILL
        BE ENCRYPTED ON SERVICE START
      },
      {
        "ProfileName": "OracleDatabase",
        "DatabaseName": "lib5",
        "DatabaseType": "Oracle",
        "DatabaseVersion": "12.1.0",
        "ConnectionString": "",
        "RootDirectory": "",
        "DataDirectory": "",
        "User": "diva",
        "Password": "<A password>"
      }
    ]
  },
  <=== CONFIGURES THE BACKUP LOCATIONS AND REPLICATION ===>
  <=== NOTE: THIS IS CONFIGURABLE VIA THE API ===>
}
```

```

"LocationSettings": {
  "Locations": [
    {
      "Name": "Primary",
      LOCATION <=== NAME OF THE
      "Primary": true,
      LOCATION CAN BE CONFIGURED <=== ONLY ONE PRIMARY
      "Enabled": true,
      LOCATION IS ACCEPTING REQUESTS <=== WHETHER THE
      "Location": "H:\\divaback",
      LOCATION <=== PATH TO THE
      "AgentUrl": "https://localhost:1878/", <=== DBAgent API
      URL
      "Type": "Local",
      (either Local or UNC) <=== TYPE OF LOCATION
      "ManagedDatabases": [
      DATABASES <== LIST OF MANAGED
        "OracleDatabase",
        "MetadataDatabase"
      ],
      <=== LIST OF DATABASES THAT ARE REPLICATED TO THIS LOCATION
      ===>
      <=== NOTE: DATABASES THAT ARE MANAGED BY A LOCATION SHOULD
      BE EXCLUDED ===>
      "BackupReplication": [],
      "SourceName": "DIVADB_Backups", <=== NAME OF THE DIVA CORE
      SOURCE FOR THIS LOCATION
      "User": "",
      <=== SYSTEM ACCOUNT TO
      CONNECT TO THIS LOCATION IF ANY
      "Password": ""
      <=== SYSTEM ACCOUNT PASSWORD
      IF ANY
    },
    {
      "Name": "Secondary",
      "Primary": false,
      "Enabled": true,
      "Location": "\\100.10.10.10\\H$\\divaback",
      "AgentUrl": "https://100.10.10.10:1878/", <=== REMOTE
      DBAgent URL
      "Type": "UNC",
      <=== TYPE UNC or
      Local
      "ManagedDatabases": [],
      "SourceName": "",
      "User": "Administrator",
      <=== UNC PATHS
      MUST HAVE CREDENTIALS
      "Password": "<A password>"
    }
  ]
},
<=== CONFIGURES THE API TIMEOUTS AND SESSION EXPIRATION ===>
"ServiceSettings": {
  "RequestExpiration": 3600,
  "RequestTimeout": 600
},
<=== CONFIGURES THE API ENDPOINTS ===>
"HttpServer": {
  "Endpoints": {

```



```
    "BackupExecutionTimeout": 120,           <=== NUMBER OF MINUTES
BEFORE BACKUP TIMEOUT
    "RestoreExecutionTimeout": 120,         <=== NUMBER OF MINUTES
BEFORE RESTORE TIMEOUT
    "StatusPollingPeriod": 3,              <=== NUMBER OF SECONDS
BETWEEN POLLING BACKUP STATUS
    "StatusReportingInterval": 1440        <=== NUMBER OF MINUTES
BETWEEN SUPPRESSING DUPLICATE ALERTS
    "ArchivePriority": "Low",              <=== ARCHIVE PRIORITY:
Min, Low, Normal, High, Max or Integer
    "ArchiveWindowStart": "00:00:00",      <=== START TIME DURING
WHICH ARCHIVE BACKUPS ALLOWED
    "ArchiveWindowEnd": "23:59:59"        <=== END TIME DURING
WHICH ARCHIVE BACKUPS ALLOWED
  }
}
```

Sample DBAgent Configuration File

The following is a sample DBAgent configuration file with field descriptions:

```
{
  <=== CONFIGURES THE LOGGING LEVEL ===>
  "Logging": {
    "LogLevel": {
      "Default": "Debug", <=== CONTROLS THE GENERAL LOGGING LEVEL
      "Microsoft": "Warning",
      "Microsoft.Hosting.Lifetime": "Information",
      "System.Net.Http.HttpClient": "None"
    }
  },
  <=== CONFIGURES THE API ENDPOINTS ===>
  "HttpServer": {
    "Endpoints": {
      "Https": {
        "Host": "localhost",
        "Port": 1878,
        "Scheme": "https",
        <=== CERT PATH FOR THE HTTPS ENDPOINT ===>
        "FilePath": "../../security/certificates/DBAgent.p12"
      }
    }
  },
  "ServiceConfiguration": {
    "BasePath": "H:\\divaback", <=== BASE PATH TO THE LOCATION
    "LogRetention": 30, <=== NUMBER OF DAYS TO PRESERVE
    STATE FILES
    "MountPointMonitors": [
      {
        "Path": "H:", <=== MOUNT POINT TO MONITOR
        "IsPercentBased": true, <=== DETERMINES WHETHER TO USE
        PERCENTAGE OR THRESHOLD VALUE
        "ErrorThreshold": 0,
        "WarnThreshold": 0,
        "ErrorPercentage": 95,
        "WarnPercentage": 85
      },
      {
        "Path": "C:",
        "IsPercentBased": true,
        "ErrorThreshold": 0,
        "WarnThreshold": 0,
        "ErrorPercentage": 95,
        "WarnPercentage": 85
      },
      {
        "Path": "E:",
        "IsPercentBased": true,
        "ErrorThreshold": 0,
        "WarnThreshold": 0,
        "ErrorPercentage": 95,
        "WarnPercentage": 85
      }
    ],
  },
}
```

```
{  
  "Path": "F:",  
  "IsPercentBased": true,  
  "ErrorThreshold": 0,  
  "WarnThreshold": 0,  
  "ErrorPercentage": 95,  
  "WarnPercentage": 85  
}  
]  
}
```

Glossary

CA (Certificate Authority)

A CA (Certificate Authority) is an issuer who receives the CSR and returns the SSL certificate with its digital signature.

CSR (Certificate Signing Request)

A CSR (Certificate Signing Request) is an encoded file that is given to a CA (Certificate Authority) when requesting an SSL certificate. It contains information that will be included in the certificate including the holder's name, serial number, expiration date and the public key. The CA returns the signed SSL certificate with its signature.

DNS (Domain Name Service)

A system for naming computers and network services that is organized into a hierarchy of domains. DNS services resolve IP addresses to host names for proper network routing.

FQDN (Fully Qualified Domain Name)

The complete domain name for a specific computer, or host, on the internet. The FQDN consists of two parts: the host name and the domain name. For example, rd-mc1-qalab.oracle.com.

Java Keystore

The JAVA Keystore is a password protected encrypted file repository containing the Key pairs, SSL certificates, and CA certificates.

Key Pair

A Key Pair consists of two uniquely related cryptographic keys; a Public Key and a Private Key (basically long random numbers).

The Public Key is what its name suggests - Public. It is made available to everyone through a publicly accessible repository or directory.

The Private Key must remain confidential to its respective owner. Because the key pair is mathematically related, whatever is encrypted with a Public Key can only be decrypted with its corresponding Private Key, and vice versa.

MPIO (Multipath I/O)

Microsoft MPIO (Multipath I/O) is a Microsoft-provided framework that allows storage providers to develop multipath solutions that contain the hardware-specific information needed to optimize connectivity with their storage arrays.

NIC Teaming

The process of combining multiple network adapter cards together for performance and redundancy reasons. Microsoft refers to this as NIC Teaming, however other vendors may refer to this as bonding, balancing, or aggregation. The process is the same regardless of which solution is used or what it is called.

OU (Organizational Unit)

An OU (Organizational Unit) is a subdivision within an Active Directory into which you can place users, Tape Groups, computers, and other organizational units. You can create organizational units to mirror your organization's functional or business structure. Each domain can implement its own organizational unit hierarchy. If your organization contains several domains, you can create organizational unit structures in each domain that are independent of the structures in the other domains.

SAS (Serial Attached SCSI)

A point-to-point serial protocol that moves data to and from computer storage devices such as hard drives and tape drives.

SSL (Secure Sockets Layer)

SSL (Secure Sockets Layer) is a standard security protocol for establishing an encrypted connection between a server and a client. Specifically, it encrypts the connection and the data transmitted along the connection. To achieve a secure connection, a service needs a Key Pair (Public Key and Private Key) and SSL Certificate.

SSL Certificate Authentication

An SSL certificate is a digital certificate that authenticates a service in network connections. To generate an SSL certificate, you must create a CSR (Certificate Signing Request) for your service Key Pairs and have it signed by your CA (Certificate Authority). An SSL certificate contains the following information:

- Certificate holder's name
- Certificate serial number and expiration date
- A copy of the certificate holder's public key
- Digital signature of the certificate issuing authority

SSL Certificate Chain

There are two types of CAs (Certificate Authorities): Root CAs and Intermediate CAs.

A certificate chain is an ordered list of certificates, containing an SSL Certificate and Certificate Authority Certificates that enable the receiver to verify that the sender and all CAs are trustworthy using its trust store. The chain (or path) begins with the SSL certificate, and each certificate in the chain is signed by the entity identified by the next certificate in the chain. Any certificate that sits between the SSL Certificate and the Root CA Certificate (last certificate in the chain) is called an Intermediate CA Certificate. The

Root CA is at the end of the chain and it signs the intermediate CA certificate, and the Intermediate CA signs the SSL certificate for the services.

For example, when a service receives its peer's SSL certificate chain that is trying to connect during the SSL handshake process, it verifies its peers SSL certificate in the chain using the Intermediate CA certificate next in the chain. It then verifies the Intermediate CA certificate by looking for the Root CA certificate that signed the intermediate CA certificate in its trust store. This verification completes the Certificate Chain. Connection is not established if the full chain verification fails.

Trust Store

A Trust Store contains the certificates of CAs (Certification Authorities) you trust. For example, when a service receives its peer's SSL certificate that is trying to connect during SSL handshake process, it verifies that its peer's SSL certificate's digital signature is signed by one of the trusted certificates in its trust store. If the certificate is not in the Trust Store, the SSL handshake fails and the connection is not established.