Telestream

Installation and Configuration Guide

Release: 9.0

Revision: 1.0

January 9, 2024

2401091235

Copyrights and Trademark Notices

Contents

Unmanaged Server Configuration Guide 171

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 Technical Support	Web Site: https://www.telestream.net/telestream-support/ 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. • Support hours for customers are Monday - Friday, 7am - 6pm local time. • P1 issues for customers are 24/7.
Telestream, LLC	Web Site: www.telestream.net Sales and Marketing Email: info@telestream.net Telestream, LLC 848 Gold Flat Road, Suite 1 Nevada City, CA USA 95959
International Distributor Support	Web Site: www.telestream.net See the Telestream Web site for your regional authorized Telestream distributor.
Telestream Technical Writers	Email: techwriter@telestream.net Share comments about this or other Telestream documents.

Preface

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

Note: The initial release of DIVA is a Windows-only release and does not currently support Linux.

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/.

Related Documents

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

https://www.telestream.net/telestream-support/.

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/.

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

Document Updates

The following table identifies updates made to this document.

Date	Update
February 2023	Removed information to create new books for the 9.0 release. Updated book with all new information from the 8.3.1 release. Removed all references to Oracle database and replaced with Postgres. Updated installation process and screenshots.
March 2023	Removed all references to Linux installation and configuration. Renamed appendixes that will remain to appropriate titles. Corrected grammar and spell checked entire book. Updated product name from DIVA Core to DIVA.
June 2023	Updated Web App term. Added section on setting up tape groups and referenced the SPM book. Added DIVA installation instructions. Added instruction to import a license. Removed DIVA View. Added ports for Analytics Windows Service Updated book with review comments
September 2023	Updated product name from Content Manager to DIVA. Publish Version 9.0 PDF.

Overview

This chapter provides you an overview of DIVA.

Topics

Release Overview

The DIVA 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 installation varies from site to site, therefore the exact configuration of your specific DIVA platform is not described in this book.

Note: The File System Interface is not released with DIVA 9.x and later and is only supported by special request.
DIVA Command has been deprecated.

Port Utilization

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

DIVA Service	Default Port	Needed by External Applications	Description and Notes
FTP	21 / TCP	Yes	Port depends on configuration
HTTP	80 / TCP	Yes	web app
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
Proxy Actor	8800 / UDP	No	Typically installed on Actor Server
Analytics Windows Service	2176/TCP 2177/TCP	No	HTTP HTTPS
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

DIVA Service	Default Port	Needed by External Applications	Description and Notes
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 | DIVA Core 9.0 and later)	8761 / TCP	Yes	Typically installed on Manager Server
REST API Data Service (DIVA Core 8.0 | DIVA Core 9.0 and later)	13443 / TCP	Yes	Typically installed on Manager Server
REST API Gateway (DIVA Core 8.0 | DIVA Core 9.0 and later)	8765 / TCP	Yes	Typically installed on Manager Server
REST API DIVA Connect Adapter (DIVA Core 8.0 | DIVA Core 9.0 and later)	17443 / TCP	No	Typically installed on Manager Server
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
Postgres Database	5432 / 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 Service	Default Port	Needed by External Applications	Description and Notes
DIVA Connect Client Adapter	9801 / TCP 7101 / TCP 7100 / 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 Release Notes in the DIVA documentation library at:

https://www.telestream.net/telestream-support/

Note: Object Names cannot begin with a dollar sign ($).

Media Storage Formats

This section describes the supported media storage formats.

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 jobs 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 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, 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 Legacy or 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 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 is written on that tape. The tape is assigned the format of the Tape Group that appears in the job. 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 job 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 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 job.

Use the Copy To Group, or Copy As New jobs 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 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 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 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 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 DIVA web app's Object Properties and Tapes panels, only a single placeholder file is shown. The DIVA API includes a command to return all of the files and folders within a complex object. See the appropriate DIVA API documentation on the DIVA Technical Support site for detailed information.

DIVA Connect does not 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 job 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 the *Unmanaged Server Configuration Guide* for detailed parameter information.

Use the following procedure to include the parameters in the Server page in web app:

1. Open the DIVA web 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 jobs 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.

Software Components

DIVA includes the software components discussed in the following subsections.

Long path names and absolute path names are supported on Windows.to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems.

Shortcuts created using the Windows operating system are not represented as symbolic Storage Links because they are treated as files.

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 Avid Connectivity
- Transcoder integration
- Tape Reading Utility

Due to degraded performance, Windows IIS and FileZilla FTP Source Servers and Destination Servers cannot be used for Complex Objects. The Windows IIS and FTP servers cannot accommodate large numbers of files.

Manager

The Manager is the main component in a DIVA system. All archive operations are controlled and handled by the Manager. Operation jobs are sent by initiator applications through the DIVA Client API. DIVA supports Main and Backup Managers.

Actor

The 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 Manager through a TLS 1.2 secure connection. Key benefits of the distributed design of the Actors are:

- Expand the archive subsystem to increase the overall bandwidth by adding more Actors to the system.
- Share SAN based disk and tape drive resources among multiple Actors.
- In combination with the Manager, multiple Actors provide scalability, load balancing, redundancy, and fail-over. Take individual Actors offline for maintenance without shutting down the DIVA 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 reports an error if an empty file or folder is discovered during the transfer.

Client APIs

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

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

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

REST API

DIVA exposes its functionality through a REST interface. It is self-contained in DIVA 9.0 and all future DIVA releases. In the 9.0 release, the API is used by the DIVA web app.

Note: Telestream recommends using the REST API rather than the previous existing APIs (that is, 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 on the DIVA Technical Support site for detailed information.

Metadata Database

DIVA stores Object metadata separately from the DIVA Database in the DIVA Metadata Database (MongoDB). The metadata database contains files stored in a file system local to the Manager. The files are located in the Metadata Database root folder. This storage method enables DIVA 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 DIVA Database, and it must be backed up at regular intervals through the DIVA 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 Windows installer starting with release 8.2.0.91. The DIVA installer identifies it as Notification Service instead of RabbitMQ because RabbitMQ is just an implementation.

The Notification Service is required for the web app to function properly.

Web App

The DIVA web app is installed as part of the DIVA installer. 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 web app connects to both the Manager and the DIVA Database. Use it to monitor, control, and supervise operations in DIVA. Multiple web app instances can be operated simultaneously from any computer that has TCP/IP connectivity to both the Manager and the DIVA Database.

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

See the DIVA Operations Guide on the DIVA Technical Support site for more information on using the interface.

The refresh rate for the web app is set in the Manager Setting page of the web app in the GUI: Dashboard Refresh Delay field.

Additional Software Components

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

Robot Manager

DIVA 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 Manager to interact with many different types of tape Managed Storage. It is connected to the Manager through TCP/IP.

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 Manager is over TCP/IP, the library does not need to be co-located near the 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 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 has its own API for communicating with the 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 API commands on computers that have TCP/IP connectivity to DIVA.

SPM (Storage Policy Manager)

The DIVA 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 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 Simple Network Management Protocol interface supports status and activity monitoring of different DIVA components. DIVA MIB (Management Information Base) is provided to third party SNMP monitoring applications. The SNMP Agent uses the Windows SNMP Service.

DIVA Connect

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

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

DIVA Connect is used in DIVA Connect installations and is the portal for multiple DIVA systems to communicate with each other. See the DIVA Connect Installation, Configuration, and Operations Guide on the DIVA Technical Support site for detailed information.

Note: DIVA Connect 4.0 is used for DIVA 9.0 and later, and is backward compatible to earlier releases of DIVA Core.

Watch Folder Monitor

The DIVA 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 Object are supported. When a new file is identified, WFM issues an archive job automatically to DIVA to archive the new file. After the files are successfully archived, they are then automatically deleted from the Source Server. Refer to the DIVA Watch Folder Monitor User's Guide for more information.

Transcoder Support

The 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. Integration to BitScream products and Vantage are supported. However, multiple transcoders are only supported for Vantage.

DIVA uses 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. Select the transcoder type "tre" from the web 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.

See the DIVA Avid Connectivity and Tools book on the DIVA Technical Support site for detailed information.

Avid DHM (Data Handler Module) Interface

The Avid DHM (Data Handler Module) interface support in DIVA 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. DHM support is implemented in DIVA TMC (Transfer Manager Communicator).

Avid DET (Dynamically Extensible Transfer) Interface

The Avid DET (Dynamically Extensible Transfer) interface support in DIVA allows storage expansion of Avid Unity infrastructures and enables editors to move native Avid content in and out of the DIVA storage system. Partially edited content stored within DIVA 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 stores these files in native Avid format. DET support is implemented in DIVA TMC.

Archive Manager Interface

An interaction between the Avid Archive Manager solution and DIVA 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 to store Avid content in its native MXF OP1 Atom format.

Utilities

The following sections describe utilities available in the DIVA system.

Web App

DIVA Command has been deprecated and is replaced by the DIVA web app for configuration and operation of DIVA. You can use it on any computer with TCP/IP connectivity to the host running the DIVA Database.

Caution: The web app is intended only for experienced users. Incorrect or incomplete changes in the web app can adversely affect DIVA 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 web app primarily connects to the DIVA 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 system configuration. Although used primarily for configuration of DIVA, some operational functions are also performed from the web app.

Robot Manager Utilities

During configuration and troubleshooting of the library and its tape drives, DIVA 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 Manager is running because this can adversely affect archive operations.

BKS (Backup Service)

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

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

See the DIVA Database and BKS Installation, Configuration and Operations book on the DIVA Technical Support site for detailed information.

Scandrive Utility

This utility 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 configuration.

Tape Reading Utility

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

This utility 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.

DIVAscript

This utility allows DIVA C++ API commands to be executed using DOS based scripts. It is designed to run automated tasks for testing rather than for any intensive uses.

RDTU (Recover Damaged Tape Utility)

The DIVA 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 system.

Metadata Database Installation and Configuration

This chapter describes installation and configuration of the Metadata Database.

Topics

Metadata Database Configuration

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

Configuring the Metadata Database

You must set the following two parameters on the Manager Setting page of the web 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 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+avg_path_file_name_size)*1.15*avg_number_component_files*number_virtual objects

When planning, enough Metadata Database disk space should be allocated to ensure expected, or unexpected, growth of your environment. You must allocated 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 being incorrectly installed, migration can be performed manually using *~/DIVA/Program/Manager/bin/DBMigrate.BAT*.

Remember the folder that contains the MDDB 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 9.0 and later.

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

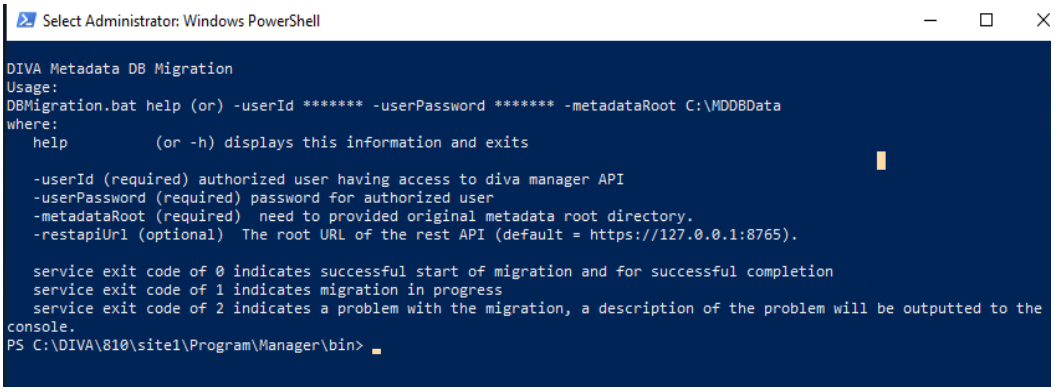

For example:

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

Migration begins the first time this is executed. Subsequent calls provide the current status until complete, as shown here:

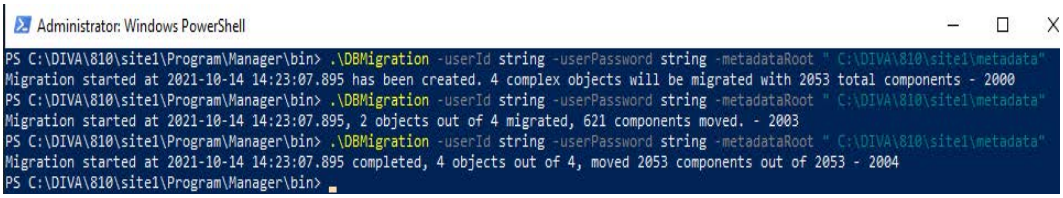

The first call created the migration and returned 2000. The next call shows two of four Objects migrated and returned 2003. This indicates the migration is still in progress. A typical migration will show this response many times. The final call shows the migration is complete with a return of 2004.

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%
```

Installing and Upgrading the Notification Service (RabbitMQ)

RabbitMQ is included in the DIVA Windows installer. It is the database used for the DIVA Core Notification Service.

Installation

The DIVA installer refers to RabbitMQ as the Notification Service instead of RabbitMQ, because RabbitMQ is simply an implementation. Use this procedure to install the Notification Service (RabbitMQ):

1. In the DIVA installer > Choose Components panel, check the Notification Service and click Next.

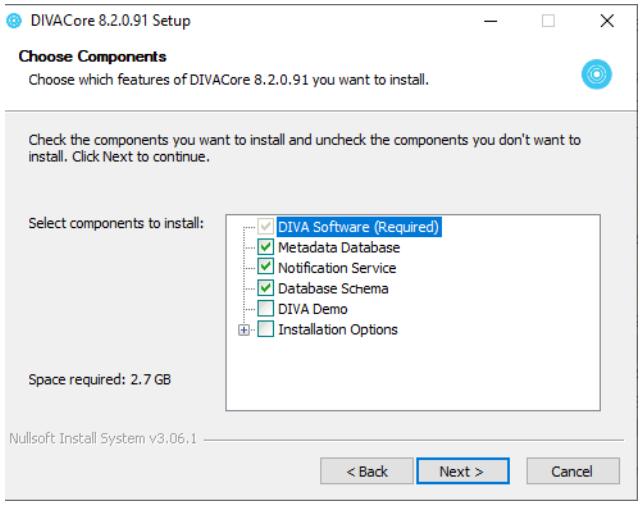

2. The Notification Service has its own database and by default is stored in the ~\DIVA\Program\rabbitmq folder. Modify as required:

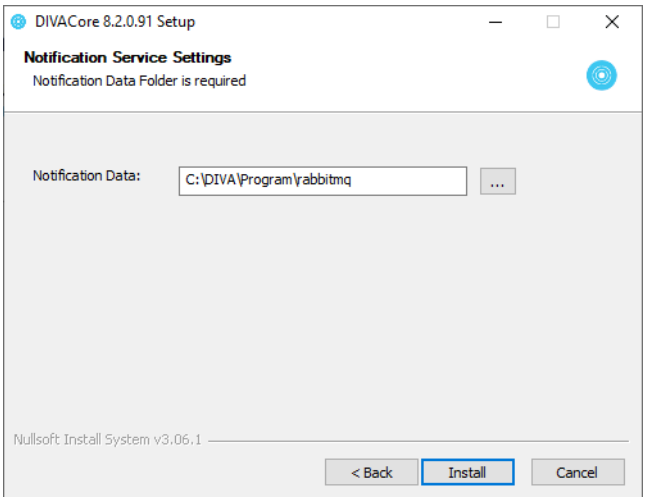

   The path to this folder is specified in a text file: *~\DIVA\rabbitmq_server\etc\NOTIFICATIONDataDir.txt*. This file is used by the installer when DIVA is upgraded into identify the path. However, the actual setting for the RabbitMQ service is in the Windows registry. Do not change the value in this file to cause RabbitMQ to use a new data directory. You must modify the registry if a different directory is desired.

3. Click Install.

Upgrading RabbitMQ Installations

When you upgrade DIVA, the installer overwriting files that already exist.

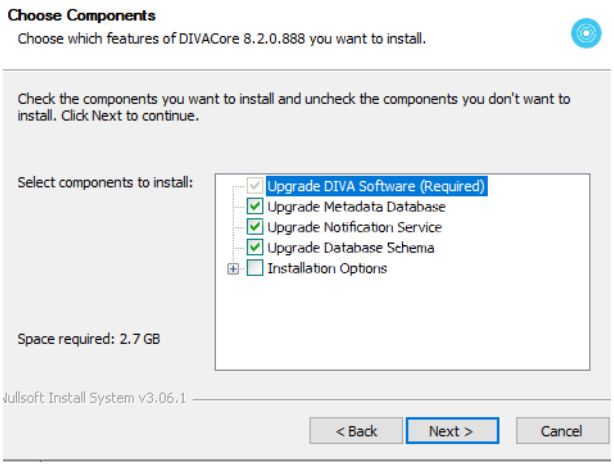

During upgrade, you can change the location of data folder but by default, installer will take the value from the DIVA\rabbitmq_server\etc\NOTIFICATIONDataDir.txt file. If the

user does not want to change it, it will be use the same folder as the previous RabbitMQ installation.

Note: 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 installer will always create a default admin user (username wsuser and password changeit): http://127.0.0.1:15672.

Troubleshooting

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

- *Metadata Database Failure Scenarios*
- *Manager Will Not Start*

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.

Note: Object Names cannot begin with a dollar sign ($).

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 system. Only Delete jobs are possible on existing Complex Objects. DIVA 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 job 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 job 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 web 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 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 jobs are possible on Complex Objects whose Metadata files are corrupted, except Delete Object jobs, until it is restored. A Metadata file modified by any external source (other than DIVA) 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 jobs 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 release; contact Technical Support for assistance.

Scenario 3: Lost or Manually Deleted Metadata Database File

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

You cannot perform any operations or jobs 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 jobs fail with the following error message:

```
Internal error: metadata database error:
Message: get: Error opening metadata for objectname/category, 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 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 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.

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.

Backup Service Will Not Start

The DIVA 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 web 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 web 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 web 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.

Installation

This chapter describes DIVA software components and system installation.

Topics

Software Component Distribution

The DIVA 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 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:

Managers

Main and Backup Manager servers

Postgres Database

Main and Backup Manager servers

Metadata Database

Main and Backup Manager servers

Backup Service

Main and Backup Manager and Actor servers

Robot Managers

Main and Backup Manager servers. Robot Managers can also be installed on a separate server when the tape library is installed a substantial distance from the Manager servers.

Storage Policy Manager

Main and Backup Manager servers

VACP Services

Main and Backup Manager servers

SNMP Agent

Main and Backup Manager servers

Connect

Main and Backup Manager servers

Actors

Actor servers

Transfer Manager Communicator (TMC)

Actor servers

Archive Manager Communicator (AMC)

Actor servers

Watch Folder Monitor

Actor servers

Installing the DIVA System

The following sections describe installation of the DIVA system. Contact Technical Support if you need assistance.

Notes: The Postgres Database must be available for DIVA before installation. See *Metadata Database Installation and Configuration*, and the DIVA Database and BKS Installation, Configuration, and Operations book.

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 web app will not be able to display the cloud array settings correctly.

The basic process to get the web 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 web app under the default URL address of https://127.0.0.1:8765.

Installation Overview

DIVA 9.0 installer includes an option to install the MDDB (Metadata Database). It is run in silent mode during installation.

The following procedure is a basic overview of the installation process on Windows. See the following operating system-specific sections for detailed instructions.

1. Install the DIVA Database user when running the DIVA installer. In Windows this is a check box.

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 web app after installation).
3. Configure the basic Manager settings to get the Manager Service operational.
4. Configure the REST API.
5. Start the web app and log in under the sysadmin account; then create a web app user. This is done so the sysadmin account is not being used to configure or view the DIVA system in the web 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 web 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 *Configuration* for configuration details.

Installed Features and Services

These Windows services and features are installed by DIVA, based on your selections:

- Backup Service
- DB Agent
- Metadata Service
- Data Service
- Discovery Service
- Rest API Gateway
- Actor
- Watch Folder Manager
- Manager
- Robot Manager
- VACP Service
- Rosetta
- Proxy Service
- Storage Plan Manager
- MongoDB
- RabbitMQ

Notes:
1. Postgres installer is not bundled and must be installed separately. Refer to the DIVA Database and BKS Installation, Configuration, and Operations Guide for detailed instructions.
2. Control GUI is removed and replaced by the web app.
3. Configuration Utility is removed and replaced by the web app.
4. Oracle Client and Database schema files are removed. Database schema files are re-factored into DBAgent/Packages folder.

Downloading the Software

You must stay current with the release of DIVA 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 software:

1. Log in to the Software Delivery Cloud and search for DIVA.
2. Select the licenses you require (for example, Actor, 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 for Windows

These figures illustrate the mandatory steps to install DIVA:

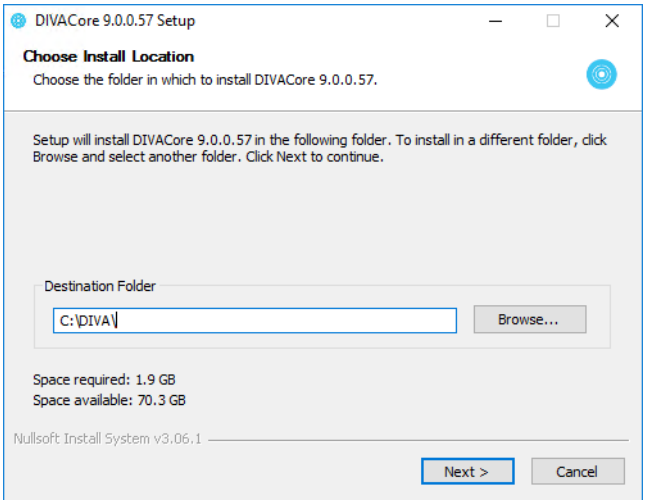

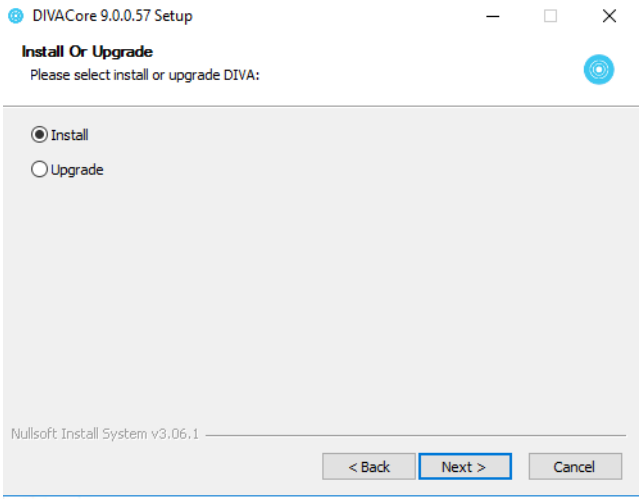

1. Choose which options to install. The DIVA Demo option can be selected to install and configure local disk storage and a library simulator. See the following DEMO section (Step 8) for more details.

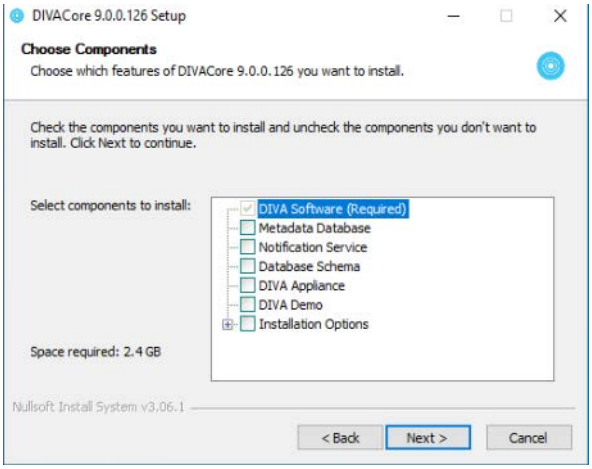

2. This option installs MongoDB and supports the following configuration options:
   - **MDDB Data Folder:** this is where MongoDB stores its data, by default it is H:\MDDBData. If H: drive does not exist it will default to C:\MDDBData.
   - **MDDB Port:** the default is 27017.

Note: If this port is changed, DIVA services that use this port will still default to 27017, so you will need to change those service configurations too).

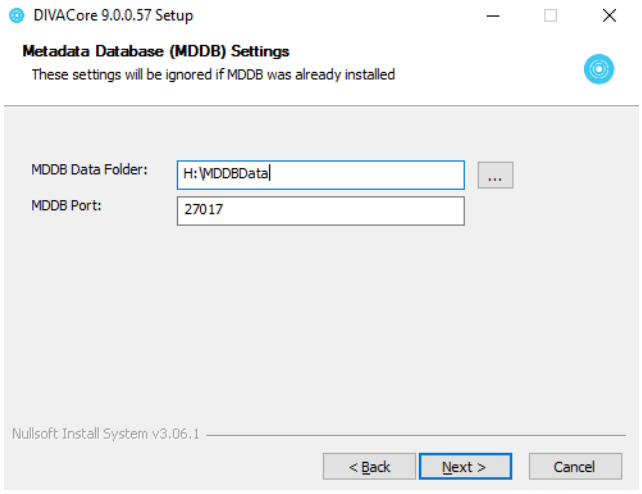

3. This option installs RabbitMQ and supports the following configuration options:
   - **Notification Data:** This is the folder where rabbitmq stores persistent notification queues. Logs and some configuration files are also stored in this folder.

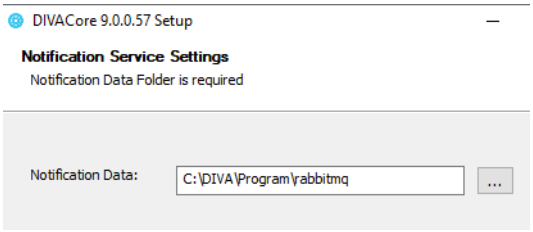

4. **Database Schema:** This option requires Postgres Database to be installed and setup properly and it will install the DIVA database schema and user to that database. It supports the following options:

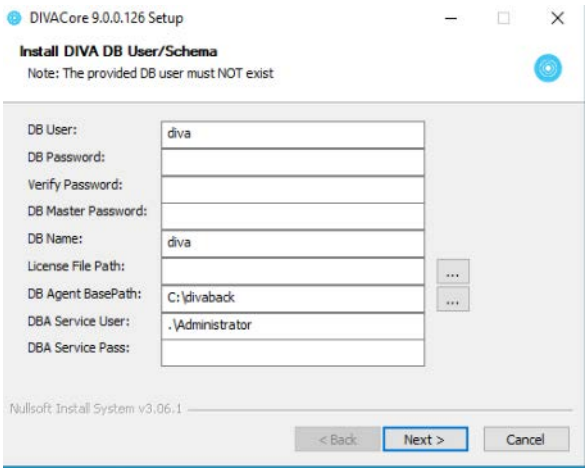

   - **DB User:** Schema User
   - **DB Password:** Schema User's password
   - **DB Master Password:** The master user (postgres)'s password when you install setup Postgres DB
   - **DB Name:** Schema's Name. Unlike Oracle, we recommend use same value as Schema's User because if you want to install more than one schema on the same Postgres DB, each schema must have unique DB Name and user
   - **License File Path:** path to the DIVA license file
   - **DB Agent Base Path:** This is the folder where DIVA will store DB Backups
   - **DBA Service User:** The user name to run DBAgent service (this needs to be an administrator user because DBA needs to access network shares to copy backup files to and from)
   - **DBA Service Pass:** The password for the user to run DBAgent service

     If the DB Agent File Path does not exist the following error message is displayed. Create the desired path to proceed.

5. **Appliance:** This option requires all all other options to be also selected except the DEMO optl1nion. It sets up a DIVA system with one Actor and an empty configuration. It is designed to install DIVA on a customer's production server, but due to the specific details of the production environment being unknown, it only

starts DIVA with an empty configuration. The DIVA web app must be used to complete the rest of the DIVA setup. It supports the following options:

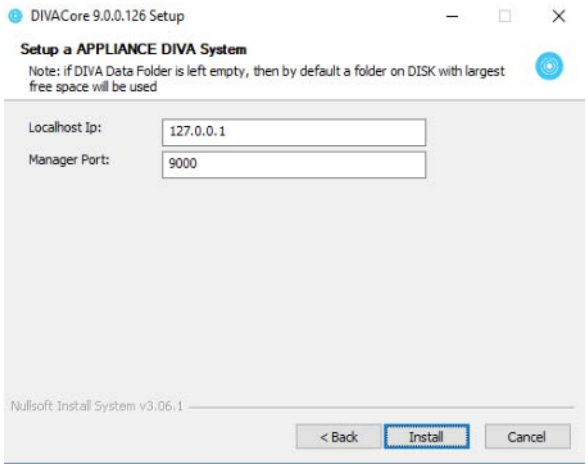

- **Localhost Ip:** It is recommended to use the actual IP of the production server instead of the loopback IP (127.0.0.1).
- **Manager Port:** Manager's legacy API port.

6. **DEMO:** This option requires all other options to be also selected except the Appliance option. It sets up a DIVA system with two Actors, local disk servers, a simulated tape library (and so on) to allow demonstrating DIVA features. This option should work on most DIVA servers, but is not designed to be flexible enough to work on any DIVA server. For example, if your DIVA server does not have a C: drive or H: drive, it may not work correctly. When the DEMO installation is complete, a fully working DIVA system should be installed and setup.

   DEMO supports the same options as Appliance, but also asks you to enter:

   - **DIVA Data Folder:** this is the folder where DIVA managed and unmanaged storages will be stored under (for example, source/destinations, disk arrays, simulated tape libraries, and so on).

The installation logs are available in:

`C:\DIVA\Program\log\diva_upgrade\diva_upgrade.trace.log`

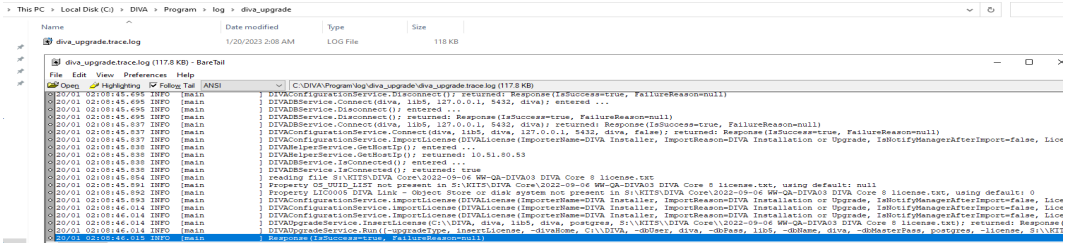

Web App Installation and Configuration

The DIVA web app is installed as part of the DIVA 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 web app. However, the Notification Service must be installed.

For the web app to function properly the following must be installed:

- Notification Service: Required for the Running Jobs page to display Job Status in real-time.
- RestAPI Data service: Required for login to DIVA web app.
- Manager: Required by everything else in DIVA web app.

By default, the Manager expects DIVA web 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 web app is accessible at https://127.0.0.1:8765/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 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 item. You can pin the full menu by clicking the hamburger button on the top banner.

Back-end Support

The DIVA web app only requires a running Manager and Data Service. You need to update api.server.port in the manager.conf file to update the web app server port. The default port is set to 8765. 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 8765 will point to a secure URL where you can access the web app interface at https://xxx.xxx.xxx.xxx:8765/DIVAWebUI where xxx.xxx.xxx.xxx is the web app server IP address. For example, https://127.0.0.1:8765/DIVAWebUI.

Importing the License

DIVA 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 Configuration > User Management > License > Import License page in the web app. In addition to enabling the Manager, the license includes a set of options that are necessary to enable the associated features in DIVA.

Use the following procedure to import a license using the web app:

1. Navigate to Configuration > User Management > License > Import License.
2. Enter the Importer's name in the Importer text box.
3. Enter the reason for importing this license in the Importer Reason text box.
4. Click Choose File under the License File Content heading and locate the license file to import.
5. Enter the Manager IP Address in the Manager Address text box.
6. Enter the Manager port number in the Manager Port text box.
7. Turn on Notify Manager After Import using the slide button.
8. Click Save at the bottom of the screen to import the license.

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.

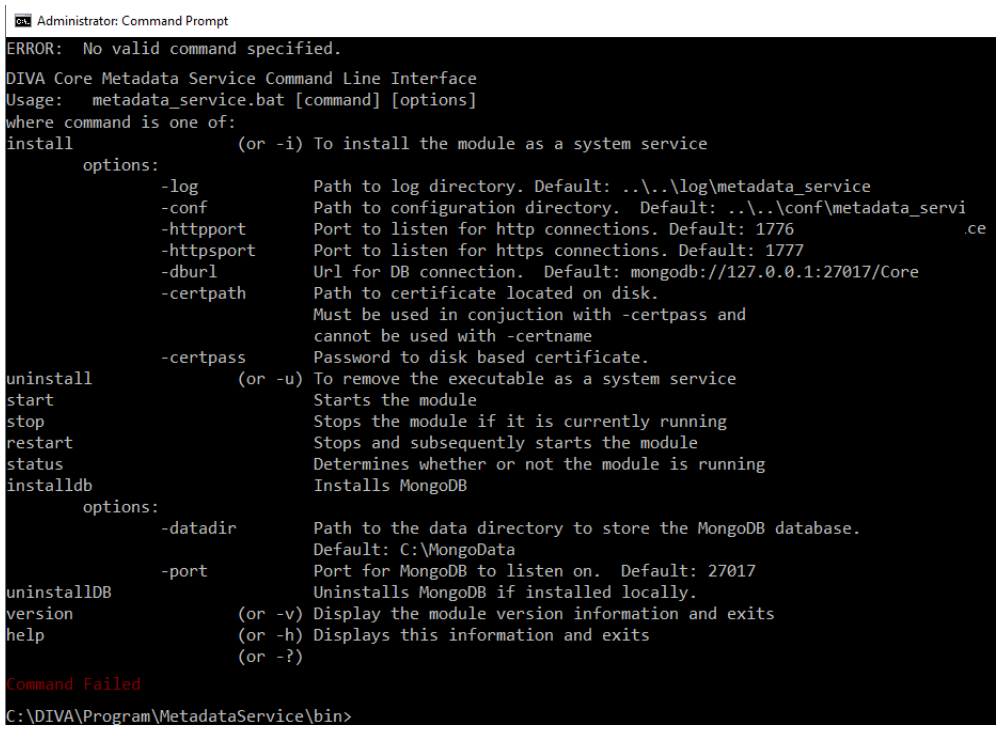

The MDDB service requires MDDB to work correctly; which must be configured in the ~\DIVA\Program\conf\metadata_service\appsettings.json file as a ConnectionString. The metadata-service.bat assumes ConnectionString = "mongodb://127.0.0.1:27017/Core" by default. However, this only works if MDDB is installed 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.

Configuration

Use the DIVA web app to configure DIVA to meet your requirements and work properly in your environment.

Topics

Configuration Overview

DIVA includes the following changes:

- Migrates configuration and operational functionality to the web app.
- Completely removes DIVA Command, DIVA Configuration Utility, and DIVA Control Panel programs.

Module Configuration Files

Each DIVA 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 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 centralizes them to a dedicated conf subfolder under the DIVA 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 will have the new releases of the .conf files appended with a .ini extension. For example, the new release of the 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 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 codependencies between some applications in the DIVA platform, so other components may also need to be restarted for changes to take effect.

Databases

See the DIVA Database and BKS Installation, Configuration and Operations Guide on the DIVA Technical Support site for detailed information on the DIVA Database. A full list of changes that can be made to the system configuration dynamically while the Manager is running is listed in *Dynamic Configuration Changes*.

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 Postgres Database. It should be backed up at regular intervals through the DIVA 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 Backup Service. Storing the Metadata Database files on a RAID array isolates the data from this type of failure.

The information stored in the Postgres 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+average_path_filename_size)*1.15*avg_num_component_files*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 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 applications on the Windows host. This particular parameter is required on any Windows computer that will run the DIVA web 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.

Prerequisites

This section describes the prerequisites necessary for successful DIVA configuration.

Database Installation and Configuration

See the DIVA Database and BKS Installation, Configuration, and Operations Guide for Database installation and configuration details.

DIVA Installer and Database Schemas

See *Installing the DIVA System* for details on the DIVA installer 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.

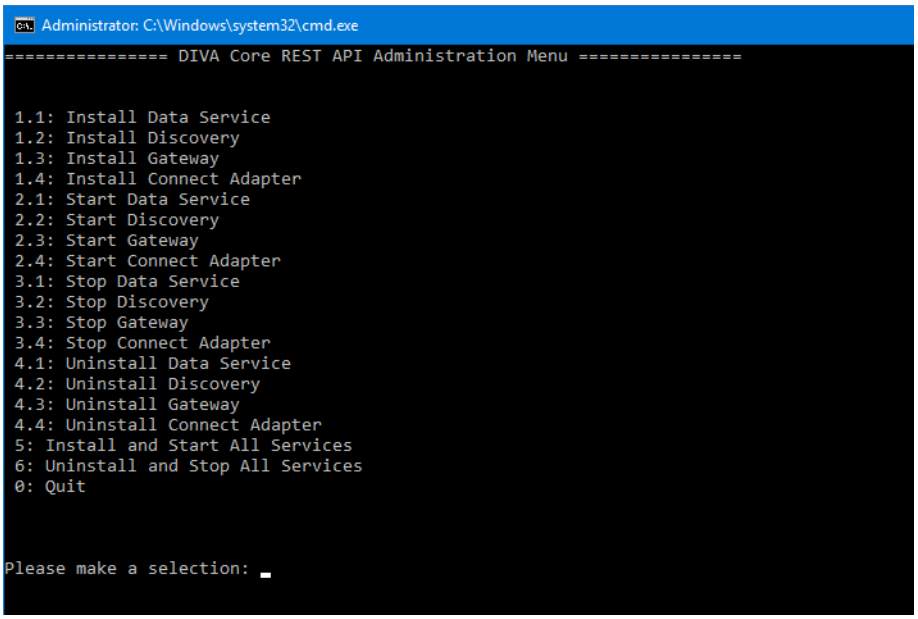

4. The listed services can be selected individually, or select option 5 to install and start all services (option 5 is recommended).
5. Open the Windows Services panel (services.msc) and confirm that the services were installed and are running.

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 web 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 web app. You log in the first time to the web app using *sysadmin* and *lib5* (the same password as the database).

Note: Log in to the web app the first time and immediately create a new user so that the sysadmin account is not being used for configuration and viewing DIVA functions.

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 Web App

Caution: The web app is intended only for experienced users. Incorrect or incomplete changes in the web app can adversely affect DIVA operations (and possibly even delete data from the archive), or prevent the Manager from running. Contact Telestream Support for assistance if you are unsure about making desired changes.

The DIVA web app primarily connects to the DIVA 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 menu on the top right of the screen as shown here:

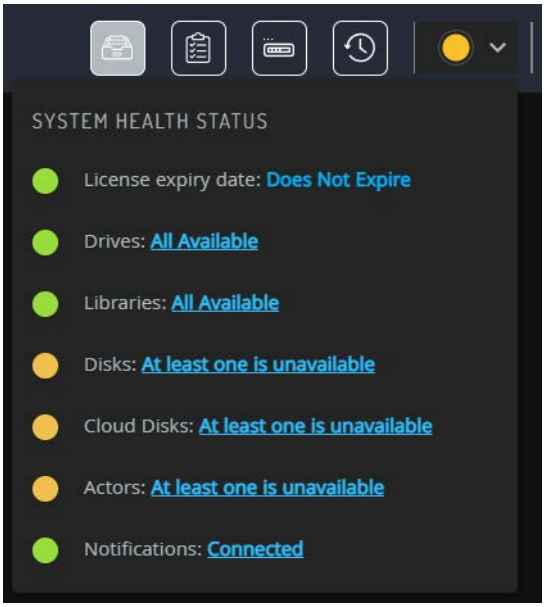

When a user logs into the web app, it automatically connects and checks the system status. Contact Telestream Support if you still cannot connect after attempting to resolve the error.

Disconnect the web app when not in use by logging out of the app and closing the browser being used.

Web App Profiles and Passwords

The DIVA web 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 change 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 web app as an Administrator or Operator. You must assign an Administrator and Operator password in the web app after DIVA installation is complete. Without an assigned password you are not permitted to switch to the respective profile in the web 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 web app the first time, it no longer matters what you use for the old password when changing passwords.

Setting Profile Passwords in the Web App

Profile passwords are initially created when the new user is created. Use the following procedure to change the profile password:

1. Open the web 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.

Web App Dashboard

The dashboard displays information at a glance when the web app starts. The dashboard home page displays four panels:

- System Info
- Utilization Stats
- Alerts
- Storage Resource

It is designed to give a quick overview of the system and it's usage.

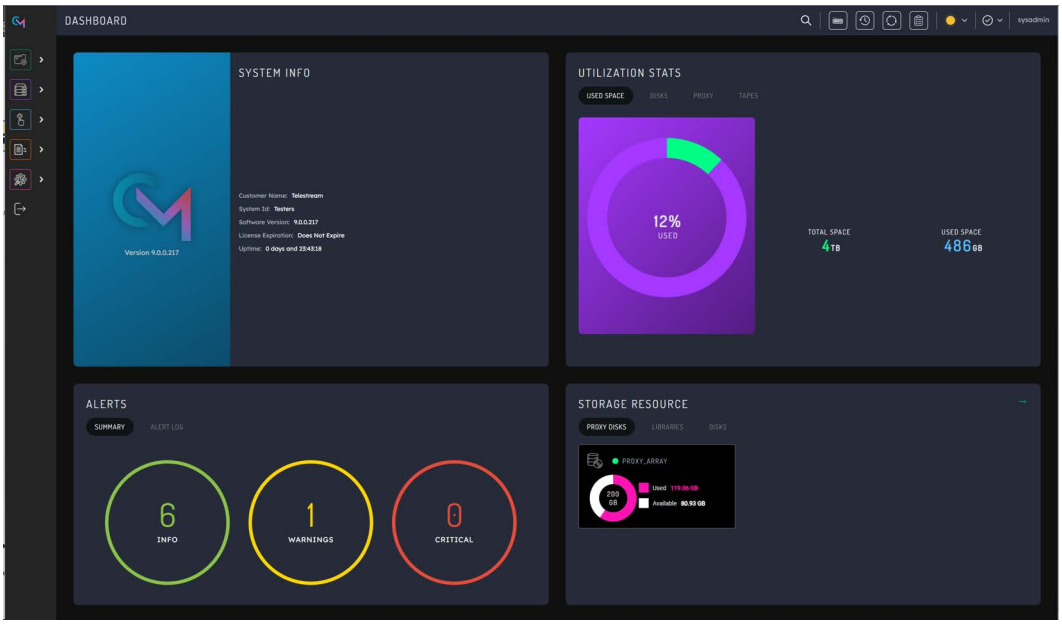

The System Info panel displays some properties from the license. For example, Customer Name, System Id, and the License Expiry Date. It also displays the software version of DIVE Core that is running, and how long the Manager has been running.

The Alerts panel displays a summary of how many of the metrics it is tracking are Info, how many are Warnings, and how many are Critical. The details of these metrics are on this panel. The metrics allow a quick view as to whether all Drives, Libraries, Disks, Cloud Disks, and Actors are available.

The Utilization Stats panel displays the total space available to the system for storage, and how much of that space is currently consumed. There is a combined view of Cloud, Disk, and Tape, and also these metrics broken out in to their individual categories.

Because Cloud storage is technically unlimited, this is omitted from the combined view unless a quota limit on the Cloud Bucket has been set in the configuration. The quota limit is only for display purposes on the dashboard. If it is not set, then the Cloud Disk will display just the amount of used space.

The Storage Resource panel displays similar information to the Utilization Stats panel. However, rather than totaling the information for all resources in a collection, it breaks them down in to individual resources.

Web App Menus

The following sections describe a general overview of each menu in the web app. Each menu includes multiple sub-menus where you configure different aspects of the system.

The web app does not have tabs like the Configuration Utility had. Instead, it contains a list of menu items on the left side of the screen that are expandable using the Start Orb on the top of the screen.

To notify the Actors of any changes in the Actor configuration in the web 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.

Note: All table views include a download button on the top right of the page. Clicking the download button will download the current table in CSV format to the local machine.

Content Management Menus

Clicking on any of the following sub-menus will display the content indicated. See the DIVA Operations Guide on the DIVA Technical Support site for more operational details.

Jobs

The Jobs view primarily displays the current jobs submitted to DIVA that are currently in, or pending, execution. jobs that have completed, canceled, or encountered warnings during execution are also displayed. This feature only applies when the web app is connected to DIVA. Completed or canceled jobs before the connection are not displayed. While connected, the number of pending, executing, completed, and canceled jobs displayed is a maximum of 300 of the most recent jobs. Clicking a job produces an informational screen pertaining to the selected job.

Job history

The Job History page displays completed jobs. Users can retry a job from this page, and in doing so the original job form will be presented pre-populated with the original parameters of the job. It is possible change these values before re-submitting the job.

Catalog Browsing

The Catalog Browsing page displays the objects in the system. Users can perform the following actions on an object or objects;

- Copy
- Copy To New
- Stage

- Restore Object
- Partial Restore
- Multiple Restore
- Delete Object
- View Object Properties

Archive Submission

The Archive Submission page less cramped than the original form, and fits the required controls. It allows the user to submit objects for an Archive Job.

SPM Actions

The SPM Actions page displays the SPM Actions in the system. The following actions can be performed on the system;

- Reschedule Action
- Mark Action Completed

Resources Management Menu

Clicking on any of the following sub-menus will display the content indicated. See the DIVA Operations Guide on the DIVA Technical Support site for more operational details.

Unmanaged Storage

The Unmanaged Storage view provides information about the sources and destinations identified in the DIVA system. This view displays the Source or Destination Server Name, Product System, Type, Address, and First Utilization Date. Clicking one of the entries displays the Entry Details dialog box.

Actors

The Actors view provides an indication of the status of each Actor defined in the web app and any current jobs. This view can be displayed by selecting Actors on the Dashboard, or by navigating to Resources Management > Actors.

Selecting one of the Actors will display the currently running job on that Actor in the window below it. If the Manager cannot establish a connection to an Actor, it is displayed as Offline. Click an Actor Panel, or access the menu from the vertical dots to display that Actor's Properties page.

The top right of the Actors page includes a button to change the view between the table view and item view. Use the button to switch views back and forth. The last view a user was using is retained and will be the initial view when the user logs out and back in to the system.

Media

This view displays information for each of the Tape Groups and Disk Arrays identified in the DIVA system. You filter the search using the lists at the top of the screen.

Disks

The Disks view displays the online status and capacities of disks configured in DIVA. The status of a disk can be set through the Disks page of the web app. If DIVA has automatically set a disk to Out of Order, the cause of the error must be investigated before setting the disk back to Working Well. When DIVA encounters an I/O error with the disk, it is set to Out of Order automatically by DIVA.

Cloud Disks

The Cloud Disks view displays the online status and capacities of disks configured in DIVA. The status of a cloud disk can be set through the Cloud Disks page of the web app. If DIVA has automatically set a disk to Out of Order, the cause of the error must be investigated before setting the disk back to Working Well. If DIVA encounters an I/O error with the disk, it is set to Out of Order automatically by DIVA.

Note: A disk linked to an OCI storage account will always report a Consumed Size of 0.

Libraries

The Libraries page displays all libraries in DIVA. Clicking the three dots on the top right of any listed library enables viewing the Library Properties page, giving additional information on number of tapes, capacity, objects, and so on.

Drives

The Drives page displays the drives in the system and a filtered table showing the Actor connections to the currently selected drive.

The top right of the Drives page includes a button to change the view between the table view and item view. Use the button to switch views back and forth. The last view a user was using is retained and will be the initial view when the user logs out and back in to the system.

Tapes

Users perform most object-based operations (Archive, Restore, Copy, Delete, and so on) from the Resource Management > Tapes page.

The Tapes view provides flexible search criteria (located at the top of the screen) to execute DIVA Database queries about the tapes managed by DIVA.

Troubleshooting Menu

Clicking on any of the following sub-menus will display the content indicated and aid in troubleshooting system issues. See the DIVA Operations Guide on the DIVA Technical Support site for more operational details.

Logged Events

The Events view is typically used with the Jobs view for troubleshooting purposes. You can filter the displayed results using the filters at the top of the screen including Dates (Start and End dates and times), Severity (Information, Warnings, Errors, and Critical), Job ID, and Description.

When a particular job fails, you can export the log of that job to a text file and send it to Telestream Support (when requested). This information can also be collected directly by the Telestream Support Engineer using the Customer Information Collection Tool.

When the query is run for the failed job's Job ID (usually retrieved from the Jobs view), it shows the same events of that job's event log. You can save this file as a text file by selecting Export.

DIVA stores a maximum of one million events in its database. When the number of logged events exceeds this value, DIVA will begin overwriting the existing events beginning with the oldest entry.

Drive Alert Logs

This view lists errors reported by tape drives. This information is vendor-specific and may vary depending on the make and model. You can filter your search using the filters at the top of the screen. For example, you can look for errors related to a particular tape.

Library Alert Logs

This view lists errors reported by direct-attached SCSI protocol Managed Storage. This information is vendor-specific and may vary depending on the library make and model. You can filter your search using the filters at the top of the screen.

Migrations Menu

Clicking on any of the following submenus will display the content indicated. See the DIVA Operations Guide on the DIVA Technical Support site for more operational details.

Start New Migration Task

This screen allows users to begin a copy or move migration job. Users follow through the wizard pages to complete the task. The wizard automatically moves to the next page when the user makes a selection on each page and continues to completion.

Migration Tasks

Navigating to the Migrations > Migration Tasks page allows the user to view the status of the current migration task, and previously executed migrations. Clicking on the ID presents more detailed information on the selected task. Clicking the Refresh icon on the top right will refresh the page and update the current status of the displayed tasks.

Configuration Menu

The Configuration Menu is where administrators configure DIVA. The following submenus enable configuring different parts of the system.

User Management

This page allows System Administrators to create, edit, or delete users. Options include assigning a Username, Password, and the Profile Role for the account.

System Settings

This page and its sub-pages allow System Administrators to configure the following:

- Productions Systems (add, edit, or delete)
  - Production systems allow dedication of a particular Actor for specific destinations.
- Sites (add, edit, or delete)
  - Additional sites are optional and may be considered by the Manager for optimal resource.
- Actors (add, edit, clone, delete, or move)
  - Define Actor host definitions and logical functions.
- Unmanaged Storage (add, edit, clone, delete, or move)
  - Define where Core archives from and restores to.
- Transcoders (add, edit, or delete)
  - Define Core transcoders and analyzers.
- Actor-Proxy Connections (add, edit, or delete)
  - Define links between Actors and Proxy Actors.

Library Storage

This page and its sub-pages allow System Administrators to configure the following:

- Robots (add, edit, delete, synchronize DB, synchronize library associations, synchronize drive & tapes list, synchronize drives list)
  - Define Robots that are to be accessible in Core.
  - Libraries (edit or delete)

- Drives (add, edit, or delete)
  - Define Drives in your tape managed storage for use with Core and its Actors.
  - Actor-Drive Connections (add or delete)
- Connected Actors (viewing only)
  - View all Actor connections to Drives.
- Media Compatibilities and Synchronize Media Compatibilities (edit or delete)
  - View drive media compatibilities.
- Drive Types and Synchronize Drive Types (delete)
  - View drive types supported by the libraries.
- Tape Types and Synchronize Tape Types (edit or delete)
  - View tape types supported by the libraries.
- Tape Groups (add, edit, clone, or delete)
  - Define Tape Groups that segment material within the tape library, or associate content with Tape Groups.
- Empty Ejected Tapes (insert or eject)
  - Manage empty ejected tapes.

Cloud Storage

This page and its sub-pages allow System Administrators to configure the following:

- Cloud Accounts (add, edit, clone, or delete)
  - Define storage accounts to connect to cloud storage.
- Cloud Buckets (add, edit, clone, or delete)
  - Define Cloud Buckets that should be accessible in DIVA.
  - Connected Actors (add, edit, clone, or delete)
- Connected Actors (viewing only)
  - View all Actor connections to Cloud Buckets.

Disk Storage

This page and its sub-pages allow System Administrators to configure the following:

- Arrays (add, edit, clone, or delete)
  - Define Disk Arrays that should be accessible in DIVA.
- Disks in DIVAgrid (add, edit, clone, or delete)
  - Define Disks that belong to the Disk Arrays accessible in DIVA.
- Connected Actors (viewing only)
  - View all Actor connections to Disk Arrays.

General Settings

This page and its sub-pages allow System Administrators to configure the following:

- Checksums
  - Configure creation of checksums by Manager.
- Media
  - Configure automated clones of media.
- Objects
  - Configure how the manager handles Complex Objects.
- SMTP Notifications
  - Configure email notifications.
- Security
  - Configure the use of encryption keys.

Local Settings

This page and its sub-pages allow System Administrators to configure the following:

- Date and Time
  - Customize how dates and times are displayed.
- Customization
  - Customize general settings for the website.
- Reset Local Config
  - Reset custom configuration for individual pages

License

- Active License
  - Displays the active license information.
- License History
  - Displays historical license information for previously imported licenses.
- Import License
  - Enables importing a new DIVA license.
- Version
  - Displays information on the current active license including DIVA Version, DIVA Web Version, DIVA API Version and Client ID.

Component Configuration

This chapter describes DIVA component configuration.

Topics

Configuration Overview

There are many interrelated components in a DIVA System. The following figure shows the basic configuration workflow.

The configuration of DIVA is hierarchical and top-level parameters such as Networks, Sites, Arrays, and Disks need to be configured before configuring other components such as Actors.

If you intend to modify an existing DIVA system, you must always start by backing up the existing DIVA installation, configuration files, and especially the DIVA and Metadata Databases.

Contact Technical Support before making any modifications to your DIVA platform if you are unsure about any steps in the procedures, or require clarification.

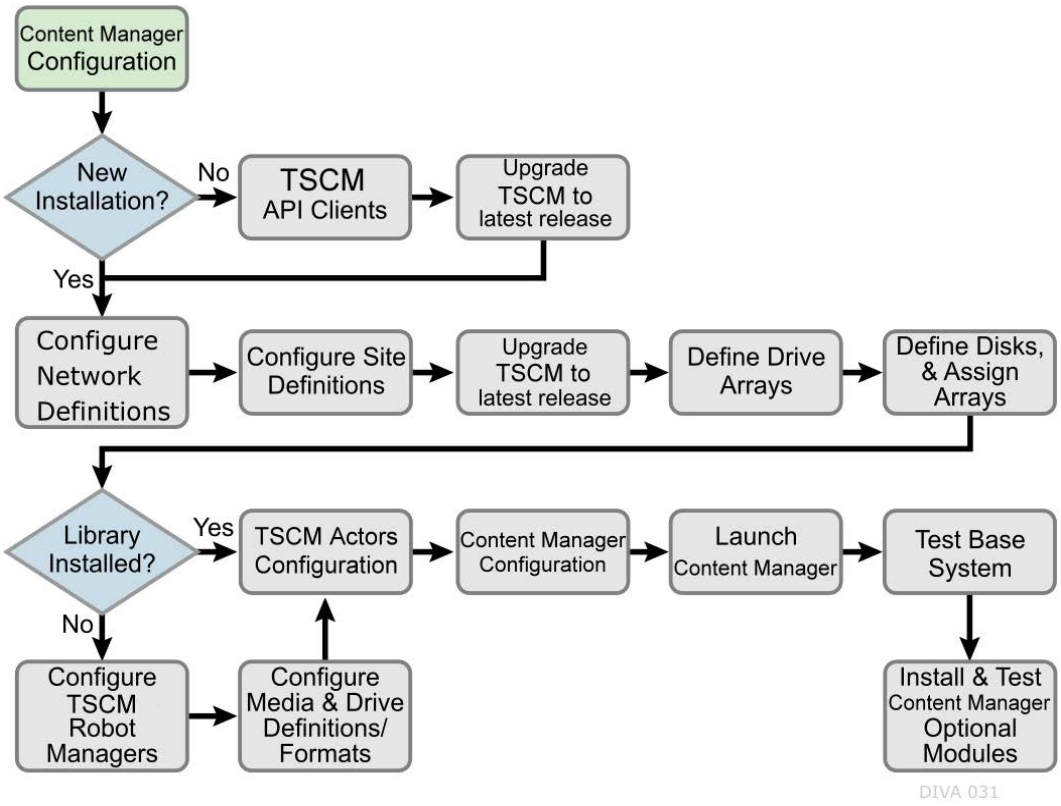

Module Configuration Files

Each DIVA 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 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 centralizes them to a dedicated conf subfolder under the DIVA 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 will have the new releases of the .conf files appended with a .ini extension. For example, the new release of the 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 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 codependencies between some applications in the DIVA 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 and runs as a Windows 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.

The following figure is the workflow for installing a Manager:

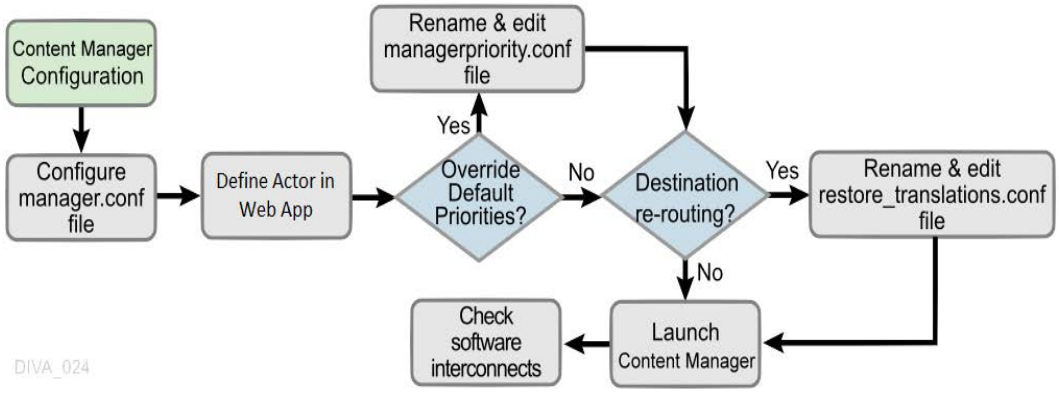

Configuring the Local Manager

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 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 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 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 software components and DIVA 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 Manager.	
DIVAMANAGER_NAME	Name	The name this Manager instance uses to identify itself to other 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 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 Services and the DIVA API.	8000

Database Settings

These parameters define the location and instance of the DIVA Database. Except for the DIVAMANAGER_TNSNAME parameter, you must define all settings in this section for the Manager to launch successfully.

The following table describes the Database settings in the manager.conf file:

Parameter	Parameter Type	Description	Default
DIVAMANAGER_TNSNAME	Name	The TNS Name of the DIVA Schema within the Postgres database. DIVA ignores this setting if the DIVAMANAGER_DBHOST and DIVAMANAGER_DBPORT settings are defined. This feature requires Postgres installed on the host running the Manager. If this setting is defined, the location of the Postgres driver 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. Example: wrapper.java.library.path=.;C:\app\postgres\product\11.1.0\BIN	
DIVAMANAGER_DBHOST	IP Address or Host Name	This specifies the Host Name or IP Address of the computer containing the DIVA Database. If using a host name, this must be present in the hosts file on the computer where the Manager is installed.	
DIVAMANAGER_DBPORT	TCP Port Number	The Postgres Listener port configured during the DIVA Database installation.	5432
DIVAMANAGER_DBSID	Name	The DIVA Database SID (Instance System Identifier) in Postgres where Manager connects.	

Parameter	Parameter Type	Description	Default
DIVAMANAGER_DBUSER	Name	The user name the Manager uses to connect to the DIVA Database. This is typically diva (case sensitive).	diva
DIVAMANAGER_DBSERVICENAME	Name	Postgres 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	Postgres 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's default behavior for task execution, resource allocation, and the number of connections it will accept from DIVA applications and DIVA API clients. These parameters are normally adjusted or fine-tuned after completing the initial installation of DIVA.

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. 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 jobs	The maximum number of jobs that can be queued for processing by DIVAMANAGER_MAX_CONCURRENT_REQUESTS processors of the Job Scheduler.	500
DIVAMANAGER_MAX_CONNECTIONS	Number of Connections	Specifies the maximum number of simultaneous client connections the Manager will accept. This includes Actors, web apps, API connections, and support tools.	200
DIVAMANAGER_MAX_SIMULTANEOUS_REQUESTS	Number of Jobs	The maximum number of jobs processed by the Manager. When this limit is reached, any further jobs 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 Jobs	Maximum number of inactive jobs that cannot find resources examined by the Job Scheduler each time it is activated.	0
DIVAMANAGER_TYPICAL_VIRTUALOBJECT_SIZE	Percentage	During operation an Actor retrieves the file size of an Object before an archive transfer. This value determines the best location on the tape for the file. Some servers do not indicate the file size of an Object before a Direct Archive. Therefore, DIVA will use this value as an estimate for tape selection. You must define this setting so that most Objects to be archived in the DIVA system are below this size.	10 (percent)
DIVAMANAGER_MAX_CONCURRENT_REQUESTS	Number of Jobs	The maximum number of concurrent jobs executed by the Manager. The maximum tested value for this setting is 16.	8

Parameter	Parameter Type	Description	Default
DIVAMANAGER_MAX_SPAN_SEGMENTS	Number	DIVA 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 Manager.	1
DIVAMANAGER_MIN_DB_CONNECTION_LIMIT	Number of Connections	The minimum number of database connections available to the Manager.	1
DIVAMANAGER_MAX_DB_CONNECTION_LIMIT	Number of Connections	The maximum number of database connections available to the 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 web 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 jobs still running or to terminate them after the Automatic Tape Repack period. If this is set to true then repack jobs 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 job.	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 periodically examines all jobs in its job queue and increments the job priority. This prevents a condition where low priority jobs might be continually superseded by higher priority jobs. 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_RS_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 Job 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 Job 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	The interval in milliseconds between Manager checks to see if the connections to its clients and services are still active (Actors, SPMs, web apps, etc.).	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 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 Manager is running.	Exported
DIVAMANAGER_MAX_RESTORE_SERVERS	Number between 2 and 200	The maximum number of servers allowed in an N-Restore job by an 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 job.	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 job.	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 job.	1000000
USE_IMPROVED_BEST_WORST_FIT_ALGORITHM	true or false	When a file was archived to tape in earlier DIVA 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 can be defined by their location. If you set this parameter to true, the Manager first tries to perform the job from the sites identified as MAIN. If unsuccessful, it retries the job with resources from all other sites. If you set this parameter to false, DIVA ignores site identification and all site resources are considered equally.	false
DIVAMANAGER_CACHE_QOS_USE_DISK	true or false	In earlier DIVA releases, a Restore job 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 to use the disk instance regardless of the QOS method specified.	true
DIVAMANAGER_PRIORITY_TIER	Number between 0 and 100	DIVA bases the execution of jobs in its job queue by the job priority number. However, there are instances where a job in the queue with lower priority uses a tape that is already mounted. Giving this job 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. If this setting is enabled, DIVA examines the job queue for lower priority jobs involving a tape that is already mounted in a drive and adds the number specified here to the job priority. For example, if the job priority is 25, and the Priority Tier value is 50, the total job priority is 75. **Note:** This feature applies only to Restore and Copy Jobs that read from tape. Archive and Copy jobs that write to tape are not supported by this feature.	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 job. 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 ignores this setting if the DIVAMANAGER_ETC_FEATURE is disabled.	50
DIVAMANAGER_OVERWRITE_POLICY	Number between 0 and 2	This value determines how DIVA handles files that already exist on a Destination Server when executing a Restore, Partial File Restore, or N-Restore job 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 job 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_TOD	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 job 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 = Math.ceil(bytes/1024)	3
DIVAMANAGER_MAX_EXCLUDED_INSTANCES	Number	Maximum number of instances excluded from a job that are logged as an event.	3

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: • 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	DIVA 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. 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 job 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 job will terminate.	0 (segments)
DIVAMANAGER_MAX_DB_CONNECTION_ATTEMPTS	Number	The maximum number of allowable attempts to connect to the database.	10000
DIVAMANAGER_MIN_DB_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 job. 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 jobs when the Destination Server is tape. Copy jobs 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 job 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 jobs.	NEARLINE_AND_DIRECT
NTH_PROGRESS_MESSAGE	Number	The number of progress messages sent to client apps. Every Nth progress message will be sent. The N=100 progress message is always sent.	5—implies send every fifth progress message to all GUIs.

Parameter	Parameter Type	Description	Default
DIVAMANAGER_TIME_TO_WAIT_FOR_GRACEFUL_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 job if it contains an empty file or folder.	false
TAPE_FULL_ON_SPAN_REJECTED	Boolean	If true, the Manager will mark a tape full when a span occurs but spanning is disabled.	false
DIVAMANAGER_RETRY_ON_SPAN_REJECTED_ALGORITHM	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: • 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 Job Priorities

Each job submitted to the Manager is placed in the Manager transfer queue. Job priorities enable DIVA to differentiate between important jobs, such as Restore jobs, over less important events. For example, tape repacks, and so on.

The job priority is a number from zero to one hundred with zero being the lowest priority and one hundred being the highest. The job priority is typically specified when you submit the job (either from the web app or the DIVA Client API). You can also alter the priority after you submit the job using the Change Priority command.

The default job priority for each job type is preset within DIVA. 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 job 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 job priority, the Manager will (by default) periodically increment the priority of every job already the job queue. This prevents a condition

where a low job priority can be continually overridden by higher priority jobs 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 to automatically override the original Destination Server specified in a Restore, Partial File Restore, or N-restore job 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;Category_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 job for this rule to apply.

Category_1

If the Object Collection of the job also matches the Destination Server will be re-routed.

TranslatedDestination_1

This is the new Destination Server for the Restore job.

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 job.

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

Manager control and management functions are performed from a command prompt on Windows platforms using the manager.bat batch file. The executable is located in the %DIVA_HOME%\Program\Manager\bin folder.

Installing and Removing the Manager Service in Windows

You must first install the Manager as a system service on new systems. You can accomplished this using the install (or -i) and uninstall (or -u) command line switches as follows:

manager install

This (or manager -i) installs the Manager service set by the SERVICE_NAME parameter defined in manager.conf. If this parameter is undefined, the service is installed as Manager.

manager uninstall

This (or manager -u) removes the Manager service set by the SERVICE_NAME parameter defined in manager.conf.

In the Windows Services applet, confirm that the 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 m,anager.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]
```

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 jobs, then it stops the 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 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) change them dynamically from the web app. See the DIVA Operations Guide on the DIVA Technical Support site for detailed information.

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 Manager has been successfully configured and launched you must check that the Manager can successfully be connected to by other DIVA clients (for example, the web 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 web app from a remote client (that is, not on the same host computer as the Manager).
2. Click the Menu Orb on the top left of the web 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 web 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 web 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 Manager can connected to each configured Core Robot Manager. This test assumes the following:

- All Core Robot Manager 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 Jobs view.
5. Double-click the Eject Tape job entry to check if an error was encountered during job 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 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 the 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, Postgres, 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.

Actor Configuration

This section describes Actor configuration and operations.

Configuration Overview

The Actor runs on Windows. Windows Actors no longer start automatically with Windows; the Actor runs as a standalone server application. The Manager connects to each Actor as a client application.

The Actor is installed in the %DIVA_HOME%\Program\Actor\bin\ folder. The Actor's configuration files are located separately in the %DIVA_HOME%\Program\conf\Actor\ folder. At the system level, the location and capabilities of each Actor are defined in the web app.

The Actor configuration parameters are located the web 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 an Actor:

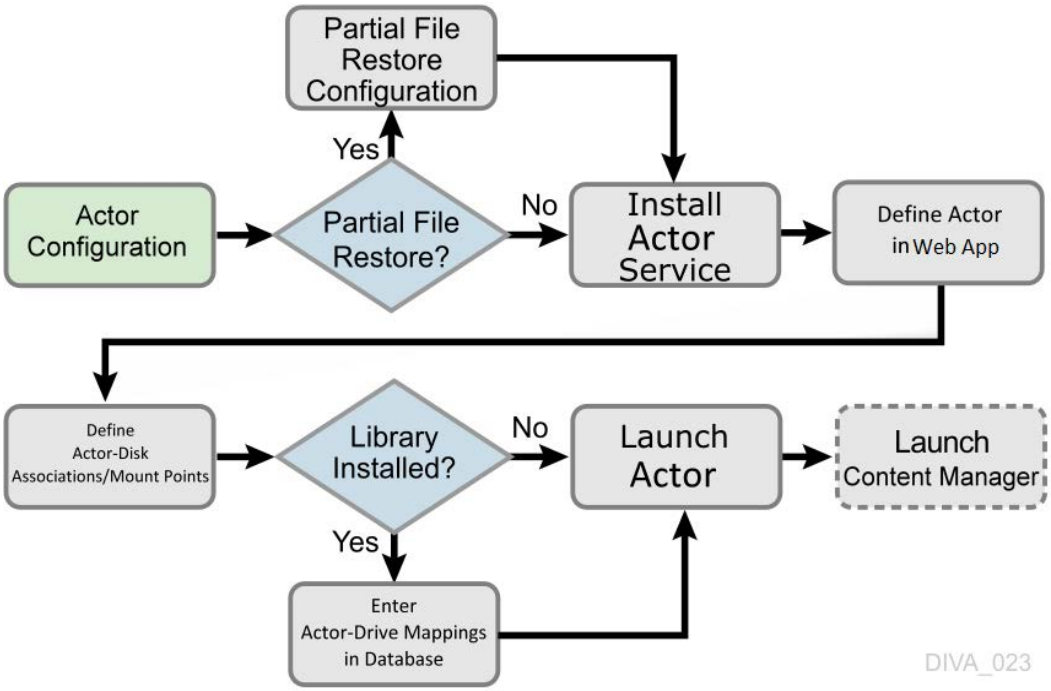

Configuring the Local Actor (actor.conf)

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	TCP Port Number for the Actor to listen on for incoming jobs. 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 Partial File Restore

The Partial File Restore parameters are located on the Partial Restore Settings page in the web 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 DIVA 9.0 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 job option available as indicated in the table that can be used when creating the job.

Parameter	Value or Type	Job 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	Specifies how the Actor handles the first frame of a QuickTime clip: • 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	Job Option	Description	Default
AVI Ignore Start Timecode	N (disabled) Y (enabled)	-PfrAviIgnoreStartTimecode	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)	-PfrEvsMxfIgnStartTimecode	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	This setting specifies how the time code SOM reference is to be derived for a GXF Partial File Restore job. The options are defined by the following values: • 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	Job 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 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	Job Option	Description	Default
MXF Partial Restore Dictionary File	Path and File Name	-PfrMxfPrDictFile	This parameter must point to the name and location of the MXF dictionary file. The dictionary is normally distributed with the Actor installation in the %DIVA_HOME%\Program\Actor\bin folder. The default dictionary file name is mxf_file.bin. Set this parameter to %DIVA_HOME%\Program\Actor\bin\mxf_file.bin. Where %DIVA_HOME% is the root path of your DIVA installation for the Actor (typically C:\Diva).	
MXF Timecode From Source Package	N (disabled) Y (enabled)	-PfrMxfTimecodeFrmSrcPkg	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.	N
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

Parameter	Value or Type	Job Option	Description	Default
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.	
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 TCOUT (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

Parameter	Value or Type	Job Option	Description	Default
MXF Use BMX Library (instead of MOG SDK)	N (disabled) Y (enabled)	-PfrMxfUseBMXLibrary	The use of either MOG SDK or BMX can be selected from the web app under Configuration > Actor Settings, by setting the Use BMX Library parameter to Y.	N
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	Job 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)	-PfrMxfUpdateTctrackOrgin	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	Job Option	Description	Default
MXF Tolerance on TCOUT	Integer between 0 and 250.	-PfrMxfTcoutTolerance	This parameter can be set to indicate a tolerance on the TCOUT supplied to a Partial File Restore job. This tolerance value is 0 by default, but it you can set it to a specific number of frames. If the supplied TCOUT is beyond the end of the clip, but not too far out (within the tolerance), DIVA will perform the Partial File Restore until the end of the clip instead of reporting and invalid TCOUT.	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 TCOUT are valid.	Y

Parameter	Value or Type	Job 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)	-PfrSeaIgnoreStartTimecode	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 TCOUT 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)	-PfrTsIgnoreStartTimecode	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 TCOUT as if it starts from 00:00:00:00.	N
MPEG2 Program Stream Ignore Start Timecode	N (disabled) Y (enabled)	-PfrPSIgnoreStartTimecode	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 TCOUT as if it starts from 00:00:00:00.	N

Defining and Declaring Actors

Each Actor must be declared in the DIVA Database. You declare the Actors in the Actors area in the web 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 jobs to and from drives that this Actor can perform. You can use this parameter to distribute jobs and bandwidth among all Actors.

Max Server Operations

This is the maximum number of simultaneous jobs to and from servers from the Servers configuration that this Actor can perform. You can use this parameter to distribute jobs 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 jobs and bandwidth among all Actors.

Max Stage Operations

This is the maximum number of staging job that an Actor is allowed to run at the same time.

Max Bridge Operations

This is the maximum number of concurrent jobs 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 Object during a transfer.

Copy To Tape Group

This parameter defines whether this Actor can be used for Copy To Tape Group jobs. You can use this option to isolate specific Actors involved in critical operations from mass Copy To Tape Group jobs, such as those from the DIVA SPM option.

Associative Copy

This parameter defines whether this Actor can be used for Associative Copy jobs.

Repack

This parameter defines whether this Actor can be used for tape repack jobs. 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 jobs 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 jobs that involve deleting DIVA Objects from a disk. You can use this option to isolate an Actor from mass deletion jobs (for example, jobs issued from the SPM option).

Direct Archive

This parameter defines whether this Actor can be used for direct Archive jobs.

Cache Archive

This parameter defines whether this Actor can be used for cache Archive jobs. You must disable this option if this Actor has no local cache storage for the temporary storage of the DIVA 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 web 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.

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 preforming 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 DIVA FTP 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 DIVA FTP 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 jobs.

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 Manager that this Actor is a Proxy by adding an Data-Proxy Actor Connection.

After configuration, Manager 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 job 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 job. The user will have to manually re-enable the Source Tape for automatic cloning by setting the corresponding Tape State in the web app.

If a write error occurs during a clone job, 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 job will select a new tape from the Clone Tape Group.

See the DIVA Operations Guide on the DIVA Technical Support site for details on tape selection, manual cloning, and automatic cloning processes.

Logging Actor Activity

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 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 job 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 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).

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 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 Actor—Host_Name.

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 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 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 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
```

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 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 Actors no longer start automatically with Windows. The Actor Services are managed through the Windows Services applet, from a Windows command line.

The Actor Service can be located 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 using the following command sequence:

```
net stop "Actor"
net start "Actor"
```

If a SERVICE_NAME is specified in the actor.conf file (for multiple Actors on a single computer), then an Actor can be restarted from a Windows command line using the following command sequence:

```
net stop "Actor -SERVICE_NAME"
net start "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.

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:

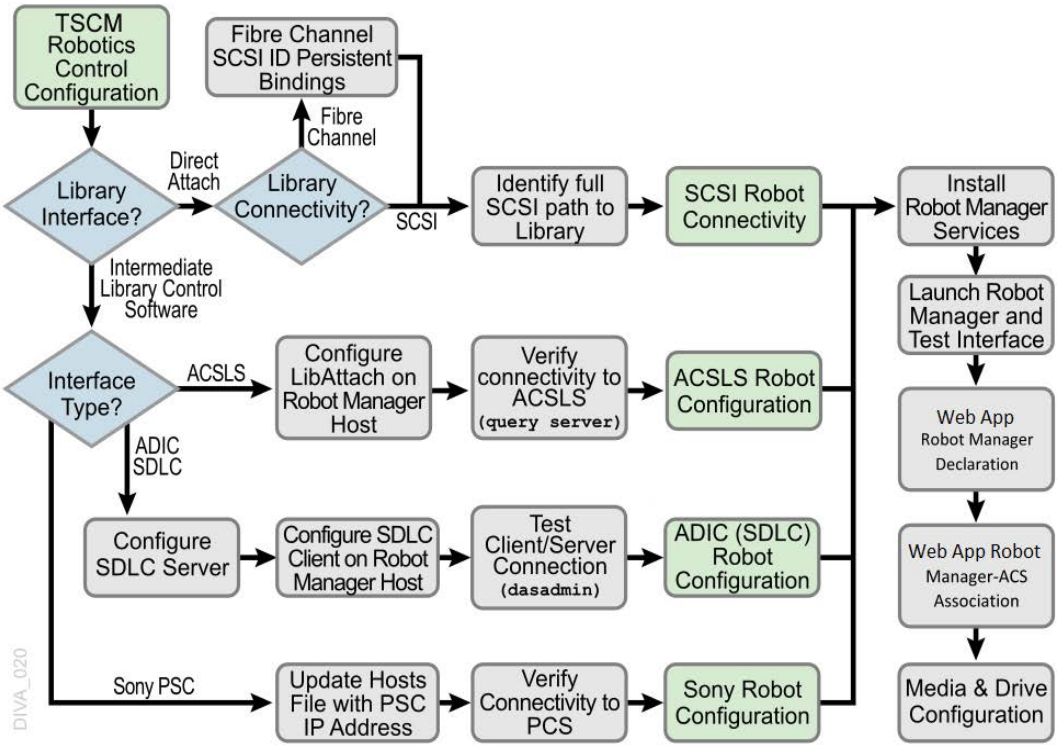

SCSI Connected Managed Storage

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 jobs. 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 job 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_SCSI_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).

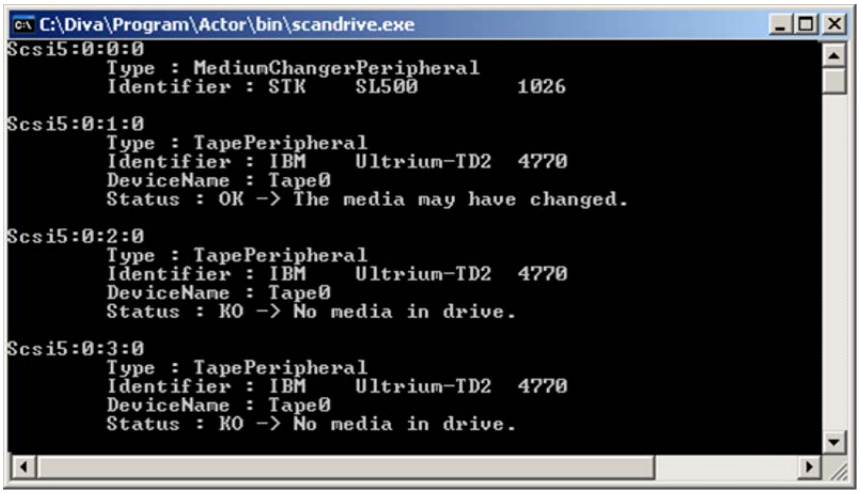

The utility reports the SCSI Device ID of the library in the format ScsiP:B:T:L (see the previous figure), where P is the port number, B is the bus number, T is the target number, and L is the Logical Unit Number.

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_SCSI_DEVICE_LSM(n) settings in the robotmanager.conf file.

Sony ODA Drives

DIVA supports the Sony new generation of ODA drives; the ODS-280F (Fiber Channel) and ODS-280U (USB). DIVA 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.

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 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 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. 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
```

Web App Settings and Information

You must configure the following settings in the DIVA web 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.

Web App Settings and Information

The Optical Drives and Discs are displayed in the DIVA web 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 provided driver (SCSI_Robot.dll) 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 is *DivaRbt-<SERVICE_NAME>*. If this variable is not set, the default Service Name is *DivaRbt*.	Uncommented
RM_PORT	TCP port number	The TCP port that the Core Robot Manager listens on for incoming jobs. 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 web 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 manages this type of library more effectively when both frames are specified. DIVA 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 DIVA 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.

Configuring ACSLS Attached Managed Storage

DIVA 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 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=ACSLS_Robot.dll		Uncomment only this line	Commented
RM_PORT	TCP Port Number	The TCP Port the Robot Manager will listen on for incoming jobs. 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	ACSLS configurations ignore this value because the ACS number is supplied from ACSLS.	Ignored
SERVICE_NAME	Name	The display name of the Robot Manager service. This variable must be set if multiple Robot Managers are installed on the same server. If this variable is used, the Service Name is *DivaRbt-<SERVICE_NAME>*. The Service Name is *DivaRbt* if 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 jobs. 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	This parameter must not be modified during production. The database may need to be patched if it is changed during production. 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. 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.	1
RM_SONY_MEDIA_TYPE_TRIMMING_LEFT	Number	This parameter must not be modified during production. The database may need to be patched if it is changed during production. 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.	0
RM_SONY_PSCSERVERNAME	IP Address or Host Name	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.	

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 retires	The 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

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 supports Alto like a tape library. A disk of the library is seen as a tape in DIVA with its UUID as label.

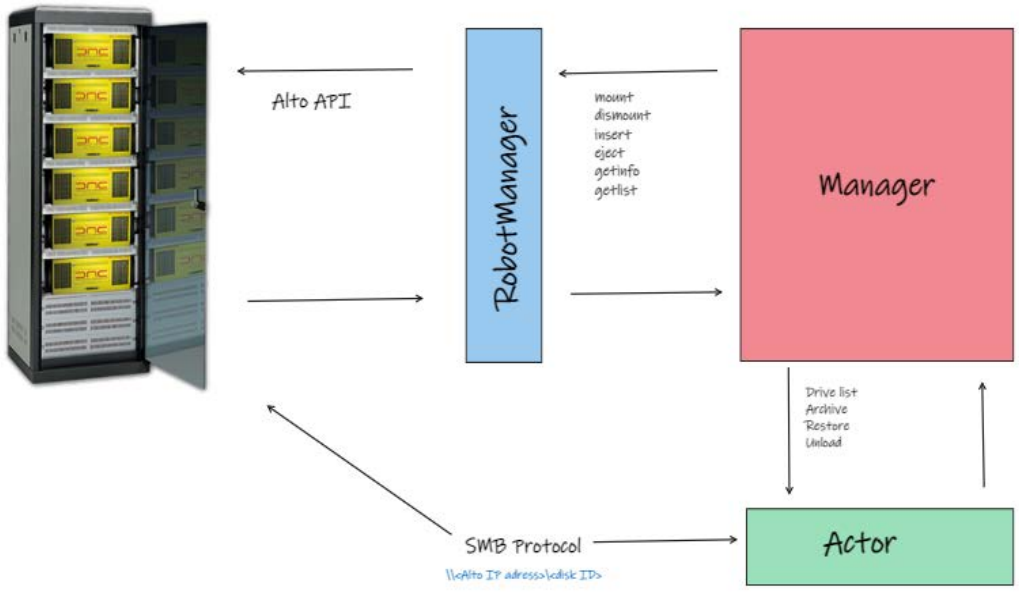

Configuration

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. 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 because tapes are mounted into drives. This number also defines the number of disks that DIVA can mount at the same time.

Actor Configuration

There is a new parameter in the configuration of Actor:

AltoVirtualDrives: This parameter contains a pattern string so that 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 Actor services must be set to log on under the same account (that is, Administrator).

Web App Support

Alto Virtual Drives value can be set for a Actor from the web app's Actors Configuration page. It is on the Advanced Settings tab, in the Disk Operations Section.

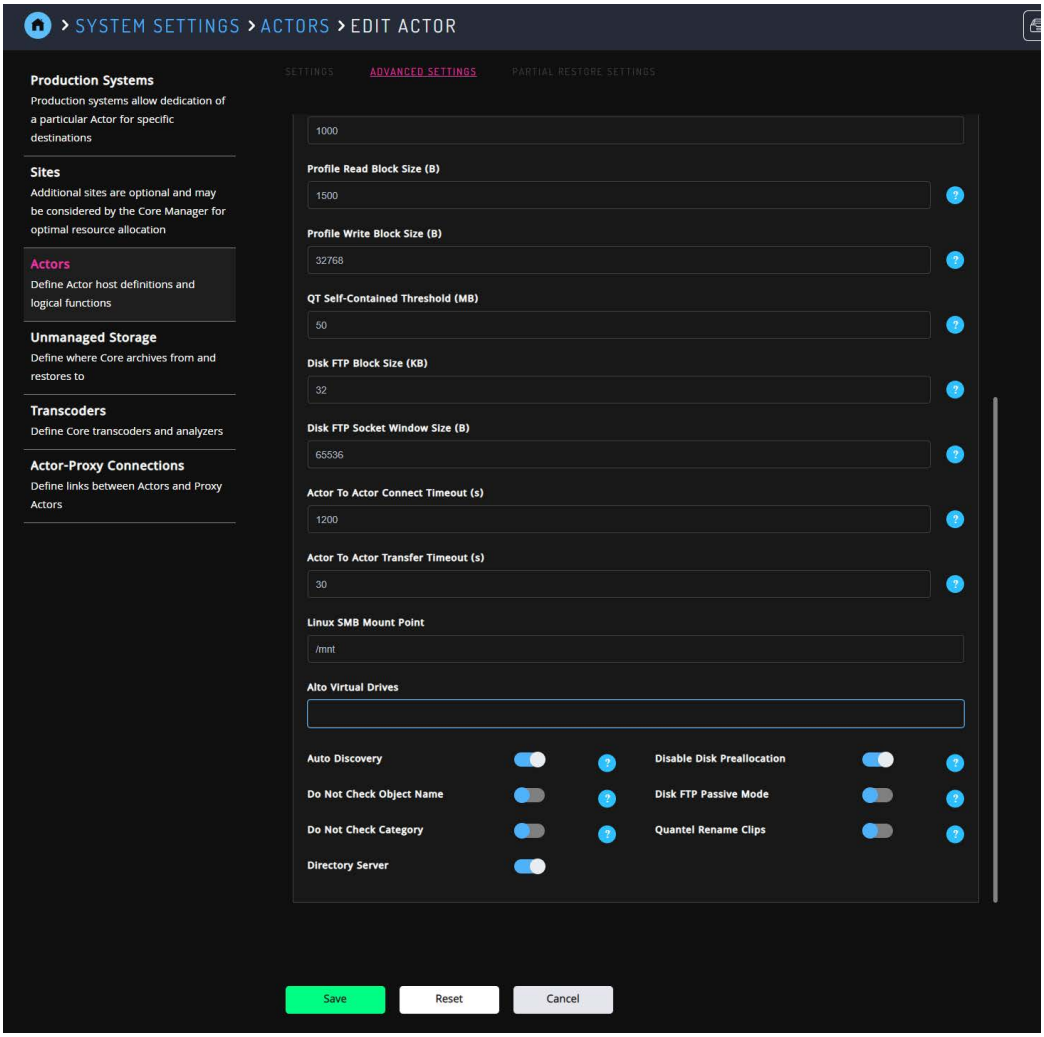

Configuring Simulated Managed Storage (for DIVA Simulators)

Simulated robots are available on Windows platforms. The settings are shown here for reference only. Refer to the DIVA Simulator Operations Guide (available to OPN partners only) for more information on installing and configuring a DIVA 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 jobs. 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 Simulator parameters:

Parameter	Parameter Type	Description	Default
RM_SIMU_BASEDIR	Directory Path	The DIVA 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.

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, 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.

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 system. You must not send commands to a Robot Manager using either of these utilities under any circumstances when the 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 Manager in the Robot Managers area of the Robots page in the web 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 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 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 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 web 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.

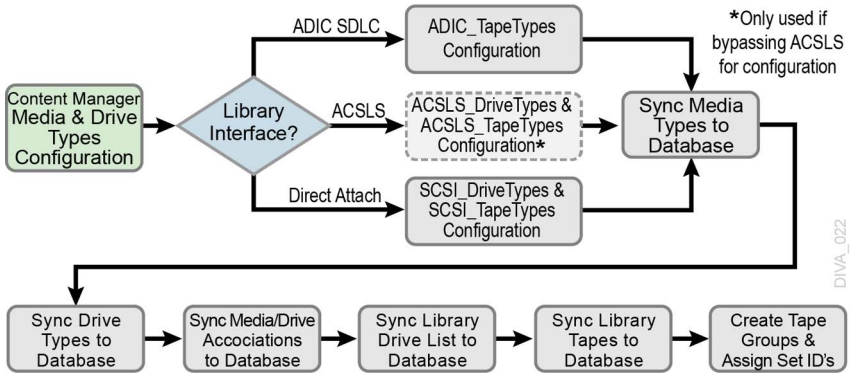

Tape Drives and their associated media types that are installed in a particular library are initially configured in the DIVA Database using static configuration files. The files are

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 web 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 DIVA 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 web 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 web 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 web 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 web 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 web 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 web app.

Configuring Tape Groups

Configuring tape groups is a different topic in DIVA. However, creating and configuring tape groups must be completed before configuring the Storage Plans. The number of Storage Plans required is usually the same as the number of tape groups configured. Tape groups enable DIVA to physically separate archived content into different tapes, and typically creating one Storage Plan per tape group is necessary. All Storage Plans are typically setup to be the same, except that the copy goes to different a tape group. However, the more tape groups DIVA uses, the less efficiently content will be stored across tapes.

Note: Refer to the DIVA Storage Policy Manager documentation for details on SPM configuration and use.

If complex objects are going to be used in the system, tape groups must be set up containing tapes configured for AXF format. Complex objects are not compatible with the Legacy formatted tapes or disks. Non-complex objects can be stored using either Legacy or AXF format.

The following example illustrates how creating too many tape groups causes fragmentation of objects across the tape group. Using fewer tape groups results in fewer Storage Plans to setup, less fragmentation across tapes, and is easier to maintain.

Having fewer tape groups will resolve the issue in the following example, and avoid fragmentation across the tape groups.

Example Configuration:

- 10 Tape Drives
- 30 Tape Groups with SetID=10
- 300 Total Tapes assigned to SetID=10

Results:

- Each tape group (in the worst case) will use at least ten tapes, and store files on each tape when an object is archived.
- After objects are archived to all 30 tape groups, all 60 tapes (total) are used.
- Over time, if any of the tape groups is heavily archived, and one of the tapes is 100% full, no more tapes are available.
- DIVA will run out of tapes, even though a lot of the tapes are still mostly empty (each containing only one object), but cannot be used because it was assigned to a different tape group.

Valid Reasons for Creating Multiple Tape Groups

Generally, multiple tape groups should only be created for a good reason. Multiple tape groups will cause tape group fragmentation, resulting in some tapes not being filled, and storage space being wasted.

The following are several valid reasons for creating multiple tape groups:

- Long and short form materials should be stored on different tape groups. If small objects are mixed with larger objects on the same tape, access to the smaller object will be delayed for an extended period until the larger object restore is complete.
- Content that is deleted regularly from the archive should be stored in a different tape group than content that will never (or rarely) be deleted. Deleting from tape will cause tape fragmentation and the fragmented space cannot be used until the tape is repacked. If the two types of content are mixed together, deleting will cause more tapes to become fragmented, and repacking will be required more frequently and take longer.
- Online and backup copies of the same content should be on two different tape groups. Backup copies are meant to be removed from the tape library and stored in an Iron Mountain type of facility as a backup. However, if both copies are mixed in the same tape group, it is impossible to determine which tape contains the backup copy. The result is being unable to remove it from the library for offline storage.
- When requirements necessitate using tapes purchased by different departments and enforcing that each department uses only the tapes they purchase themselves, tape groups must be created for each department with a different set of tapes.
- Different storage formats must be assigned to different groups. For example, one group would be Legacy Format while another group is AXF Format. Complex

objects are not compatible with legacy format, and must be processed to and from AXF formatted medium. Non-complex objects are compatible with both Legacy and AXF formats.

Invalid Reasons for Creating Multiple Tape Groups

Having too many tape groups will cause DIVA to work less efficiently, result in tape group fragmentation, and are more difficult to configure and maintain. Unless one of the previously described valid reasons is necessary, the recommended best practice for creating tape groups is to use as few tape groups as possible. See *Configuring Tape Groups* for recommended best practices.

The following are several invalid reasons for creating multiple tape groups:

- To store different content in different tape groups because you think it is easier to manage. DIVA manages the tapes. When restoring an object on a tape, DIVA automatically knows which tape the object is on, and mounts that tape, or notifies the user to insert the necessary tape into the library. It will not require the user to figure out which tape is needed.
- Creating many tape groups because it is easier to search. DIVA performs searches using object metadata stored in the database, not tape groups. The tape group only makes sure content is physically separated, and does not assist in searching functions.

Creating multiple tape groups for cataloging. Users migrating from an analog tape environment tend to label what is recorded on each analog tape directly on the tape itself. A group of those tapes are then stored on different sections of a shelf. DIVA does not work efficiently this way, and this method should not be used.

Additional Functionality

This chapter describes DIVA additional functionality.

Topics

Checksum Support Configuration

You configure the Checksum Support functions through the web 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 web 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. 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 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 web 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 web 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. 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 web 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 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 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 job with -axf in the Job Options.

TEXT Genuine Checksum Mode

This mode enables DIVA 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 jobs generated by DIVA.
- 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.

Web App Settings

Use this procedure to configure AXF Genuine Checksum Mode in the web 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 job 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 Windows to a maximum of 4000 characters. Relative path names are limited to 256 characters on Windows systems.
- 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.

Web App Settings

Use the following procedure to configure TEXT Genuine Checksum Mode in the web 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 job 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 job 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 Job (File Path Root)	Archive Job (Files)
D:\	Data\Video\NewTitle	*
D:\Data	Video\NewTitle	*
D:\	Data\	Video\NewTitle*
D:\		Data\Video\NewTitle*

Server (Root Path)	Archive Job (File Path Root)	Archive Job (Files)
/	Data/Video/NewTitle	*
/Data	Video/NewTitle	*
/	Data/	Video/NewTitle/*
/		Data/Video/NewTitle/*

Transcoder Installation and Configuration

The following instructions are directed toward servers running the Windows Server 2016 SP1 operating system.

Upgrading from 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 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.

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 and Transcoders

The following instructions identify the configuration of DIVA 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.
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.

- 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.

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 web app. with the following settings:
   - Transcoder Type: *vantage*
   - Working directory: *M:\VantageCache*
   - Leave the remaining options at the default settings.
3. Open the web app.
4. Navigate to the Transcoders area on the System page.
5. Ensure that the DIVA 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 web app.
9. Navigate to the Transcoders area on the System page.
10. Add the transcoder to the web 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 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 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 to use the workflow. In this example, the workflow should look like the following figure.

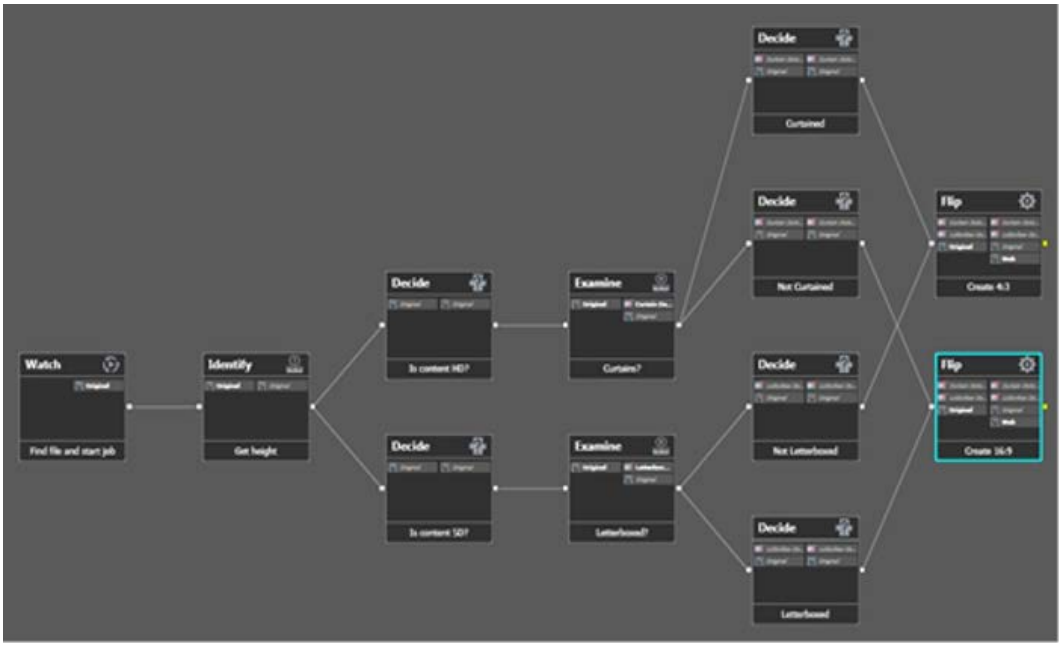

Configuring Transcoders

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 web 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
```

Object Auto-Indexing

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. For example:

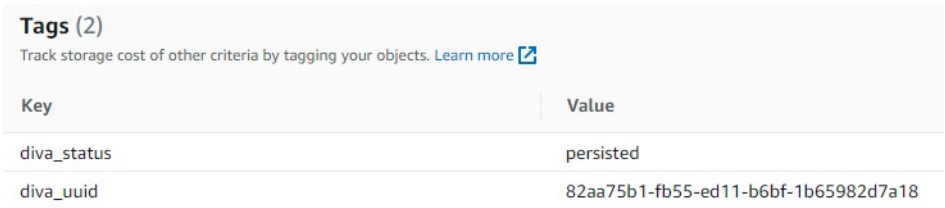

Here is how to clean this up:

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.

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	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). Example with these parameters set: • Whatever the configuration of Object Path Prefix • Bucket name is set to `MyBucket/MyVideoFiles` • SubPathFilter is set to `ABC` Auto-indexing will only consider the cloud objects starting from `MyVideoFiles/ABC` inside `MyBucket`.
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. 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.

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.

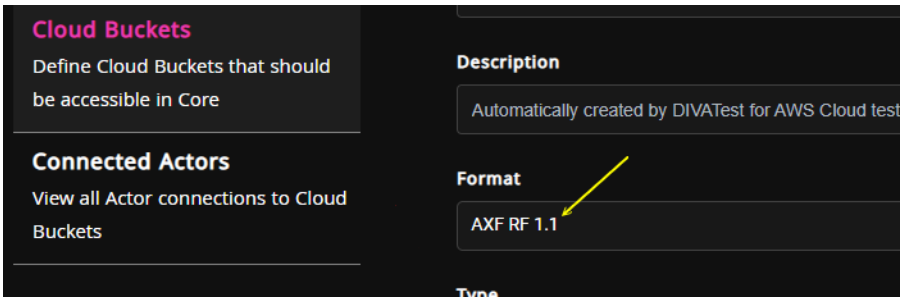

- Bucket Name

  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.

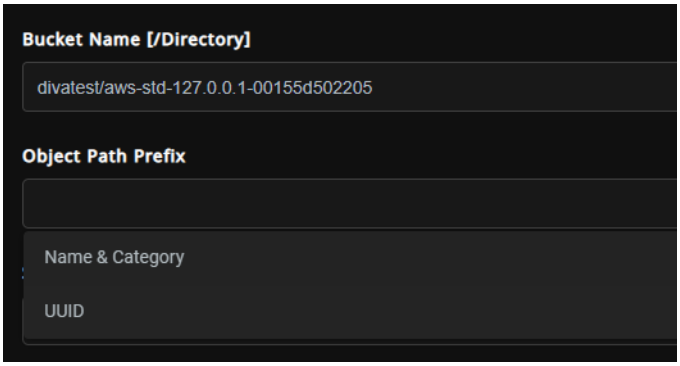

The following table describes how auto-indexing will set Objectname, Collection, and UUID depending on the Object Path Prefix selected:

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.	ObjectName and Collection are extracted from the path using a regular expression containing three patterns in parenthesis: `([^/])/([^/])/?(.*)` • The first is the collection. • The second is the Objectname. • The third is the relative path of the file within the DIVA object. The file paths not matching the regular expression are ignored.
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	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]+)/?(.*)` The file paths not matching the regular expression are ignored.

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-INDEXING). The name of this file is the name of the array with a .usage extension. For example:

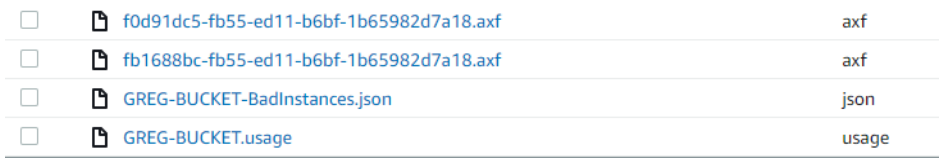

If DIVA is trying to use the same array for different purposes, the operation will fail.

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-INDEXING 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 job 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-INDEXING.
```

Web App Object Auto-Indexing Support

Object Auto-Indexing support has been combined with Cloud Bucket Scanning support. 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.

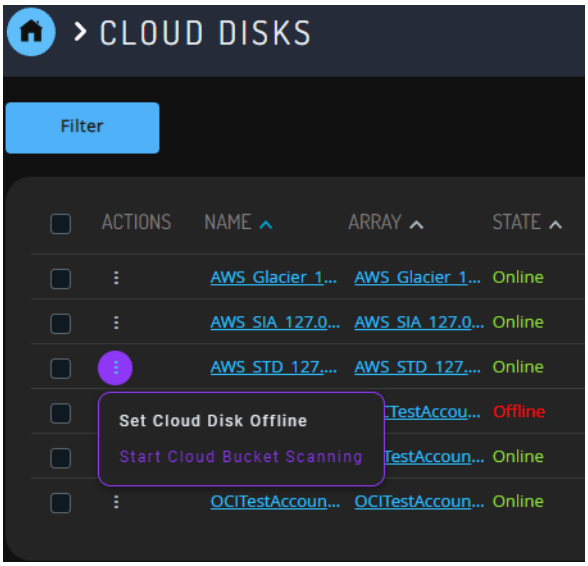

When Start Cloud Bucket Scanning is selected, a form is presented to the user. This is the same form for starting a cloud bucket scan as before, but a Scan Type selector has been added to choose the type of scan to perform. AXF Discovery is the type of scan that was supported in DIVA 8.2.

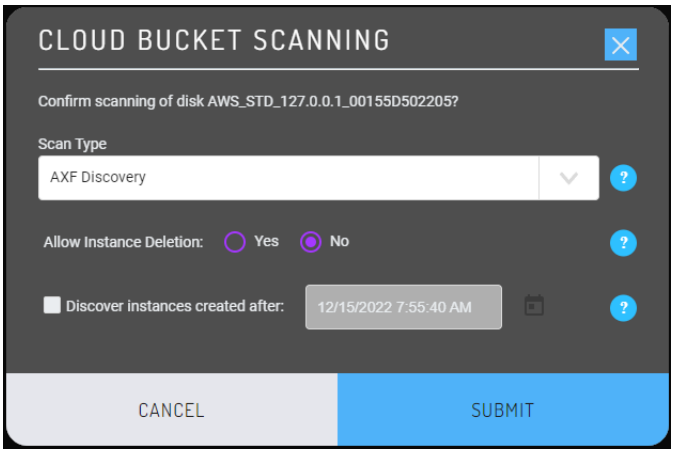

Content Indexing is the name for the Object Auto-Indexing scan. This option is only available if the Cloud Bucket format for the disk is set to `AXF RF 1.1`.

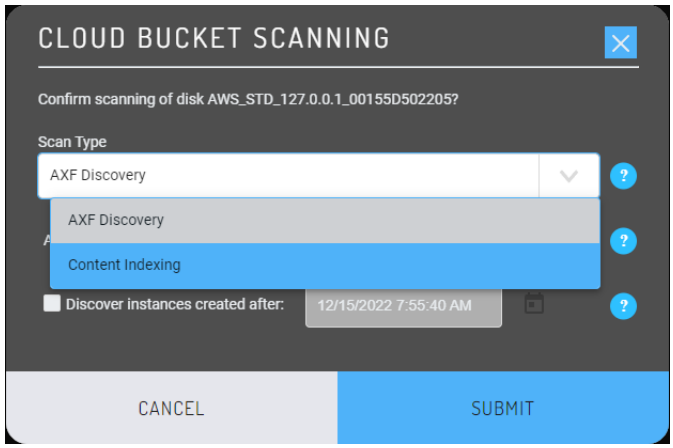

Selecting Content Indexing as the scan type displays the dialog for the options pertinent to that type of scan as shown here:

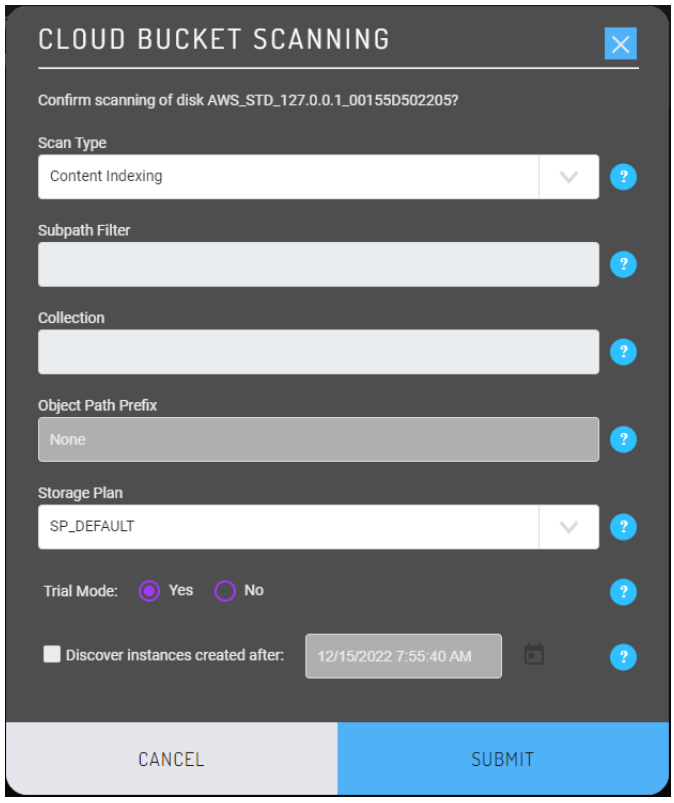

Viewing the status of the indexing is the same as previously done, and scan events will appear in the Scan Events table on the Property page of the cloud disk.

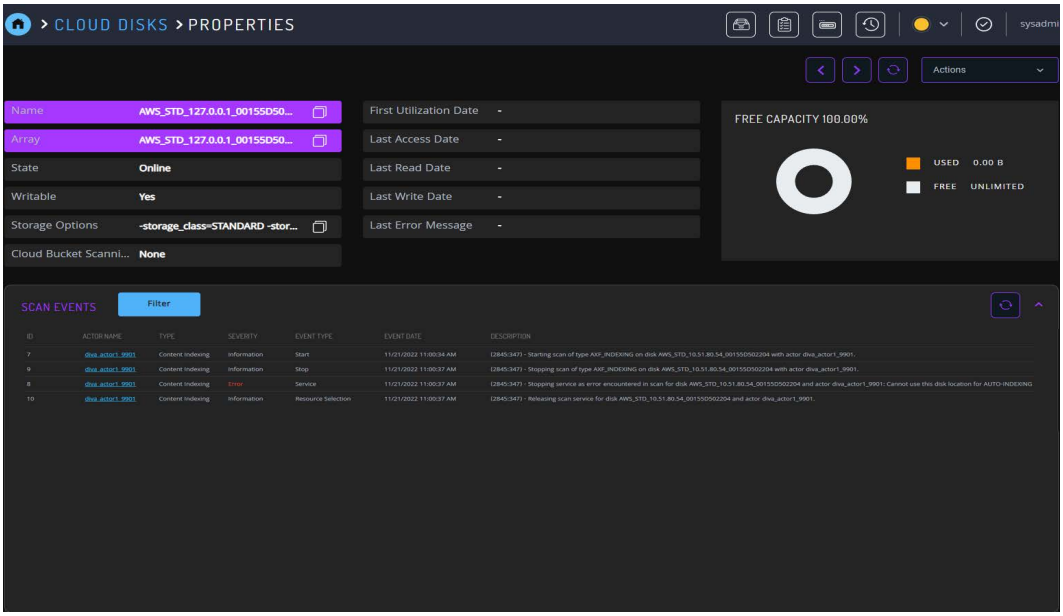

It is not possible to run both types of scan against a bucket. The default view of the properties table is to show all scan events. However, the Scan Events table can be filtered by scan type.

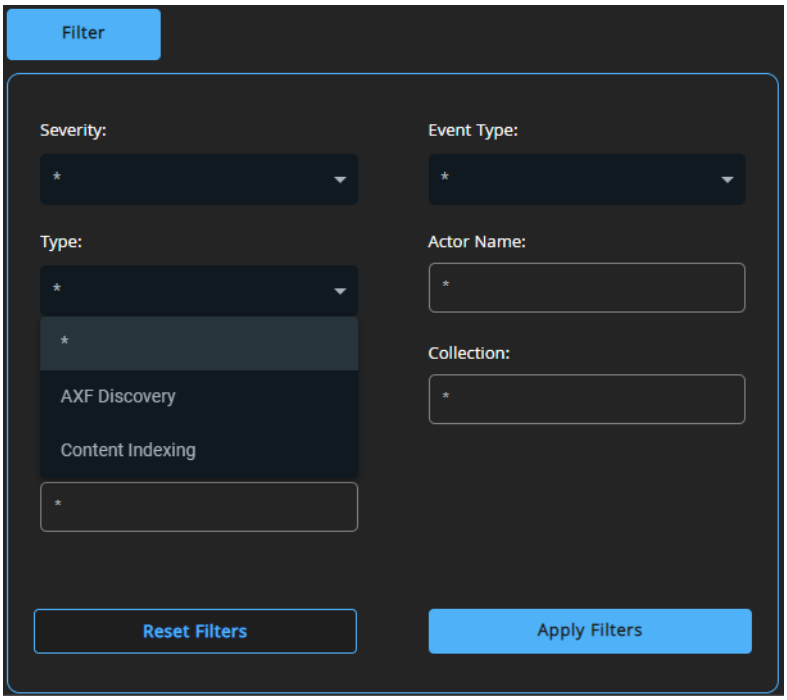

DIVA Frequently Asked Questions

Contact Technical Support for any additional questions not covered here.

What Happened to DIVA Command?

DIVA Command has been replaced by the DIVA web app.

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 operating systems be kept up to date with optional updates?

Optional operating system updates are not necessary in the DIVA 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 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 Operations Guide on the DIVA Technical Support site 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.

Where are vendor-specific logs located?

The vendor-specific log files are located in the %DIVA_HOME%\Program\log folder.

How far back in time do the logs go?

The log file retention period is configurable in the DIVA configuration file. Contact Technical Support for more information. The log files are retained as follows by default:

- DIVA 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 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 solution (for example, loss of a RAID disk, failover to the backup manager, and so on).

Are there 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.

How do I recover when a Complex Object's Metadata file is corrupted in the Main Manager System?

The DIVA 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 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 release. Contact Technical Support for assistance.

How do I locate a Complex Object's Metadata inside the Metadata Database?

Contact Technical Support for assistance.

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 web app.

What information is stored in the Metadata Database file?

All file details including file names, folder names, location, size, checksums, etc.

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 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 jobs 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 jobs. 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 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.

Can the location of the Metadata Database backups be configured?

Yes, you can configure the backup location. See *Metadata Database Installation and Configuration* for DIVA Backup Service installation and configuration procedures.

Unmanaged Server Configuration Guide

This appendix describes Source and Destination Server configuration guidelines for each type of DIVA supported content server. See the DIVA Supported Environments Guide on the DIVA Technical Support site for detailed and up-to-date lists of supported content servers, formats, and related DIVA platforms.

Topics

- 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 jobs. 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 jobs.

For Partial File Restore jobs, the file names on the destination are those specified when archiving. If no Files Path Root is entered, DIVA 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 jobs, unless the path specified in a job is an absolute path.

This approach provides better Server abstraction. You specify the server directories used by DIVA at the configuration level, not at the job level. They can be changed at any time without requiring a change to DIVA 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 Windows 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 ignores the parameter. However, if you do specify a Root Path its value is combined with the Files Path Root you specified in a job 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**	**Job 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	Job 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 (Job)	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 (Job)	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 jobs. This section also specifies, for each Connect Option, whether it can be superseded at the job 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 job option.

If the archive initiator sets the QualityOfService to something other than DEFAULT, DIVA 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 are preferred, but cache transfers will occur if processing the job in direct mode is not possible.

CACHE_AND_DIRECT

Cache transfers from (or to) a Server to (or from) DIVA are preferred, but direct transfers will occur if processing the job 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 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 job 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 job 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 job 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 job 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 job can use the Delete on Source QOS option, and is in the format -allow_delete_on_source.

The Archive job optional parameter delete_on_source instructs DIVA 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 job and the Source Server Type is not LOCAL, DISK or CIFS, DIVA automatically terminates the job.

This parameter applies to the FTP_STANDARD Source Server Type. you can change this behavior so that jobs will not fail when delete_on_source is specified in an Archive job.

If the -allow_delete_on_source option is specified, and the delete_on_source parameter is specified in an Archive job, DIVA will attempt to delete the asset from the Source Server after the archive has been completed successfully.

This option cannot be superseded by the job 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 jobs. 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 web app. The configuration can be superseded by the job 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 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 web app except the Service Name and Port parameters. DIVA 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>
```

This statement will process the file names as follows:

Input file Name	Output File Name
Star Wars	Star Wars.gxf
Readme.txt	Readme.txt
Jaws.gxf	Jaws.gxf

Example Two

Use the following statement to remove the .gxf extension (if any) at archive time:

```
-arch_renaming <.*\.gxf:(.*)\.(.*):\1>
```

This statement will process the file names as follows:

Input File Name	Output File Name
Star Wars.gxf	Star Wars
Readme.txt	Readme.txt
Jaws.avi	Jaws.avi

Example Three

When Flip Factory transcodes clip FOO to Pinnacle MSS, the resulting files are named FOO.MSS.header, FOO.MSS.ft, FOO.MSS.info, and FOO.MSS. These names are not those expected by Pinnacle MSS Servers, and this option fixes these discrepancies. Use the following statement to achieve this:

```
-arch_renaming
<.*\.header:(.*):header><.*\.ft:(.*):ft><.*\.info:(.*):info><.*MSS
:(.*):std>
```

This option will process the file names as follows:

Input File Name	Output File Name
FOO.MSS.header	header
FOO.MSS.ft	ft
FOO.MSS.info	info
FOO.MSS	std

Example Four

A variation of the following statement can be used to re-parent some files under a different relative directory:

```
arch_path_renaming <media:(.*):media.dir>
```

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 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 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 jobs. 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 job will be complete.

During an N-Restore, if multiple renaming rules are defined, DIVA 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 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 job 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 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 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 job option unless used in a TranscodeArchived job.

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 job option, unless used in a TranscodeArchived job.

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 according to the transcoders defined in the DIVA 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 Transcoder defined in the Transcoders area of the Systems page of the web app.

If this option is not present, DIVA will select one of the transcoders defined in the DIVA 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 job. 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 job.

The metadata file format is compliant with the DIVA 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 job 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 job constantly fails.

This option applies to any Server type and cannot be superseded by the job 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 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 job 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 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 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 releases 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 jobs and instructs the Manager to restore the original asset into an AXF File. Instead of purely restoring the content of an Object to the destination, DIVA 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

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 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.

MSS with Independent Storage

One record is created for each MSS that DIVA must move data to and from.

Attribute	Value	Example
IP Address	MSS IP address	10.80.114.21

Attribute	Value	Example
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 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 must move data to and from.

Attribute	Value	Example
IP Address	IP Address of the gateway through which DIVA 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 has to move data to and from.

Attribute	Value	Example
IP Address	IP Address of gateway through which DIVA accesses 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
FilesPathRoot	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 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 Avid Connectivity and Tools book on the DIVA Technical Support site for detailed information.

One record is created for each video server DIVA 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 jobs 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 Restore or Partial File Restore job, 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 Avid Connectivity and Tools book on the DIVA Technical Support site for detailed information.

One record is created for each video server DIVA 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 jobs 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 using FTP. 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 configuration. In this case, DIVA 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 configuration (see the following table). In this case, all nodes of the BMC must be declared in the DIVA configuration.

You can also add a Metasource that encompasses all nodes of the cluster to enable load balancing and failover from within DIVA.

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 using the FTP protocol.

DIVA 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

Attribute	Value	Example
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_SEACHANGECHECKDELAY=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
FilesPathRoot	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 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 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 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 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 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 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 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_SEACHANGECHECKDELAY=1000

-ftp or -cifs

One of these two options must be specified. Otherwise, Streaming API protocol is assumed, which is not supported by DIVA for BMLe and BMLex. This option cannot be superseded by the job 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 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 job option.

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.	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 depreciated. 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 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 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 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 Actors is an IP/FC network.

You must create one record for each video server DIVA 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 web 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
FilesPathRoot	/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). 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 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 job 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_re-naming options. This option can be superseded by the job 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
FilesPathRoot	/explodedFile/disk:/folder	/explodedFile/INT1:/default
FileNames	Enter the name of the clip in this field.	MyClip

Grass Valley K2 Servers

From DIVA's perspective, 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 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 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.

When this option is not set, DIVA 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 job 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 job 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 job 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 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
FilesPathRoot	/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.

One record must be created for each MAV70 server DIVA 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.

DIVA 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 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 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 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 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 job will terminate if this option is not included in the job.

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	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.

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 gets access to the media through a MediaGrid ContentDirector. DIVA 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 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 are properly shared for a given Windows user.
- The Actor service is configured to run under this user account.
- The user has local administrative rights on the Actor.

One record must be created for each ContentDirector DIVA 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
FilesPathRoot	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 searches for files matching the format clipname.mov or clipname.MOV. DIVA automatically retrieves and processes all of the potentially referenced files.

In cases where the material is wrapped as MXF, DIVA 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 web 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 jobs 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 transfers files called 8152.D10 and 8152.WAV to Power Portal.
- If DIVAACTOR_QUANTELRENAMECLIPS is set to 1 (enabled), DIVA 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
FilesPathRoot	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 Actor captures the new FID after the transfer and forwards it to the Manager.

DIVA_GetRequestInfo must be called to obtain the new FID using the DIVA API. If the job is completed, the new FID is in the job'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 job 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 jobs. 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 job:

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 configuration

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 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
FilesPathRoot	Leave this field empty.	
FileNames	Enter the Clip ID in this field.	1444247

Standard FTP and SFTP Servers

DIVA 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.

Create one record for each video server that Content Manager transfers data with.

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 jobs as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	Full path to files Partial path to files No path entry	/my_videos/movies /my_videos
FileNames	Names of files Partial path and names of files Full path and names of files	maniolia, matrix movies/maniolia, movies/matrix /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 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 jobs as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	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 FilesPathRoot section (that is, use scheme 1 or 2 in the previous table). Including them in the FileNames section prevents the job 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 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 jobs as described in the following table:

DIVA_archiveVirtual Object Parameter	Value	Example
FilesPathRoot	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 FilesPathRoot section (that is, use scheme 1 or 2 in the previous table). Including them in the FileNames section prevents the job 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 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 jobs being processed because it can lead to job 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 job 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 job 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 job 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 will try a maximum of four combinations; most likely a0-sd1, a0-sd2, a1-sd3, a1-sd1.

This option cannot be superseded by the job 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 deletes the content of C:\sourceserver\root recursively.
C:\sourceserver\root	*	-r -delete_fpr	DIVA deletes the content of C:\sourceserver\root recursively, and then deletes root.
C:\sourceserver\root	*		DIVA deletes the content of C:\sourceserver\root (first level only).
C:\sourceserver\root	*	-delete_fpr	DIVA 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 deletes the content of C:\sourceserver\root\obj recursively, and then deletes C:\sourceserver\root\obj.
C:\sourceserver\root	obj*	-r -delete_fpr	DIVA 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.

FilesPathRoot	Files	Options	Result
C:\sourceserver\root	obj1* obj2*	-r	DIVA 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 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 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 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 web app, or specifying the relative path in the FilesPathRoot field of the GUI Manager or API job for the archive or restore operation.

To set up a default home location so that an API job 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 job 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 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 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.

Dynamic Configuration Changes

This appendix lists the currently supported changes to DIVA configuration that become effective while the Manager is running, and those that require a software component or the Manager to be restarted.

Topics

Updates in the Manager Configuration

If a parameter in the 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 Web 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 jobs are complete.

- Add
- Delete (Notify Manager and after jobs complete)
- Source Server Name (Notify Manager and after jobs complete)
- IP Address (Notify Manager and after jobs complete)
- Source Server Type (Notify Manager and after jobs complete)
- Network (Notify Manager and after jobs complete)
- Site (Notify Manager and after jobs complete)
- Connect Options (Notify Manager and after jobs complete)
- Root Path (Notify Manager and after jobs complete)
- Max Throughput (Notify Manager and after jobs complete)

- Max Accesses (Notify Manager and after jobs complete). You must not make changes to this parameter while there are active job because it could lead to the job being terminated.
- Max Read Accesses (Notify Manager and after jobs complete). You must not make changes to this parameter while there are active job because it could lead to the job being terminated.
- Max Write Accesses (Notify Manager and after jobs complete). You must not make changes to this parameter while there are active job because it could lead to the job 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 jobs 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 jobs complete)
- Site (Notify Manager and after jobs complete)
- Max Drive Operations (Notify Manager and after jobs complete)
- Max Server Operations (Notify Manager and after jobs complete)
- Max Disk Operations (Notify Manager and after jobs complete)
- Direct Restore (Notify Manager and after jobs complete)
- Cache Restore (Notify Manager and after jobs complete)
- Copy to Tape Group (Notify Manager and after jobs complete)
- Associative Copy (Notify Manager and after jobs complete)
- Repack (Notify Manager and after jobs complete)
- Delete (Notify Manager and after jobs complete)
- Direct Archive (Notify Manager and after jobs complete)
- Cache Archive (Notify Manager and after jobs 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 Web 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 Web 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 Web 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 Web 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 Web 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 Web 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 Web 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 Web 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

	Event Type	Tape Type	Tape Barcode	Drive Type	Drive Name	Disk Name	Drive Serial Number	Library Serial Number	SD Name	Actor Name
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
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									

	Event Type	Tape Type	Tape Barcode	Drive Type	Drive Name	Disk Name	Drive Serial Number	Library Serial Number	SD Name	Actor Name
COPY_AS_REQUEST (to new)	Yes									
CREATE_INSTANCE	Yes									
RESTORE and PARTIAL_RESTORE	Yes								Yes	
DELETE_VIRTUALOBJECT	Yes									
DELETE_INSTANCE	Yes									
TRANSCODE_END	Yes									Yes
TRANSCODE_ERR	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 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 job
- 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 job

- Tape positioning before a tape write (Archive job)
- End Of Tape (EOT exception) encountered during an Archive job

	Object Name[1]	Object Category[1]	Object Instance[1]	Media (*Tape Group or array*)	Job ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
TAPE_INSERT						Yes	Yes		
TAPE_INSERT_ERR				Yes		Yes			
TAPE_MOUNT				Yes		Yes	Yes		
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

	Object Name[1]	Object Category[1]	Object Instance[1]	Media (*Tape Group or array*)	Job ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
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
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	

	Object Name[1]	Object Category[1]	Object Instance[1]	Media (*Tape Group or array*)	Job ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
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	
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	
TRANSCODE_END	Yes	Yes	Yes		Yes	Yes	Yes	Yes	Yes
TRANSCODE_ERR	Yes	Yes	Yes		Yes	Yes			
STOPPED_ON_CANCEL	Yes	Yes			Yes	Yes			
CHECKSUM_ERROR_TAPE	Yes	Yes	Optional	Yes	Yes	Yes			

	Object Name[1]	Object Category[1]	Object Instance[1]	Media (*Tape Group or array*)	Job ID	Event End Time	Event Duration	Transfer Size	Transfer Rate
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 jobs.
2. The transcoder work directory is not a DIVA 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]						

	Transfer Error Rate	Error Code	Error Message	Transcoder or Analyzer Name	Number of Archive Operations	Data Size
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			
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						

	Transfer Error Rate	Error Code	Error Message	Transcoder or Analyzer Name	Number of Archive Operations	Data Size
CHECKSUM_ERROR_DISK						
CHECKSUM_ERROR_SD						
TAPE_IMPORT						Yes
TAPE_EXPORT						Yes

1. The transcoder work directory is not a DIVA disk. No DISK READ or DISK WRITE events are created when accessing this directory.

Metrics Definitions

The following table identifies DIVA 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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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_AVG_TRANSFER_RATE_READ	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Lifetime
DISK_AVG_TRANSFER_RATE_READ_DAY	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Day
DISK_AVG_TRANSFER_RATE_READ_MONTH	DISK: average transfer rate of READ	DISK_READ	Average	Null	Transfer Rate	Disk Name	Month
DISK_AVG_TRANSFER_RATE_WRITE	DISK: average transfer rate of WRITE	DISK_WRITE	Average	Null	Transfer Rate	Disk Name	Lifetime
DISK_AVG_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 System: number of active Archive jobs	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Lifetime
DIVA_SYSTEM_ACTIVE_ARCHIVE_NUMBER_DAY	DIVA System: number of active Archive jobs	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Day
DIVA_SYSTEM_ACTIVE_ARCHIVE_NUMBER_MONTH	DIVA System: number of active Archive jobs	ARCHIVE_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Month
DIVA_SYSTEM_ACTIVE_COPY_AS_NUMBER	DIVA System: number of active Copy As New Object jobs	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_ACTIVE_COPY_AS_NUMBER_DAY	DIVA System: number of active Copy As New Object jobs	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Day
DIVA_SYSTEM_ACTIVE_COPY_AS_NUMBER_MONTH	DIVA System: number of active Copy As New Object jobs	COPY_AS_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Month
DIVA_SYSTEM_ACTIVE_COPY_NUMBER	DIVA System: number of active Copy jobs	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Lifetime
DIVA_SYSTEM_ACTIVE_COPY_NUMBER_DAY	DIVA System: number of active Copy jobs	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Day
DIVA_SYSTEM_ACTIVE_COPY_NUMBER_MONTH	DIVA System: number of active Copy jobs	COPY_REQUEST	Maximum	Null	Number of Operations	Local DIVA System	Month
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER	DIVA System: number of active Restore jobs	RESTORE	Maximum	Null	Number of Operations	Local DIVA System	Lifetime
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER_DAY	DIVA System: number of active Restore jobs	RESTORE	Maximum	Null	Number of Operations	Local DIVA System	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_ACTIVE_RESTORE_NUMBER_MONTH	DIVA System: number of active Restore jobs	RESTORE	Maximum	Null	Number of Operations	Local DIVA System	Month
DIVA_SYSTEM_AVG_READ_WRITE	DIVA 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 System	Lifetime
DIVA_SYSTEM_AVG_READ_WRITE_DAY	DIVA 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 System	Day
DIVA_SYSTEM_AVG_READ_WRITE_MONTH	DIVA 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 System	Month
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_ARCHIVE	DIVA System: number of Objects archived	ARCHIVE_REQUEST	Count	Null	Transfer Size	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_ARCHIVE_DAY	DIVA System: number of Objects archived	ARCHIVE_REQUEST	Count	Null	Transfer Size	Local DIVA System	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_ARCHIVE_MONTH	DIVA System: number of Objects archived	ARCHIVE_REQUEST	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_CREATED	DIVA System: number of Objects created	ARCHIVE_REQUEST COPY_AS_REQUEST	Count	Null	Event ID	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_CREATED_DAY	DIVA System: number of Objects created	ARCHIVE_REQUEST COPY_AS_REQUEST	Count	Null	Event ID	Local DIVA System	Day
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_CREATED_MONTH	DIVA System: number of Objects created	ARCHIVE_REQUEST COPY_AS_REQUEST	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_DELETED	DIVA System: number of Objects deleted	DELETE_VIRTUALOBJECT	Count	Null	Event ID	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_DELETED_DAY	DIVA System: number of Objects deleted	DELETE_VIRTUALOBJECT	Count	Null	Event ID	Local DIVA System	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_NUMBER_VIRTUALOBJECT_DELETED_MONTH	DIVA System: number of Objects deleted	DELETE_VIRTUALOBJECT	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_NUMBER_VIRTUALOBJECT_INSTANCE_COPY	DIVA System: number of Object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUALOBJECT_INSTANCE_COPY_DAY	DIVA System: number of Object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA System	Day
DIVA_SYSTEM_NUMBER_VIRTUALOBJECT_INSTANCE_COPY_MONTH	DIVA System: number of Object instances copied	COPY_REQUEST	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_NUMBER_VIRTUALOBJECT_INSTANCE_CREATED	DIVA System: number of Object instances created	CREATE_INSTANCE	Count	Null	Event ID	Local DIVA System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_INSTANCE_CREATED_DAY	DIVA System: number of Object instances created	CREATE_INSTANCE	Count	Null	Event ID	Local DIVA System	Day
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_INSTANCE_CREATED_MONTH	DIVA System: number of Object instances created	CREATE_INSTANCE	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_NUMBERVIRTUAL_OBJECT_INSTANCE_DELETED	DIVA System: number of Object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_INSTANCE_DELETED_DAY	DIVA System: number of Object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Local DIVA System	Day
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_INSTANCE_DELETED_MONTH	DIVA System: number of Object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Local DIVA System	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_RESTORE	DIVA System: number of Objects restored	RESTORE	Count	Null	Event ID	Local DIVA System	Lifetime
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_RESTORE_DAY	DIVA System: number of Objects restored	RESTORE	Count	Null	Event ID	Local DIVA System	Day
DIVA_SYSTEM_NUMBER_VIRTUAL OBJECT_RESTORE_MONTH	DIVA System: number of Objects restored	RESTORE	Count	Null	Event ID	Local DIVA System	Month
DIVA_SYSTEM_READ_WRITE	DIVA 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 System	Lifetime
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER	DIVA 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 System	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER_DAY	DIVA 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 System	Day
DIVA_SYSTEM_READ_WRITE_ABORTED_NUMBER_MONTH	DIVA 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 System	Month
DIVA_SYSTEM_READ_WRITE_DAY	DIVA 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 System	Day
DIVA_SYSTEM_READ_WRITE_MONTH	DIVA 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 System	Month

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
DIVA_SYSTEM_READ_WRITE_NUMBER	DIVA 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 System	Lifetime
DIVA_SYSTEM_READ_WRITE_NUMBER_DAY	DIVA 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 System	Day
DIVA_SYSTEM_READ_WRITE_NUMBER_MONTH	DIVA 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 System	Month
MEDIA_ARCHIVED_VIRTUALOBJECT_DATASIZE_DAY	Media: data size of all Objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	Media Name	Day
MEDIA_ARCHIVED_VIRTUALOBJECT_DATASIZE_MONTH	Media: data size of all Objects archived	ARCHIVE_REQUEST	Sum	Null	Transfer Size	Media Name	Month
MEDIA_VIRTUALOBJECT_INSTANCE_CREATE	Media: number of Object instances created	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
MEDIA_VIRTUALOBJECT_INSTANCE_CREATE_DAY	Media: number of Object instances created	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Day
MEDIA_VIRTUALOBJECT_INSTANCE_CREATE_MONTH	Media: number of Object instances created and deleted	CREATE_INSTANCE	Count	Null	Event ID	Media Name	Month
MEDIA_VIRTUALOBJECT_INSTANCE_DELETE	Media: number of Object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Media Name	Lifetime
MEDIA_VIRTUALOBJECT_INSTANCE_DELETE_DAY	Media: number of Object instances deleted	DELETE_INSTANCE	Count	Null	Event ID	Media Name	Day
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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_DISMOUNT_ABORTED	Tape Drive: number of terminated MOUNT and DISMOUNT operations (*together*)	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 (together)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Drive Serial Number	Lifetime
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_NUMBER_DISMOUNT_ABORTED	Tape Library: total number of terminated DISMOUNT operations	TAPE_DISMOUNT_ERR	Count	Null	Event ID	Library Serial Number	Lifetime
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
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
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_NUMBER_WRITE	Tape Library: total number of WRITE operations	TAPE_WRITE TAPE_WRITE_ERR	Count	Null	Event ID	Library Serial Number	Lifetime
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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_LIBRARY_WRITE_DAY	Tape Library: total amount of data WRITE	TAPE_WRITE	Sum	Null	Transfer Size	Library Serial Number	Day
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 (*together*)	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 (*together*)	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 (*together*)	TAPE_READ_ERR TAPE_WRITE_ERR	Count	Null	Event ID	Tape Barcode	Day
TAPE_READ_WRITE_NUMBER	Tape: number of READ and WRITE operations (*together*)	TAPE_READ TAPE_WRITE	Count	Null	Event Id	Tape Barcode	Lifetime

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TAPE_READ_WRITE_NUMBER_DAY	Tape: number of READ and WRITE operations	TAPE_READ TAPE_WRITE	Count	Null	Event ID	Tape Barcode	Day
TAPE_REPACK_NUMBER	Tape: number of REPACK, REUSE, and REFORMAT operations (together)	TAPE_REPACK	Count	Null	Event ID	Local DIVA 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

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TRANSCODE_AVG_THROUGHPUT_DAY	Transcoder: average transcoding throughput	TRANSCODE_END	Average	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day
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_MAX_THROUGHPUT	Transcoder: maximum transcoding throughput	TRANSCODE_END	Maximum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_MAX_THROUGHPUT_DAY	Transcoder: maximum transcoding throughput	TRANSCODE_END	Maximum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day
TRANSCODE_MIN_THROUGHPUT	Transcoder: minimum transcoding throughput	TRANSCODE_END	Minimum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Lifetime
TRANSCODE_MIN_THROUGHPUT_DAY	Transcoder: minimum transcoding throughput	TRANSCODE_END	Minimum	Null	Transfer Rate	Transcoder Name or Analyzer Name	Day

Metric Name	Description	Events	Operation	Weight Factor	Collection Field	Aggregation Field	Collection Interval
TRANSCODE_NUMBER	Transcoder: number TRANSCODE operations	TRANSCODE_END	Count	Null	Event ID	Transcoder Name or Analyzer Name	Lifetime
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

Parameter	Default	Values
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