

SGI® OpenGL Multipipe™
User's Guide

Version 2.4.1



007-4318-015



CONTRIBUTORS

Written by Ken Jones and Jenn Byrnes

Illustrated by Chrystie Danzer

Production by Karen Jacobson

Engineering contributions by Ye Cong, Craig Dunwoody, Bill Feth, Alpana Kaulgud, Claude Knaus, Ravid Na'ali, Jeffrey Ungar, Christophe Winkler, Guy Zadicario, and Hansong Zhang

COPYRIGHT

© 2000–2004 Silicon Graphics, Inc. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, InfiniteReality, IRIS, IRIX, Onyx, Onyx2, OpenGL, and Reality Center are registered trademarks and GL, InfinitePerformance, InfiniteReality2, IRIS GL, Octane2, Onyx4, Open Inventor, the OpenGL logo, OpenGL Multipipe, OpenGL Performer, Power Onyx, SGI Advanced Linux, SGI Propack, Silicon Graphics Prism, Tezro, and UltimateVision are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

GNOME is a trademark of the GNOME Foundation. KDE is a trademark of KDE e.V, Incorporated. Linux is a registered trademark of Linus Torvalds. MIPS and R10000 are registered trademarks of MIPS Technologies, Inc. used under license by Silicon Graphics, Inc. Netscape is a trademark of Netscape Communications Corporation. XFree86 is a trademark of The XFree86 Project, Inc. Xinerama, X Window System, and the X device are trademarks of The Open Group. All other trademarks mentioned herein are the property of their respective owners.

New Features in This Release

OpenGL Multipipe 2.4.1 contains the following enhancements:

- Improved daemon processing

Slave render servers and other OpenGL Multipipe daemons are now launched by the `inetd` or `xinetd` meta daemon on IRIX and Linux, respectively.

The `ompslave`, `ompcull`, and `ompswapready` processes will only be running while an application is actively running under OpenGL Multipipe.

- Improved support for static tiling modes of hardware compositors
- Enhanced application compatibility

Record of Revision

Version	Description
001	August 2000 Beta release.
002	November 2000 Updated for release 1.0 of the OpenGL Multipipe product.
003	February 2001 Updated for release 1.1 of the OpenGL Multipipe product. New features: <ul style="list-style-type: none">- Increased overall performance- Support for overlapping screens, as in SGI Reality Center facilities
004	May 2001 Updated for release 1.2 of the OpenGL Multipipe product. New features: <ul style="list-style-type: none">- Transparent OpenGL Pipe Management- Subset of multipipe applications made aware of Xinerama
005	August 2001 Updated for release 1.3 of the OpenGL Multipipe product. New features: <ul style="list-style-type: none">- Enhanced Support for Multithreaded Applications- Enhanced tgl Script
006	November 2001 Updated for release 1.4 of the OpenGL Multipipe product.

Bugfixes:

- Enhanced GLX conformance for context manipulation
- Support for pixmaps, pbuffers, and GLXWindows

Beta features:

- Curved Screen Support

This allows you to run applications on a non-planar Reality Center in immersive mode by adapting the 3D projections to the display layout.

- Window Manager Support for Aware Windows

All applications started in aware mode can now be under window manager control by using the customized window manager included with this release.

007

February 2002

Updated for release 1.4.1 of the OpenGL Multipipe product.

Broader application support

Bugfixes:

- Enhanced OpenGL conformance for applications using `glCallList()` within another display list
- Stability improvements to (beta) aware window manager

008

April 2002

Updated for release 1.4.2 of the OpenGL Multipipe product.

- Broader application support
- Stability improvements to (beta) aware window manager

009

October 2002

Updated for release 2.1 of the OpenGL Multipipe product.

Features:

Replacement of the Transparent OpenGL (TGL) layer with a proxy render library and render servers

Support for additional servers. The list of supported servers now includes the following:

- Silicon Graphics Onyx
- Silicon Graphics Onyx2
- SGI Onyx 3000, InfiniteReality

- SGI Onyx 3000, InfinitePerformance
- Silicon Graphics Octane2

010

May 2003

Updated for release 2.1.2 of the OpenGL Multipipe product.

Features:

- Performance enhancements over the OpenGL Multipipe 2.1.1 and 2.1 releases
- Increased application compatibility for vertex array applications
- Option of running in master render mode or slave-only mode
- Support for compositor-based systems
- Seamless cursor movement across overlapped or composited screen regions (IRIX 6.5.20 or later required)
- Support for SGI-SCREEN-CAPTURE and ReadDisplay X extensions in SGI Xinerama mode (IRIX 6.5.20 or later required)
- An API introduced for integration of multipipe applications with SGI Xinerama
- Hardware swap-synchronization option (Swap Ready, Genlock)
- Support for additional platforms. The list of supported visualization systems now includes the following:
 - o SGI Onyx 3000 series with InfinitePerformance graphics
 - o SGI Onyx 3000 series with InfiniteReality graphics
 - o SGI Onyx 350
 - o Silicon Graphics Octane2
 - o Silicon Graphics Onyx
 - o Silicon Graphics Onyx2

011

July 2003

Updated for release 2.2 of the OpenGL Multipipe product.

Features:

- Performance enhancements over the 2.1.2 release including the following:
 - o Pixel drawing enhancements
 - o Geometry culling, which can improve application performance for some applications beyond what can be achieved on a single pipe
- Support for the DMX (Distributed Multihead X) meta display server

- Support for DMX and SGI Xinerama meta display servers by the MPC API for multipipe-aware applications
- Support for additional platforms. The list of supported visualization systems now includes the following:
 - o SGI Onyx 3000 series with InfinitePerformance graphics
 - o SGI Onyx 3000 series with InfiniteReality graphics
 - o SGI Onyx 350
 - o Silicon Graphics Octane2
 - o Silicon Graphics Onyx
 - o Silicon Graphics Onyx2
 - o Silicon Graphics Onyx4 UltimateVision

012

December 2003

Updated for release 2.3 of the OpenGL Multipipe product.

Features:

- Performance enhancement options including (disabled by default):
 - o Geometry culling to optional, additional OpenGL clip planes
 - o OpenGL viewport clipping and geometry culling to improve fill and geometry performance when using a compositor
 - o Spatial partitioning of large display lists for efficient geometry culling
- Baseline performance enhancements over OpenGL Multipipe 2.2
 - o Swap synchronization and frame latency control improvements between slave rendering processes
 - o Improved performance when using geometry culling
- Support for Onyx4 OpenGL extensions
- Support for additional platforms. The list of supported visualization systems now includes the following:
 - o SGI Onyx 3000 series with InfinitePerformance graphics
 - o SGI Onyx 3000 series with InfiniteReality graphics
 - o SGI Onyx 350
 - o Silicon Graphics Octane2
 - o Silicon Graphics Onyx
 - o Silicon Graphics Onyx2
 - o Silicon Graphics Onyx4 UltimateVision
 - o Silicon Graphics Tezro

- 013 April 2004
Updated for release 2.3.1 of the OpenGL Multipipe product.
Enhancements over OpenGL Multipipe 2.3:
- Capability to expand the OpenGL viewport size beyond the single-pipe limit
 - Support for overlay visuals in the DMX proxy X server if supported by underlying X servers
 - Support for applications that use GLX pbuffers or GLX pixmaps
 - Improved application compatibility
- 014 October 2004
Updated for release 2.4 of the OpenGL Multipipe product.
Enhancements over OpenGL Multipipe 2.3.1:
- Support for most OpenGL Multipipe features on Silicon Graphics Vizualization System for Linux 64-bit platforms
 - Small object culling (small relative to screen space)
 - Drawpixels culling
 - Improved application compatibility
- 015 November 2004
Updated for release 2.4.1 of the OpenGL Multipipe product.
Enhancements over OpenGL Multipipe 2.4:
- Improved daemon processing
 - Improved support for static tiling modes of hardware compositors
 - Enhanced application compatibility

Contents

	New Features in This Release	iii
	Record of Revision	v
	About This Guide.	xv
	Related Publications	xv
	Obtaining Publications	xvi
	Conventions	xvi
	Reader Comments.	xvii
1.	OpenGL Multipipe Overview	1
	What OpenGL Multipipe Provides	1
	Architecture of OpenGL Multipipe	4
	Components of OpenGL Multipipe	5
	The X Proxy Layer	5
	The SGI Xinerama Extension	5
	The DMX Proxy Server	6
	The 3D (Master) Proxy Render Library.	6
	3D (Slave) Render Servers	7
	Supported Platforms	8
2.	Installing OpenGL Multipipe	9
	Installing on IRIX	9
	Installing on 64-Bit Linux.	11

3. Using OpenGL Multipipe 13
Setting up the OpenGL Multipipe Environment 13
Configuring OpenGL Multipipe with DMX as the X Proxy Layer 14
Initializing DMX 14
Creating DMX Configuration Files 18
Configuring OpenGL Multipipe with SGI Xinerama as the X Proxy Layer 20
Initializing SGI Xinerama 20
Configuring Overlapping Screens with SGI Xinerama 21
Verifying That the OpenGL Multipipe Environment is Enabled 21
Disabling the OpenGL Multipipe Environment 22
Running Applications with OpenGL Multipipe 22
Running OpenGL Single-Pipe Applications 24
Running Pure X Applications 24
Running IRIS GL Applications 24
Running IRIX o32 Applications 25
Running Multipipe Applications in Multipipe-Aware Mode 26
Performance Features 27
Viewport Clipping 27
Geometry Culling 28
Small Object Culling 29
Display List Partitioning 29
Master Rendering Modes 29
-mstrmode off Mode 30
-mstrmode track Mode 31
-mstrmode render Mode 32
Frame Latency Control 33
Buffer Swap Synchronization 34
Software Swap Synchronization 34
Hardware Buffer Swap Synchronization 34

Using SGI Scalable Graphics Hardware with OpenGL Multipipe	35
Configuring Compositing Screens with SGI Xinerama	36
Configuring Compositing Screens with DMX	36
Specifying Static Compositing Regions with <code>ompstartdmx</code>	37
Specifying Static Compositing Regions for Fully Overlapped Screens	38
Enabling Duplicate Cursor Images in Overlap Regions	38
Under SGI Xinerama.	39
Under DMX	39
Managing Windows for Aware Applications	40
Starting an Aware Window Manager	41
Exiting an Aware Window Manager	43
Setting an Aware Window Manager as the Default	43
4. Limitations.	45
Performance Enhancement	46
X Extensions	46
The Multipipe-Aware Window Manager	46
OpenGL Window Size Constraints	46
Processor Requirements	47
SGI ProPack and OpenGL Multipipe Versions.	47
Overlay Windows Support in DMX.	47
SGI Xinerama Not Supported on Onyx4 or Silicon Graphics Prism Platforms.	47
Graphics Pipe Requirements	48
Some Features Not Available on Linux	48
5. Troubleshooting	49
Problems Enabling SGI Xinerama	50
Problems Starting DMX	50
Problems Starting Applications with <code>omprun</code>	51
Problems Running IRIS GL Applications	51
Problems Running o32 Applications	52

Graphics Do Not Display Correctly on All Screens.52
Coding Problem in the Application52
You Did Not Use the <code>omprun</code> Script53
A User-Defined Script Invokes an IRIS GL or o32 Application53
You Are Using the Aware Window Manager54
Set-User-ID (“s-bit”) Applications54
Mouse Behavior Offset by a Screen55
Problems Running Inherently Multipipe Applications55
Multipipe-Aware Applications Fail to Receive Events on Screen 056
Nothing Displays or the Graphic Stalls or Hangs56
Coding Problem in the Application56
Window Exceeds Maximum OpenGL Window Size57
Improperly Wired Genlock or Swap Ready Cables57
OpenGL Graphics Render Slowly57
X Applications Are Not Behaving Correctly or Fail to Start58
X Application Uses Unsupported X Extension.58
SGI Xinerama Client or Server Uses Nonstandard Protocol59
Application Window Disappears59
Application Explicitly Opens a Display Connection to <code>:0.0</code>60
Simultaneously Running X Servers with and without SGI Xinerama Enabled60
Tiled Background Image61
Flickering Grey Rubberband During Window Movement61
Mouse Disappears on Composited or Edge-Blended Display.62
Problems Running Multithreaded Applications62
Problems with Aware Window Management62
Windows of Some Aware Applications are Not Managed63
Problems with Desktop Background Images63
Mouse Events Sometimes Register on the Wrong Screen.63
Ghost Windows Appear In Overlap Regions on Edge-Blended Displays64
Applications Do Not Behave Correctly in Aware Mode64

About This Guide

This guide describes the OpenGL Multipipe product, which allows you to run single-pipe applications in a multipipe environment without modification. You can seamlessly move single-pipe application windows across the single logical display that OpenGL Multipipe creates from multiple pipes. Both multipipe applications and single-pipe applications run concurrently.

Related Publications

The following SGI documents contain additional information that may be helpful:

- *InfiniteReality Video Format Combiner User's Guide*
- *POWER Onyx and Onyx Rackmount Owner's Guide*
- *SGI InfinitePerformance: Scalable Graphics Compositor User's Guide*
- *IRIX Admin: Software Installation and Licensing*

These books might also be helpful:

- Dave Shreiner, OpenGL Architecture Review Board, Mason Woo, Jackie Neider and Tom Davis. *OpenGL Programming Guide: The Official Guide to Learning OpenGL, Version 1.4*. Reading, MA: Addison Wesley Longman Inc., 2003. ISBN 0-321-17348-1.
- Nye, Adrian, *Volume One: Xlib Programming Manual*. Sebastopol, California: O'Reilly & Associates, Inc., 1992.

Obtaining Publications

You can obtain SGI documentation in the following way:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. With an IRIX system, select **Help** from the Toolchest, and then select **InfoSearch**. Or you can type `infosearch` on a command line.
- On IRIX, you can also view release notes by typing either `grelnotes` or `relnotes` on a command line.
- You can also view man pages by typing `man< title>` on a command line.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
interface	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists. Functions are also denoted in bold with following parentheses.
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.

Right angle brackets (>) These brackets indicate a path through menus to a menu option. For example, “**File > Open**” means “Under the **File** menu, choose the **Open** option.”

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library Web page:
<http://docs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, CA 94043-1351

SGI values your comments and will respond to them promptly.

OpenGL Multipipe Overview

This overview of OpenGL Multipipe consists of the following sections:

- “What OpenGL Multipipe Provides”
- “Architecture of OpenGL Multipipe”
- “Components of OpenGL Multipipe”
- “Supported Platforms”

What OpenGL Multipipe Provides

SGI has always been focused on high-end graphics solutions. The Onyx family of scalable visualization supercomputers allows you to have multiple graphics pipes on one single-system-image machine in order to reach new visualization performances. These multipipe systems are commonly used to drive expanded visualization systems such as SGI Reality Center facilities. OpenGL Multipipe extends the use of these powerful supercomputers to a broad spectrum of graphics applications without the requirement of modifying the applications.

Many existing graphics applications—such as Netscape or applications based on Open Inventor, for example—are constrained to run on a single pipe. On these single-pipe applications, you can choose the pipe on which to open the application’s windows, but the windows cannot be dragged from one pipe to another. The main reason is that the graphics pipes are separate logical units and are handled by an X server as different, unconnected screens. This means that the X server does not provide any functionality to group multiple screens into a single logical display. A second reason is that OpenGL applications connect directly to a specified graphics pipe and bypass the X protocol layer.

In the past, displaying an application on multiple screens required you to explicitly write the application for that purpose. You had to use tools like the OpenGL Performer or OpenGL Multipipe SDK libraries to help you create these multipipe applications. These tools allow you to explicitly open windows on different screens and to draw into them

using OpenGL. However, this solution lacks consistency. In fact, all of the windows on the different pipes are independent; hence, moving or iconifying one window on one screen will not handle the other windows accordingly.

OpenGL Multipipe has been designed to overcome these difficulties. The goal is to group pipes managed by the X server in order to create a consistent, single virtual screen as shown in Figure 1-1. This means that the applications are unaware of the underlying hardware configuration. Rather, they only know about a single display and behave accordingly.




Figure 1-1 OpenGL Multipipe with Non-Overlapping Screens

In contrast to Figure 1-1, if you have screens that overlap each other (such as in an SGI Reality Center wall display with edge blending), OpenGL Multipipe allows you to

specify the amount of this overlap. Figure 1-2 shows the image blended on overlapping screens.

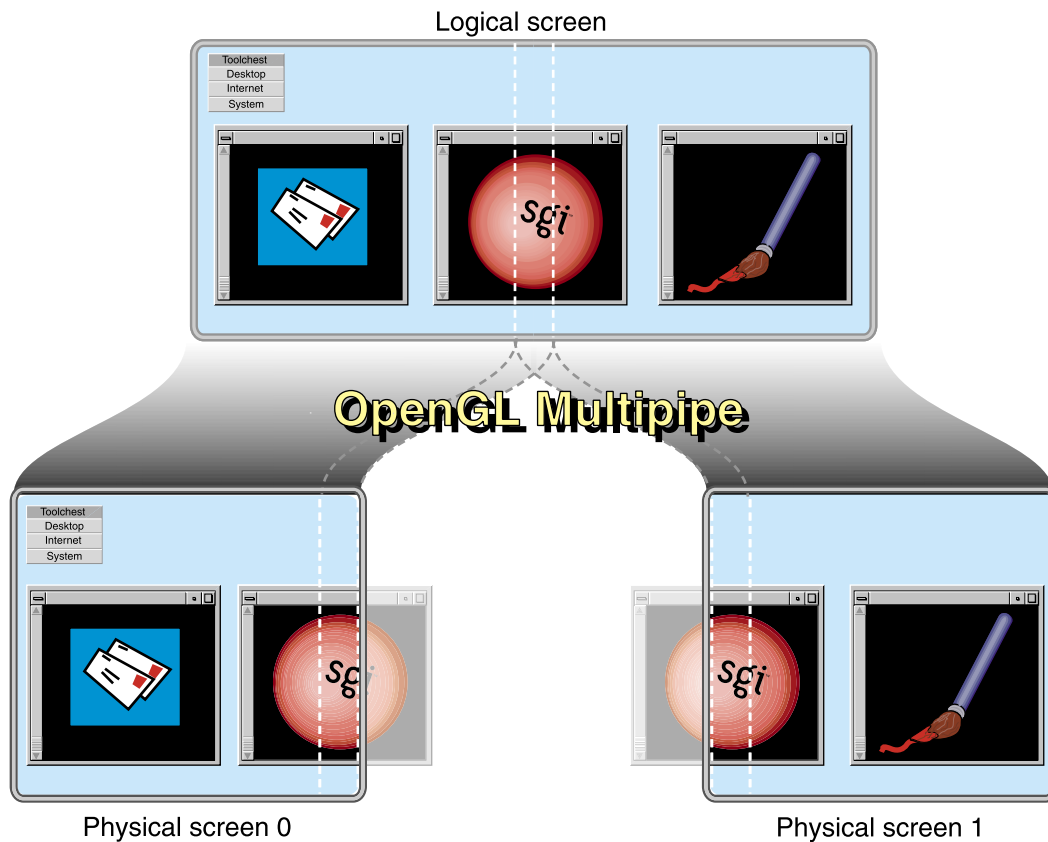


Figure 1-2 OpenGL Multipipe with Overlapping Screens

Note: OpenGL Multipipe does not require you to modify or recompile your application.

Architecture of OpenGL Multipipe

OpenGL Multipipe provides the illusion that an application is rendering 2D (X perspective) and 3D (OpenGL perspective) on a single local pipe when it is actually using one or more pipes. In this regard, a single logical display, OpenGL Multipipe 2 is similar to the first-generation product OpenGL Multipipe 1.

Unlike OpenGL Multipipe 1, however, the architecture of OpenGL Multipipe 2 allows the application processing and the rendering to occur in separate processes instead of separate threads of the same process. The separation improves application compatibility and is a step toward providing better scalability.

Like OpenGL Multipipe 1, OpenGL Multipipe 2 is a set of protocols and proxies coupled with clients and servers. Both versions of the product use an X proxy layer to hide the physical screen layout. This X proxy layer presents a single logical pipe or “meta screen” to all applications and allows their windows to be freely moved across or to span any set of pipes. In OpenGL Multipipe 1, the X proxy layer is the SGI Xinerama X extension. In OpenGL Multipipe 2 (OpenGL Multipipe 2.2 and later), the X proxy layer may be SGI Xinerama or the new Distributed Multihead X (DMX) proxy server.

OpenGL Multipipe 1 uses the Transparent OpenGL (TGL) library to send OpenGL calls to each real pipe. In contrast, OpenGL Multipipe 2 uses a 3D proxy library and render servers for this purpose.

Table 1-1 charts the primary interfaces of OpenGL Multipipe 1 and OpenGL Multipipe 2.

Table 1-1 Primary OpenGL Multipipe Interfaces

Product	X Server Interface	Interface with OpenGL and Graphics Pipes
OpenGL Multipipe 1	SGI Xinerama extension	Transparent OpenGL (TGL)
OpenGL Multipipe 2	SGI Xinerama extension or Distributed Multihead X (DMX) proxy server	Proxy render library and render servers

Components of OpenGL Multipipe

OpenGL Multipipe has the following components:

- An X proxy layer (the SGI Xinerama extension or the DMX proxy server)
- A 3D proxy render library
- 3D render servers

The X Proxy Layer

For pure X applications—that is, applications that do not use other graphics libraries (such as OpenGL) to draw into their windows—the X proxy layer is all that is needed to enable such applications to run transparently over multiple pipes. This means that windows of applications that are based on the X protocol and that use X extensions can be dragged from one pipe to another and even span multiple pipes. The applications behave as if they are running on a single, large virtual pipe. The X proxy layer hides the real screens from the client applications connecting to it. It distributes to all pipes the X requests from the clients but only sends the clients information about the large virtual display.

The X proxy layer can be either the SGI Xinerama extension or the DMX proxy server. In the case of the SGI Xinerama extension, the X proxy layer is a part of the X server. However, the DMX proxy server is a separate entity. The following sections describe the two options.

The SGI Xinerama Extension

SGI Xinerama is an enhanced version of the standard Xinerama X extension, which groups all screens managed by the X server into one logical screen that it exposes to applications. You must have administrative privileges to enable or disable the SGI Xinerama extension. For more information about the SGI Xinerama extension, see the `xinerama(3X11)` man page.

The DMX Proxy Server

The DMX proxy server, unlike SGI Xinerama, is not part of the X server; it is an X application that behaves like an X server to other X applications. DMX is more flexible than SGI Xinerama both in its supported display geometries and in its ability to act as a proxy for many different X servers. DMX also has built-in support for OpenGL applications through its support of the GLX X extension. This means that DMX will enable X and OpenGL applications to run transparently across multiple pipes. However, DMX's GLX extension is limited in performance. Hence, it is best to run graphics-intensive applications under the full OpenGL Multipipe environment.

Unlike SGI Xinerama, administrative privileges are not required to start and stop DMX. For more information about the DMX proxy server, see the `Xdmx(1)` man page, which is installed in `/usr/share/omp/doc/user/Xdmx.1.html`.

The 3D (Master) Proxy Render Library

OpenGL applications are X applications that use another graphics library (namely the OpenGL library) to draw into their windows. OpenGL applications open a direct connection to a graphics pipe. This means that the application is bypassing the X protocol (and the X proxy layer, which replicates the the X protocol stream to each pipe) in order to draw in the windows through this direct connection. The X proxy layer, which accounts only for the X protocol, is unable to handle this case. In OpenGL Multipipe 1, the Transparent OpenGL (TGL) library was designed to handle the OpenGL side of any application.

The *master* proxy library corresponds to the TGL layer in OpenGL Multipipe 1. It intercepts OpenGL calls to enable distribution to multiple pipes, but it does this by sending an OpenGL wire protocol stream to separate *slave* render processes rather than by spawning threads within the application to render to local pipes as does TGL. Figure 1-3 illustrates the functions of the master proxy library.




Figure 1-3 Master Proxy Library Functions

In addition to sending an OpenGL stream to each slave, the master also has the capability of rendering directly to a single local pipe in place of a single slave render process (for faster GL state queries), or it may use a local pipe only to track OpenGL state while a slave process renders to that pipe (for improved parallelism).

3D (Slave) Render Servers

An application running under the master render library of OpenGL Multipipe communicates its OpenGL commands to slave render server processes. Each slave process translates the OpenGL wire protocol stream into OpenGL commands that are executed on the application's behalf. For each rendering application thread of execution, one slave process exists for each screen of the display.

Supported Platforms

Table 1-2 shows the requirements for OpenGL Multipipe 2.4.1.

Table 1-2 Supported Platforms

Server	Processor	Operating System
Silicon Graphics Onyx	MIPS R10000 or later	IRIX 6.5.15 or later
Silicon Graphics Onyx2	MIPS R10000 or later	IRIX 6.5.15 or later
SGI Onyx 3000 with InfiniteReality graphics	MIPS R10000 or later	IRIX 6.5.15 or later
SGI Onyx 3000 with InfinitePerformance graphics	MIPS R10000 or later	IRIX 6.5.15 or later
SGI Onyx 350	MIPS R10000 or later	IRIX 6.5.15 or later
SGI Onyx4 UltimateVision	MIPS R10000 or later	IRIX 6.5.24 or IRIX 6.5.22 or later plus patch 5448
Silicon Graphics Octane2	MIPS R10000 or later	IRIX 6.5.15 or later
Silicon Graphics Tezro	MIPS R10000 or later	IRIX 6.5.15 or later
Silicon Graphics Prism	Intel Itanium 2	SGI ProPack 3 Service Pack 3

Installing OpenGL Multipipe

This chapter describes how you install OpenGL Multipipe on the following two platforms:

- “Installing on IRIX” on page 9
- “Installing on 64-Bit Linux” on page 11

Installing on IRIX

This section lists information supplemental to the guide *IRIX Admin: Software Installation and Licensing*. The information listed here is product-specific; use it with the installation guide to install this product.

The following are the prerequisites for installing OpenGL Multipipe on your system:

- Hardware: an Octane2, Onyx, Onyx2, Onyx 3000, Onyx 350, Onyx4, or Tezro system with MIPS R10000 or later processors
- Software:
 - IRIX 6.5.15 or later, generally
 - IRIX 6.5.24 or IRIX 6.5.22 or later plus patch 5448 for Onyx4 platforms
 - C++ Standard Execution Environment (c++_eoe)

Note: IRIX 6.5.18 or later is required to use the aware window management feature under SGI Xinerama. IRIX 6.5.20 or later is required to enable duplicate cursor images in screen overlap regions under SGI Xinerama.

To install OpenGL Multipipe, follow these steps:

1. Go to the following URL:
<http://www.sgi.com/software/multipipe/>
2. Click on the **Download** button and follow the instructions to download OpenGL Multipipe.

This installer includes the OpenGL Multipipe libraries and tools described in Table 2-1.

3. Use `inst` or `swmgr` to install OpenGL Multipipe.

The libraries are provided in two versions:

`n32` The new 32-bit libraries. Located in the `/usr/lib32` directory. Usable only on IRIX 6.2 and later operating system releases. The `n32` libraries operate at increased efficiency in many situations.

`64` The 64-bit libraries. Located in the `/usr/lib64` directory.

This step installs the file subsystems shown in Table 2-1.

Table 2-1 File Subsystems for OpenGL Multipipe 2.4.1

Subsystem	Description
<code>omp_eoe.sw.base</code>	Contains the actual OpenGL Multipipe binaries as well as the script for starting OpenGL applications in OpenGL Multipipe mode.
<code>omp_eoe.sw64.base</code>	Contains the OpenGL Multipipe 64-bit software.
<code>omp_eoe.man.relnotes</code>	Contains the release notes, located in the following directory: <code>/usr/share/omp/release_notes/user</code>
<code>omp_eoe.man.base</code>	Contains man pages and other documentation located in the <code>/usr/share/omp/doc/user</code> directory.
<code>omp_eoe.sw.wm</code>	Contains the aware window manager for multipipe-aware window support under SGI Xinerama.
<code>omp_eoe.sw.wmp</code>	Contains the aware window manager proxy for multipipe-aware window support under DMX.
<code>omp_eoe.sw.xdmx</code>	Contains the DMX proxy server, <code>Xdmx</code> , configuration utilities, and documentation related to the DMX X proxy layer.

Table 2-1 File Subsystems for OpenGL Multipipe 2.4.1 (continued)

Subsystem	Description
<code>omp_eoe.sw.mpc</code>	Contains libraries, header files, and example code for integrating multipipe applications with the X proxy layer.
<code>omp_eoe.sw64.mpc</code>	Contains 64-bit libraries for integrating multipipe applications with the X proxy layer.
<code>omp_eoe.man.mpc</code>	Contains documentation related to the OpenGL Multipipe MPC API, which is located in the directory <code>/usr/share/MPC/doc/developer</code> .

For any critical updated information, you can check the release notes, `/usr/share/omp/release_notes/user/relnotes.html`, and the OpenGL Multipipe website, <http://www.sgi.com/software/multipipe/>.

Installing on 64-Bit Linux

The following are the prerequisites for installing OpenGL Multipipe on your system:

- Hardware: Silicon Graphics Prism visualization system
- Software: SGI Advanced Linux Environment with SGI ProPack 3 Service Pack 3

OpenGL Multipipe for Silicon Graphics Prism systems is installed by default when installing the SGI ProPack package. OpenGL Multipipe contains the following file subsystem:

```
sgi-omp-2.x-yyyy.ia64.rpm
```

Detailed documentation of the `rpm` utility for managing software installation on Linux is available in the `rpm(8)` man page.

For any critical updated information, you can check the release notes, `/usr/share/omp/release_notes/user/relnotes.html`, and the OpenGL Multipipe website, <http://www.sgi.com/software/multipipe/>.

Using OpenGL Multipipe

As described in Chapter 1, OpenGL Multipipe consists of three main components: an X proxy layer, a proxy 3D render library, and 3D render servers. This chapter describes how to effectively use these components with your graphics applications. The following sections describe the pertinent tasks:

- “Setting up the OpenGL Multipipe Environment” on page 13
- “Running Applications with OpenGL Multipipe” on page 22
- “Performance Features” on page 27
- “Using SGI Scalable Graphics Hardware with OpenGL Multipipe” on page 35
- “Managing Windows for Aware Applications” on page 40

For information about other features of OpenGL Multipipe specific to this release, see the release notes in the following file:

`/usr/share/omp/release_notes/user/relnotes.html`

Setting up the OpenGL Multipipe Environment

To begin using OpenGL Multipipe, you must enable an X proxy layer. This will cause all applications to see a single logical pipe. To deactivate OpenGL Multipipe, just disable the X proxy layer. Some of the steps required to enable or disable OpenGL Multipipe may require `root` access. This section notes this requirement in the applicable steps.

This section describes the following tasks:

- “Configuring OpenGL Multipipe with DMX as the X Proxy Layer” on page 14
- “Configuring OpenGL Multipipe with SGI Xinerama as the X Proxy Layer” on page 20
- “Verifying That the OpenGL Multipipe Environment is Enabled” on page 21
- “Disabling the OpenGL Multipipe Environment” on page 22

Configuring OpenGL Multipipe with DMX as the X Proxy Layer

DMX will group multiple screens into a logical display. This section describes how you initialize DMX and how to create DMX configuration files.

Initializing DMX

To initialize DMX, do the following:

1. If your system supports SGI Xinerama, ensure that it is not currently enabled.

To determine if SGI Xinerama is enabled, see the section “Verifying That the OpenGL Multipipe Environment is Enabled” on page 21. If enabled, disable SGI Xinerama and restart the X server.

Enter the following to disable SGI Xinerama:

```
# chkconfig xinerama off
```

Enter the following as `root` in a command shell to restart the X server:

```
# (/usr/gfx/stopgfx; /usr/gfx/gfxinit; /usr/gfx/startgfx) &
```

The X server has to be restarted for the `chkconfig` change to take effect. With SGI Xinerama disabled, the X server will manage pipes as separate screens.

2. Run DMX on top of the existing X server(s).

You may do this manually after logging into your desktop or you may configure an `.xsession` script to run DMX immediately upon login.

To manually initialize DMX, enter the following (`root` access not needed) in a command shell:

```
$ ompstartdmx
```


You can use the flag `-help` for more information about the starting options. If you specify no flags, DMX starts on top of the existing X server and will configure a single large screen that overlays the existing n screens such that screen 0 will be the leftmost and screen $n-1$ will be the rightmost. To use a different configuration, such as a vertical configuration, you must provide a DMX configuration file. The following section “Creating DMX Configuration Files” describes how to create such a configuration file.

By default on IRIX, `ompstartdmx` will run the following clients unless a session script has been specified: `4Dwm`, `toolchest`, and `winterm`. On Linux, only `xterm` will be started as the session. For more details, see the later subsection “Differences between the IRIX and Linux Versions of `ompstartdmx`”.

You can configure DMX to run automatically upon login. How you do so depends on your operating system.

Configuring DMX to Run Automatically on IRIX

To configure DMX to run automatically upon login, you need to start with an `.xsession` file in your `$HOME` directory. If you do not already have one, you may copy one of the example `.xsession` files provided from the `/usr/share/omp/examples/X11` directory, or you may copy the main system `Xsession` file from `/var/X11/xdm/Xsession` by entering the following:

```
$ cp /var/X11/xdm/Xsession $HOME/.xsession
$ chmod +w $HOME/.xsession
```

The second command ensures the file is writeable. If not already present, add the following lines at the beginning of your `$HOME/.xsession` file:

```
if [ -x /usr/bin/ompstartdmx -a -z "$SGIOMP_META_DISPLAY" ]; then
    /usr/bin/ompstartdmx -wm none -session /var/X11/xdm/Xsession
    exit
fi
```

Note that the `if` clause is necessary to prevent infinite recursion. Also, note that `-wm none` is only needed if your `.xsession` script starts a window manager by itself (which is the case if you copied the system `Xsession` file). Only one window manager can be started on a display, and without the `-wm none` argument, `ompstartdmx` would try to start a window manager by default, which would result in an error.

For more information about `.xsession` files, see the `X(1)` and `xdm(1)` man pages.

Configuring DMX to Run Automatically on Linux

This following describes how you configure DMX to run automatically upon login using GDM, KDM, or XDM.

Using GDM

At the login screen, choose one of the following from the the **Session** menu:

GNOME-OpenGL-Multipipe

KDE-OpenGL-Multipipe

Choosing one of those will start DMX with a GNOME or KDE session, respectively.

The session will use the default DMX configuration. If the display is managing only one screen, an appropriate message will be displayed.

If you want to customize the `ompstartdmx` options (for example, to use a special DMX configuration file), modify the script(s) in the directory `/etc/X11/gdm/Sessions/GNOME-OpenGL-Multipipe` or `/etc/X11/gdm/Sessions/KDE-OpenGL-Multipipe`. You can also copy and add your own session scripts in the `/etc/X11/gdm/Sessions` directory.

Using KDM

You will need to configure KDM to add more session type options to the **Session type** menu of the KDM login screen. OpenGL Multipipe does not provide any scripts or automation for doing this. You must also configure the new session type to start `ompstartdmx` or to use a session script provided by OpenGL Multipipe in directory `/etc/X11/gdm/Sessions`.

The following steps describe one way to add a new session type, to set the associated session script, and to write that script:

1. In file `/etc/X11/xdm/kdmrc`, add the label for your new session type to the comma-separated list for the variable `SessionTypes`.

The label (for example, `KDE-OpenGL-Multipipe`) will be displayed in the option menu on the login screen and passed as the first argument to the session script.

2. In file `/etc/X11/xdm/kdmrc`, set your own session script by replacing the following line with the file path to your script:

```
Session=/usr/share/config/kdm/Xsession
```

For example, you might choose the following:

```
Session=omp_kdm_session
```

3. Write your session script to look something like the following:

```
#!/bin/bash
if [ "$1" = KDE-OpenGL-Multipipe ]; then
    exec /etc/X11/gdm/Sessions/KDE-OpenGL-Multipipe
elif [ "$1" = GNOME-OpenGL-Multipipe ]; then
    exec /etc/X11/gdm/Sessions/GNOME-OpenGL-Multipipe
else
    exec /usr/share/config/kdm/Xsession $1
fi
```

Using XDM

You can add a few lines to your `~/ .xsession` script, see `ompstartdmx -help` for more details. The Linux procedure is slightly different than that of IRIX.

After DMX has initialized, you will see a new session covering all the screens. At this point, you can start using OpenGL Multipipe by running applications with `omprun` (see “Running Applications with OpenGL Multipipe” on page 22).

Differences between the IRIX and Linux Versions of `ompstartdmx`

This section lists differences (and one noteworthy similarity) in the behavior and command-line options of the `ompstartdmx` script between IRIX and Linux:

- On Linux, there is no `-wm` option. The window manager is the default window manager for GNOME or KDE. Use their standard tools to set your preferred window manager.
- On Linux, the `ompstartdmx` script will return control only when the session ends. On IRIX, it returns control once the session starts.
- On Linux, the valid options for the `-session` flag are as follows:

```
[ -session {gnome | kde | noaware | scriptname} ]
```

Option	Behavior
<code>-session unspecified</code>	Starts a single xterm as the session. The session ends when this xterm dies.
<code>gnome</code>	Starts a GNOME session with aware support if GNOME is installed; otherwise, xterm will be used.
<code>kde</code>	Starts a KDE session with aware support if KDE is installed; otherwise, xterm will be used.

<code>noaware</code>	Launches without aware window management the default session that would run with the <code>startx</code> command.
<code>scriptname</code>	Runs the script as the session. When the script returns control, the session ends. Note: The <code>ompstartdmx</code> script does not launch any window manager; the session script must do so.

- The way to end the session remains the same: logout or `Ctrl-Alt-q`.

DMX Limitations on Linux

When using a GNOME session with an aware window manager, the session will not be saved; the default session will be used at each login. This happens because the `gnome-wm` script does not pass the session ID parameter to any aware window manager, only to those which are supported by the GNOME environment (`sawfish`, `metacity`, `twm`, and their like). When using a KDE or GNOME session without aware window management, the session will be saved and restored at the next login.

Creating DMX Configuration Files

A DMX configuration file is simply a text file that describes the configuration of a virtual display, the real displays it manages, and the geometry of the virtual screen. This section provides some short examples of configuration files. These and other example configuration files may be found in the directory `/usr/share/omp/examples/dmx`.

To start DMX with one of these configurations, do the following:

1. Save the configuration to a text file with any name—for example, `updown.dmx`.
2. Invoke `ompstartdmx` with the option `-cfgfile`, as shown in the following entry:

```
$ ompstartdmx -cfgfile updown.dmx
```

It is also possible to place many configurations in a single file. In this case, you can choose one configuration from the file by specifying both the `-cfgfile` and `-cfgname` options, as shown in the following:

```
$ ompstartdmx -cfgfile allmyconfigs.dmx -cfgname updown
```

The following example configuration file specifies a vertical layout:

```
virtual updown 1280x2048 {
    display :0.0 1280x1024;
    display :0.1 1280x1024 @0x1024;
}
```

This configuration file defines a virtual screen configuration named `updown` of size 1280x2048. The virtual screen includes the following two real back-end displays:

`Display :0.0` It has a size of 1280x1024 and is located at location 0x0 in the virtual screen space.

`Display :0.1` It has a size of 1280x1024 and is located at 0x1024 in the virtual screen space.

You may also define some overlap between each of the screens, as in the following horizontal layout:

```
virtual overlap 2460x1024 {
    display :0.0 1280x1024;
    display :0.1 1280x1024 @1180x0;
}
```

This configuration file defines two screens of 1280x1024, each with 100 pixels of overlap, resulting in a virtual screen size of 2460x1024.

The display value specified can be any valid display value, including a display value that specifies a remote machine, as in the following example:

```
virtual remote 2560x1024 {
    display localhost:0.0 1280x1024;
    display remotehost:0.1 1280x1024 @1280x0;
}
```

There is also a graphical tool to create and edit configuration files. You can find documentation for this tool in `/usr/share/omp/doc/user/xdmxconfig.1.html`. The tool is installed in `/usr/share/omp/X11/bin/xdmxconfig`.

More information about the configuration file format can be found in the file `/usr/share/omp/doc/user/Xdmx.1.html`.

Configuring OpenGL Multipipe with SGI Xinerama as the X Proxy Layer

Note: SGI Xinerama is not supported on Onyx4 or Silicon Graphics Prism servers. Only the DMX proxy layer is supported on these servers.

SGI Xinerama groups multiple screens of a single X server into a logical screen. This section describes how you initialize and configure SGI Xinerama.

Initializing SGI Xinerama

To configure OpenGL Multipipe with SGI Xinerama as the X proxy layer, perform the following steps:

1. Enable SGI Xinerama.

If you are enabling SGI Xinerama on your system for the first time, enter the following as `root` in an IRIX shell:

```
# chkconfig -f xinerama on
```

Otherwise, enter the following to enable SGI Xinerama:

```
# chkconfig xinerama on
```

On systems having only one graphics pipe or in the case where the X server is directed to handle only one pipe (see the `Xsgi(1)` man page), enabling SGI Xinerama has no effect. In these cases, SGI Xinerama will be disabled, regardless of the value of the `xinerama` flag supplied on the `chkconfig` command.

2. Restart the X server.

The X server has to be restarted for the `chkconfig` change to take effect. Enter the following as `root` in an IRIX shell to restart the X server:

```
# (/usr/gfx/stopgfx; /usr/gfx/gfxinit; /usr/gfx/startgfx) &
```

After the X server is started with SGI Xinerama enabled and you have logged in to the system, you can start using OpenGL Multipipe by running applications with `omprun` (see “Running Applications with OpenGL Multipipe” on page 22).

Configuring Overlapping Screens with SGI Xinerama

Reality Center environments with multiple projectors and multiple graphics pipes often have overlapping screens. To allow seamless alignment of these screens, projectors typically have edge blending capability.

You control the amount of overlap or gap space between screens by specifying the `xoffset` and `yoffset` arguments (in units of pixels) of the `-hw` parameters in the file `/var/X11/xdm/Xservers`. For example, if you have three pipes, each with a resolution of 1280 x 1024 and a horizontal overlap of 256 pixels between each screen, you would place the following `-hw` arguments together on one line in the `/var/X11/xdm/Xservers` file:

```
:0 secure /usr/bin/X11/X
-boards 4,5,6
-hw board=4,right=1
-hw board=5,left=0,right=2,xoffset=-256
-hw board=6,left=1,xoffset=-256
... other X server arguments ...
```

Notes:

- The `-boards` numbers are physical pipe numbers, but the indexes given to the `right`, `left`, `above`, and `below` parameters refer to the logical order of the `-hw` parameters.
- The lines are separated in the example only for readability.

See the `Xsgi(1)` and `xdm(1)` man pages for more information.

Verifying That the OpenGL Multipipe Environment is Enabled

The OpenGL Multipipe environment is enabled if an X proxy layer is enabled (DMX or SGI Xinerama).

To verify that an X proxy layer is enabled, ensure that your `DISPLAY` environment variable is pointing to the correct display (usually `:0.0` for SGI Xinerama or `:1.0` for DMX) and enter the following commands in a command shell:

```
$ xdpinfo | grep SGI-XINERAMA
```

If `SGI-XINERAMA` appears as the result of the prior commands, SGI Xinerama is enabled.

```
$ xdpinfo | grep DMX
```

If DMX appears as the result of the prior commands, DMX is enabled.

Disabling the OpenGL Multipipe Environment

To disable the OpenGL Multipipe environment, do the following:

- If enabled, disable DMX.

To end the DMX session, select **Desktop > Logout** from the Toolchest or run `/usr/bin/X11/endsession`. You may also force the DMX server to exit by pressing `Ctrl+Alt+q`.

After the session has ended, you will be returned to your regular X session.

If you have configured DMX to run automatically upon login and you end the DMX session, it will automatically end the regular X session as well and you will return to the login screen. To permanently disable DMX from starting upon login, revert the DMX-related changes you made to your `.xsession` file, or delete or rename your `$HOME/.xsession` file.

- If enabled, disable SGI Xinerama and restart the X server.

Enter the following to disable SGI Xinerama:

```
# chkconfig xinerama off
```

Enter the following as `root` in an IRIX shell to restart the X server:

```
# (/usr/gfx/stopgfx; /usr/gfx/gfxinit; /usr/gfx/startgfx) &
```

The X server has to be restarted for the `chkconfig` change to take effect.

Running Applications with OpenGL Multipipe

Plain X applications will generally run under an X proxy layer without assistance. OpenGL (3D) applications need to use the OpenGL Multipipe 3D proxy library to handle 3D rendering correctly and efficiently on all screens.

To use the OpenGL Multipipe 3D proxy library with OpenGL applications, just run the program using the `omprun` script:

```
$ omprun app_name app_args
```


This will preload the OpenGL Multipipe proxy libraries to intercept OpenGL calls.

The following is an example:

```
$ omprun ivview /usr/share/data/models/Banana.iv
```

Note: Failure to use the `omprun` command under DMX will cause the application to use the slower GLX indirect rendering support in DMX to draw OpenGL to all screens.

The `omprun` script causes an OpenGL application to use the intermediate 3D proxy libraries instead of the normal OpenGL library. This enables the OpenGL application to behave correctly when its windows are moved across parts of the logical screen. Such an application is considered to be started in multipipe-unaware mode (or simply, *unaware* mode), since it is not aware of the underlying graphics hardware structure.

Technically, the `omprun` script sets the link editor environment variable (`_RLDN32_LIST` or `_RLD64_LIST` on IRIX and `LD_PRELOAD` on Linux) to use the `libOMPmaster.so` library of matching format prior to using the `libGL.so` library.

For more information on using `omprun`, see the `omprun(1)` man page or use the `-help` command-line option of `omprun` as follows:

```
$ omprun -help
```

The following sections describe how to best run different types of graphics applications:

- “Running OpenGL Single-Pipe Applications”
- “Running Pure X Applications”
- “Running IRIS GL Applications”
- “Running IRIX o32 Applications”
- “Running Multipipe Applications in Multipipe-Aware Mode”

For more information on running applications with OpenGL Multipipe, see the release notes, `/usr/share/omp/release_notes/user/relnotes.html`.

Running OpenGL Single-Pipe Applications

OpenGL single-pipe applications are the targeted applications for OpenGL Multipipe. To run such applications, simply enable the OpenGL Multipipe environment and invoke the application using the `omprun` script.

Under SGI Xinerama, any OpenGL application started without the `omprun` script will not behave correctly. In that case, OpenGL drawings will appear only in the part of the window overlapping screen 0. On the other screens, the application will display a random image that corresponds to the current content of the framebuffer on that pipe.

Under DMX, OpenGL applications started without the `omprun` script will display correctly on all screens, using the GLX indirect rendering support in DMX. However, using the `omprun` script will provide better performance for OpenGL applications.

Hint: For an easy way to run a number of single-pipe OpenGL applications under OpenGL Multipipe without the need to always explicitly invoke `omprun`, start a command shell under `omprun`, as shown in the following :

```
$ omprun xwsh
<omprun xwsh>$ app_name app_args
```

Any application started from this new command shell will be started automatically in transparent multipipe mode.

Running Pure X Applications

As noted in Chapter 1, “OpenGL Multipipe Overview”, the X proxy layer enables pure X applications to run transparently over multiple pipes. To run pure X applications, simply enable SGI Xinerama or DMX before invoking them and they will run correctly. You do not need to use the `omprun` script for these applications.

Running IRIS GL Applications

On IRIX, there are applications that use the older IRIS GL graphics library instead of that of OpenGL. OpenGL Multipipe does not support IRIS GL. To check whether your current application is attempting to use IRIS GL, enter the following:

```
$ elfdump -D1 app_name | grep libgl.so
```

The following is an example:

```
$ elfdump -D1 /usr/sbin/showcase | grep libgl.so
[11]   Jun  6 22:31:51 1996   0xe9155925 ----- libgl.so   sgi1.0
```

The `omprun` script does this check and prevents OpenGL Multipipe from executing IRIS GL applications.

If your system supports IRIS GL, you can still run IRIS GL applications, but not using the OpenGL Multipipe `omprun` layer. Under SGI Xinerama, they will run correctly only on screen 0. Also, IRIS GL applications will run correctly in multipipe-aware mode, which is described in the subsequent section “Running Multipipe Applications in Multipipe-Aware Mode”.

Only when DMX is used as the X proxy layer, will IRIS GL applications run correctly on all screens without using the `omprun` script. This happens through the GLX indirect rendering support in DMX. Consequently, performance decreases.

Running IRIX o32 Applications

On IRIX, there are applications that use the older o32 application binary interface (ABI) instead of the newer n32 or 64-bit ABIs. OpenGL Multipipe does not support applications that were built using the o32 ABI. To check whether your current application was built with the o32 ABI, enter the following:

```
$ file app_name | grep 32-bit
```

If the text `ELF 32-bit ...` is printed as a result, it is an o32 application.

The following is an example:

```
$ file /usr/sbin/showcase | grep 32-bit
/usr/sbin/showcase: ELF 32-bit MSB mips-2 dynamic executable MIPS -
version 1
```

The `omprun` script does this check and prevents OpenGL Multipipe from executing o32 applications.

If your system supports the o32 ABI, you can still run o32 applications, but not using the OpenGL Multipipe `omprun` layer. Under SGI Xinerama, they will run correctly only on screen 0. Also, o32 applications will run correctly in multipipe-aware mode, which is

described in the subsequent section “Running Multipipe Applications in Multipipe-Aware Mode”.

Only when you use DMX as the X proxy layer, will o32 applications run correctly on all screens without using the `omprun` script. This happens through the GLX indirect rendering support in DMX. Consequently, performance decreases.

Running Multipipe Applications in Multipipe-Aware Mode

Multipipe applications are intentionally written to take advantage of systems that have multiple graphics pipes. If they know about the underlying graphics hardware, they can explicitly address and take advantage of the individual graphics pipes. Typically, multipipe applications are written using OpenGL Performer or OpenGL Multipipe SDK.

It is best not to run such applications under OpenGL Multipipe, which hides the hardware configuration of the system from the applications. To run at full potential, these applications should be aware of the different graphics pipes in the system. To allow such applications to see the real graphics hardware does not require you to disable OpenGL Multipipe.

The OpenGL Multipipe layer may be bypassed on a per-application basis. This allows a multipipe application to be fully aware of the multipipe environment while other applications, aware of only a single large pipe, continue to run under OpenGL Multipipe. Applications that bypass the OpenGL Multipipe layer are said to run in multipipe-aware mode (or simply, *aware mode*), because they are aware of the different graphics pipes handled by the X server.

To run a multipipe application in aware mode while other single-pipe applications run concurrently in unaware mode, set the `DISPLAY` environment variable to point to the desired backend display that is managed by the X proxy layer (for example, `:0.1`), and then start the multipipe application with the `-aware` command-line option of the `omprun` script, as in the following example:

```
$ setenv DISPLAY :0.0
$ omprun -aware perfly
```

Under DMX, it is especially important to set the `DISPLAY` environment variable because the DMX meta display has a completely different display name than its component backend displays. By default, the DMX display is `:1` and the backend aware display is `:0`. In the case of SGI Xinerama, the SGI Xinerama meta display and its component backend displays are referred to by the same display name (for example, `:0`).

Note: Under SGI Xinerama, applications started in aware mode will be under window manager control only with `omp4Dwm`. Under DMX, other window managers may be used. See the subsequent section “Managing Windows for Aware Applications” for more information about aware window management.

Hint: For an easy way to run a number of commands in aware mode, start a command shell in aware mode.

```
$ setenv DISPLAY :0.0
$ omprun -aware xwsh
<aware xwsh>$ app_name app_args
```

Any application started from this new command shell will be started automatically in aware mode.

Performance Features

OpenGL Multipipe has several methods of optimizing OpenGL applications to work correctly and efficiently across multiple pipes. This section provides an overview of the following performance features and some guidelines for their use:

- “Viewport Clipping” on page 27
- “Geometry Culling” on page 28
- “Small Object Culling” on page 29
- “Display List Partitioning” on page 29
- “Master Rendering Modes” on page 29
- “Frame Latency Control” on page 33
- “Buffer Swap Synchronization” on page 34

The release notes provide a more technical discussion of each of these features.

Viewport Clipping

Applications that draw large polygons with complex texturing or shading procedures are likely to be fill-limited—that is, the rasterization stage of the graphics pipeline is the

bottleneck to improving performance. If slower performance results in proportion to an increase in the OpenGL window size, this is an indicator that fill performance could be the problem.

To eliminate the pixel fill bottleneck, polygon rasterization work can be divided among multiple graphics pipes. Using OpenGL Multipipe, this can be accomplished by simply positioning a window so that it spans multiple graphics pipes and each pipe performs an equal fraction of the rasterization work. On each component screen of the logical display, OpenGL Multipipe automatically clips the OpenGL viewport to the physical screen boundaries.

Viewport clipping is enabled by default. It can be disabled with the `omprun -novpclip` option.

The fill performance benefits of viewport clipping can be more fully realized by using an SGI Scalable Graphics Compositor and specifying additional parameters to OpenGL Multipipe. For information on these parameters, see the section “Specifying Static Composited Regions for Fully Overlapped Screens” on page 38.

Geometry Culling

Applications that render large amounts of geometry in display lists can sometimes reach the limit of the polygon processing capabilities of the graphics hardware. Such an application is said to be transform-limited or geometry-limited—that is, the geometry transformation stage of the graphics pipeline is the bottleneck to improving performance. The geometry limit varies for each graphics architecture. For example, an Onyx 3000 pipe can handle millions of polygons per second; an Onyx4 pipe can handle hundreds of millions of polygons per second. If performance remains the same when lighting or textures are disabled by the application, these are indicators that geometry performance is the limiting factor.

To eliminate the geometry transformation bottleneck, OpenGL Multipipe can divide geometry among multiple pipes. By default, OpenGL Multipipe renders all geometry on each graphics pipe, even if not all of the geometry is visible on a given pipe. When geometry culling is enabled, each OpenGL Multipipe slave process renders only the geometry from display lists, vertex arrays, and immediate mode commands that is visible on its pipe. It is also possible for OpenGL Multipipe to cull geometry to user-specified OpenGL clip planes.

This feature is enabled with the `omprun -cull` command-line option. Other options related to geometry culling are described in the `omprun -help` text and in the Usage Tips and Tricks section of the release notes.

Small Object Culling

OpenGL Multipipe can cull objects whose screen-space bounding box is less than a threshold number of pixels. You can specify the threshold with the `-minpixels` command-line option for the `omprun` script. Geometry culling must be turned on in order for this feature to take effect.

This feature can increase the rendering speed in applications that have many features that are too small to discern and, therefore, may not be of immediate interest—for example, screws in CAD/CAM applications. With a proper `minpixels` culling threshold, OpenGL Multipipe culls such small features and rendering performance is improved as a result.

Display List Partitioning

When rendering display lists with OpenGL Multipipe's geometry culling option enabled, OpenGL Multipipe culls or renders an entire display list as a unit. When the original display list has a large amount of geometry and spans large areas of the scene, performance scalability can suffer.

OpenGL Multipipe can optionally spatially divide user display lists to break them into smaller ones. After the division, the smaller, more spatially coherent pieces are more friendly to load balancing and display list culling.

This feature is enabled with the `omprun -dlsplit` command-line option. Other options related to display list partitioning are described in the `omprun -help` text and in the Usage Tips and Tricks section of the release notes.

Master Rendering Modes

As cited earlier, the OpenGL proxy layer of OpenGL Multipipe has two components: a master render library that intercepts OpenGL calls made by the application and render slave processes that each receive a stream of OpenGL calls from the master and perform OpenGL rendering on the application's behalf.

The master render library functions as part of the application process (that is, the “master process”) and can have additional responsibilities besides intercepting, packing, and distributing OpenGL calls to the slave processes. The master process may also render directly to a single local pipe in place of a single slave process, or it may use a local pipe only to track OpenGL state while a slave process renders to that pipe.

The `omprun -mstrmode` option allows you to specify what functions the master component performs on the single local reference pipe. The master’s mode may improve or hinder OpenGL performance depending upon the behavior of a particular application. Therefore, it is important to understand the implications of each mode.

The following master modes are available:

- `-mstrmode off`
- `-mstrmode track`
- `-mstrmode render`

`-mstrmode off` **Mode**

The `-mstrmode off` mode is most efficient for applications that do not perform `glGetxxx()` calls (GL state queries) in every frame, because in this mode querying GL state requires round-trip communication to a slave. The master process does not render; it only packs and distributes GL calls to all slave processes. A slave process renders to each pipe and one of the slaves is designated to track state for any occasional `glGetxxx()` or `glIsxxx()` queries. Figure 3-1 illustrates running in `-mstrmode off` mode.

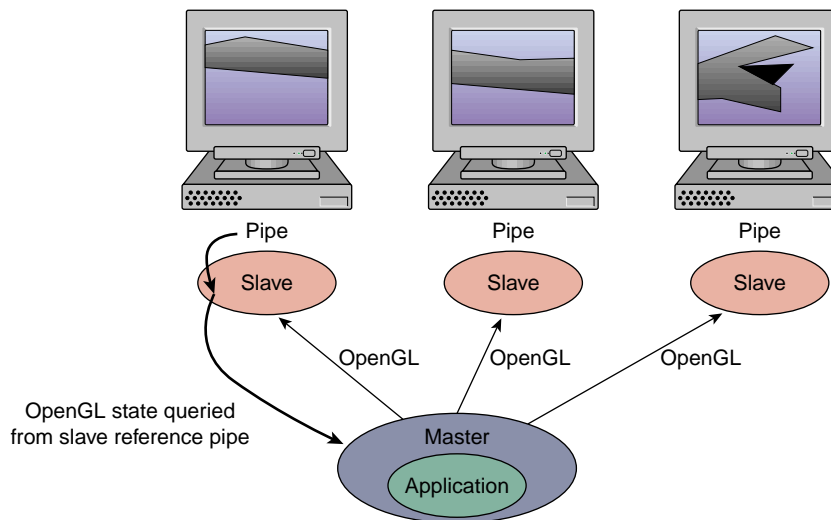


Figure 3-1 Running in `-mstrmode off` Mode

`-mstrmode track` Mode

The `-mstrmode track` mode is most efficient for applications that frequently query GL state. The master process does not render but, in addition to packing and sending GL calls to all slave processes, it tracks the GL state on a local reference pipe. Slave processes render on all pipes. Figure 3-2 illustrates running in `-mstrmode track` mode.

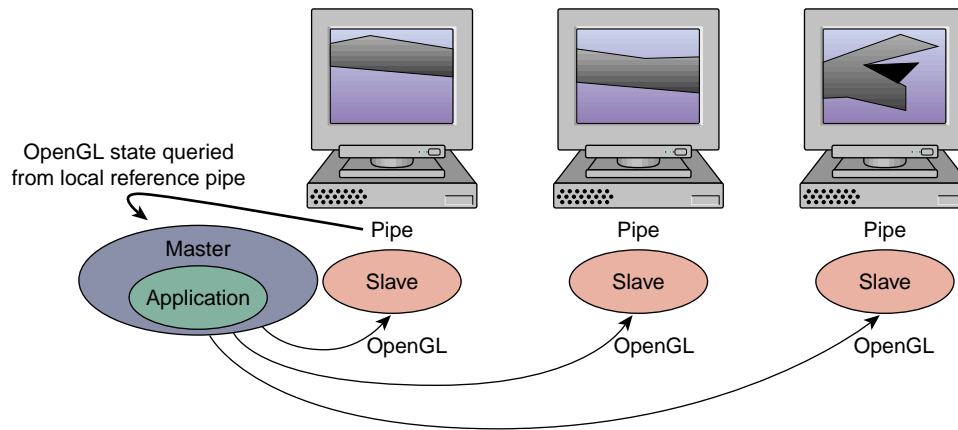


Figure 3-2 Running in `-mstrmode track` Mode

`-mstrmode render` Mode

The `-mstrmode render` mode may yield slightly better performance for applications that do not use display lists and that run on systems with only two graphics pipes or with a limited number of processors. The master process renders and tracks GL state on a local reference pipe. One less slave process is needed because the application (master) process renders itself. State queries again are made to the master's local reference pipe. Figure 3-3 illustrates running in `-mstrmode render` mode.

Note: Some of the performance features described in this section, including geometry culling, are not available in `-mstrmode render` mode.

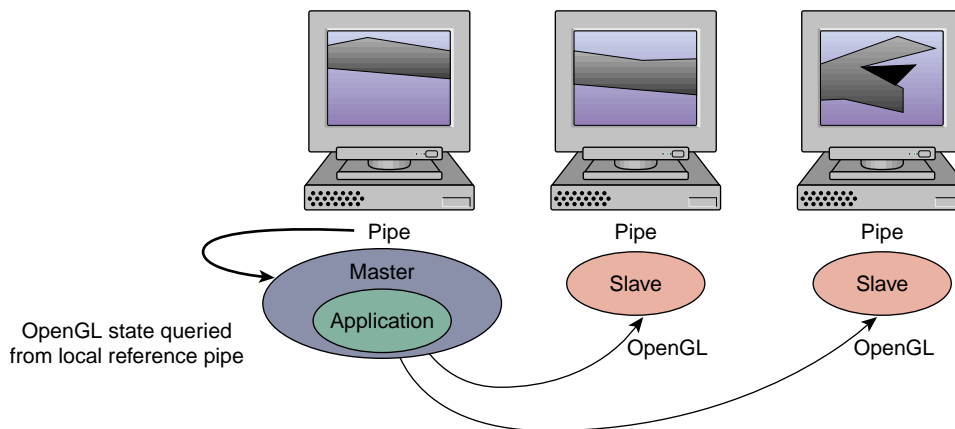


Figure 3-3 Running in `-mstrmode` render Mode

Frame Latency Control

OpenGL Multipipe uses a shared-memory buffer in between the application and the drawing slaves. This buffer can introduce latency—that is, multiple frames can be buffered to be consumed gradually by the slave. If the application does not call `glFinish()` by itself, then OpenGL Multipipe allows the number of buffered frames to reach a small preset limit.

The latency helps smooth out application and drawing speed differences, and thereby increase throughput. However, if the amount of latency is beyond what you can accept, it can be limited by using the `omprun -latency` command-line argument to specify the maximum latency in number of frames. For more information related to frame latency control, see the `omprun -help` text and the Usage Tips and Tricks section of the release notes.

The following command forces the master and slaves to have zero frames of latency between them:

```
omprun -latency 0 app_name app_args
```

The master will wait for the slaves to finish executing enough of the previous frames so that the latency is below the threshold before beginning to pack data into the buffer for the next frame.

Buffer Swap Synchronization

Variations in pixel fill, geometry load, and many other factors can lead to an unbalanced load among the graphics pipes. Some pipes will render their parts of the scene faster or more slowly than the rest. Synchronization among the pipes is required to prevent one pipe from rendering faster or slower than another, which in some cases can present visible “tearing” in the output image.

This section describes the following synchronization schemes:

- “Software Swap Synchronization”
- “Hardware Buffer Swap Synchronization”

Software Swap Synchronization

By default, OpenGL Multipipe performs a software synchronization among the slave processes to ensure that they issue their respective swap-buffer commands at the same time. The software synchronization approximates a synchronized swap in hardware.

Software swap synchronization is enabled by default. It can be disabled with the `omprun -nosync` command-line option. Note that this also disables any meaningful sense of frame latency.

Hardware Buffer Swap Synchronization

Note: OpenGL Multipipe does not support this feature on Onyx4 and Silicon Graphics Prism servers.

OpenGL Multipipe supports hardware synchronization of `glXSwapBuffers()` across all pipes. Normally, when an application that is run with `omprun` makes a call to `glXSwapBuffers()`, OpenGL Multipipe sends swap-buffer requests to **all** pipes since the application window might be visible on all pipes. When multiple pipes are used to drive a large logical screen (that is, a wall display), ensuring that the actual buffer swaps happen at exactly the same time on every pipe improves the perception that the display is a single logical pipe.

To use Swap Ready hardware synchronization, follow these steps:

1. Before using the `-swapready` option, make sure you have properly connected Genlock and Swap Ready cables to all pipes and that the pipes are configured to be genlocked.

Running any OpenGL application that attempts to use Swap Ready hardware without proper configuration can cause a serious graphics failure. This includes all applications started with `omprun -swapready`. For information about wiring the Genlock and Swap Ready cables, see the *POWER Onyx and Onyx Rackmount Owner's Guide*. Also, see the `genlock(1)` man page.

2. When running the application with `omprun`, use the `omprun -swapready` flag:

```
$ omprun -swapready app_name app_args
```

Using Swap Ready implicitly disables software swap synchronization. If you are not using Swap Ready hardware to synchronize `glXSwapBuffers()` calls, then OpenGL Multipipe uses a software synchronization that is less accurate.

Note: More than one application may be started with hardware swap synchronization using the `omprun -swapready` option. However, multipipe applications that support Swap Ready natively will conflict with OpenGL Multipipe if an application that was started with `omprun -swapready` is running. Likewise, if a multipipe application is already using the Swap Ready line, the `omprun -swapready` option will revert to software swap synchronization.

Using SGI Scalable Graphics Hardware with OpenGL Multipipe

You may configure SGI Scalable Graphics Compositor hardware for use with OpenGL Multipipe to improve geometry and fill performance for an application. This requires no changes to the application. Using the following topics, this section describes how to configure SGI Xinerama or DMX for this purpose as well as settings for OpenGL Multipipe to improve performance and usability in composited logical screen mode:

- “Configuring Composited Screens with SGI Xinerama” on page 36
- “Configuring Composited Screens with DMX” on page 36
- “Specifying Static Composited Regions with `ompstartdmx`” on page 37
- “Specifying Static Composited Regions for Fully Overlapped Screens” on page 38

- “Enabling Duplicate Cursor Images in Overlap Regions” on page 38

Configuring Composited Screens with SGI Xinerama

Composited screens are a special case of overlapped screens in which each compositor input screen completely overlaps the other compositor input screens. That is, each compositor input pipelet displays the same area of the logical screen.

To configure totally overlapped pipes, specify negative `xoffset` and `yoffset` parameters that equal the width (and height) of the screens you are overlapping. For example, if you have four pipelets, each with a resolution of 1280 x 1024, connected to a single compositor, you would place the following `-hw` arguments together on one line in the `/var/X11/xdm/Xservers` file:

```
:0 secure /usr/bin/X11/X
-hw boards 4,5,6,7
-hw board=4,right=1
-hw board=5,left=0,right=2,xoffset=-1280
-hw board=6,left=1,right=3,xoffset=-1280
-hw board=7,left=2,xoffset=-1280
... other X server arguments ...
```

Notes:

- The `-boards` numbers are physical pipe numbers, but the indexes given to the `right`, `left`, `above`, and `below` parameters refer to the logical order of the `-hw` parameters.
- The lines are separated in the example only for readability.

Configuring Composited Screens with DMX

To configure completely overlapped screens under DMX, simply create a DMX configuration file to manage the screens of the backend X server in the desired order. Do not specify an offset or a `0x0` offset after the screen specifications. The following is an example configuration file:

```
virtual totaloverlap 1280x1024 {
    display :0.0 1280x1024;
    display :0.1 1280x1024;
    display :0.2 1280x1024 @0x0; # "@0x0" is optional
    display :0.3 1280x1024;
```

```
}
```

For more information on DMX configuration files, see section “Creating DMX Configuration Files” on page 18.

Specifying Static Composited Regions with `ompstartdmx`

You can specify a static tiling mode for a hardware compositor during DMX startup. The `ompstartdmx` script will start both the hardware compositor and DMX with the specified tiling mode.

The `-compositor` option on the `ompstartdmx` script allows you to specify the desired static tiling mode, as shown in the following example:

```
$ ompstartdmx -compositor mode
```

The following values are available for *mode*:

Mode	Description
<code>quad</code>	A quad tiling
<code>vert2</code>	Two vertical stripes
<code>vert3</code>	Three vertical stripes
<code>vert4</code>	Four vertical stripes
<code>horiz2</code>	Two horizontal stripes
<code>horiz3</code>	Three horizontal stripes
<code>horiz4</code>	Four horizontal stripes

You can also set up the video format using the `-compositor` option, as shown in the following entry:

```
$ ompstartdmx -compositor quad,1280x1024_75
```

The video format must be supported by both the hardware compositor and the pipes, and all pipes must have the same video format. Otherwise, `ompstartdmx` exits with an error. The default video format is the current one for the pipes.

Specifying Static Composited Regions for Fully Overlapped Screens

To establish static composited regions, do the following:

1. Set up fully overlapping screens with SGI Xinerama or DMX.

The sections “Configuring Composited Screens with SGI Xinerama” on page 36 and “Configuring Composited Screens with DMX” on page 36 describe how you do this.

2. Configure the SGI Scalable Graphics Compositor hardware and OpenGL Multipipe to constrain drawing to areas of the physical pipes that will provide the best static load balancing for your application with `sgcombine`.

Note that `sgcombine` must be run in multipipe-aware mode:

```
$ env DISPLAY=:0.0 omprun -aware /usr/gfx/sgcombine
```

For more information about setting up composited screens with `sgcombine`, see the *SGI InfinitePerformance: Scalable Graphics Compositor User's Guide*.

3. Configure OpenGL Multipipe to match the static compositor configuration set with `sgcombine`.

For example, if you specify a compositor tiling of four input rectangles of 640x512 pixels in `sgcombine`, setting the environment variable `SGIOMP_SCREEN_RECTS` to the following string provides the matching settings to OpenGL Multipipe:

```
"640x512+0+0 640x512+640+0 640x512+0+512 640x512+640+512"
```

OpenGL Multipipe will then clip drawing to these subregions on backend screens `:0.0`, `:0.1`, `:0.2`, and `:0.3` (pipes 4, 5, 6, and 7), respectively, if DMX or SGI Xinerama are configured as in the previous sections. The rectangles are specified in pipe-relative coordinates, one per backend X screen, using the format described in the `XParseGeometry(3X11)` man page. Extra rectangles are ignored and screens for which a rectangle is not specified will have clipping performed at the screen borders.

Enabling the `omprun -cull` option will cause geometry to be culled to these areas as well.

Enabling Duplicate Cursor Images in Overlap Regions

When an X proxy layer is used to overlap screen regions on an edge-blended display or a compositor-based system, the cursor will seem to disappear when it enters the overlapped or uncomposited regions of the display. In X, a mouse belongs to one screen

of the X server at a time. Therefore, it is normally not possible to draw the mouse multiple times (on different screens) in the overlap region.

To prevent the cursor from disappearing in these cases, you need to create additional or *duplicate cursor* images (not real cursors) where two or more screens overlap. There is still just one real cursor position on the display. The way you do this depends on your X proxy layer, SGI Xinerama or DMX.

Under SGI Xinerama

In IRIX 6.5.20 or later, you enable duplicate cursors by adding the `-phantomcursors` flag to the X server command line in the `/var/X11/xdm/Xservers` file. For more information about the `-phantomcursors` option, see the `Xsgi(1)` man page.

Under DMX

There are two ways to enable duplicate cursors under DMX:

- Managing an appropriate subarea of each composited input pipe instead of the whole screen area
- Managing multiple backend X servers

Managing Screen Subregions with DMX

You can use a DMX configuration file to cause DMX to manage less than a whole backend screen. When the subregions of each screen are visually assembled using a congruent static compositor tiling, the resulting logical screen will look and behave identically to the DMX or SGI Xinerama displays with fully overlapping screens described previously. The following configuration file entry illustrates this feature:

```
virtual quadtilesubregions 1280x1024 {
    display :0.0 640x512+0+0      @0x0;
    display :0.1 640x512+640+0   @640x0;
    display :0.2 640x512+0+512   @0x512;
    display :0.3 640x512+640+512 @640x512;
}
```

Note: The `SGIOMP_SCREEN_RECTS` environment variable, described in section “Specifying Static Composited Regions for Fully Overlapped Screens” on page 38, is not necessary here and should not be used with a DMX configuration that manages screen subregions.

Managing Multiple Backend X Servers with DMX

The X mouse cursor can be made visible in the partially or fully overlapped screen regions needed for edge-blended or composited displays. As long as the logically overlapping areas of the DMX display are composed by screens of different backend X servers, DMX can utilize the mouse cursor that is available on each backend X server.

Instead of having DMX manage multiple screens of a single backend X server, you can start multiple backend X servers and specify a DMX configuration that logically overlaps areas from different backend X servers, as shown in the following:

```
virtual totaloverlapmulti 1280x1024 {
    display :0.0 1280x1024;
    display :1.0 1280x1024;
    display :2.0 1280x1024;
    display :3.0 1280x1024;
}
```

Note: The multiserver technique can be combined with the technique of managing screen subregions to create a logical display on which mouse cursors are visible in all overlapped and composited regions.

Managing Windows for Aware Applications

To allow window manager support for applications started in aware mode, OpenGL Multipipe provides aware window management. Under SGI Xinerama, the window manager `omp4Dwm`, a specialized version of the SGI standard window manager (`4Dwm`), is used to manage aware windows. Under DMX, a window manager wrapper is provided that is much like `omprun` for GL applications.

When `omp4Dwm` or the window manager wrapper is not used, applications started in aware mode (using `omprun -aware app_name`) will bypass the window manager. This

means that their windows cannot be moved, resized, iconified, or otherwise managed. This includes the regular decoration provided by the window manager. The windows will not have this decoration. This occurs because an unaware window manager does not know about all of the real screens that the X proxy layer is managing. It cannot manage aware windows that were not created through X proxy layer.

If you are using DMX as your X proxy layer and you invoke `ompstartdmx` with the standard arguments (or if DMX is configured to start automatically when you log in), a window manager of your choice will be automatically started with the window manager wrapper so that it is able to manage aware windows. You can change the DMX default window manager by using the `ompstartdmx -wm` option. You may follow the steps in the following subsections if you want to run a different window manager to manage aware windows.

This section describes the following topics:

- “Starting an Aware Window Manager”
- “Exiting an Aware Window Manager”
- “Setting an Aware Window Manager as the Default”

Note: The multipipe-aware window manager is currently not supported for compositor-based systems.

Starting an Aware Window Manager

To start an aware window manager, perform the following steps:

1. Exit or kill any unaware window manager that is currently managing the display.

If you are using `4Dwm` (the default window manager on IRIX), enter the following in a command shell:

```
$ tellwm quit
```

Otherwise, if possible, exit your window manager without logging out. One way to do this is to find the process number for your window manager and kill it manually, as the following illustrates:

```
$ ps -e | grep my_window_manager
 23878 ? 0:42 my_window_manager
$ kill 23878
```

Some window managers may not allow you to exit the window manager and remain logged in. If this is the case, you will need to start the aware window manager from a `.xsession` file. See the section “Setting an Aware Window Manager as the Default” on page 43 for more information.

2. Start the specialized window manager.

Under SGI Xinerama, enter the following:

```
$ start_ompwm
```

Under DMX, enter the following:

```
$ ompwrapwm twm
```

The `start_ompwm` script starts `omp4Dwm` after first checking if the display server supports SGI Xinerama. The `ompwrapwm` script starts the window manager `twm` in aware window management mode under DMX. If the display server is determined to be compatible, the script starts the window manager with aware window management support enabled. If the display server is not compatible, the script will exit. The script can be made to start the window manager in unaware mode (with aware window management disabled) as a contingency.

For more information on using the `start_ompwm` and `ompwrapwm` scripts, see their man pages or use the `-help` command-line option of the scripts as follows:

```
$ start_ompwm -help
```

or

```
$ ompwrapwm -help
```

Notes:

- You can use any window manager with the `ompwrapwm` script or with the `ompstartdmx -wm` option on IRIX. Only `4Dwm` and `twm` are supported on IRIX. GNOME is recommended on Linux platforms.
- Starting an application in aware mode and then starting the window manager will result in the application’s windows being unmanaged by the aware window manager. An aware window manager must be started prior to running an application in aware mode in order for its windows to be managed.

Exiting an Aware Window Manager

To exit an aware window manager, simply log out and log back in. The default window manager will again manage your display.

If you are running `omp4Dwm` under SGI Xinerama, you may also exit `omp4Dwm` by entering the following:

```
$ tellwm quit
```

Then start your original window manager.

Setting an Aware Window Manager as the Default

An alternate way to run an aware window manager is to invoke the `start_ompwm` or `ompwrapwm my_window_mgr` script in your `$HOME/.xsession` file. Which script should be invoked (`start_ompwm` or `ompwrapwm`) depends respectively on whether SGI Xinerama or DMX is running.

For IRIX users, the directory `/usr/share/omp/examples/X11/` contains some example `.xsession` files. For more information about `.xsession` files, see the man pages `X(1)` and `xdm(1)`.

For Linux users, you must set up a default aware window manager under DMX. See section “Initializing DMX” on page 14.

Limitations

OpenGL Multipipe allows single-pipe applications to run in a multipipe environment without any modification and without the need to recompile the application. It also allows single-pipe and multipipe applications to run concurrently on the same X server. However, OpenGL Multipipe has limitations and the following sections describe them:

- “Performance Enhancement” on page 46
- “X Extensions” on page 46
- “The Multipipe-Aware Window Manager” on page 46
- “OpenGL Window Size Constraints” on page 46
- “Processor Requirements” on page 47
- “SGI ProPack and OpenGL Multipipe Versions” on page 47
- “Overlay Windows Support in DMX” on page 47
- “SGI Xinerama Not Supported on Onyx4 or Silicon Graphics Prism Platforms” on page 47
- “Graphics Pipe Requirements” on page 48
- “Some Features Not Available on Linux” on page 48

For release-dependent limitations, see the OpenGL Multipipe release notes. For supported platforms and configurations, see “Supported Platforms” in Chapter 1.

Performance Enhancement

OpenGL Multipipe does not replace performance-focused multipipe APIs—such as OpenGL Performer or OpenGL Multipipe SDK—or any other custom multipipe solution. Using OpenGL Multipipe results in some minimal overhead (performance loss) for traditional single-pipe applications. This is due to the inherent cost of distributing the X and OpenGL commands among the graphics pipes.

It should be noted that process placement—that is, the way in which processes are assigned to processors—can significantly affect performance. For the affect of process placement on performance, see the OpenGL Multipipe release notes.

X Extensions

Some X extensions are not supported by SGI Xinerama, and others are not supported by DMX. For example, SGI Xinerama does not support the `SGI-VIDEO-CONTROL` extension, which permits control of video operations on the base graphics hardware, and `SCREEN-SAVER`, which is used by some screen saver programs. Applications using these X extensions may not function properly. The behavior of these applications started in unaware mode is undefined, though they will generally behave correctly on screen 0 or in aware mode. To determine whether a particular extension is supported, see the section “X Application Uses Unsupported X Extension” on page 58.

The Multipipe-Aware Window Manager

Due to the nature of the screen overlapping required for composited displays, the aware window manager is currently limited to managing aware windows on noncomposited displays only. Unaware windows will be managed properly on composited displays.

OpenGL Window Size Constraints

The hardware graphics pipes have a hardware-dependent limit on the size of the region into which an OpenGL application renders. The consequence is that an OpenGL application is constrained to draw into a limited area. Under SGI Xinerama, enlarging an OpenGL window beyond this size limit results in undefined behavior. An OpenGL

window may be placed anywhere within the the total area managed by the X server. Only the size of the region into which OpenGL renders is restricted.

Under DMX, however, OpenGL Multipipe allows OpenGL rendering to be unaffected by the window size limit of the graphics hardware. If the OpenGL application is invoked using `omprun`, the application may render into windows of any size.

Processor Requirements

On IRIX, OpenGL Multipipe requires that you use a MIPS R10000 processor or later. The following example shows how you check for the processor type:

```
$ hinv -t cpu  
CPU: MIPS R16000 Processor Chip Revision: 2.1
```

SGI ProPack and OpenGL Multipipe Versions

Versions of OpenGL Multipipe that ship with SGI Propack for Linux are only supported on the version of SGI Propack with which they ship due to backwards compatibility issues on Linux.

Overlay Windows Support in DMX

DMX supports overlay windows—that is, windows that use *overlay visuals*—only if overlay visuals are available on each of the underlying X servers that DMX manages. This is the case on most SGI systems. On Onyx4 and Silicon Graphics Prism platforms, you must explicitly enable support of overlay visuals in the configuration file(s) of the underlying XFree86 server(s).

SGI Xinerama Not Supported on Onyx4 or Silicon Graphics Prism Platforms

Due to changes in the X server on Onyx4 and Silicon Graphics Prism platforms, SGI Xinerama will not be supported on these platforms. Only the DMX proxy layer is

supported on these graphics platforms. DMX and SGI Xinerama are available on all other platforms supported by OpenGL Multipipe.

Graphics Pipe Requirements

Each of the graphics pipes managed by SGI Xinerama or DMX must have identical capabilities. Each pipe must provide the same set of X visuals and GLX FBConfigs, have the same amount of texture memory, and so on.

For example, if a system has three InfiniteReality3 graphics pipes where two pipes have two raster managers (RMs) and one pipe has four RMs, the pipe with four RMs must be configured to look as if it has two RMs. This can be done by ensuring the maximum pixel depth setting on each pipe is consistent. You can inspect the pixel depth by using the command `/usr/gfx/gfxinfo`. Depending on your system type, you can adjust the pixel depth on a particular pipe with one of the following commands: `/usr/gfx/ircombine`, `/usr/gfx/sgcombine`, or `xsetmon`. See the man pages for `gfxinfo(1)`, `ircombine(1)`, `sgcombine(1)`, and `xsetmon(1)` for more information.

Some Features Not Available on Linux

All OpenGL Multipipe features are available on IRIX platforms. However, some of these features are not available on Linux platforms. For a list of such features, see the OpenGL Multipipe release notes.

Troubleshooting

This chapter describes some problems you might encounter and what to do to solve them. For additional considerations, see the OpenGL Multipipe release notes, /usr/share/omp/release_notes/user/relnotes.html.

This chapter documents the following problems:

- “Problems Enabling SGI Xinerama” on page 50
- “Problems Starting DMX” on page 50
- “Problems Starting Applications with omprun” on page 51
- “Problems Running IRIS GL Applications” on page 51
- “Problems Running o32 Applications” on page 52
- “Graphics Do Not Display Correctly on All Screens” on page 52
- “Mouse Behavior Offset by a Screen” on page 55
- “Problems Running Inherently Multipipe Applications” on page 55
- “Multipipe-Aware Applications Fail to Receive Events on Screen 0” on page 56
- “Nothing Displays or the Graphic Stalls or Hangs” on page 56
- “OpenGL Graphics Render Slowly” on page 57
- “X Applications Are Not Behaving Correctly or Fail to Start” on page 58
- “Simultaneously Running X Servers with and without SGI Xinerama Enabled” on page 60
- “Tiled Background Image” on page 61
- “Flickering Grey Rubberband During Window Movement” on page 61
- “Mouse Disappears on Composited or Edge-Blended Display” on page 62
- “Problems Running Multithreaded Applications” on page 62
- “Problems with Aware Window Management” on page 62
- “Applications Do Not Behave Correctly in Aware Mode” on page 64

Problems Enabling SGI Xinerama

On systems having only one graphics pipe or in the case where the X server is directed to handle only one pipe (see the `Xsgi(1)` man page), enabling SGI Xinerama has no effect. In these cases, SGI Xinerama will be disabled, regardless of the value of the `xinerama` flag supplied on the `chkconfig` command.

Note: SGI Xinerama is not available on Onyx4 and Silicon Graphics Prism platforms. Only the DMX proxy layer is supported on these platforms.

Problems Starting DMX

If there is a problem starting the DMX proxy server, you may see output such as the following after running `ompstartdmx`:

```
ompstartdmx fatal: An error occurred when starting Xdmx
Check the Xdmx log file for details: /tmp/Xdmx.log.xxxxx
```

This can result from a number of conditions, some of which have workarounds that are described in the following paragraphs. Inspect the `Xdmx.log.xxxxx` file, especially toward the end of the log, for messages that indicate one of the following conditions:

- Incompatible screens, no common visuals
DMX will not create a logical display from graphics pipes with differing graphics capabilities. If the DMX proxy server detects that there are no common X visuals on the backend screens it tries to manage, DMX will abort with an error to this effect.

- Only one screen on display

On systems having only one graphics pipe or in the case where the X server is directed to handle only one pipe, `ompstartdmx` will exit with an error such as the following:

```
ompstartdmx fatal: Display :0.0 has only one screen.
DMX was not started
```

In these cases, it does not make sense to start DMX since there is only one pipe. However, specifying a configuration file with the `ompstartdmx -cfgfile` option will not prevent DMX from running on a single backend screen. Use the `ompstartdmx -help` option for more information.

Problems Starting Applications with `omprun`

If an application will not start when using the `omprun` command, one likely scenario is that the `DISPLAY` environment variable does not point to a meta display. If the `DISPLAY` environment variable does not point to a meta display, ensure that the following conditions true (check them in the order listed):

1. The `DISPLAY` environment variable points to the correct display.
2. An X proxy layer is enabled.

See the section “Verifying That the OpenGL Multipipe Environment is Enabled” on page 21. An X proxy layer must be enabled or when you invoke an application with `omprun`, the application will exit with the following error:

```
SGIomp fatal: ":0.0" is not a meta display
```

3. The application was not run from a shell that was started with the `omprun -aware` command or from a script that used the `omprun -aware` command to start the application.

The `omprun -aware` command effectively disables the X proxy layer for any programs it invokes.

4. Your application does not use either the OpenGL Multipipe SDK or OpenGL Performer multipipe API.

Recent versions of these APIs may have integrated with SGI Xinerama or DMX and may not run under OpenGL Multipipe. The solution is to run these applications as is or to simply ensure that they are run in aware mode (with `omprun -aware`). Another alternative is to use older versions of these APIs that do not contain the X proxy aware code.

Problems Running IRIS GL Applications

OpenGL Multipipe does not support IRIS GL applications. In some cases (when the application started with the `omprun` script is an executable and not a script), `omprun` can determine if the application is based on IRIS GL. In such a case, a warning message is generated and the application will not be started, as shown in the following example:

```
$ omprun clock  
omprun warning: clock is an IRIS GL application  
OMP Library does not support IRIS GL applications
```

For information about how to run an IRIS GL application when an X proxy layer is enabled, see the section “Running IRIS GL Applications” on page 24.

Problems Running o32 Applications

OpenGL Multipipe does not support o32 applications. In some cases (when the application started with the `omprun` script is an executable and not a script), `omprun` can determine if the application was built using the o32 application binary interface (ABI). In such a case, a warning message is generated and the application will not be started, as shown in the following example:

```
$ omprun showcase
omprun warning: showcase is an O32 application
OMP Library does not support O32 applications
```

For information about how to run an o32 application when an X proxy layer is enabled, see the section “Running IRIX o32 Applications” on page 25.

Graphics Do Not Display Correctly on All Screens

If a graphics window displays correctly on some screens only, there are several likely scenarios, which are described in the following subsections:

- “Coding Problem in the Application”
- “You Did Not Use the `omprun` Script”
- “A User-Defined Script Invokes an IRIS GL or o32 Application”
- “You Are Using the Aware Window Manager”
- “Set-User-ID (“s-bit”) Applications”

Coding Problem in the Application

If you are using the `omprun -cull` feature and you resize or move the application window to different screens, some applications may not draw an image properly on all screens. This can occur if an application does not call `glClear()` at the beginning of each frame (that is, it “builds up” an image, relying on a sort of rendering history from past frames). When culling is enabled, applications that do not call `glClear()` at the beginning

of each new frame may have unusual rendering artifacts when they are moved from their initial window position. The culling feature by nature eliminates drawing commands that would otherwise be rendered into an off-screen region. To avoid this behavior, do not use the `-cull` option.

You Did Not Use the `omprun` Script

Note: This section pertains to you only if you use the SGI Xinerama X proxy layer.

If a graphics window displays correctly on one screen only (usually screen 0), ensure that you start the application with the `omprun` script. If the same behavior persists when you invoke the application using the `omprun` script, ensure that one of the other conditions described in the following subsections does not exist.

A User-Defined Script Invokes an IRIS GL or o32 Application

Note: This section pertains to you only if you use the SGI Xinerama X proxy layer.

The `omprun` script cannot detect IRIS GL or o32 applications if it starts another script that in turn starts the target application. The following shell session illustrates this case:

```
$ cd /usr/demos/General_Demos/atlantis
$ omprun ./atlantis
omprun warning: ./atlantis is an IRIS GL application
OMP Library does not support IRIS GL applications
$ omprun ./RUN
```

In the preceding session, `RUN` is a script that invokes Atlantis. `RUN` does start the application, but it will be displayed correctly on one screen only.

If you start an application by using a user-defined script, ensure that the application is not an IRIS GL or o32 application. The following session shows how to test for an application that uses IRIS GL:

```
$ elfdump -Dl /usr/sbin/clock | grep libgl.so
[1] Oct 20 20:39:53 2000 0xe5383809 ----- libgl.so sgi1.0
```

If there is no output, the application does not use IRIS GL.

The following demonstrates how to test for an application that uses the o32 ABI:

```
$ file /usr/sbin/iconsmith | grep '32-bit'  
/usr/sbin/iconsmith: ELF 32-bit MSB mips-2 dynamic executable MIPS -  
version 1
```

If there is no output, the application does not use the o32 ABI.

You Are Using the Aware Window Manager

If you started an application in aware mode (that is, by running the script `omprun -aware app_name...`), the application running in aware mode only draws to one screen. If you are running an aware window manager, it is possible that the window manager may position the window on a screen other than the one on which the application is drawing. To see the window rendered correctly, move the application's window to the correct screen.

Set-User-ID (“s-bit”) Applications

OpenGL Multipipe cannot override the OpenGL calls of set-user-ID applications because the `omprun` script makes use of the following environment variable(s) to force an application to load the OpenGL Multipipe library `libOMPmaster.so` instead of the standard OpenGL library, `libGL.so`:

- `_RLDN32_LIST` and `_RLD64_LIST` (on IRIX)
- `LD_PRELOAD` (on Linux)

For security reasons, IRIX and Linux may ignore these environment variables for set-user-ID programs. Therefore, OpenGL Multipipe is not able to intercept and distribute OpenGL calls to all pipes. As a workaround, you may run the application under the DMX proxy layer to use its native support for GLX indirect rendering. To do so, simply run the application under DMX without using the `omprun` command. If this is not feasible, the application should run under `omprun` if you invoke the application while logged in as the user that owns or created the executable (or as `root` if administrative access is available).

Mouse Behavior Offset by a Screen

Note: This section pertains to you only if you use the SGI Xinerama X proxy layer.

If logical pipe 0 is not in the top left screen position, mouse events (such as clicks) are offset by one screen. Logical pipe 0 can be any physical pipe; it is the physical pipe specified by the first `-hw` argument in the X server configuration file, `/var/X11/xdm/Xservers`.

To work around this problem, list the graphics pipe of the monitor that is in the top left position first in the list of `-hw` arguments in the `Xservers` file. For example, in a configuration where pipes 5, 3, and 4 are in a linear array in that order, you would use the following `-boards` and `-hw` parameters together on one line in the `/var/X11/xdm/Xservers` file:

```
:0 secure /usr/bin/X11/X
-hw board=5,3,4
-hw board=5,right=1
-hw board=3,left=0,right=2
-hw board=4,left=1
... other X server args ....
```

Notes:

- The `-boards` numbers are physical pipe numbers, but the indexes given to the `right`, `left`, `above`, and `below` parameters refer to the logical order of the `-hw` parameters.
- The first `-hw` parameter is that of the top, leftmost pipe and should never have a left or top neighbor.
- The lines are separated in the example only for readability.

See the `Xsgi(1)` and `xdm(1)` man pages for more information about the `-hw` options and the `Xservers` file.

Problems Running Inherently Multipipe Applications

Applications—such as `ircombine`, `sgcombine`, `xsetmon`, `setmon`, and `gamma`—which are designed to manage graphics hardware pipes individually, are inherently

multipipe applications. Therefore, they should be started in multipipe-aware mode, as shown in the following examples:

```
$ setenv DISPLAY :0.0
$ omprun -aware gamma

$ setenv DISPLAY :0.0
$ omprun -aware sgcombine
```

This allows the application to run as if the X proxy layer were disabled. Applications with a GUI—such as `sgcombine`, `ircombine`, and `xsetmon`—will not be under window manager control unless an aware window manager is active.

Multipipe-Aware Applications Fail to Receive Events on Screen 0

Windows of applications that are run in aware mode are not handled by ordinary window managers. This can cause some problems on screen 0 for keyboard events.

Moving away all the windows that are overlapping the aware window (even if these windows are displayed behind the aware window) will set the correct focus. The aware window will then receive the keyboard events.

Alternately, running the aware window manager will also fix the focus problem.

Nothing Displays or the Graphic Stalls or Hangs

If you start an OpenGL application with `omprun` and it does not display anything or the graphic stalls or even hangs, the source of the problem might be one of the following:

- “Coding Problem in the Application”
- “Window Exceeds Maximum OpenGL Window Size”
- “Improperly Wired Genlock or Swap Ready Cables”

Coding Problem in the Application

You may see a blank display or experience stalls or hangs for OpenGL applications that do not call `glFlush()`, `glFinish()`, or `glXSwapBuffers()` at the end of each frame. This

causes OpenGL Multipipe to draw only when its internal buffer overflows. It can happen that the buffer never fills, in the case of an event-driven application—that is, the application draws one frame and waits for an event before drawing the next frame. Unfortunately, there is no workaround for applications that are not frame-based because OpenGL Multipipe relies on regular calls to the functions just cited.

Window Exceeds Maximum OpenGL Window Size

Note: This section pertains to you only if you use the SGI Xinerama X proxy layer.

If the size of a window with OpenGL content is larger than the maximum width or height allowed by the graphics hardware, undefined drawing behavior will result. Although this limit can vary depending on the graphics configuration, it is typically around 4000 pixels. Logical screen configurations with a width or height larger than the maximum OpenGL window size are particularly susceptible to this behavior. As long as an OpenGL window is smaller than the maximum size, it may be placed anywhere within the total area managed by the X server. Only the size of the region into which OpenGL renders is restricted. For more information, see “OpenGL Window Size Constraints” on page 46.

Improperly Wired Genlock or Swap Ready Cables

If you are experiencing long delays between frames of an OpenGL application (whether or not it was started with `omprun`), the condition may have resulted from using the `omprun -swapready` option with improperly configured or improperly wired Genlock or Swap Ready cables.

For more information about this problem and a workaround, see the OpenGL Multipipe release notes.

OpenGL Graphics Render Slowly

While there could be many reasons for slow rendering, ensure that you used the `omprun` command to start your OpenGL application, in particular, if you are using the DMX X proxy layer. Failure to use the `omprun` command under DMX will cause the application

to draw to all screens using the slower GLX indirect rendering support in DMX instead of the OpenGL direct rendering provided through the `omprun` command. Use the `omprun` command to achieve the best rendering performance for single pipe (“multipipe-unaware”) applications under DMX.

For other performance optimization suggestions when using the `omprun` command, see the section “Performance Features” on page 27.

To achieve the best performance with multipipe applications, see the section “Running Multipipe Applications in Multipipe-Aware Mode” on page 26.

X Applications Are Not Behaving Correctly or Fail to Start

If X applications are not behaving correctly or fail to start, consider the cases described in the following subsections:

- “X Application Uses Unsupported X Extension”
- “SGI Xinerama Client or Server Uses Nonstandard Protocol”
- “Application Window Disappears”
- “Application Explicitly Opens a Display Connection to :0.0”

X Application Uses Unsupported X Extension

Verify that the application is not using unsupported X extensions. Unfortunately, there is no way to accurately list the extensions an application uses. However, the following examples using the `nm` command give some hints about the extensions used. If an application is using an X extension, this can usually be detected by searching for occurrences of the string `extension` or for the name of a particular extension. The `xdpyinfo` command lists the names of extensions supported by the X server.

Indicating the use of the DOUBLE-BUFFER extension (DBE), the following example shows that command `gmemusage` calls `XdbeQueryExtension`:

```
# nm /usr/sbin/gmemusage | grep -i extension
[116] |2143299120| 436|FUNC |GLOB |DEFAULT |UNDEF| XdbeQueryExtension
```

For a list of X extensions supported by SGI Xinerama, see the `Xinerama(3X11)` man page. For a list of X extensions supported by DMX, see the `Xdmx(1)` man page in the

directory `/usr/share/omp/doc/user`. Applications that use unsupported X extensions may be run in aware mode by running them with the `omprun -aware` option so that they bypass the X proxy layer.

SGI Xinerama Client or Server Uses Nonstandard Protocol

The SGI Xinerama versions in IRIX 6.5.11 and earlier use a protocol that is incompatible with versions of SGI Xinerama released in IRIX 6.5.12 and later. If an application links (dynamically at run time or statically at compile time) with X client libraries that came with IRIX 6.5.11 and earlier and then attempts to make SGI Xinerama calls to an X server from IRIX 6.5.12 or later, the behavior will be undefined. Similarly, linking with X client libraries from IRIX 6.5.12 or later and connecting to an X server from IRIX 6.5.11 or earlier will also yield undefined behavior.

Since the OpenGL Multipipe layer calls `XineramaQueryVersion`, it is able to reliably detect and report this server-client version incompatibility.

If you encounter this protocol incompatibility, the workaround is to use a client library and server that both support the same SGI Xinerama version—that is, use a client library and X server from the same IRIX version. More simply stated, if you are connecting to a remote display that has SGI Xinerama enabled, ensure that both the local and remote hosts are running the same version of IRIX.

See the `Xinerama(3X11)` and `XineramaQueryVersion(3X11)` man pages for more details.

Application Window Disappears

Note: This section pertains to you only if you use the DMX X proxy layer.

This can occur if you start an application on one of DMX's backend displays without using `omprun -aware`—for example, if your `DISPLAY` environment variable is set to `:0.0` and the DMX proxy is managing display `:1`. The application will run on `:0.0`, a backend (“aware”) display, but because it was not started properly in aware mode (using `omprun -aware`), the window will not be managed by the window manager running under DMX. Consequently, it may be “pushed” behind the DMX root window. Exiting the DMX session will reveal the application window. To avoid this problem, open the application in aware mode using `omprun -aware app_name`.

Application Explicitly Opens a Display Connection to :0.0

Note: This section pertains to you only if you use the DMX X proxy layer.

This problem can manifest itself in many ways:

- An apparent X error may appear in the output of the application.
- A window may suddenly disappear behind the DMX root window.
- The command `omprun` may output the following message:

```
DISPLAY does not point to a meta display.
```

This problem originates from a program explicitly requesting an X display connection to :0.0 even if the `DISPLAY` environment variable contains a different default display string, which is usually the case when running DMX.

The workaround is for DMX to run as `display :0` and to have the backend X servers use other display numbers so that when the application tries to open :0.0, it correctly connects to the DMX meta display on :0 instead of a backend display by accident. To have other X servers use other display numbers so that :0 is available for DMX, do the following:

1. Open file `/var/x11/xdm/xservers`.
2. Replace the instance of :0 with another display number (for example, :1 or :5).
3. Restart graphics.
4. Remove the `/tmp/.X11-unix/X0` Unix socket file if it exists.

The existence of this file will cause `ompstartdmx` to start DMX on a display other than :0.

Now when you run `ompstartdmx`, by default it will use the first available X display number, which should now be :0.

Simultaneously Running X Servers with and without SGI Xinerama Enabled

To run X servers with SGI Xinerama enabled simultaneously with regular X servers (that is, with SGI Xinerama disabled) on the same machine, add `+xinerama` or `-xinerama`

to the existing arguments in the file `/var/X11/xdm/Xservers`. This allows you to override the `chkconfig xinerama` flag. Refer to the `Xsgi(1)` and `xdm(1)` man pages for more information.

Tiled Background Image

Note: This section pertains to you only if you use the IRIX window manager `4Dwm`.

Setting a large image as the root window background image will result in having a tile image displayed across the screens. You can overcome this problem by the using `4Dwm` desktop features as follows:

- Set the following line in the `$HOME/.Sgiresources` file:
`4Dwm*SG_useBackgrounds: True`
- Create the background image in the `xpm` (X Pixmap) file format. The fewer colors used in that image, the less impact it will have on the colormaps used by other applications.
- Copy the `/usr/lib/X11/system.backgrounds` file to `$HOME/.backgrounds`.
- Edit `$HOME/.backgrounds` and, using the syntax of `/usr/lib/X11/system.backgrounds` as a template, add a new setting for your image.
- Select your background from the GUI background program.

See the `4Dwm(1X)` man page for more information.

Flickering Grey Rubberband During Window Movement

Note: This section pertains to you only if you use the DMX X proxy layer.

This occurs as a result of the lack of overlay visual support in DMX. See “Overlay Windows Support in DMX” on page 47.

Mouse Disappears on Composited or Edge-Blended Display

When an X proxy layer is used to overlap screen regions on an edge-blended display or a compositor-based system, the cursor will seem to disappear when it enters the overlapped or uncomposited regions of the display. In X, a mouse belongs to one screen of the X server at a time. Therefore, it is normally not possible to draw the mouse multiple times (on different screens) in the overlap region.

To prevent the cursor from disappearing in these cases, you need to create additional or *duplicate cursor* images (not real cursors) where two or more screens overlap. There is still just one real cursor position on the display. The way you do this depends on your X proxy layer, SGI Xinerama or DMX. See the section “Enabling Duplicate Cursor Images in Overlap Regions” on page 38.

Problems Running Multithreaded Applications

If the application supports the use of POSIX threads (*pthreads*), use the pthread threading model with OpenGL Multipipe.

To force the use of the pthread threading model, use the `-pthread` option when starting the application, as shown in the following:

```
$ omprun -pthread app_name
```

Problems with Aware Window Management

The following subsections detail problems with aware window management and workarounds:

- “Windows of Some Aware Applications are Not Managed”
- “Problems with Desktop Background Images”
- “Mouse Events Sometimes Register on the Wrong Screen”
- “Ghost Windows Appear In Overlap Regions on Edge-Blended Displays”

Windows of Some Aware Applications are Not Managed

For windows of aware applications to be managed, first start the aware window manager, then start the desired application in aware mode. If the reverse is done, the windows will not be able to be manipulated.

Problems with Desktop Background Images

Note: This section pertains to you only if you use the IRIX window manager 4Dwm or its Xinerama-aware variant omp4Dwm.

Under SGI Xinerama, if desktop background images do not appear when the `start_ompwm` script is used to start `omp4Dwm`, you must enable the `SG_UseBackgrounds` resource for 4Dwm. This can be done through an X resource file or on the `start_ompwm` command line, as shown in the following command entries:

```
$ tellwm quit
$ start_ompwm -xrm "*SG_UseBackgrounds: True"
```

If you started DMX as the X proxy layer using `ompstartdmx`, 4Dwm is started by default. If you need to explicitly start 4Dwm with proper background images under DMX, you may do so by entering the following:

```
$ tellwm quit
$ ompwrapwm 4Dwm -xrm "*SG_UseBackgrounds: True"
```

Mouse Events Sometimes Register on the Wrong Screen

Note: This section pertains to you only if you use the SGI Xinerama X proxy layer.

A number of known mouse pointer and keyboard focus issues arise with aware windows that reside on screens other than physical screen 0—that is, on any of the physical screens 1.. n that SGI Xinerama is managing. Events (mouse clicks, keyboard strokes, etc.) are sometimes generated as though the mouse were on screen 0 instead of on the correct screen (1.. n). This frequently occurs when a popup menu of an aware window is open and the mouse is clicked outside the menu.

To avoid this particular problem, close the popup menu first (for example, by using a keyboard shortcut to close the menu) then click the mouse.

Ghost Windows Appear In Overlap Regions on Edge-Blended Displays

Aware windows bypass the X proxy layer and are only created on one physical screen, but when an aware window manager manages an aware window, it creates window frames on each screen. The frames are multipipe-transparent—that is, drawn on every screen. However, the application window within the frame is multipipe-aware—that is, drawn only on one screen.

Since an aware application window is not drawn on every screen, the multipipe-transparent frame behind the application window will “show through” in screen-overlap regions on an edge-blended display.

To work around this problem, you may want to run your application in a window of a size and position such that it does not overlap any of the screen-overlap regions of the display. Alternatively, you may want to temporarily quit the aware window manager.

Applications Do Not Behave Correctly in Aware Mode

If you are using DMX as your proxy X layer, note that the `ompstartdmx` script sets the environment variable `SGIOMP_META_DISPLAY`, which is important for running applications in aware mode under DMX. If you are starting an application in aware mode from a remote shell or other login shell, this variable will not be set. If unexpected results occur when you try to run an application in aware mode, ensure that the `SGIOMP_META_DISPLAY` variable is set to the DMX display (often `:1`).