



***Command Dispatch System  
Administrator's Guide Version 2020  
22 December 2020***

**Important Note:** You must be a GeoCue System Administrator to set all advanced features in the Command Dispatch System as well as to modify Priority or to use the "Launch Now" feature of Dispatch Manager.

It is important to note that the default configuration of the Command Dispatch System will run all dispatchable GeoCue commands. Thus the information in this guide is for advanced system configuration.

GeoCue Group, Inc.  
9668 Madison Blvd.  
Suite 202  
Madison, AL 35758  
1-256-461-8289  
[www.geocue.com](http://www.geocue.com)



ISV/Software Solutions

# NOTICES

The material in **GeoCue Group, Inc.** documents is protected by United States Copyright laws.

You may make as many copies of this material for use internal to your company as you desire. Please do not distribute this material outside of your company without first discussing with us.

## ***Trademarks, Service Marks***

- *Windows* and *.NET* are trademarks of Microsoft Corporation.
- *MicroStation* is a trademark of Bentley Systems Incorporated.
- *TerraScan* is a trademark of TerraSolid Oy.
- GeoCue and CuePac are registered trademarks of GeoCue Group, Inc.
- *Intergraph* and *DMC* are trademarks of Intergraph Corporation.

# Getting Help

We are sure that you will experience different problems with GeoCue that range from installation issues to defects that made it through our testing undetected. We hope that you will immediately contact us with any problems or questions and have the patience to work with us through a successful GeoCue deployment.

Please contact us via email for assistance with or comments about GeoCue or our various CuePacs.

email:

[support@geocue.com](mailto:support@geocue.com)

Phone:

**1-256-461-8289**

Just ask for GeoCue Support and you will get connected with someone who can assist you. There is usually someone in the office between the hours of 0600 and 1800 CDT, USA on weekdays. Weekends are sort of hit or miss.

Fax (always on):

**1-256-461-8249**

## *About this Document*

Welcome to the GeoCue Release 2020 Command Dispatch System Administrator's Guide. This document supplements the GeoCue User Guide and contains advanced information for the GeoCue Command Dispatch System (CDS). This guide is aimed at advanced users of the Command Dispatch System. Note that most information in this guide is not required for routine dispatched and distributed processing.

# Contents

1	Introduction .....	1-1
1.1	Installation .....	1-2
1.2	Licensing.....	1-2
2	Terminology.....	2-1
2.1	Dispatch vs. Distributed .....	2-1
2.2	Tasks and Subtasks.....	2-1
2.3	Machine Resources.....	2-2
2.4	Command Resources .....	2-4
2.5	Processing Priority .....	2-4
2.6	Processing Groups.....	2-5
2.7	Clouds (Machine Clusters) .....	2-5
2.8	Priority Weights .....	2-7
2.9	Machine Restriction Time .....	2-9
2.10	Processing Group Restriction Time.....	2-9
3	Dispatching and Monitoring Tasks (Commands).....	3-1
3.1	Dispatch Invocation.....	3-1
3.2	The Dispatch Dialog .....	3-1
3.2.1	Schedule .....	3-2
3.2.2	Machine Mode .....	3-2
3.2.3	Cloud Mode.....	3-4
3.3	Machine/Cloud Tables .....	3-6
3.3.1	Machine Table Fields.....	3-6
3.3.2	Cloud Table Fields.....	3-8

- 3.3.3 Override Auto-Start Next Step ..... 3-9
- 3.4 Command Dispatch System Tool Tray Tool ..... 3-11
- 3.5 Monitoring Dispatched Tasks..... 3-14
  - 3.5.1 Tasks Table ..... 3-15
  - 3.5.2 Sub Task Dispatch Status Summary ..... 3-18
  - 3.5.3 Subtasks Table ..... 3-19
- 3.6 Clearing Rows from the Tasks Pane ..... 3-22
- 3.7 Rescheduling a Task ..... 3-23
- 3.8 Pausing a Task ..... 3-23
- 3.9 Terminating a Task ..... 3-23
- 3.10 Modifying a Subtask (GeoCue Administrators Only) ..... 3-24
- 3.11 Creating Entity Queues ..... 3-24
- 3.12 Monitoring Machine Activity ..... 3-26
- 4 Configuring Machines ..... 4-1
  - 4.1 Adding Machines to the CDS..... 4-1
  - 4.2 Machine Resources ..... 4-4
  - 4.3 Setting Machine Configuration Parameters..... 4-5
  - 4.4 Machine Schedules ..... 4-8
  - 4.5 Machine Loading ..... 4-10
  - 4.6 Machine Group/Cloud Membership ..... 4-11
- 5 Configuring Commands ..... 5-12
  - 5.1 Enabling Dispatch..... 5-15
  - 5.2 Restrict to Single Instance..... 5-16
  - 5.3 Auto Kill Size ..... 5-16
  - 5.4 Caching..... 5-17
  - 5.5 Machines ..... 5-18

5.6	Maximum Failures Allowed.....	5-19
5.7	Predicted Sizes .....	5-19
5.8	Equilibrium Time .....	5-19
6	Configuring Priority.....	6-1
6.1	Priority Algorithm.....	6-1
6.2	Setting Priority Weights .....	6-3
6.2.1	Entity Priority Weight.....	6-3
6.2.2	Project and Processing Group Priority Weight, Clipping.....	6-6
6.3	Setting Relative Priority .....	6-7
6.3.1	Entity Relative Priority.....	6-7
6.3.2	Project Relative Priority.....	6-8
6.3.3	Processing Group Relative Priority.....	6-11
7	Configuring Processing Groups.....	7-1
7.1	Processing Groups Overview .....	7-1
7.2	Creating Groups .....	7-4
7.3	Configuring Resource Groups .....	7-5
7.3.1	Setting Resource Allocations (non-scheduled).....	7-9
7.3.2	Reserving Resources.....	7-10
7.3.3	Scheduling Resources.....	7-11
7.4	The RESERVED Group.....	7-11
7.5	Processing Groups Considerations.....	7-12
8	Reserving Machine Resources for non-CDS Use.....	8-1
8.1	Reserving Machine Resources .....	8-1
8.2	Scheduling the RESERVED Group.....	8-2
8.3	An Example .....	8-3
9	Configuring Clouds.....	9-1

- 9.1 Configuring Clouds ..... 9-4
- 9.2 Processing Groups and Cloud ..... 9-5
- 9.3 The Universal Group/Cloud..... 9-5
- 9.4 Summary ..... 9-6
- 10 Configuring Commands within Processing Groups/Clouds ..... 10-1
  - 10.1 Allocating Resources at the Command Level..... 10-2
  - 10.2 Reserving Command Resources..... 10-3
  - 10.3 Scheduling Command Resources ..... 10-4
- 11 Troubleshooting..... 11-1
  - 11.1 Tasks Remain in a Pending State..... 11-1
- 12 Command Dispatch System Internals..... 12-1
  - 12.1 Processing Granularity ..... 12-1
  - 12.2 Machine Assignment..... 12-2
    - 12.2.1 Machine Dispatch Mode ..... 12-3
    - 12.2.2 Cloud Dispatch Mode ..... 12-4
  - 12.3 Subtask Queuing ..... 12-4
  - 12.4 Machine (Node) Selection..... 12-5
  - 12.5 Subtask Execution ..... 12-7
  - 12.6 How the CDS Dispatches With Processing Groups ..... 12-7
    - 12.6.1 Dispatch Queue Entry..... 12-8
    - 12.6.2 Processing Groups and Dispatch ..... 12-8
    - 12.6.3 Dispatch Pass One ..... 12-9
    - 12.6.4 Dispatch Pass Two ..... 12-10
    - 12.6.5 Starvation Prevention..... 12-11
  - 12.7 Subtask Completion Processing ..... 12-11
- 13 Concluding Remarks ..... 13-1



14 Acronyms ..... 14-1

# 1 Introduction

The Command Dispatch System (CDS) was introduced in GeoCue version 3.0. This system allows you to run certain commands on remote nodes (such as Populate Working Segments in LIDAR 1 and DEM CuePac) and to *distribute* certain commands across multiple workstations (such as Generate Images and distribute TerraScan Macros in LIDAR 1 CuePac, DMC Post-Processing in

the DMC PPS CuePac and so on). The CDS contains many advanced distributed processing management features. Among these are:

- Machine aggregation into local “Clouds”
- Machine Resource Allocations
- Command Resource Allocations
- Multi-level Priority System (Entity, Project, Cloud)
- Machine Priority
- Advanced Resource Scheduling
- Support for UNIX/LINUX Remote Nodes<sup>1</sup>

Most users of the Command Dispatch System will not need to use these advanced features. Thus this Administrator's Guide is aimed primarily at users who need to set up advanced processing scenarios that use priorities, resource grouping, resource scheduling and cloud computing.

## **1.1 Installation**

The Command Dispatch System is an integral component of the Departmental Server version of GeoCue (as well as several OEM versions). Please refer to the separately supplied GeoCue Installation Guide for GeoCue Departmental Server.

## **1.2 Licensing**

The Command Dispatch System is available as an integral component of the Departmental level of GeoCue Server. It is not included in other versions of GeoCue Server.

The fundamental software required for a distributed processing system includes:

- GeoCue Departmental Server
- A minimum of 1 GeoCue Client
- A Dispatched Subtask License (DSL) per dispatched subtask

For example, suppose you wish to dispatch an orthorectification task to a batch processing server containing 8 compute cores. If you desired to run 8 parallel tasks, your software requirements would be:

1 GeoCue Departmental Server

1 GeoCue Client (for task management and initiation)

8 Dispatched Subtask Licenses

Note that the licensing requirements of the actual processing software are independent of the Command Dispatch System. In the above example, if the rectification software was licensed *per machine* regardless of the number of *instantiations*, you would need a single rectify license. On the other hand, if the software were licensed such that one license were required per instance, you would need 8 licenses.

NOTE: The Command Dispatch System *never* circumvents the licensing policy of the ultimate processing software. In general, you will need a minimum of 1 license per machine for end-use software and often 1 license per *instance*. You should keep this in mind as you design your distributed processing architecture.

## 2 Terminology

The features of the Command Dispatch System (CDS) use terminology that may be unique to GeoCue. This chapter provides an introduction to the terms that we use in configuring the CDS. The actual use of the terms introduced in this chapter is presented in later chapters of this guide.

### 2.1 *Dispatch vs. Distributed*

*Dispatch* is a task that can be run on a remote machine via the CDS but it cannot be distributed across multiple machines. This option is configurable by the system administrator and is usually a limitation imposed by the type of processing that needs to occur. For example, in the GeoCue LIDAR 1 CuePac, populating working segments from LIDAR segments is performed by selecting the tiles to be populated and then running the populate step. No matter how many tiles are simultaneously selected for this operation, it can be dispatched only to a single remote machine. This is due to the algorithm that we implemented for the population operation.

*Distributed* processing, on the other hand, means that the items that you send off to the CDS in a single operation (i.e. by selecting into the working set and pressing the In Progress button) will be *Queued* to a dispatch queue and sent off to multiple processing nodes. An example is the generation of LIDAR orthos in the GeoCue LIDAR 1 CuePac. If you select 20 LIDAR orthos and have 5 machines enabled in your CDS for LIDAR ortho generation, these five machines will be fed orthos from the queue until all 20 have been processed.

### 2.2 *Tasks and Subtasks*

We use the terminology Task and Subtask when referencing items that can participate in processing via the CDS. These terms are interpreted differently depending on the processing mode; Dispatch or Distributed.

When a task will be *Dispatched*, all the entities participating in the operation are sent to the selected recipient machine and thus a *Task* and *Subtask* are synonymous.

When a task can be *Distributed*, the task is split into *Subtasks* (usually one subtask per entity) and these *subtasks* are launched to the nodes that are set up to process the particular command. The subtasks remain associated with the original launching task.

## 2.3 Machine Resources

Resources available per machine are specified in terms of:

- Compute Power
- Memory
- I/O Bandwidth

We use units called *Standard x Units* where x is Processor, Memory or I(/O) for the three controlling attributes:

- Standard Processor Units (SPU)
- Standard Memory Units (SMU)
- Standard I/O Units (SIU)

These metrics really only have meaning in the context of a command. Each command that can be executed via the CDS can have its resource requirements assigned in terms of these three metrics (SPU, SMU, SIU).

While these units are completely arbitrary (meaning you can set your own scale) we recommend scaling in terms of the values in Table 2-1.

**Table 2-1: Recommended Standard Units**

<b>Metric</b>	<b>Reference Unit</b>
1 SPU	1 GHz Xeon, Single Core
1 SMU	1 GB RAM
1 SIU	1 MB/second of I/O bandwidth

Thus if you used our recommended scaling for a computer with a single 2 GHz quad core Xeon machine, 2 GB of memory and a disk I/O bandwidth of 80 MB/sec, your settings would be as listed in Table 2-2

**Table 2-2: Settings for the example Machine**

<b>Metric</b>	<b>Reference Unit</b>
SPU	8.0 (4 x 2 GHz)
SMU	2.0 (2 MB / 1 MB/SMU)
SIU	80.0 (80 MB/sec divided by 1 MB/SIU)

As will be explained in a later section, you can have any parameter except SPU ignored in a dispatch computation by setting the corresponding resource to a value of -1 (or *ignore* if this setting is available in the dialog). These values are set using the *Machine* tab of Dispatch Manager (as will be described in a later chapter).

## 2.4 Command Resources

The complement to Machine Resources is Command Resources. Machine Resources specify, in terms of the Standard Units, the processing capabilities of a particular machine (or, synonymously, node). Command Resources specify the amount of Standard Resources necessary to effectively execute a particular command. This is, of course, rather arbitrary since you do not generally know these values (e.g. how many SPUs, SIUs, SMUs does Intergraph's OrthoPro require?). We generally advise that you make educated guesses and then, using data available from Dispatch Manager, *tune* these values over time.

These values, as will be described in a subsequent chapter, are set using the *Command* tab of *Environment Builder*.

## 2.5 Processing Priority

Processing Priority (or simply *priority*) determines the processing order of subtasks that could otherwise process in any order. Priority in the CDS is a floating point number ranging from 0.0 to MAX\_FLOAT. Higher values are considered higher priority. For example, assume that two LIDAR Orthos are in the CDS dispatch queue and that these ortho subtasks have priorities of 1.0 and 7.3. Also assume that only one machine is available for processing. In this simple scenario, the ortho with priority 7.3 will be processed prior to the ortho with priority 1.0.

Subtask priority is computed in GeoCue as subtasks are entered into the dispatch queue. Priority is set at three levels in GeoCue:

- Entity
- Project
- Processing Group

These priorities are combined in a priority weighting equation to arrive at a single final priority value.



## 2.6 Processing Groups

*Processing Group* is a concept in the Command Dispatch System that allows partitioning of machine resources based on assigned *groups*. Groups can be used to modify the processing priority of subtasks. It is typically used to control how resources are used on multiprocessor machines where you do not want a particular project or production function (for example, your ortho department) to dominate the use of this shared resource.

Processing Groups are simply names that you create and assign via the **Groups** tab of Dispatch Manager. For example, you might create the groups ORTHO and LIDAR to control the resource usage of projects assigned to these groups. Note that Groups are also used to designate Clouds.

## 2.7 Clouds (Machine Clusters)

*Cloud Computing* (which used to be called "Clustering") allows you to select a named collection of processing workstation when you dispatch a task rather than the individual machines. Clouds are created using the Groups/Clouds tab of Dispatch Manager.

In order to use Cloud dispatch mode, it must be enabled on the Configuration tab of Dispatch Manager. Once enabled (and assuming that you have defined at least one machine cloud), you will be presented a Machine/Cloud option on the Dispatch dialog. Switching to the Cloud mode will present a list of your defined Clouds as opposed to the normal list of machines.

Cloud processing is an extremely flexible processing mode in that it allows the individual machines that comprise the cloud to be managed (e.g. taken off-line and brought back on-line) while the Cloud is processing a Task. Thus, for example, individual machines in the cloud can be taken off-line automatically (a feature of the Command Dispatch System), a backup automatically executed and then brought back on-line without disturbing the processing of the subtasks that comprise the executing Task(s).

NOTE: Cloud and Cluster are used interchangeable throughout this manual. Gradually we are deprecating the term Cluster in favor of Cloud.

## 2.8 Priority Weights

The ultimate priority of a subtask in the CDS is computed based on the individual priority settings for Entities, Project and Processing Group. The computed priority is *scaled* by *priority weights* that are set in Environment Builder and Dispatch Manager, **Configuration, Priority** tab. The parameters that are used in ultimate priority computation are given in Table 2-2

**Table 2-3: Machine Settings**

<b>Parameter</b>	<b>Description</b>	<b>Notes</b>
Entity Relative Priority (ERP)	The priority assigned to an individual entity The valid range is 0.0 to MAX_FLOAT.	The system default is 1.0. Can be interactively assigned using Entity Manager.
Project Relative Priority (PRP)	The priority assigned to an individual <i>project</i> . The valid range is 0.0 to MAX_FLOAT.	The system default is 1.0. Assigned using Project Properties or Project Utilities.
Group Relative Priority (GRP)	The priority assigned to a <i>Processing Group</i> (Cloud). The valid range is 0.0 to MAX_FLOAT.	The system default is 1.0. Assigned using the Configuration tab of Dispatch Manager, Processing Groups sub-tab.
Entity Priority Weight (EPW)	A weighting factor applied to the Entity Relative Priority. The valid range is 0.0 to MAX_FLOAT.	The default value is 1.0. Assigned using the Entity tab of Environment Builder, Entity Types section, Modify option.
Project Priority Weight (PPW)	A weighting factor applied to the Project Relative Priority. The valid range is 0.0 to MAX_FLOAT.	Assigned using the Configuration tab of Dispatch Manager, Priority sub-tab.
Group Priority Weight (GPW)	A weighting factor applied to the Group Relative Priority. The valid range is 0.0 to MAX_FLOAT.	Assigned using the Configuration tab of Dispatch Manager, Priority sub-tab.

It is important to note that the Command Dispatch System will order the processing of subtasks based on the final computed priority of the subtask. Thus, depending on the priority parameters and weights, a subtask from a low priority project may ultimately process prior to a subtask from a higher priority subtask.

If clipping is not being used, the ultimate priority of a subtask is given by the formula of Eq: 1.

$$\text{Priority}^2 = (\text{ERP} \times \text{EPW}) + (\text{PRP} \times \text{PPW}) + (\text{GRP} \times \text{GPW}) \quad \text{Eq: 1}$$

*Priority Range Clipping* are values that you can set to *clip* the maximum value of each of the partial priorities (entity, project, group) above. This provides more flexibility in setting up system priority schemes.

NOTE: Priority is global within the Command Dispatch System. For example, assume task A comprising 100 subtasks, each having priority 10 is dispatched. Sometime later, Task B comprising subtasks of priority 11 is launched. All of the Task B subtasks will move ahead of the remaining Task A subtasks due to their higher priority values.

---

<sup>2</sup> Entity Priority can be dynamically modified based on the GeoCue *Linking* system. This is discussed in a separate section.

## **2.9 Machine Restriction Time**

You can set up schedules for processing nodes (machines) that indicate when the node is not available for processing. For example, if you have a server that runs backups from 2 AM to 4 AM on Tuesdays and Thursdays each week, you can configure the CDS to *not* use this machine during those time intervals. If the machine is actually processing a subtask at the time it is scheduled to go off-line, it will complete the current subtask first.

## **2.10 Processing Group Restriction Time**

You can restrict the times when nodes will process subtasks from a particular Processing Group using a scheduling system similar to the node time restriction schedule. For example, you might want to allow a particular machine to process light weight tasks such as generating LIDAR orthos at any time but restrict processing of TerraScan macros to nights only. This type of command specific scheduling can be accomplished by using Processing Groups and Process Group restriction times.

## 3 Dispatching and Monitoring Tasks (Commands)

In this chapter we provide a brief overview of dispatching commands from the user's perspective. This information is covered in more detail in the GeoCue User Guide.

### 3.1 Dispatch Invocation

Tasks (that are capable of being dispatched) can be vectored to the Command Dispatch System via three different mechanisms:

1. by enabling *Dispatch* from the checklist toolbar controls (this is the most common mechanism)
2. by enabling *Dispatch* on dialogs that support dispatching (such as the Import LIDAR sources dialog of the LIDAR 1 CuePac)
3. Programmatically (some GeoCue commands make direct access to the Command Dispatch System. End-users can access the Command Dispatch System in this manner by using the GeoCue Software Developer's Kit – SDK)

### 3.2 The Dispatch Dialog

Following invocation by one of the methods described in the previous section, the Dispatch dialog (Figure 3-1) is displayed. Each of the sections of this dialog is discussed in the following subsections.

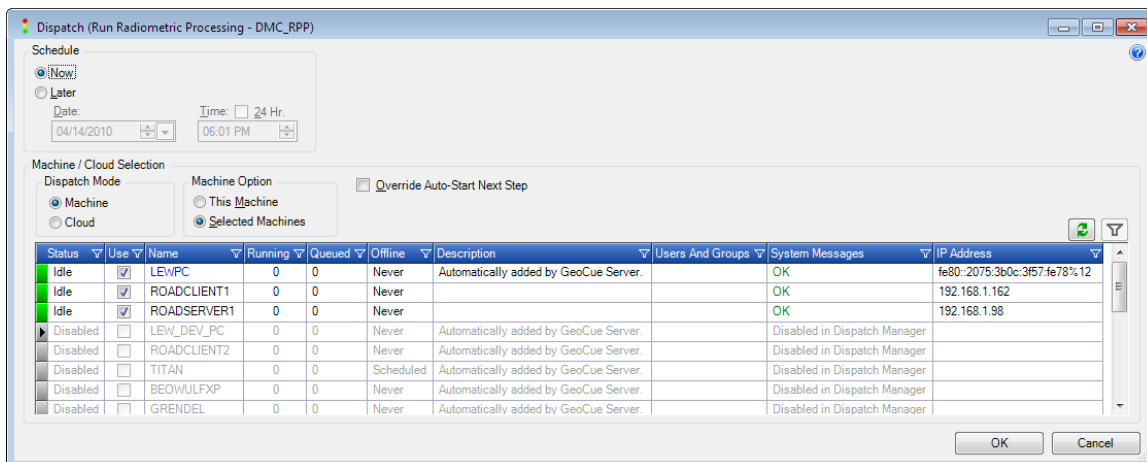


Figure 3-1: The GeoCue Dispatch dialog

### 3.2.1 Schedule

The default is Now which means that the Task will be submitted to the dispatch queue for immediate execution. Of course, when the Task actually executes is determined by the ultimate subtask priority, the number of other subtasks in the dispatch queue and the available of machines on which this Task can execute.

If you select *Later* in the Schedule radio button control, you can set a date and a time for processing this Task. Again, the exact time of execution will depend on the overall constellation status at the schedule time. However, you will be assured that the Task will not execute prior to the scheduled date/time.

### 3.2.2 Machine Mode

This is the default for launching tasks. When this mode is active, you can select either “this Machine” or “Selected Machines” in the Machine Option section of the dialog. If Selected Machines is enabled, you can select the desired machines from the Machine list table (Figure 3-2). Note that machines may not be selectable for several different reasons:

1. The machine has failed a Command Dispatch System validity check (it may be shut down or off the network)
2. The machine is Disabled because it has been set to “Disabled for Dispatched Processing” either via Dispatch Manager or via the tool tray icon on the disabled machine’s desktop (see the Configuring Machines chapter)
3. The Command that is associated with the Task being dispatched is not configured to run on the machine (see the Configuring Commands chapter).

If a Task can be dispatched to a remote machine but cannot be distributed across multiple machines (e.g. Populate LAS Working Segments in the LIDAR 1 CuePac), the machine selection checkboxes will exhibit “radio button” behavior (meaning you can select only one machine).

To dispatch the Task, select the desired machine(s) and press the OK button.

NOTE: Effective with GeoCue Version 5.0, you can *dispatch* to machines that are *off-line*. These machines will pick up subtasks for processing if they come back on-line while subtasks are still pending.



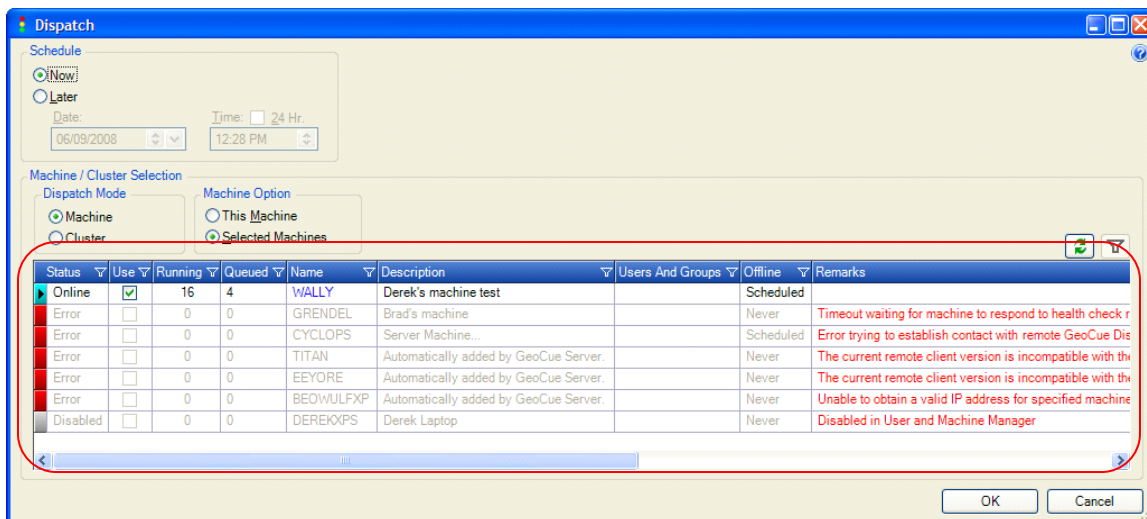


Figure 3-2: The Machine table

### 3.2.3 Cloud Mode

NOTE: You must have "Cloud" mode enabled on the Options sub-tab of the Configuration tab of Dispatch Manager to be able to select the Cloud mode during dispatch. The Cloud table will be empty if you have not defined Clouds using the Group/Cloud tab of the Dispatch Manager.

If Cloud mode is selected in the Dispatch Mode section of the Dispatch dialog, the Machine Options section of the dialog becomes disabled and the Cloud Table is displayed in the lower section of the dialog. An example of two configured clouds is shown in Figure 3-3.

Note that Cloud selection always has *radio button* behavior – that is, you can select only one Cloud as the target for the Task.

NOTE: Only Clouds that include machines that can support the command associated with the Task being dispatched will be listed in the Cloud table. This eligibility criterion is established on the Command tab of Environment Builder.

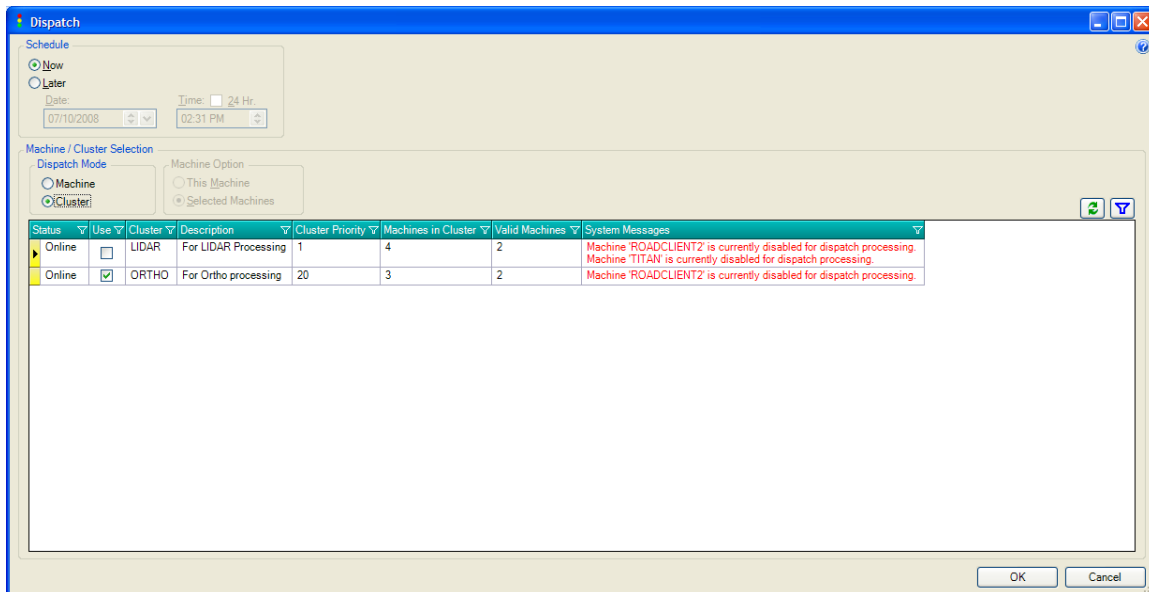


Figure 3-3: The Dispatch dialog displaying a Cloud table

To Dispatch the Task, select the desired Cloud and press OK.

### 3.3 Machine/Cloud Tables

The machine and cloud tables provide a variety of diagnostic information. An overview of the fields is provided in the follow subsections.

#### 3.3.1 Machine Table Fields

The fields of the Machine table and their descriptions are listed in Table 2-1.

**Table 3-1: Machine Dispatch Table**

<i>Field</i>	<i>Description</i>	<i>Notes</i>
Status	Idle (Green) Processing (Cyan) Disabled (Grey) Offline (Yellow) Error (Red)	The System Messages section of the dialog lists the reason for Errors. You can dispatch to machines in the states: <ul style="list-style-type: none"> <li>• Idle</li> <li>• Processing</li> <li>• Offline</li> </ul>
Use	Selector field for choosing machines	If the Task can be Dispatched but not Distributed, this field behaves in radio button fashion
Name	The name of the machine	This corresponds to the Machine name in the Machine tab of the Dispatch Manager
Running	This is the number of subtasks that are currently running on the machine	

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Queued	This is the number of subtasks that are queued that could potentially run on this machine	The actual machine on which a subtask will be hosted is not known until the subtask is actually dispatched.
Offline	Indicates if the machine has an offline schedule set (this is set in the Machine tab of Dispatch Manager)	“Never” means a schedule has not be set for this machine. If the field indicates “scheduled”, you can hover the cursor over the field to see a tool tip of the scheduled offline times.
Description	The description for the machine as entered in the Machine tab of Dispatch Manager	
Users and Groups	This lists the Users/Groups who can access this machine	
System Messages	This is system added information. It is usually a detailed message that relates to the reason for error conditions.	
Address	The Internet Protocol (IP) address of the machine	This is useful diagnostic information. If this IP address differs from the true IP address of the machine, CDS will not be able to communicate with this node.

### 3.3.2 Cloud Table Fields

Table 3-2: Cloud Dispatch Table

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Status	Ready (Green, Yellow if not all machines are available)  Offline (yellow)  Error (Red)	The Remarks section of the dialog listed the reason for Errors or non-available machines.
Use	Selector field for choosing Clouds	This field <i>always</i> behaves in radio button mode (you can dispatch only to a single cloud)
Cloud	The cloud name as defined in the Groups/Clouds tab of Dispatch Manager	
Description	The cloud description as defined in the Groups/Clouds tab of Dispatch Manager	
Cloud Priority	The cloud priority as defined in the Groups/Clouds tab of Dispatch Manager	
Machines in Cloud	The total number of machines that are members of this cloud	A machine is a member of a cloud if it has SPU resources assigned in the Group/Cloud resource allocation tables of Dispatch Manager

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Valid Machines	The number of machines in the cloud that could process the current Task	<p>A machine that is a member of a cloud may not be able to process the current command because:</p> <ul style="list-style-type: none"> <li>• It is offline</li> <li>• It is not configured for the command associated with the Task (this definition performed using the Command tab of Environment Builder)</li> </ul>
System Messages	This field contains detailed system status messages	Examine this field for information about why not all machines are valid.

### 3.3.3 Override Auto-Start Next Step

The GeoCue processing checklist allows you to configure steps such that if you execute a particular step and that step completes successfully, the next step will execute automatically. This capability is graphically indicated in the checklist by a green triangle on the step type indicator (see Figure 3-4).

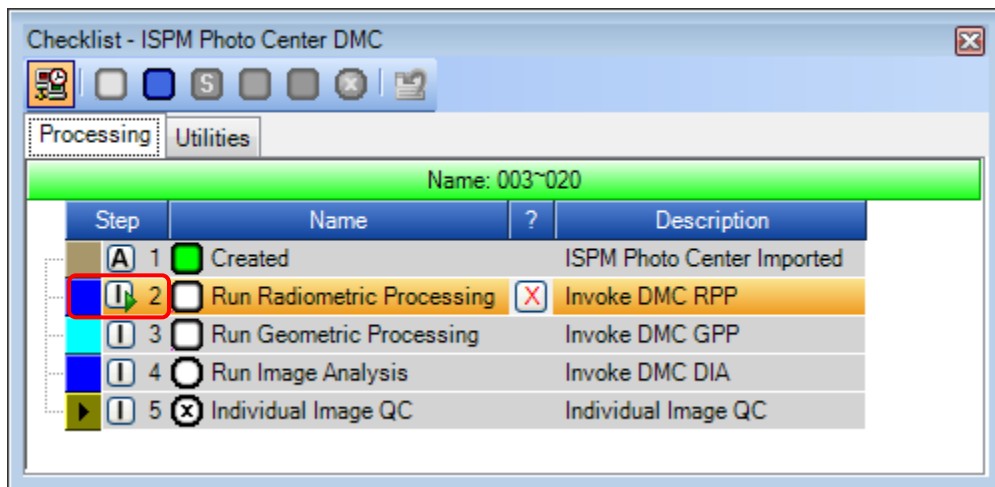


Figure 3-4: Auto-run next step indicator

You can temporarily disable this function for the set of subtasks currently being dispatched by checking the box on the Dispatch Dialog (see Figure 3-5)

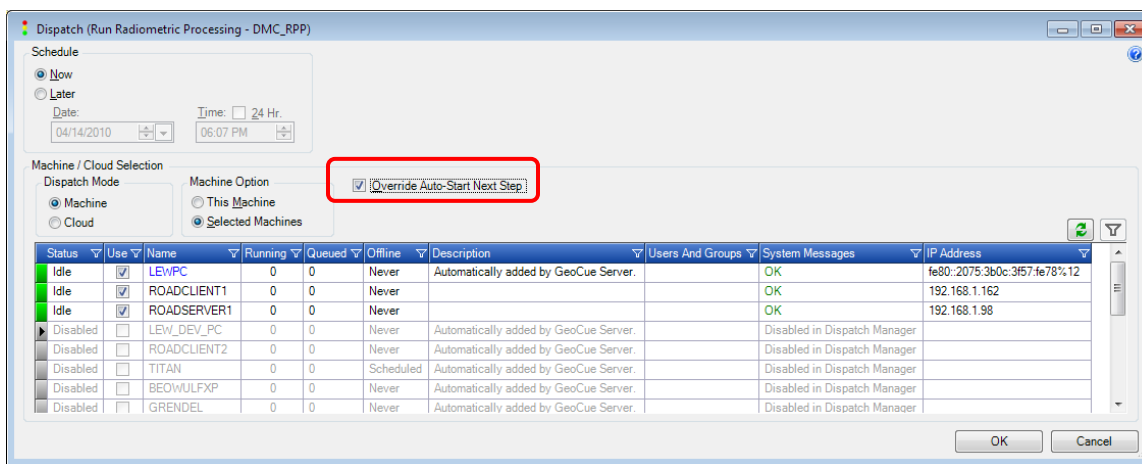


Figure 3-5: Override Auto Start Next Step

### 3.4 Command Dispatch System Tool Tray Tool

Workstations on which GeoCue Client has been installed display a Command Dispatch System icon in the system tray (the lower right area of your desktop) – see Figure 3-6. The primary purpose of this tool is to allow a user to easily block their machine from being used by the Command Dispatch System. This is useful if you run a large dispatch job that includes a user's workstation and they need the full resources of their machine. Note that effective with the Version 5.0 release of CDS, this tool sets a machine to "offline" rather than "disabled."

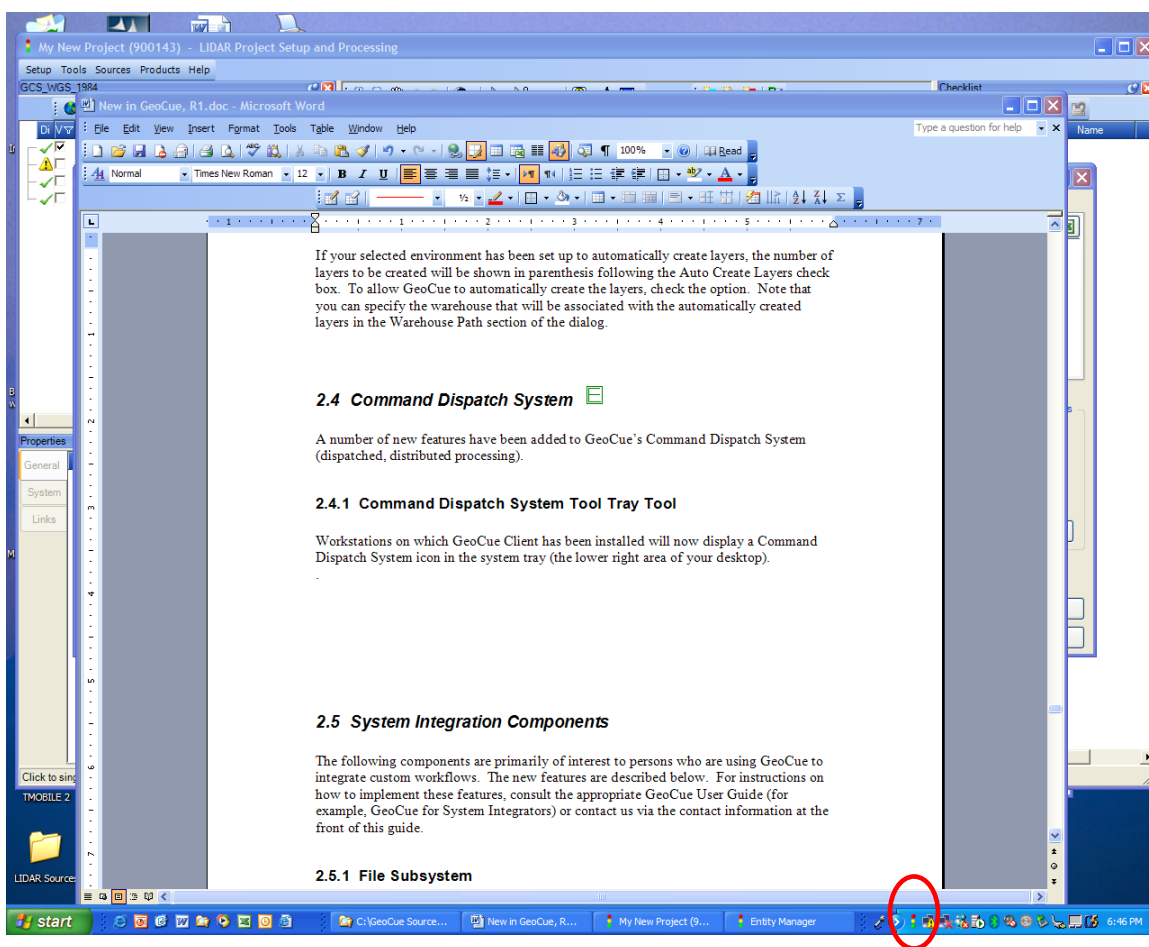


Figure 3-6 GeoCue Command Dispatch System tool



Hovering your cursor over this icon will tell you the status of any dispatch jobs that happen to be running on your Client machine as a tool tip (Figure 3-7). If your client is running a dispatch job in the background, this tip will inform you of the number of current tasks running on your node. If the Dispatch tool tip indicates "Idle", this is only for your machine; Command Dispatch System may be running tasks on other machines. You can always examine the full status by using Dispatch Monitor.



Figure 3-7 Dispatch tool tip

When GeoCue is processing a task on your node, the icon will change as indicated in Figure 3-8.



Figure 3-8 Dispatch In Progress icon

Hovering your cursor over the icon during dispatch processing will provide a tool tip status (Figure 3-9) as to the number of tasks currently executing on your node. This is typically the product of Cores X Processors as set up for your machine in User/Machine manager.



Figure 3-9 Tool Tip for a dispatch task In-Progress

Right Clicking the tool will provide you a menu with three selections:

1. Dispatch Monitor...
2. Set Offline/Online
3. Exit

Select **Take Offline** to block your workstation from accepting Dispatched commands. A machine that is offline for dispatch will display the icon of Figure 3-10.




Figure 3-10 GeoCue Dispatch disabled for this node

If you have taken your machine off-line, the selection on the tool changes to “Bring Online.” Simply press this selection to re-enable this node for dispatching.

Selecting **Exit** removes the tool from your tool tray. To restore the tool, simply exit and restart GeoCue Client.

NOTE: Setting a machine to “offline” using the tool tray Dispatch tool will not cause the machine to immediately cease processing any currently executing subtasks. The current subtasks will always complete. This is necessary in order to maintain project integrity. For example, if a set of LIDAR Stereo models have been dispatched to a number of machines for processing and you take your machine

offline, the model that is currently being generated on your machine will complete prior to your machine being removed from the list of eligible processing nodes.

### ***3.5 Monitoring Dispatched Tasks***

Running or queued tasks are monitored using the Tasks tab of Dispatch Manager. Dispatch Manager can be invoked either from the GeoCue tool tray icon on a desktop workstation, from the Tools menu of a GeoCue Client or from the Windows Start menu, GeoCue program group. Figure 3-11 shows the Task tab of Dispatch Manager. There are three major informational sections on this tab of the Dispatch Manager.

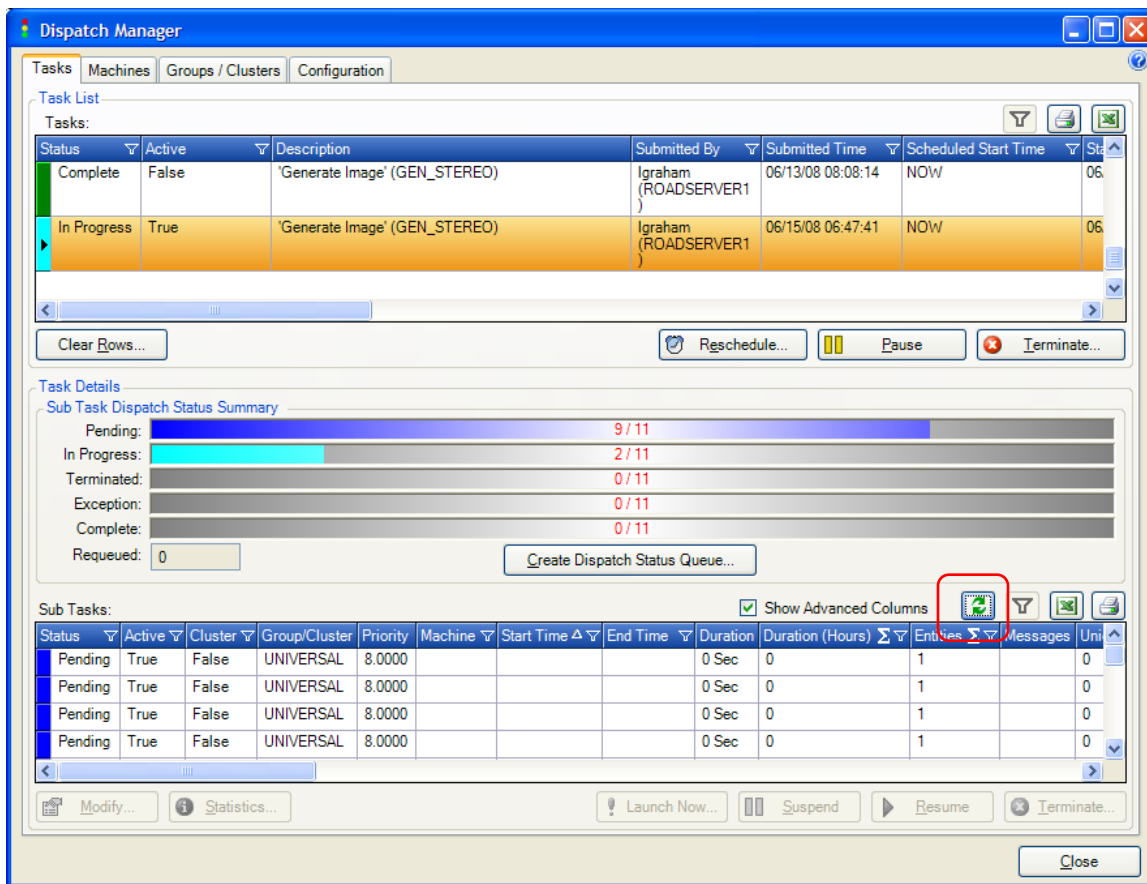


Figure 3-11: The Task tab of Dispatch Manager

### 3.5.1 Tasks Table

The upper section of the dialog shows dispatched tasks. You can filter this list to show only *Active* tasks (tasks in the Pending or Running state) by setting the filter on the Active column to True. You can delete rows by *selecting* the row and pressing the *Clear Rows...* button. Note that you cannot clear rows of Pending or In Progress Tasks.

The information contained in the Task table is detailed in

**Table 3-3: Task Table**

<b>Item</b>	<b>Description</b>	<b>Notes</b>
Status	The current status of this Dispatch Task	One of: Pending In Progress Exception Terminated (meaning a user explicitly terminated the command)
Active	True for Pending and In-Progress Tasks	Use this column to filter the Task dialog to show only active/pending tasks
Description	This is the Checklist Step name with the actual command from Environment Builder parenthetically indicated.	
Submitted by	The User: Machine that submitted this Task	
Submitted Time	Date/Time the Task was submitted	
Scheduled Start Time	Date/Time the Task is scheduled to start if the user elected Scheduling from the Dispatch dialog	NOW indicates the Task was submitted for immediate execution.
Start Time	The Date/Time the first subtask was dispatched to a machine	
End Time	The Date/Time the last subtask completed	For <i>dispatched</i> tasks there will be only one subtask
Duration	The time, in hours, minutes and seconds that the subtasks associated with the Task spent in execution.	

<b>Item</b>	<b>Description</b>	<b>Notes</b>
Duration (Hours)	The same as Duration but in decimal hours	Use this field to sort on time
Machines	The machines to which this Task could be dispatched	
Cloud	True if the dispatch mode was Cloud	
Group/Cloud	The name of the Group (or Cloud if the previous field was True) for this Task.	If you have not set any advanced dispatch modes, this field will indicate the UNIVERSAL group.
Project	The Project from which this Task was dispatched	
Environment	The Environment that was set in the Project from which this Task was dispatched	
Entities	The number of entities associated with this Task	
Messages	System messages (if any) are displayed in this column	
Group Access Code	This is a system code used by the <i>stochastic</i> re-queuing system	This is used for diagnostic purposes by GeoCue technical staff
User Params	Any additional parameters entered by the user in the Checklist Step configuration dialog	
Override Working Directory	If not blank, the target machine's working directory will be set to this path.	Primarily used on UNIX systems
ID	A unique integer ID associated with this Task	

### 3.5.2 Sub Task Dispatch Status Summary

The progress bars provide a visual indication of activity in the overall dispatch system. These bars reflect the information listed in Table 3-4.

**Table 3-4: Task monitor progress bars**

<b>Bar</b>	<b>Description</b>	<b>Notes</b>
Pending	This is the total number of currently pending subtasks in the CDS.	A Dispatched task results in a single subtask (such as Populate Working Segments). A distributable task results in <i>n</i> subtasks where <i>n</i> is typically the number of entities associated with the command.
In-Progress	The number of subtasks that are currently dispatched to machines and running	
Terminated	The number of subtasks that have been terminated by a user.	
Exception	The number of subtasks that have caused an error condition	
Complete	The number of completed subtasks	

The Requeued field indicates the number of Subtasks that have been Requeued. This occurs when one subtask needs the resources of another subtask but the other subtask is currently processing (an example is processing TerraScan macros with an over edge requirement). In this case, the subtask will be aborted and placed back in to the Pending queue.

You can create a GeoCue Name Queue associated with the various statuses of subtasks. Select the Task and then press the Create Dispatch Status Queue. This feature is very useful for actions such as locating all of the subtask entities that resulted in an error in the GeoCue Map View.

### 3.5.3 Subtasks Table

The lower section of the dialog shows Subtasks. The lower section is populated by selecting one or more tasks in the upper section of the dialog and pressing the Refresh button at the upper right of the Sub Task pane. Note that you must press the refresh button each time you want to review the list – it is not automatically updated.

Recall that if a Task will be *distributed*, it will spawn  $n$  subtasks where  $n$  is generally the number of entities associated with the Task. If a Task cannot be distributed but only *dispatched*, it will result in the creation of a single subtask.

The various fields of the Subtask Table are described in Table 3-5. The fields highlighted in light blue display only when the “Show Advanced Columns” box is checked.

Table 3-5: Subtask Table

<i>Item</i>	<i>Description</i>	<i>Notes</i>
Status	This is the total number of currently pending subtasks in the CDS.	A Dispatched task results in a single subtask (such as Populate Working Segments). A distributable task results in $n$ subtasks where $n$ is typically the number of entities associated with the command.



<b>Item</b>	<b>Description</b>	<b>Notes</b>
Active	True for Pending and In-Progress Subtasks	Use this column to filter the Subtask table to show only active/pending tasks
Cloud	True if the dispatch mode was Cloud	
Group/Cloud	The name of the Group (or Cloud if the previous field was True) for this Task.	If you have not set any advanced dispatch modes, this field will indicate the UNIVERSAL group.
Priority	The subtask priority	See the section on the Priority Computation Algorithm (PCA)
Machine	The machine to which the subtask was dispatched	Blank if a subtask is pending since, in general, the machine on which a subtask will run is not known until it is actually dispatched
Start Time	The Date/Time the subtask was dispatched to a machine	
End Time	The Date/Time the subtask completed	
Duration	The time, in hours, minutes and seconds that the subtask spent in execution.	
Duration (Hours)	The same as Duration but in decimal hours	Use this field to sort on time
Entities	The number of entities associated with this subtask	Generally one for <i>Distributable</i> commands and <i>n</i> for <i>Dispatched</i> commands
Messages	System generated status messages	
Unique Processing ID	This is a system ID used by the stochastic processing subsystem	

<b>Item</b>	<b>Description</b>	<b>Notes</b>
Task Description	A description of the Task associated with this subtask	This description is formed by combining the name of the processing step used to launch this subtask with the name of the command from the command table – Example: "Generate Image (GEN_STEREO)"
Task ID	The ID of the Task to which this subtask belongs	The same as the Task ID in the Task table
Subtask ID	A unique integrate ID associated with this subtask	
Command	The command associated with this Task/subtask	The command from the Command table of Environment Builder associated with this subtask
Machine Process ID	This is the Process ID (PID) of the actual process on the machine that is running this command. For example, this is the Process ID that you would see in Windows Task Manager for this subtask execution.	This value is only displayed for processing machines equipped with a Dynamic Resource Interrogation Module (DRIM)
Accumulated CPU Time	The amount of CPU time used by this subtask on the target machine	DRIM only
Command Line Arguments	The command line arguments passed with the subtask	This field is only populated after the subtask is dispatched (it is blank for <i>pending</i> commands)
Messages	System messages (if any) are displayed in this column	

<b>Item</b>	<b>Description</b>	<b>Notes</b>
Priority Override	"True" if a user has overridden the priority using Dispatch Manager.	Priority override can only be performed by a GeoCue Administrator.
Size	The current amount of memory being used by this subtask	DRIM only
Predicted Size	The predicted amount of memory this subtask will use	Set in Environment Builder
Size Equilibrium Time	The time from the start of the subtask until its memory usage reached the predicted size	Set in Environment Builder
Auto Kill Size	The amount of memory usage that will trigger the autokill system.	DRIM only
CPU Usage	The percent of the host machine's CPU resources that are being used by this subtask	DRIM only
Predicted CPU Usage	The predicted percentage of CPU resources that will be consumed by this subtask	Set in Environment Builder
CPU Usage Equilibrium Time	The length of time from the start of the subtask until it is consuming the predicted CPU percentage	Set in Environment Builder
Reserve Resources	The Command associated with this Task has Reserved Resources	
Statistics Sampling	The frequency with which the DRIM for this subtask will gather statistics	DRIM Only

### 3.6 Clearing Rows from the Tasks Pane

You can clear rows from the Tasks pane of the Tasks tab of Dispatch Manager by selecting the rows that you wish to clear and pressing the **Clear Rows...** button at the lower left of the pane. Note that once a row(s) is cleared, it cannot be restored. Clearing a row has no affect on the processing history of the entity or entities associated with the Task.

Note that you cannot clear a Pending Task.

### **3.7 Rescheduling a Task**

You can reschedule a task by selecting the row of the Task in the Tasks pane of the Tasks tab of Dispatch Manager and pressing the **Reschedule...** button. Note that you cannot reschedule a task that has already started, completed or that has been terminated.

### **3.8 Pausing a Task**

You can *pause* a task by selecting the row of the Task in the Tasks pane of the Tasks tab of Dispatch Manager and pressing the **Pause...** button. Note that you cannot pause a task that is scheduled, completed or that has been terminated.

### **3.9 Terminating a Task**

You can *terminate* a task by selecting the row of the Task in the Tasks pane of the Tasks tab of Dispatch Manager and pressing the **Terminate...** button. Note that you cannot terminate a task that is complete. If you terminate a task that is in progress, the subtasks that are currently executing will complete; only the subtasks that have not yet started will be terminated. Note that since a Dispatched task executes all subtasks on the same machine, a Dispatched task cannot be terminated once it has started.

### 3.10 Modifying a Subtask (GeoCue Administrators Only)

You can modify a subtask(s) by selecting the subtask rows in the Subtask pane of the Task tab of Dispatch Manager and pressing the **Modify...** button. You must be a GeoCue Administrator to access the Modify dialog). This will invoke the dialog of Figure 3-12.

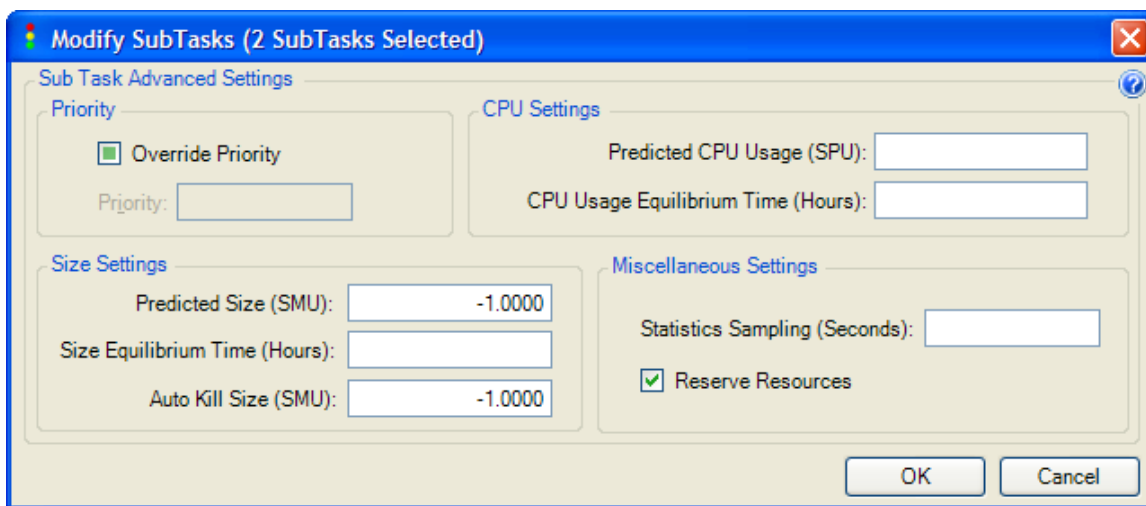


Figure 3-12: Modify Subtask dialog

As with all multi-select dialogs in GeoCue, if a field is blank, the selected objects do not currently share a common value.

The most common use of the override is to modify the subtask(s) priority.

### 3.11 Creating Entity Queues

You can create *queues* of entities associated with a Task. This is very useful for operations such as adding *failed* subtasks back to the working set for reprocessing or for examining the error status code for these entities.

The queue creation commands are on the Task tab of Dispatch Manager under the Sub Task Dispatch Status Summary progress bars.

Select the row of the Task for which you wish to create a queue (you cannot create queues for multiple Tasks simultaneously). Press the **Create Dispatch Status Queue...** button. This displays the dialog of Figure 3-13. Fill in a name for the queue and select the type of status using the Entity Dispatch Status drop-down selector. In the example of Figure 3-13 we are creating a queue of the subtasks that caused an exception. Press OK to create the queue.

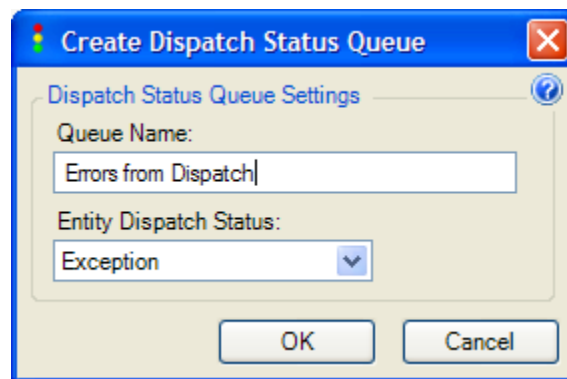


Figure 3-13: Create Dispatch Status Queue dialog

This queue will now be available in GeoCue Clients that are accessing the *project* from which this Task was dispatched. The example queue is depicted in Figure 3-14.

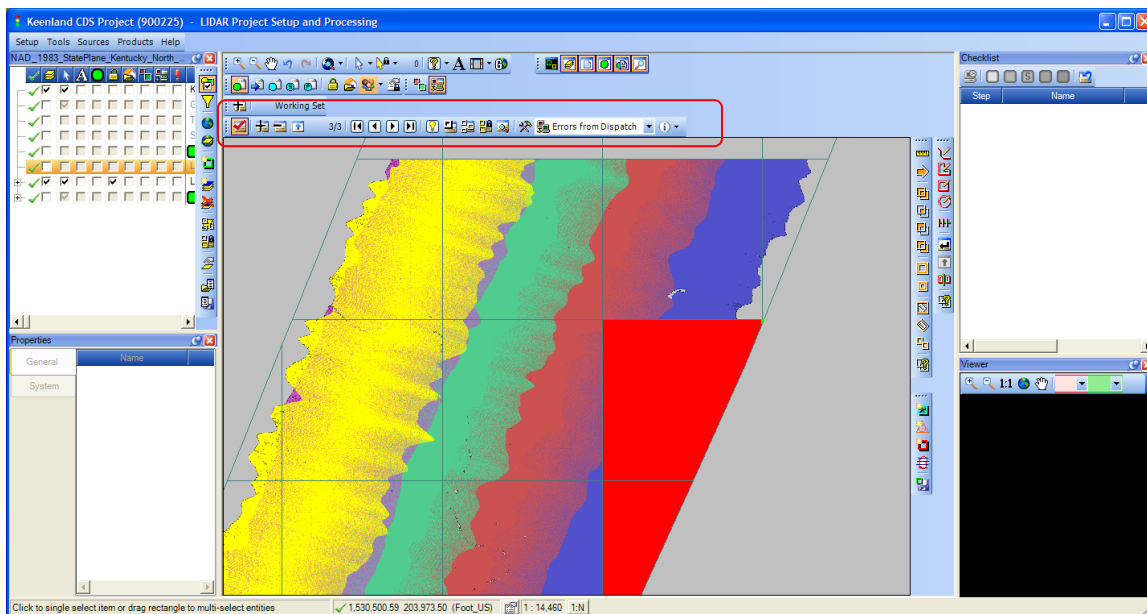


Figure 3-14: The Named Queue accessed from GeoCue Client

### 3.12 Monitoring Machine Activity

You can monitor machine activity using the Machine tab of Dispatch Manager.

NOTE: If you are not a GeoCue Administrator (set in the GeoCue User Manager dialogs), the Groups/Clouds and Configuration tabs of Dispatch Manager will be disabled. Additionally, all of the modification capabilities of the Machines tab will be disabled.

Figure 3-15 depicts the Available Dispatch Machines pane of the Machines tab of Dispatch Manager. Note that you must press the *refresh* button to see the latest status of Machines.

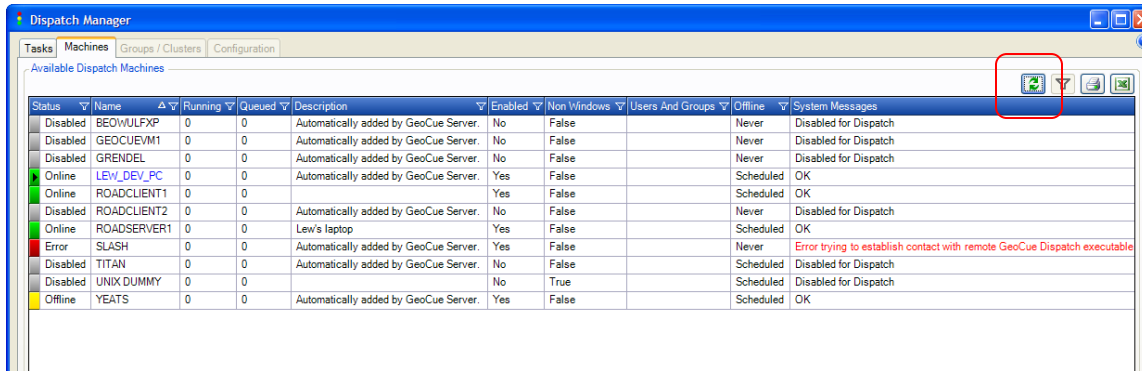


Figure 3-15: Available Dispatch Machines pane of Machines tab of Dispatch Manager

The fields germane to machine monitoring as well as their descriptions are listed in Table 3-6.

Table 3-6: Machine Monitoring Fields

Field	Description	Notes
Status	Idle (Green) Processing (Cyan) Disabled (Grey) Offline (Yellow) Error (Red)	The System Messages section of the dialog lists the reason for Errors.
Name	The name of the machine	This corresponds to the Machine name in the Machine tab of the Dispatch Manager
Running	This is the number of subtasks that are currently running on the machine	This is an instantaneous snapshot so it can momentarily fluctuate.



<b>Field</b>	<b>Description</b>	<b>Notes</b>
Queued	This is the number of subtasks that are queued that could potentially run on this machine	The actual machine on which a subtask will be hosted is not known until the subtask is actually dispatched.
Description	The description for the machine as entered in the Machine tab of Dispatch Manager	
Enabled	Indicates if the machine has been disabled	A machine can be disabled via this tab of Dispatch Manager (by a GeoCue Administrator) or via the tool tray icon on the machine itself.
Non-Windows	True if this is a LINUX/UNIX machine	
Users and Groups	This lists the Users/Groups who can access this machine	
Offline	"Never" of "Scheduled"	Hover the mouse cursor over "scheduled" fields to display the schedule as a tool tip.
System Messages	This is system added information. It is usually a detailed message that relates to the reason for error conditions.	
Address	The Internet Protocol (IP) address of the machine	This is useful diagnostic information. If this IP address differs from the true IP address of the machine, CDS will not be able to communicate with this node.

## 4 Configuring Machines

In this chapter we discuss the various attributes that can be configured for processing machines that will participate in your GeoCue Command Dispatch System constellation. Note that configuring Processing Groups is discussed in the next chapter.

### 4.1 Adding Machines to the CDS

GeoCue dispatch licenses are installed on machines as follows:

- Windows Machines – Simply install GeoCue Client on all machines that you would like to participate in the Command Dispatch System.
- UNIX Machines – see the UNIX installation guide.

Windows machines are automatically added to the Machines tab of Dispatch Manager. UNIX machines must be explicitly added via the Machines tab of Dispatch Manager, **Add** button.

You can explicitly add a machine to your CDS constellation using the **Machines** tab of Dispatch Manager. Dispatch Manager is invoked either from the GeoCue program group on the Windows Start menu or by right clicking the GeoCue tool tray icon and selecting Dispatch Manager.

Select the Machines tab (Figure 4-1).

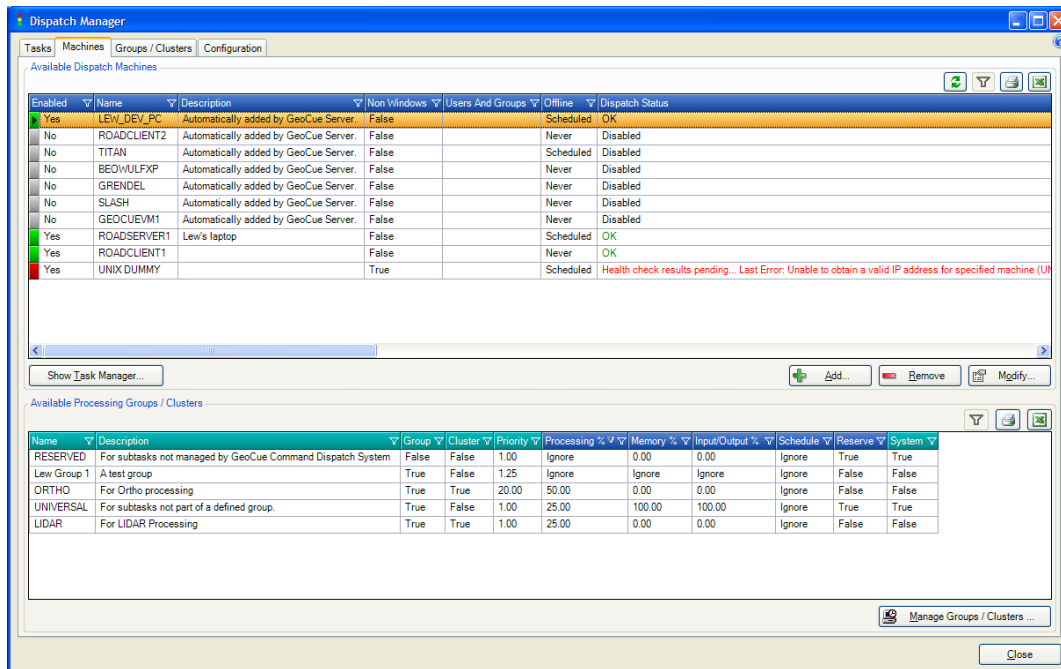


Figure 4-1: Machines tab of Dispatch Manager

If a machine is not listed, it can be explicitly added by pressing the **Add ...** button. This invokes the dialog of Figure 4-2.

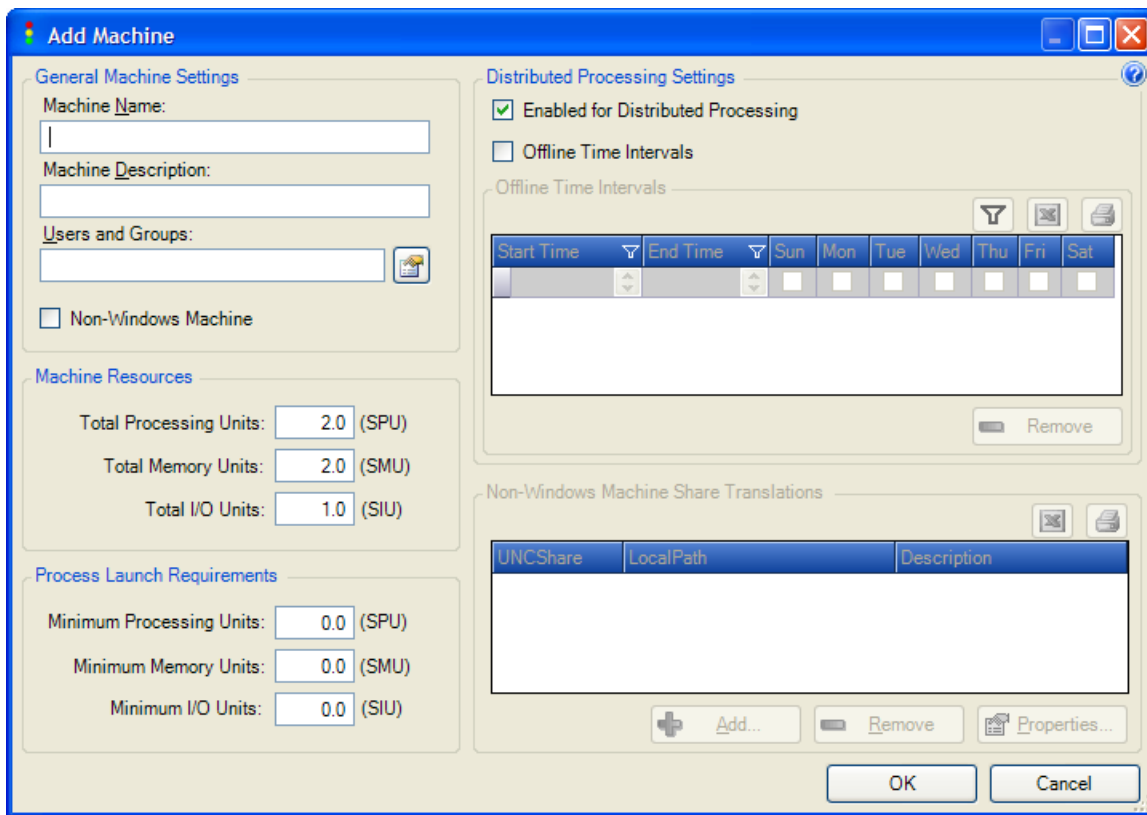


Figure 4-2: Add Machine dialog

Add the name of the machine in the Machine Name field. Note that for Windows machines, you can discover a machine's name by viewing the "Computer Name" tab of the My Computer, Properties dialog in Windows (right click My Computer from the Start menu and select *properties* from the right click menu).

The remainder of the fields on the Add Machine dialog will be discussed in subsequent sections.

## 4.2 Machine Resources

You can specify three aspects of the resources of a particular machine:

- Processor “power” (Standard Processor Units, SPU)
- Available memory (Standard Memory Units, SMU)
- Input/Output “capacity” (Standard Input/Output Units, SIU)

Each of these is characterized in terms of a “Standard Unit.” The definition of Standard Units is entirely up to you, the CDS Systems Administrator. We recommend using the normalization values listed in the Terminology chapter and repeated in Table 4-1 below. Note that commands included in GeoCue CuePacs are normalized using these values!

**Table 4-1: Recommended Standard Units**

<i>Metric</i>	<i>Reference Unit</i>
1 SPU	1 GHz Xeon, Single Core
1 SMU	1 GB RAM
1 SIU	1 MB/second of I/O bandwidth

These values are used in determining the possible loading of a machine by the CDS. Each Command within GeoCue is assigned values for their resource requirements in terms of SPU, SMU and SIU.

Obviously, once you have decided upon your own normalization scheme, you will have to be consistent throughout your GeoCue constellation including changing the load factors of Commands shipped with GeoCue (and third party) CuePacs.

**IMPORTANT:** If you change the GeoCue defined loading factors of Table 4-1, you *must* redefine the load factors of all commands that you use in the Command Table of Environment Builder. Failure to do this will result in either too many instances of commands executing on a machine or blocking of the command from executing at all.

### 4.3 Setting Machine Configuration Parameters

Machines are configured by selecting the machines that you wish to configure in the **Machines** tab of Dispatch Manager, Available Dispatch Machines pane and pressing the **Modify...** button. This displays the dialog of Figure 4-2. Note that you can select multiple machines if you wish to set properties for more than one machine simultaneously. The fields and their descriptions are listed in Table 4-2.

**Table 4-2: Machine Settings**

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Machine Name	Windows Operating System - The Windows name of the machine as shown in My Computer.  Unix Operating System – The machine name as configured in your UNIX system	Note – Once set, this field cannot be modified.
Machine Description	A text field used to add information about the machine	

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Users and Groups	Set values in this field to restrict the machines to a particular set of users and/or groups <sup>3</sup> .	Blank means that the machine can be used by any user or group.
Non-Windows Machine	Check this box if the machine is a UNIX system	Checking this box will enable the "Non-Windows Machine Share Translations" section of the dialog.
Enabled for Distributed Processing	When checked, the machine is enabled for the Command Dispatch System	You can uncheck this box to take a machine off-line. All other setting will be preserved so you will not lose the machine schedule.
Offline	When checked, the machine is taken off-line. When unchecked, the machine is brought back on line so long as it is not in a scheduled off-line period.	Useful for immediately taking a machine off-line or restoring a machine to on-line that has been taken off-line by a local machine tool tray tool.
Offline Time Intervals	This section allows you to set a calendar of times in which this machine will <i>not</i> be available to the CDS.	This is useful for allowing the CDS to access machines in low usage times (evenings, for example) that are otherwise heavily loaded. You can also use this feature to prevent CDS from attempting to access machines during backup times.
Total Processing Units	The total available processing capacity of the machine in terms of Standard Processor Units (SPU)	The minimum value for this field is 0.1

<sup>3</sup> Note that Permission Groups (set in the User/Group Administration tool of GeoCue) are quite different from Processing Groups (set in the Configuration tab of Dispatch Manager)

<b>Field</b>	<b>Description</b>	<b>Notes</b>
Total Memory Units	The total available memory capacity of the machine in terms of Standard Memory Units (SMU)	-1.0 means that memory will not be considered when dispatching to this machine.
Total I/O Units	The total Input/Output bandwidth of the machine in terms of Standard I/O Units (SIU)	-1.0 means that Input/Output capacity will not be considered when dispatching to this machine.
Minimum Processing Units	The machine must have at least this much SPUs remaining for a task to be launched.	This value is computed assuming the subtask was launched. For example, if a Machine has 2.0 SPUs available <sup>4</sup> and a minimum launch value of 0.5, a command that required 1.75 SPU could not launch on the machine.
Minimum Memory Units	The same as Minimum Processing Units but for Memory	
Minimum Input/Output Units	The same as Minimum Processing Units but for Input/Output	

Share Translation settings (Figure 4-3) are used to translate a GeoCue Windows Share (for example, a warehouse path) into the share path as recognized on the UNIX machine. For example, when a Windows share is mounted via an application such as Samba, the path specification is different than that seen by Windows.

---

<sup>4</sup> Note that the available Processor, Memory and I/O are dynamically determined based on the current CDS load of a machine.



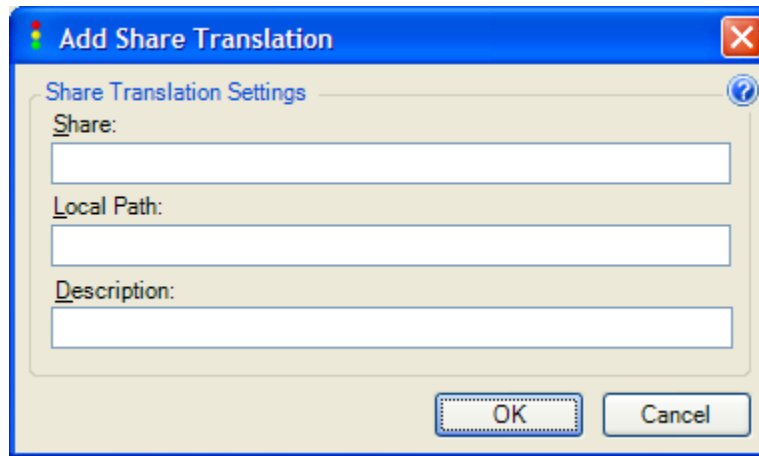


Figure 4-3: Share Translation

## 4.4 Machine Schedules

Machine Schedules are set via the Offline Time Intervals (Figure 4-4). You can set as many independent schedules as desired in this section. Note that if you selected multiple machines when launching the Properties dialog, the same schedule will be set for each.

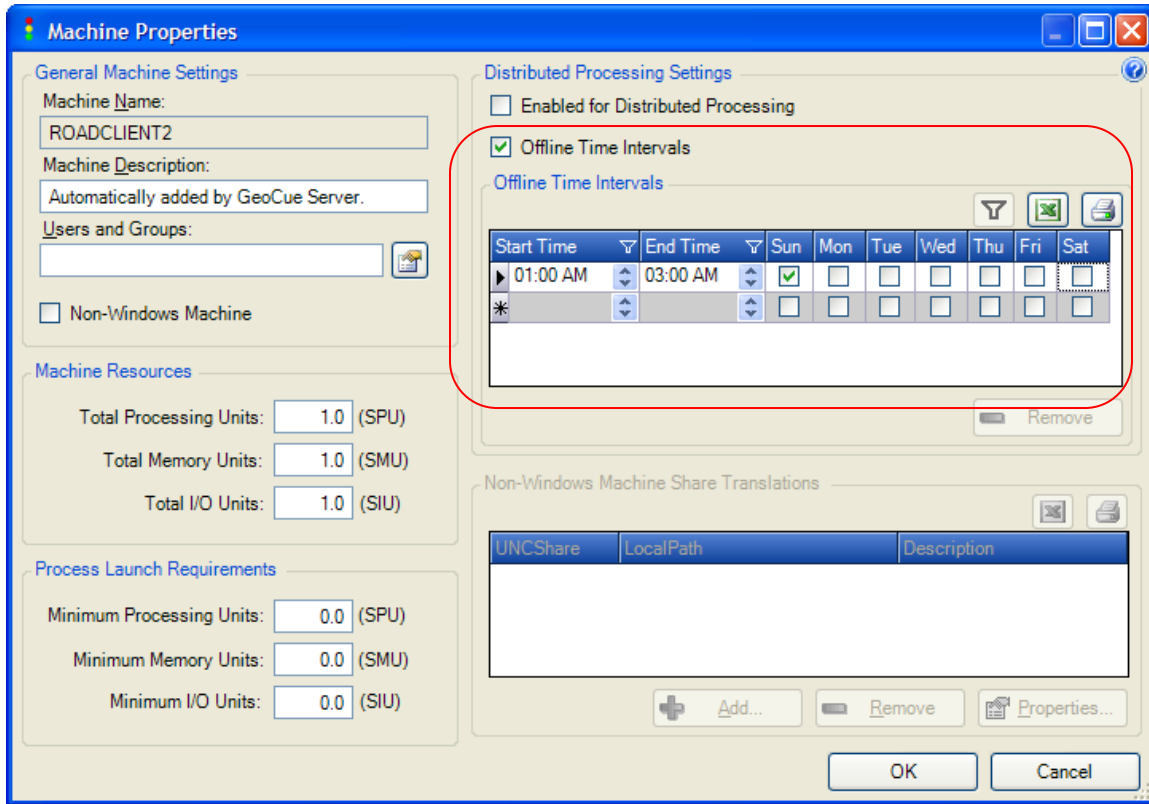


Figure 4-4: Setting Off-line times

Set the start and end times (the meaning is the start and end of the off-line time) and check the days of the week applicable for this time.

Remove a time by selecting the row containing the schedule you wish to remove and pressing the **Remove** button. Note that you can multi-select rows when removing.

## 4.5 Machine Loading

The Command Dispatch System now supports the idea of Machine Priority. The normal mode of dispatch scheduling is a "Round Robin" algorithm under which machines are loaded with subtasks in round robin fashion. We refer to this as **Symmetric** loading.

You can also elect to load machines in an Asymmetric fashion. This is achieved by setting a machine to **Asymmetric Loading** and specifying a priority (see Figure 4-5).

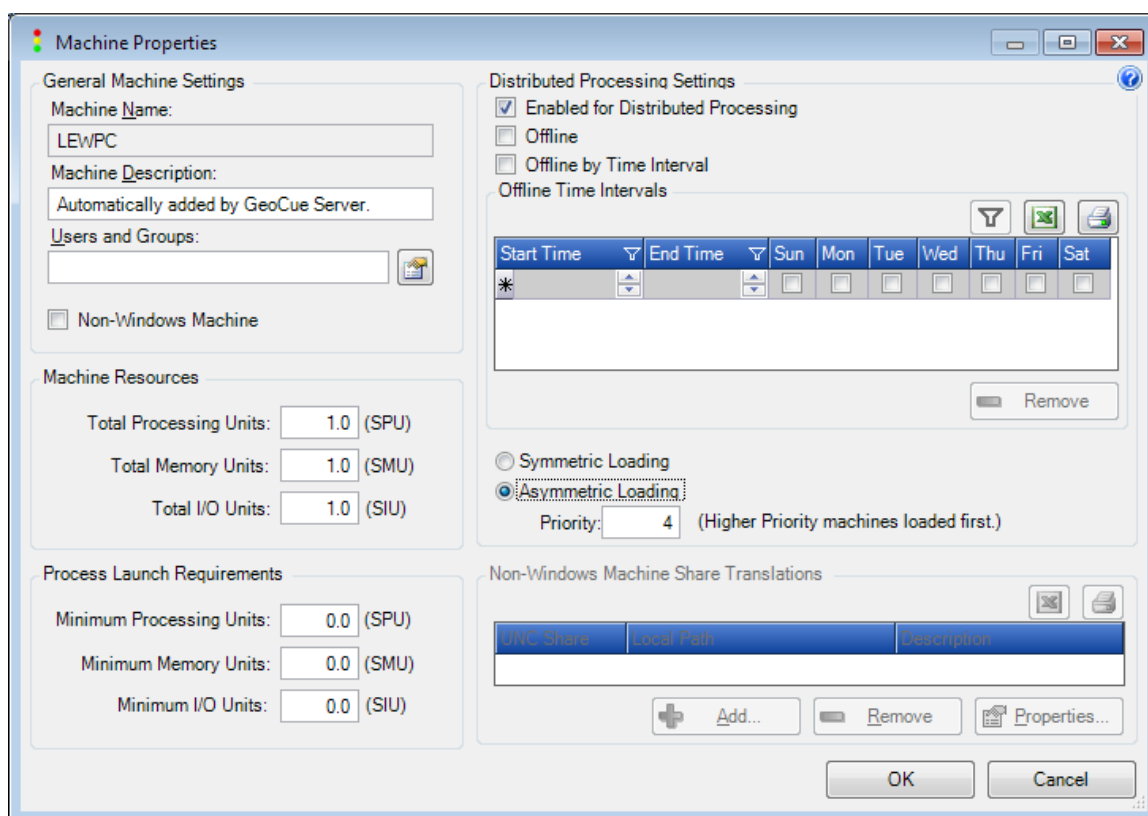


Figure 4-5: Asymmetric Machine Loading

In the asymmetric mode, machines with higher priority (higher numerical value) will be loaded with subtasks prior to those with lower number. The priority of all machines set to Symmetric mode is 1 and thus they are at the lowest priority.

Consider the following example, A DMC PPS Cloud has a server machine with a specified Standard Processing Units (SPU) of 4. The technician has a workstation with a specified SPU of 2. The user would like to tasks to always run on the server until it is “maxed out” prior to running on the workstation. This could be accomplished by setting the Asymmetric Loading check for the server and setting its priority to a value of 2 or more. This will cause the server to load with 4 subtasks (assuming the executed subtask is set to requiring 1 SPU) before any subtasks are dispatched to the workstation.

Note that Asymmetric loading has little effect when a large number of subtasks are queued since, even though the machines will load in priority order, they all eventually become fully loaded. Thus this mode is most useful for causing “casual” dispatches to be sent to preferential servers.

**IMPORTANT:** The Command Dispatch System contains tools to allow you to specify the machines on which specific commands can run. Thus if you would like for Command A to be able to run on both a particular workstation and a Server, but command B to run only on the Server, this can be specified. See the C section of this manual.

## 4.6 Machine Group/Cloud Membership

The lower pane of the Machines tab of Dispatch Manager lists the Groups and Clouds that have been configured in your Command Dispatch System. When you select a machine or machines in the top pane, the columns in the lower pane will populate for the selected machine(s). If you have selected more than one machine and a field in the lower pane is blank, it means that this value is different between the selected machines. Note that the groups RESERVED and UNIVERSAL are predefined factory groups that are present in all CDS deployments. You can modify the membership characteristics of the selected machines by pressing the **Manage Groups/Clouds** button at the lower right side of the Available Processing Groups/Clouds pane. The use of Groups and Clouds are explained in subsequent chapters of this manual.

## 5 Configuring Commands

Other than their assignment to Processing Groups (as described in a subsequent chapter), Commands are configured in the **Command** tab of Environment Builder. Invoke Environment Builder (from the Setup ► Administration menu of a GeoCue Client). Select the **Command** tab, **Command** sub-tab. This will present a list of all commands in your GeoCue installation (Figure 5-1). Note that you can filter and sort the columns to facilitate locating commands in the table.

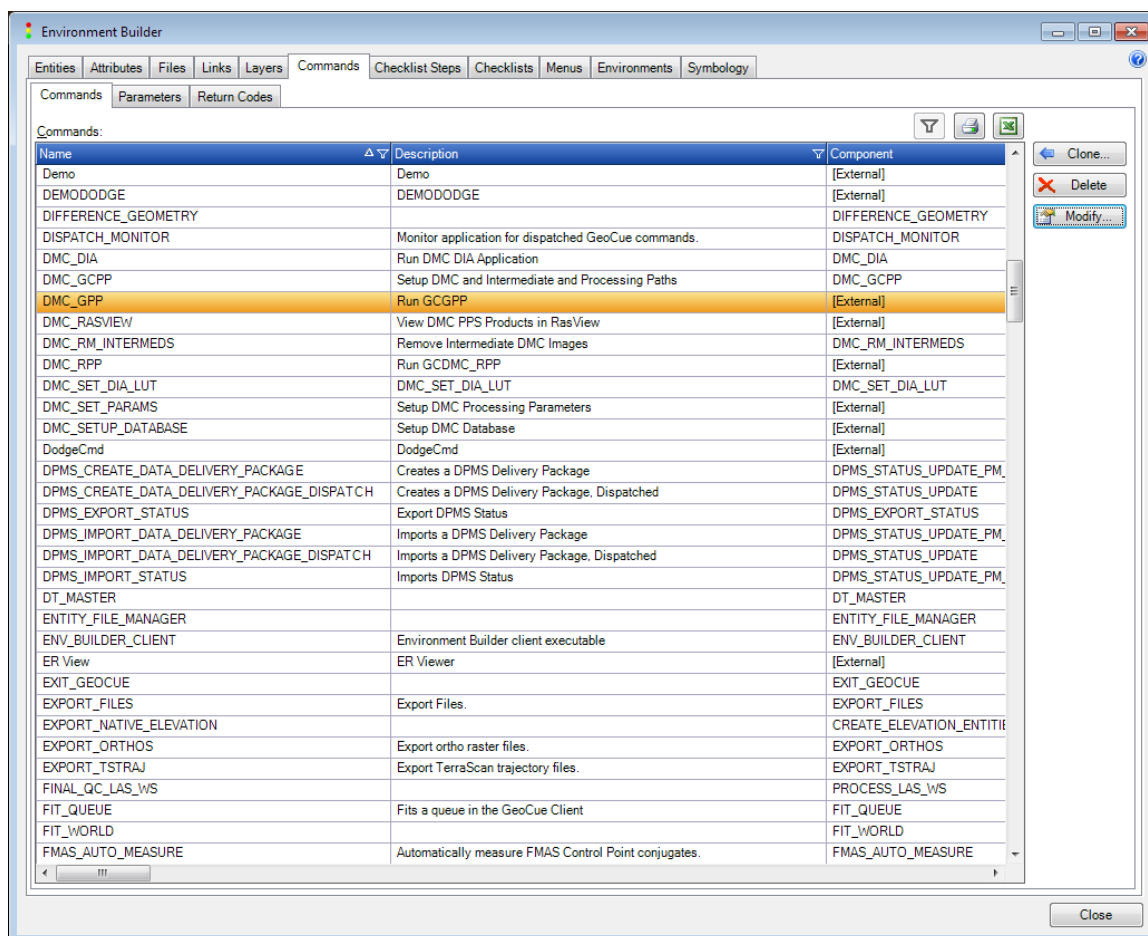


Figure 5-1: Invoking the Command functions in Environment Builder

Select the command that you wish to modify and press the **Modify** button. This displays the dialog of Figure 5-2.

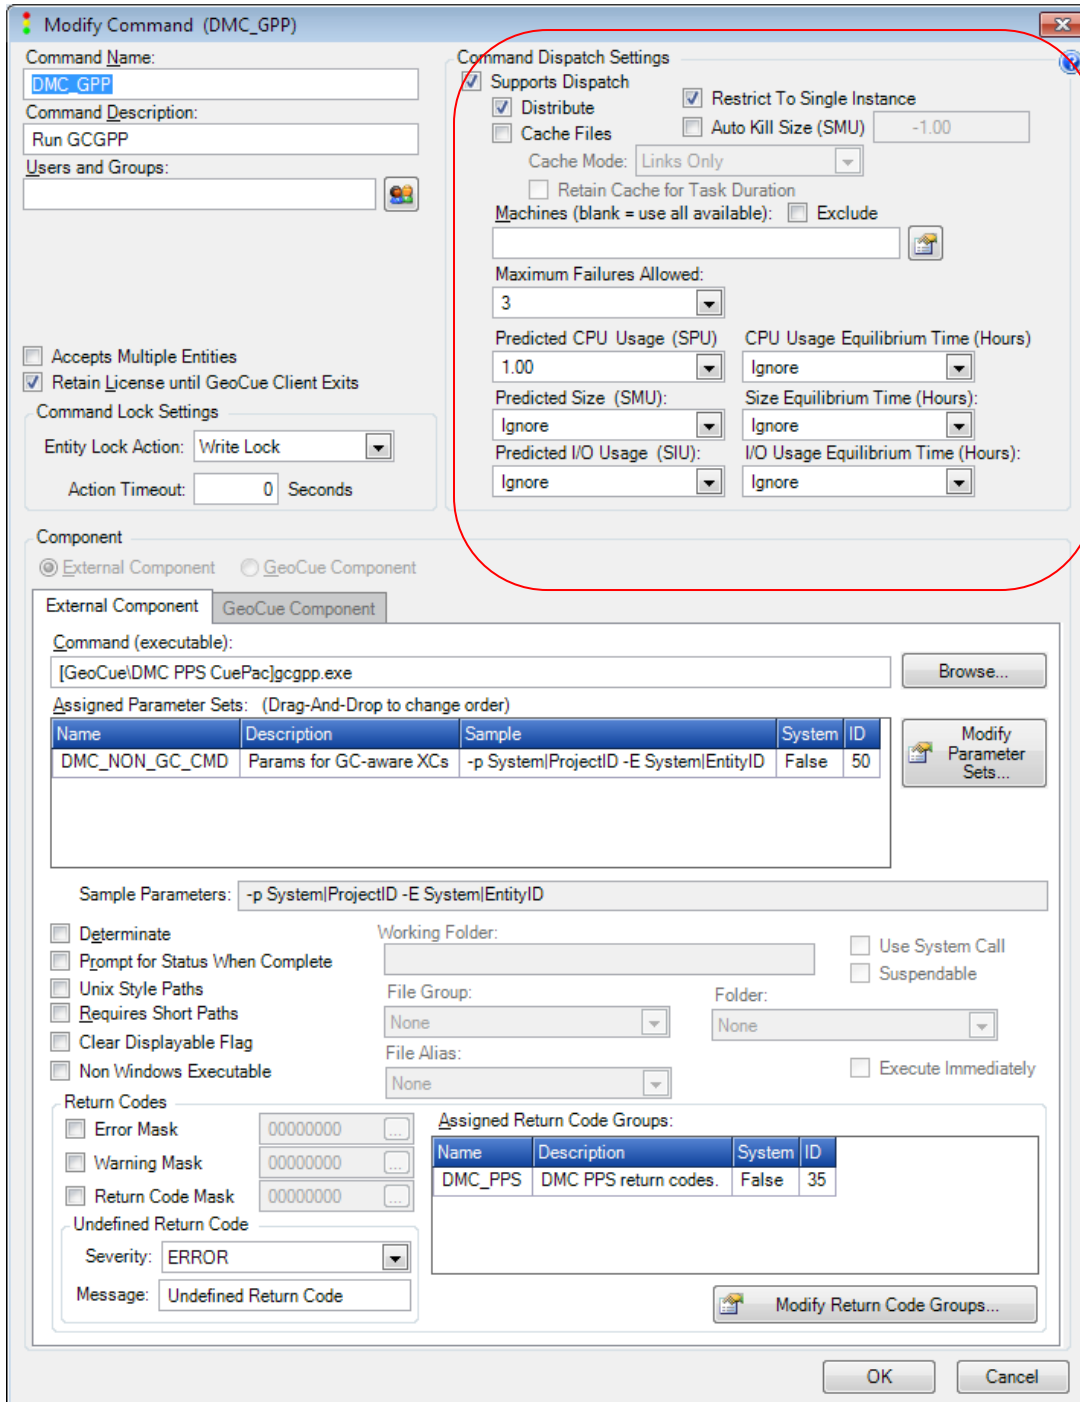


Figure 5-2: Modify Command dialog of Environment Builder

The portion of this dialog relevant to the Command Dispatch System appears in the upper right of the dialog and is presented in an expanded view in Figure 5-3. All of these values will have default settings if the command that you are modifying is a GeoCue supplied command.

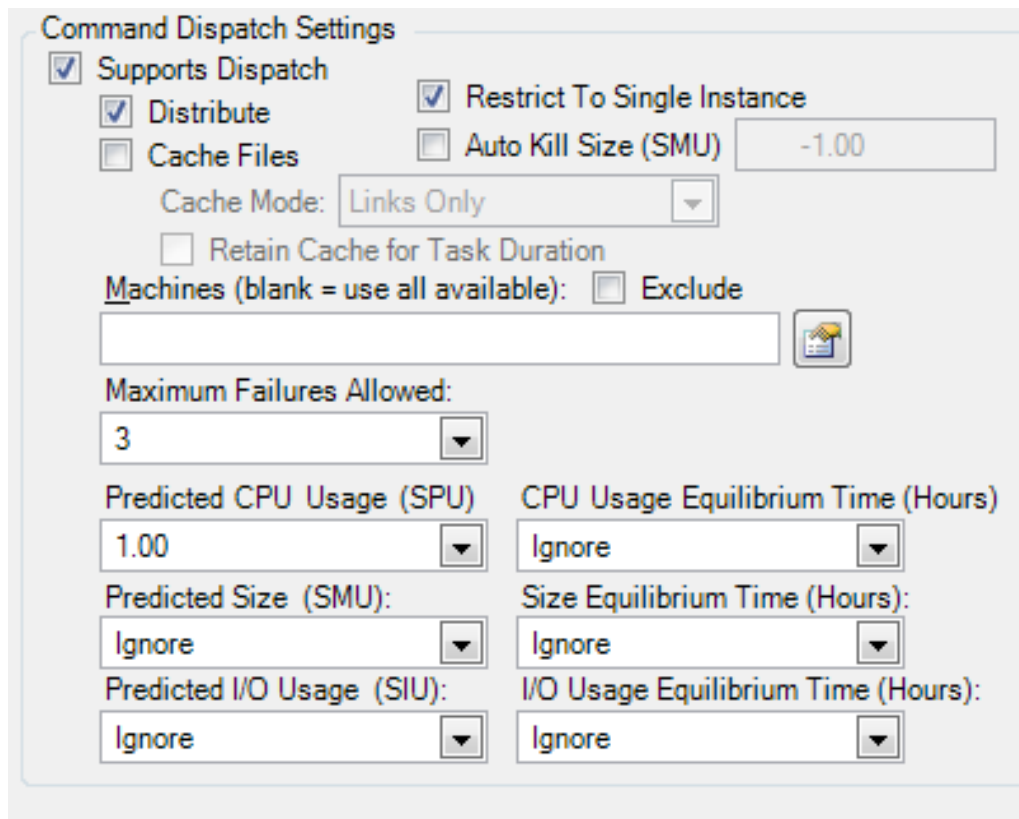


Figure 5-3: Dispatch section of the Command settings

GeoCue supports three types of commands. These types are listed in Table 5-1.

**Table 5-1: GeoCue Command Types**

<b>Type</b>	<b>Description</b>	<b>Notes</b>
Intrinsic Commands (IC)	These are commands that execute entirely within the context of GeoCue Clients.	These cannot be written by third party developers. An Intrinsic Command cannot be used in the CDS.
Extrinsic Commands (XC)	These are commands that execute outside the context of the GeoCue Client but have access to GeoCue internal functions via the GeoCue Repository API.	These commands can be written by third party developers. Authoring an XC requires the GeoCue Repository Software Developer Kit (SDK). These commands are the only ones that can support group dispatch (all entities sent to a single remote machine).
Foreign Commands (FC)	These are command line executable programs that have no knowledge of GeoCue. These are the most commonly integrated third party commands.	These commands can be executed as either local commands (Dispatch disabled) or as Distributed commands.

The following subsections describe the various options of the Command Table in Environment Builder. Refer to Figure 5-3 for the location of the various options.

## 5.1 Enabling Dispatch

Check the **Supports Dispatch** checkbox if the command supports dispatching. If the command also supports *distribution*, then check the **Distribute** box. As a reminder:



- Dispatch – The command can be run on a remote machine. If more than one entity is selected, the entities must be sent to a single remote machine as a group (*Populate Working Segments* in the LIDAR 1 CuePac is an example of the Dispatch mode of operation). NOTE: Only Extrinsic Commands (XC) can support Dispatch without distribution. An XC can only be created using the GeoCue Repository SDK<sup>5</sup>.
- Distribute – The command supports each selected entity running on a different machine. In this case, the subtasks (usually there is a subtask per entity when a command can be distributed) are queued by the CDS and sent to machines for processing as those machines become available. An example is the *Generate Image* step in the LIDAR 1 CuePac.

## 5.2 Restrict to Single Instance

This option allows only a single instance of a particular command to execute on a node at a given time. This is used for situations in which a machine is physically capable of running many instances of a command but some limitation such as software licensing or a software limitation prevents more than a single instance of the command from executing simultaneously. An example of this is the Geometric Post-Processing (GPP) command of the Intergraph DMC PPS workflow.

Prior to the addition of this option, the command had to be limited via one of the resource parameters (SPU, SMU, SUI) but this, of course, would impact other parameters of operation.

## 5.3 Auto Kill Size

This field allows you to put a “runaway memory” limit on the executable. If checked, fill in the size (in standard memory units, SMU) for the kill trigger. If an instance of the executable grows

---

<sup>5</sup> Dispatched commands that are not distributed receive Entity information in a GeoCue Processing Queue. The GeoCue Repository API must be used to extract information from this queue.

in memory beyond this size, the Command Dispatch System will automatically terminate the instance. The reason for the termination will be posted to the System field of the process step tracking history. Note that this action will only be taken for operating systems that support a system level call to explicitly kill an executable.

NOTE: Certain actions such as Auto Kill require active monitoring of the target machine by the Command Dispatch System. This is accomplished by a component called a Dynamic Resource Interrogation Module (DRIM). DRIMs are not available for all machine types so if feature such as Auto Kill are not enabled for your configuration, check for the presences of a DRIM.

## 5.4 Caching

The Command Dispatch System supports *caching* of the files associated with a processing task/subtask. The meaning of these settings is described in Table 5-2.

**Table 5-2: Command Cache Settings**

<b>Option</b>	<b>Meaning</b>
Cache Files	Enables caching for the command
Cache Mode – Entity Only	Cache only the entity's files
Cache Mode – Links Only	Cache only the files associated with the entities that are linked to this entity but not the entity's own files
Cache Mode – Entity and Links	Cache all files
Retain Cache for Task Duration	Do not flush the cache until all subtasks associated with the dispatched task complete.

The Retain cache option is mainly useful for linked files when all entities are linked to the same file. An example of this occurs in the UltraCam processing CuePac. Each L0 image is linked to the same calibration file set. The file set is about 500MB so we want to retain these files on the remote dispatch nodes until all images that were submitted as a task are complete. This prevents a target processing node from reloading calibration files from the server for each new image.

## 5.5 Machines

Note – Prior to using this portion of the dialog, you should ensure that all of the machines in your GeoCue constellation are registered with Dispatch Manager. Check this by viewing the Machines tab of Dispatch Manager.

The machines field allows you to limit the command to running only on selected machines. This is useful for preventing distributed commands from loading down certain machines (such as interactive workstations). It is particularly useful for managing executables that require a license (other than GeoCue software where licensing is integral to the Command Dispatch System). For example, if you were dispatching the Rectify portion of Leica's LPS software and you had licenses (and the software itself, of course) only installed on BIG\_SERVER and BIGGER\_SERVER, you would set this field to these two machine names.

Note that by checking the Exclude box, the command will run on all machines registered in Dispatch Manager *except* the listed machines.

Leaving this field blank means that your command can run on all machines within your GeoCue constellation (assuming you have not placed restrictions on the machines using Dispatch Manager).

## 5.6 Maximum Failures Allowed

This option is useful to prevent *runaway* dispatching to a failed node. If you key in a number, the CDS will stop dispatching subtasks, *of this task*, to a specific node after the specified number of failures. The failed node will be returned to the dispatch table following the completion of the current Task.

## 5.7 Predicted Sizes

The predicted size of the executable, in terms of Standard Units, is entered into the three predicted fields. A value of -1.0 tells the system to ignore this parameter in dispatch management computations. Note that you cannot set the Standard Processor Units field to -1.0 since this would cause (when no other restrictions were in place) all of the subtasks to be sent to a single machine. For this reason, the minimum value for SPU is 0.10.

The I/O limitation is useful for applications that are not all that compute intensive but are very I/O intensive. An example of this is a Rectify program. Generally machines running a rectification process run out of I/O bandwidth prior to maxing out processor usage.

The default configuration for a command is 1.0 for SPU and ignore for SMU and SIU.

## 5.8 Equilibrium Time

The three settings for Equilibrium Time are predictions of how long after the start of execution the resource usage predicted in the fields described in the previous section will be met. These default to zero which means that the CDS should assume that the executable will immediately use the predicted resources.

These fields are used by an advanced scheduling algorithm of the CDS that uses real time machine monitoring. You should not set these to values other than zero unless you have received explicit instructions from GeoCue Corp.

## 6 Configuring Priority

In this chapter we present a description of how priority works in GeoCue and a brief description of how priority is set. The details of dialogs related to priority are presented in the respective section of other chapters.

GeoCue dispatches subtasks to machines based on the priority of the subtask (there are, of course, many other factors that come into play such as the Schedule of a machine, the Processing Group of the subtask, the Processing Group parameters of the target machine and so forth). Given two subtasks with non-competing Command and Processing Group settings (both of these factors are discussed in later chapters), the subtask with the higher priority will be dispatched first.

NOTE: You must have GeoCue Administrator privilege to set or change any aspect of the priority system.

### 6.1 Priority Algorithm

Priority in the GeoCue Command Dispatch System is a floating point value ranging from 0.0 to MAX\_FLOAT. A higher numerical value is a higher priority. Thus zero is the lowest possible subtask priority in the GeoCue CDS. It has no special significance other than being the lowest possible priority.

The dispatch subtask priority is computed based on a Priority Computation Algorithm (PCA). This algorithm uses *priority factors* and *priority weights*.

*Relative Priority* (with default value noted in parenthesis) is set at three levels in GeoCue:

- ERP - Entity Relative Priority ( 1.0)
- PRP - Project Relative Priority (1.0)
- GRP - Group Relative Priority (1.0)

Three *Priority Weights* are used in the PCA. These weights are:

- EPW - Entity Priority Weight (1.0)
- PPW - Project Priority Weight (1.0)
- GPW - Group Priority Weight (1.0)

*Intermediate* Priorities are computed by multiplying the Relative Priority by the associated Priority Weight:

Entity Priority = ERP x EPW

Project Priority = PRP x PPW

Group Priority = GRP x GPW

If desired, the maximum value of each type of priority can be controlled. This is called *priority clipping*. For example, if your priority scheme needs a constraint that the Entity Priority never exceed 100.0, you can set this as the Entity Priority clipping value (the Clipping values are set on the Priority tab of the Configuration tab of Dispatch Manager. The three settings are:

- Entity Priority Clip (EPC)
- Project Priority Clip (PPC)
- Group Priority Clip (GPC)

The ultimate *computed* priority of a subtask, where *priority linking* is not in effect, is given by the Priority Computation Algorithm of of Eq: 1 and repeated below for convenience:

$$\begin{aligned} \text{Priority}^6 = & \text{EPC}[(\text{ERP} \times \text{EPW})] + & \text{Eq: 1} \\ & \text{PPC}[(\text{PRP} \times \text{PPW})] + \\ & \text{GPC}[(\text{GRP} \times \text{GPW})] \end{aligned}$$

---

<sup>6</sup> Entity Priority can be dynamically modified based on the GeoCue *Linking* system. This is discussed in a separate section.

Note that since all of the six priority parameters in the PCA default to 1.0 and clipping is, by default, off, the default priority of a subtask in the factory configuration of GeoCue is 3.0. Thus you can think of 3.0 as being the average priority in the GeoCue system.

The PCA has nine user adjustable parameters. Therefore, there are nine ways you can affect the priority of a subtask.

NOTE: You can manually override the computed priority of a subtask using the subtask section of the Task tab in Dispatch Manager.

It is important to note that the PCA implemented in GeoCue will allow the priority of a subtask from a low priority project to exceed the priority of a subtask from a higher priority project unless you have set appropriate clipping values.

## **6.2 Setting Priority Weights**

You can change the *weights* used in the PCA. We recommend that you do not do this unless you have carefully devised your priority computation approach. This is not the best way to do things such as raise the priority of certain project relative to others (Use Project Relative Priority instead).

### **6.2.1 Entity Priority Weight**

Entity Priority Weight (EPW) is set using the Entity tab of Environment Builder. Select the entity whose EPW you wish to modify and press the *Modify* button. This is illustrated in Figure 6-1 for a LIDAR\_Ortho\_Raster entity type.

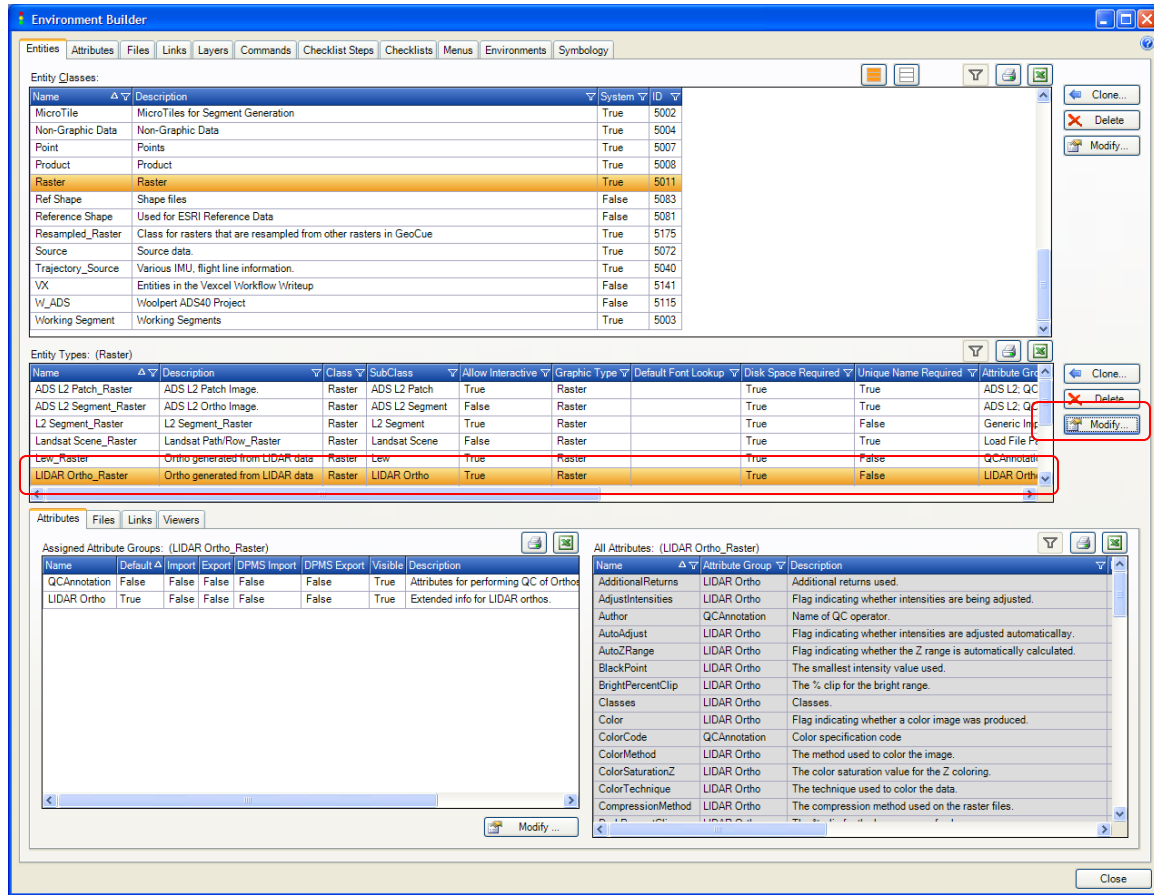


Figure 6-1: Selecting the LIDAR\_Ortho\_Raster entity type

After pressing the Modify button, you will be presented the Modify Entity dialog of Figure 6-2. Key in the new Entity Priority Weight. The valid range is 0.0 to MAX\_FLOAT.



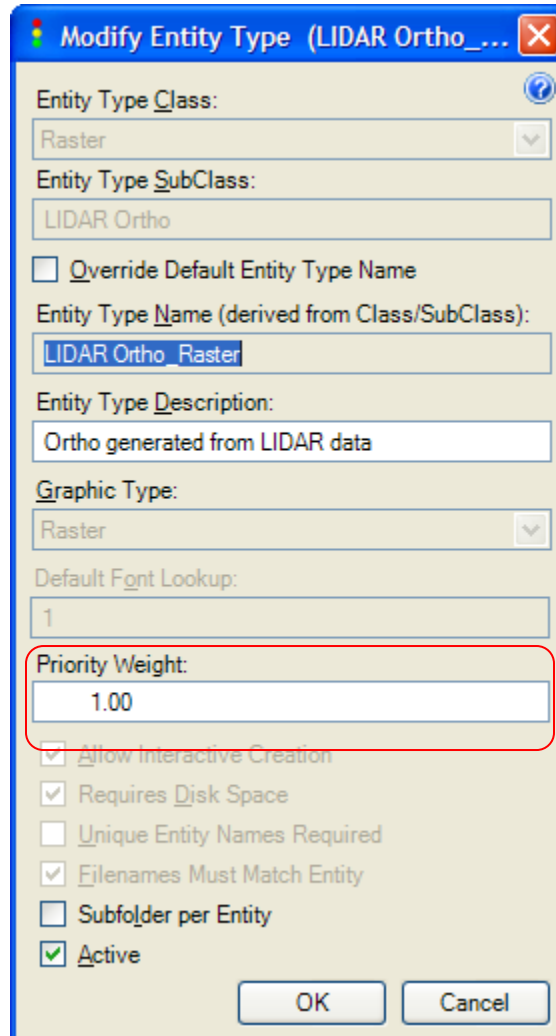


Figure 6-2: The Modify Entity dialog

You can use the EPW to force certain types of entities to process prior to others during dispatch. For example, suppose you were using a combined LIDAR/Ortho project and wanted generation of your LIDAR Orthos to always receive priority over rectification of your image orthos. You could achieve this by setting the EPW of the LIDAR\_Ortho\_Raster entity type to 2.0. This would cause all LIDAR ortho generation subtasks to have a computed priority of 4.0 (see the Priority Computation Algorithm) while all image orthos would maintain their system default subtask priority of 3.0. Since the LIDAR\_Ortho\_Raster priority is higher, these subtasks will process first (given all other factors are the same).

## 6.2.2 Project and Processing Group Priority Weight, Clipping

Both Project Priority Weight (PPW) and Group Priority Weight (GPW) are set using the Configuration tab, Options sub-tab (see Figure 6-3). Priority clipping is also set in this dialog if you opt to use clipping. Simply key in the value you desire and press Apply.

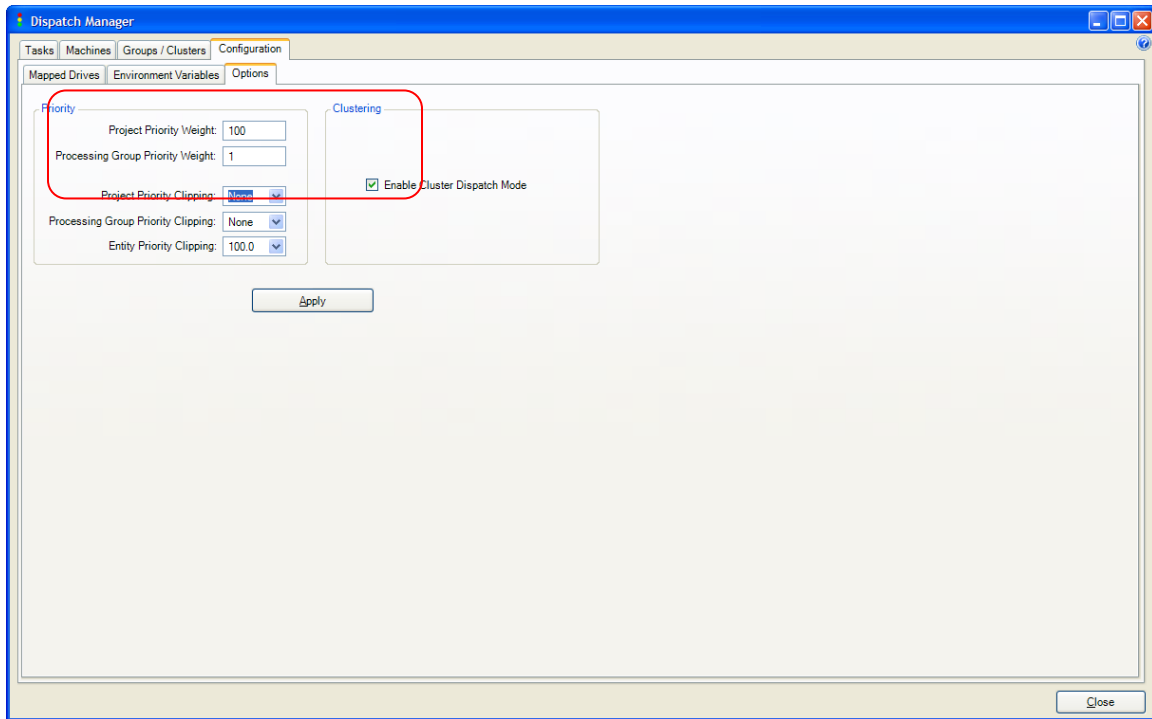


Figure 6-3: Setting Project and Group Priority Weights

These two weights are typically only set to values other than the default values of 1.0 when you are devising your own priority schemes. For example, if you have set Relative Entity Priority and Entity Priority Weight such that the computed Entity Priority (ERP $\times$ EPW) can never exceed 100.0, you could ensure that projects with RPP of 2.0 will always execute prior to Projects with RPP of 1.0 by setting the Project Priority Weight to 100.0.

Note that in the above discussion, there is no guarantee that Entity Priority (ERP x EPW) cannot exceed 100.0 since a user can set entity relative priority using Entity Manager. This can be controlled by setting a Clipping value for Entity Priority to 100.0. When clipping values are set, the PCA uses this value to limit each priority factor prior to adding the three priorities. Thus for the above scheme, you would set the Entity Priority Clipping value to 100.0. The default value for Clipping is "None", meaning that no clipping will be applied. You can type into the combo box any clipping value greater than or equal to 0.00. Be aware (see the PCA) that there is no numeric value that has the same effect as "none" so be certain to reselect "none" if you do not want clipping applied.

## **6.3 Setting Relative Priority**

As discussed in a previous section, there are three relative priority settings that are used in computing the final subtask priority via the Priority Computation Algorithm.

### **6.3.1 Entity Relative Priority**

All entities have a default Entity Relative Priority (ERP) of 1.0. You can change ERP using Entity Manager. The general procedure is to invoke Entity Manager (from the GeoCue Client, Tools menu) and add the entities for which you wish to change the priority to the Working Set. Figure 6-4 depicts a set of LIDAR Orthos filtered in Entity Manager, set to the Priority tab.

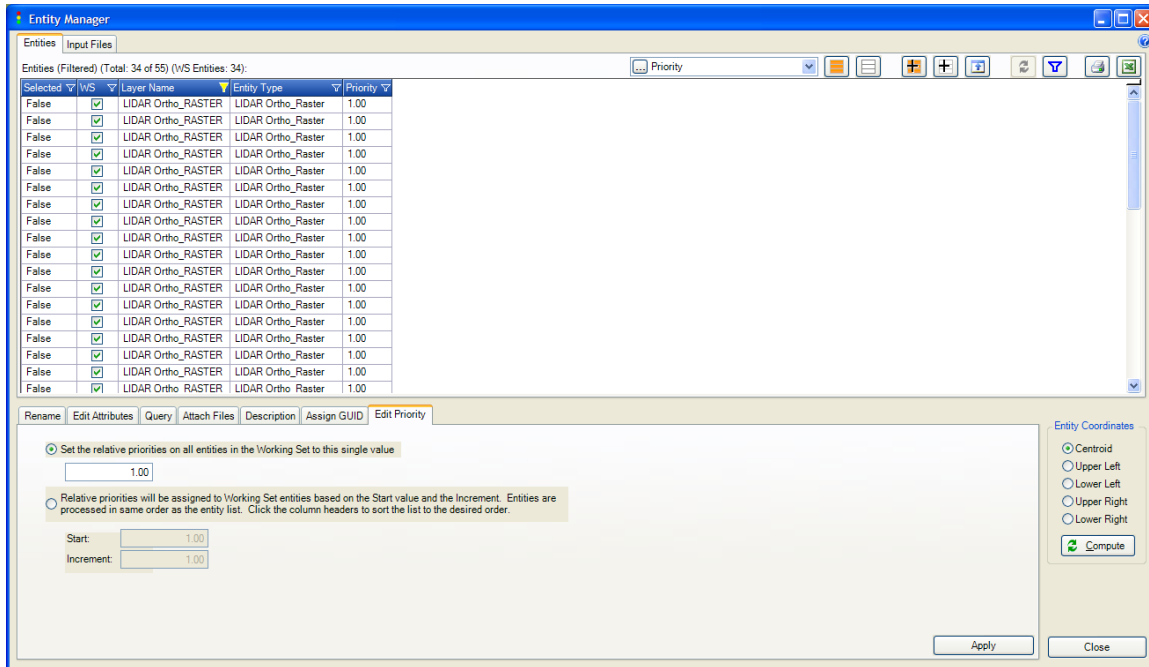


Figure 6-4: Setting Entity Relative Priority

Entity Manager provides two methods of setting ERP:

- Set each entity to the same, user entered ERP
- Set a start and increment value for ERP. This method applies the increment in the order in which the entities are listed in Entity Manager.

Remember that you must add the entities to the Working Set prior to changing the ERP value.

### 6.3.2 Project Relative Priority

Project Relative Priority (PRP) can be set with either the Project Properties dialog (Figure 6-5) or via the Project Utilities dialog, Priority tab (Figure 6-6). When using the Project Utilities dialog, remember to press *Apply* after changing values.

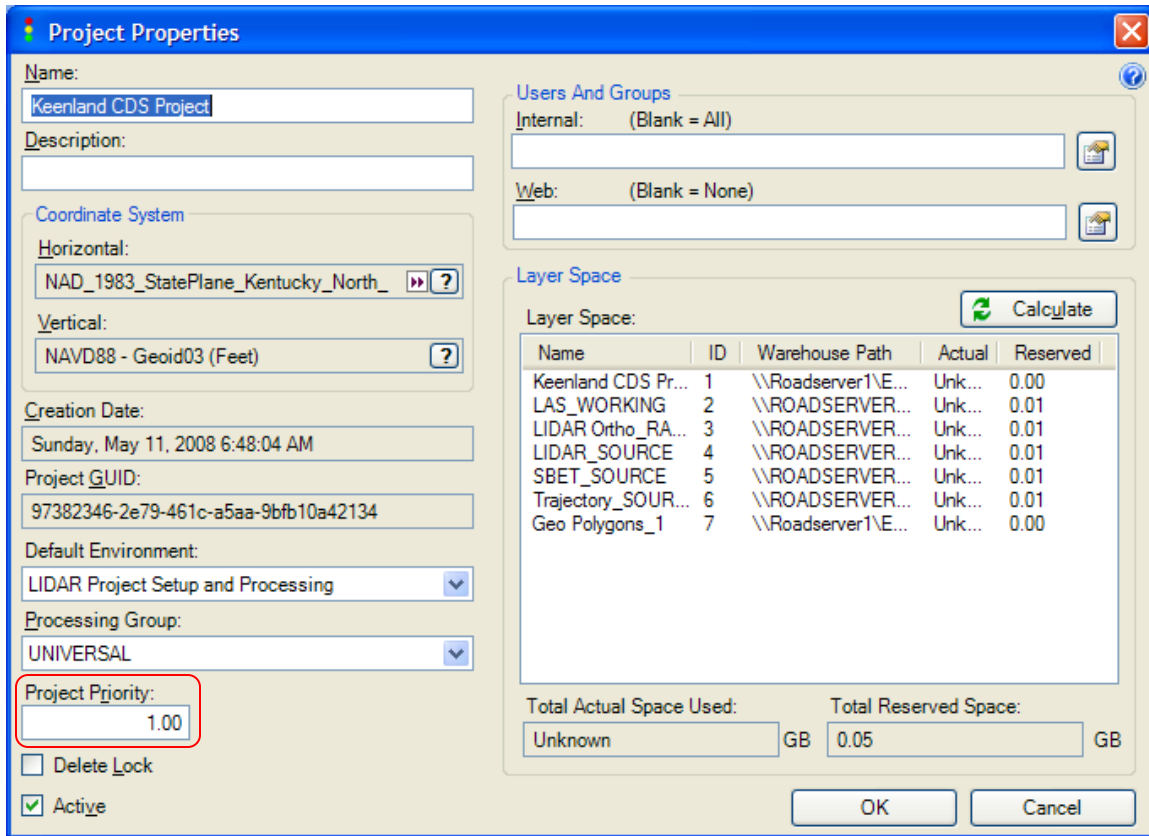


Figure 6-5: Project Relative Priority set from the Project Properties dialog

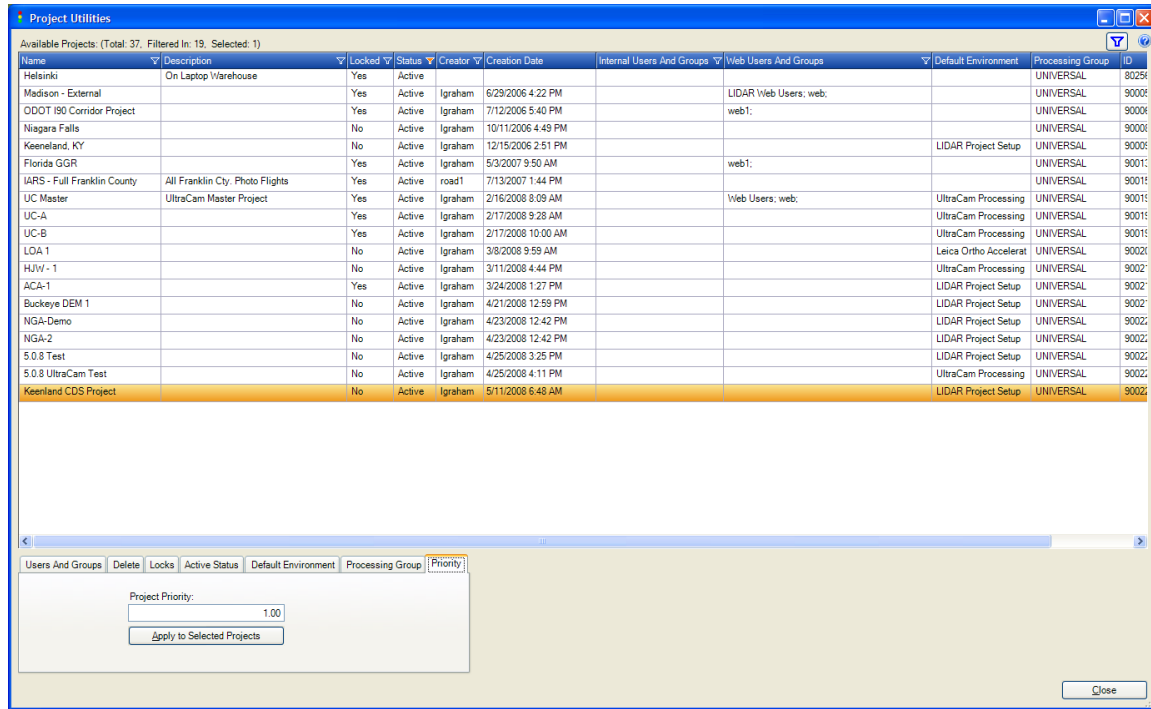


Figure 6-6: Setting Project Relative Priority via the Project Utilities dialog

### 6.3.3 Processing Group Relative Priority

The Processing Group Relative Priority (GRP) is set using Dispatch Manager, Groups/Clusters tab (Figure 6-7). Select the Group whose priority you wish to modify and press Modify. Note that you cannot modify the RESERVED group.

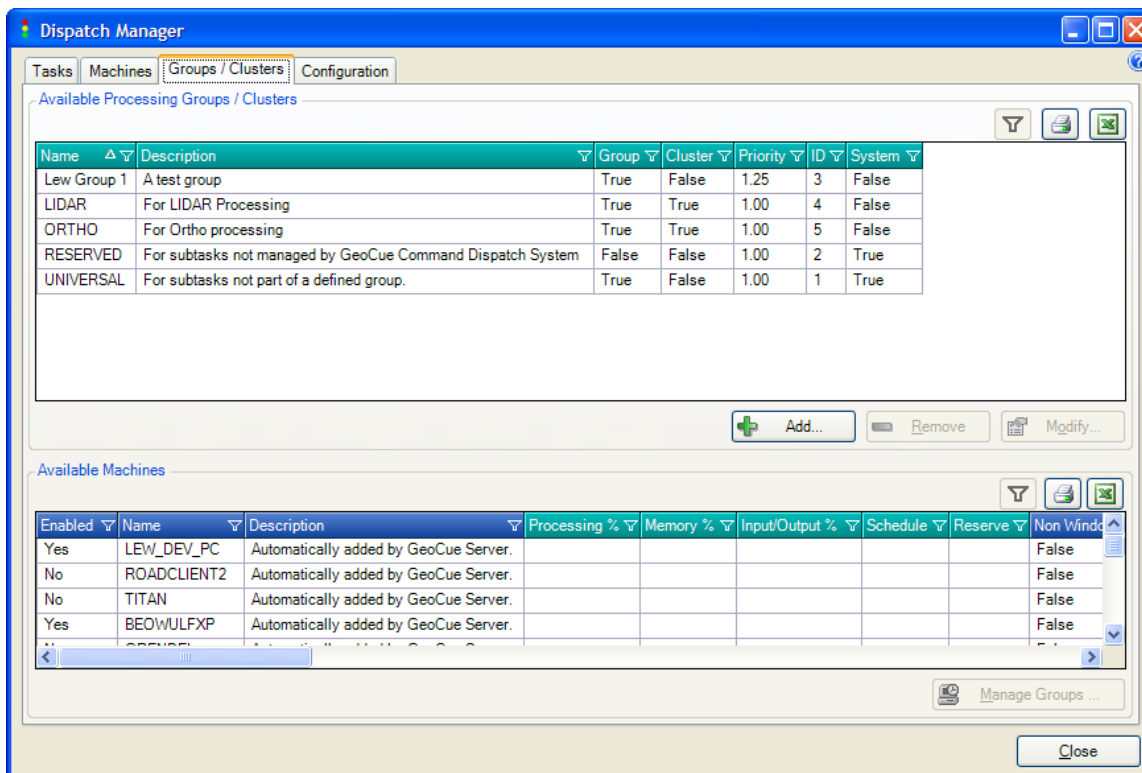


Figure 6-7: Setting Processing Group Relative Priority via the Dispatch Manager dialog

## 7 Configuring Processing Groups

*Processing Groups* are used to control the resources of machines based on the membership of a subtask in a user definable category called a *Processing Group* (*Processing Groups* are also used to define machine *Clouds* - discussed in a subsequent chapter). This is an advanced capability of the Command Dispatch System that is typically only used for controlling machines with substantial resources that need to be balanced across the processing needs of different production activities. It was originally developed for allocating the resources of a 128 processor machine where it was undesirable to just let tasks load up the machine based only on a first come, first served algorithm.

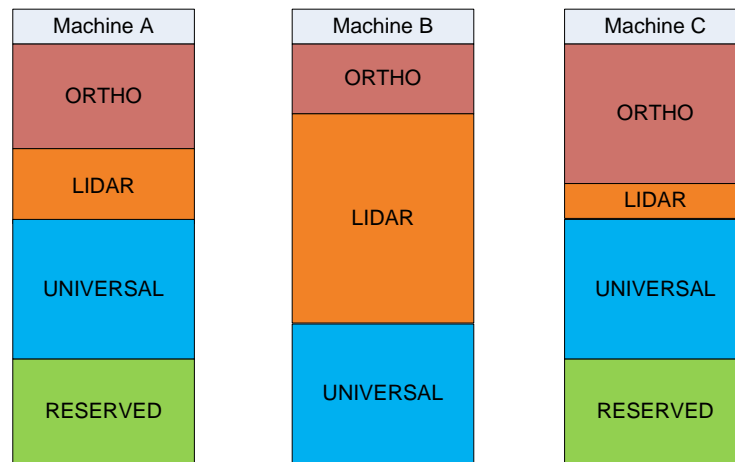
Note – *Processing Groups* are needed only for more advanced installations. You should not activate *Processing Groups* in the Command Dispatch System until you thoroughly understand the material in this section. The inappropriate use of *Processing Groups* can render a machine “unreachable” by the CDS.

The concept of *Processing Groups* is very similar to *Cloud* except that *Processing Groups* extends the concept of machine *Clouds* to the Command level (described in the next chapter). The use of *Processing Groups* for Machine *Clouds* is explained in a later chapter.

### 7.1 *Processing Groups Overview*

*Processing Groups* are simply names that you make up for categorizing the resources of machines (and, optionally, commands). For example, you might name your processing groups LIDAR and ORTHO. You would then categorize your GeoCue projects into these groups (leaving projects that were not in these two groups set to their default values). This is graphically depicted in Figure 7-1.





**Figure 7-1: Processing Groups**

GeoCue reserves two names for special processing groups; RESERVED and UNIVERSAL. Projects that have not been explicitly added to groups are, by default, members of the UNIVERSAL group. Machines that have not been assigned to processing groups have all of their resources assigned to the UNIVERSAL group. Thus if no processing groups are established (and machine resources have not been RESERVED), each machine will have all of its resources available for general command processing.

The configuration of Processing Groups is accomplished via several controls:

- Processing Groups are created via the Groups/Clouds tab of the Dispatch Manager (Figure 7-2).
- Projects are assigned to Processing Groups via the Properties dialog of a Project or via the Group tab in Project Utilities.
- Commands are assigned to Processing Groups via the Modify button on the Groups/Clouds tab of Dispatch Manager (Command assignment to Processing Groups is described in a subsequent chapter).

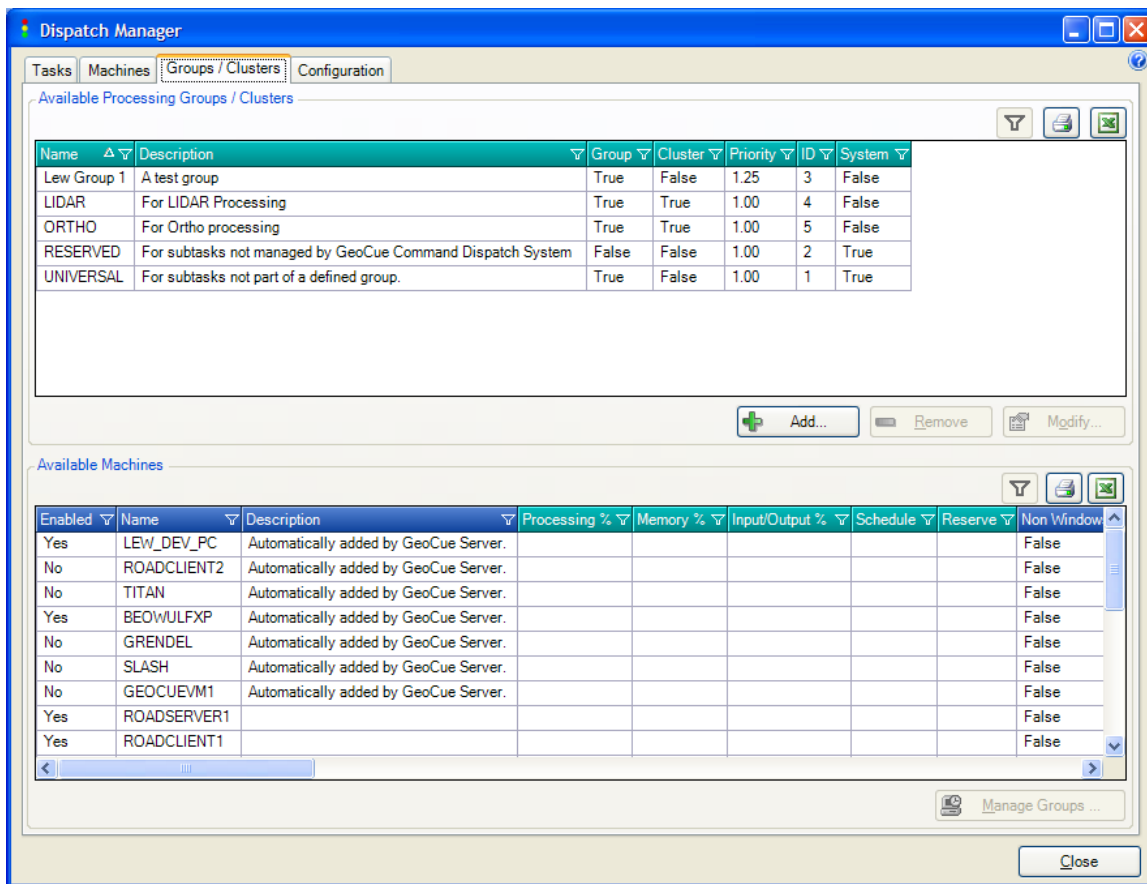


Figure 7-2: The Groups/Clouds tab of Dispatch Manager

A portion of the total resources of a machine (Processor, Memory and Input/Output) can be allocated to Processing Groups. By default, the resources of a machine are allocated 100% to the UNIVERSAL group. This means that subtasks have full access to machine resources on a first come, first served basis (of course subtasks arrive at machines in priority order).

The Command Dispatch System includes a system group called RESERVED. This group allows you to reserve resources of a machine for processing that occurs outside the control of the CDS. For example, you could set a particular server's reserved allocation to 50%, thus limiting the GeoCue CDS use of this machine to half of its processing resources. As discussed in a later

section of this chapter, you can apply a schedule to the resource allocations of machines. This is useful when the required machine loading is dynamic with respect to time.

## 7.2 Creating Groups

Groups are created using the Groups/Cloud tab of Dispatch Manager (Figure 7-3). Create a new Group by pressing the **Add** button (note – as with many GeoCue dialogs, if you have a row selected, the **Add** button will change to a **Clone** button). Modify a Group by *selecting* the group that you wish to modify and press the modify button. Note that (as with all GeoCue dialogs) you can *deselect* by holding down the Control key while pressing the left mouse button. Finally you can clone an existing group by selecting the group you wish to clone and pressing “Clone.”

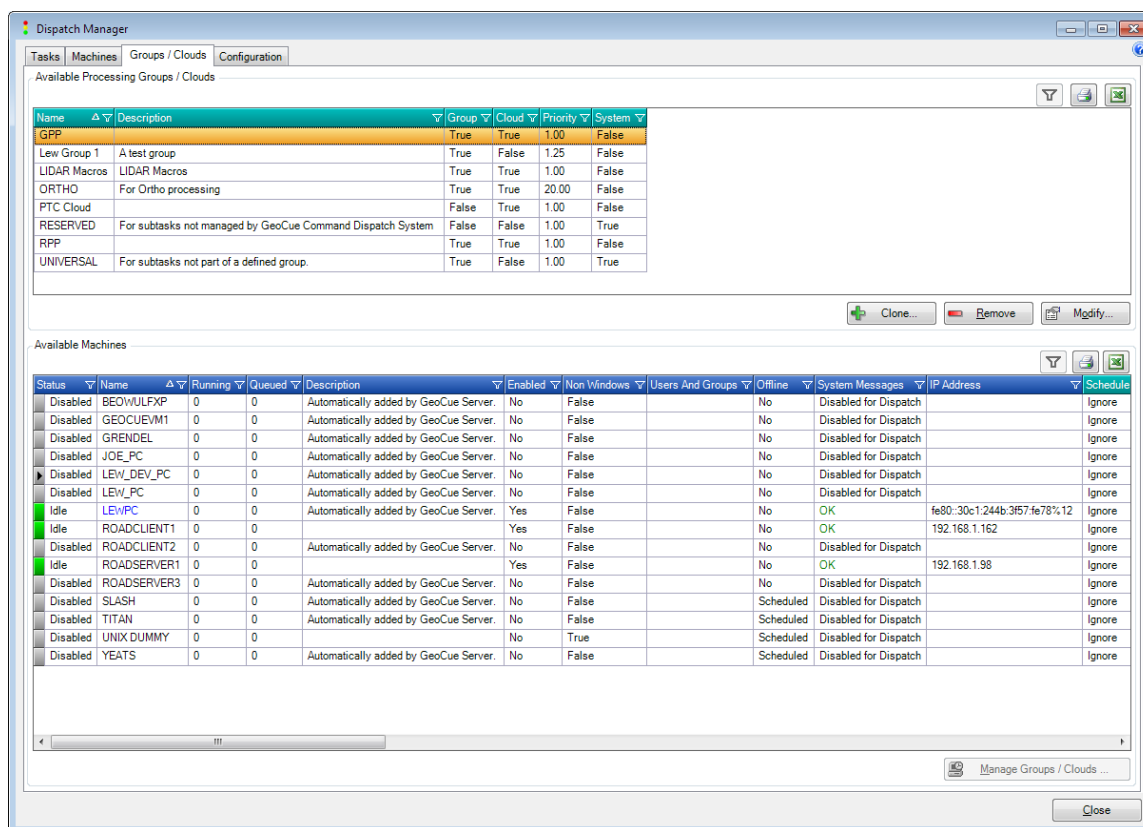


Figure 7-3: The Groups/Clouds dialog

You can use any name (except RESERVED and UNIVERSAL) you desire for a Processing Group so long as it is unique among existing processing group names. RESERVED and UNIVERSAL are system reserved names that cannot be used.

If you are not performing advanced scheduling, simply accept the system default priority of 1.0. Otherwise, set the priority to the value you intend to use in your priority scheme.

### **7.3 Configuring Resource Groups**

Resource Groups are modified on a machine basis. The Resource settings made in each session will be homogeneous across the selected machines. Thus if you have 3 machines that you wish to allocate across processing groups at the same allocation levels (e.g. 40% for ORTHO, 25% for LIDAR, 20% RESERVED) you would simultaneously select these machines and set their resources. If you wish to set different resources for the various machines, it will be necessary to edit their resources individually.

The configuration starts from either the Machines tab or the Groups/Clouds tab of Dispatch Manager. If you are using the Machines tab, first *Select* the machines that you wish to configure into Processing Groups from the Machines pane of the tab. Use the normal GeoCue interface for selecting (Control-left mouse click for individual selection and Shift-left mouse click to select multiple rows). In the Available Processing Groups/Clouds pane, select the group that you wish to configure. Press the **Manage Groups/Clouds...** button at the bottom right of the dialog to invoke the Machine Processing Groups/Clouds dialog Resources dialog. An example of 3 machines being selected for Group modification from the Machine tab is depicted in Figure 7-4.

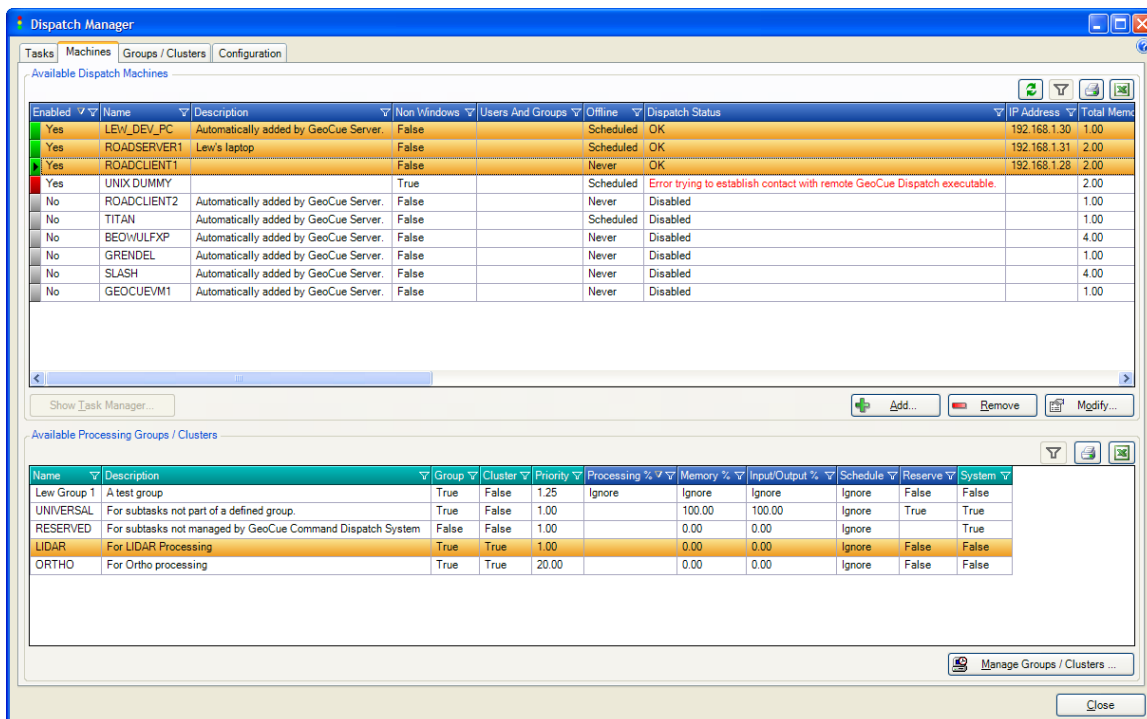


Figure 7-4: Preparing to configure Groups from the Machines tab

The selection process is reversed if you are starting from the Groups/Clouds tab of Dispatch Manager. Select the Group/Cloud you wish to modify in the upper pane (you can only single select) and then select the Machines from the lower tab. An example of selecting the same group and machines as before is depicted in Figure 7-5.

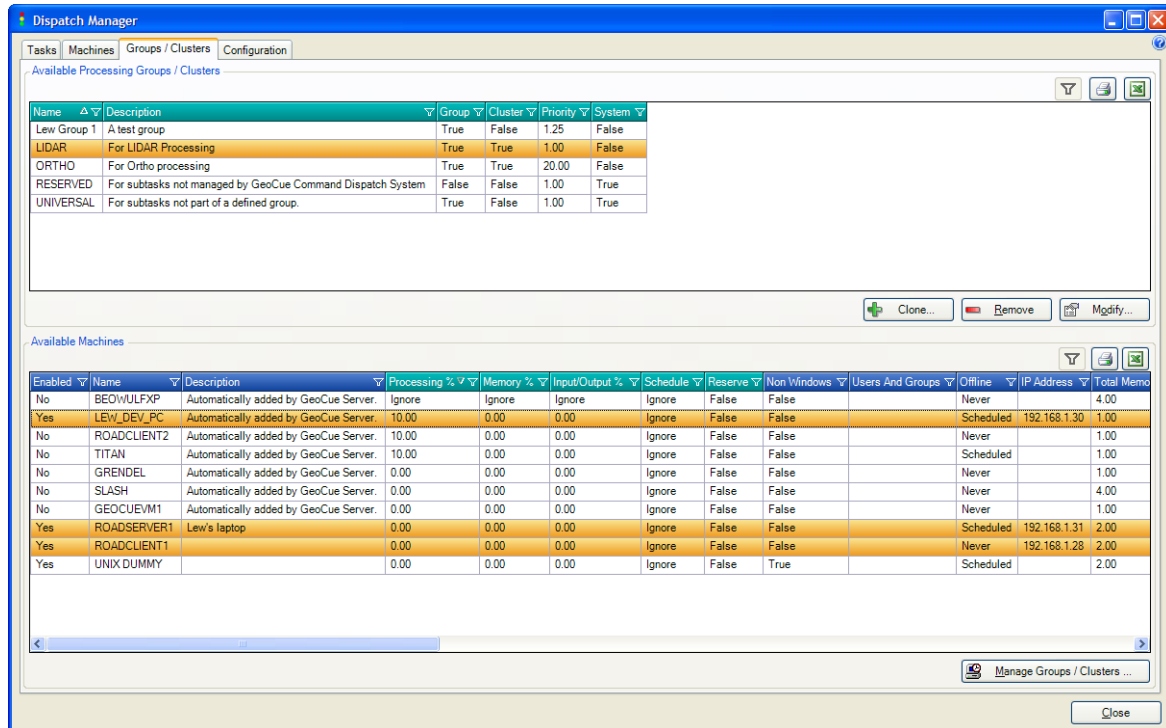


Figure 7-5: Preparing to configure Groups from the Groups/Clouds tab

After selecting the desired machines/groups and pressing **Manage Group/Clouds**, the dialog of Figure 7-6 will be displayed. The values displayed in this dialog apply *only* to the machines that were *selected* when the dialog was invoked.

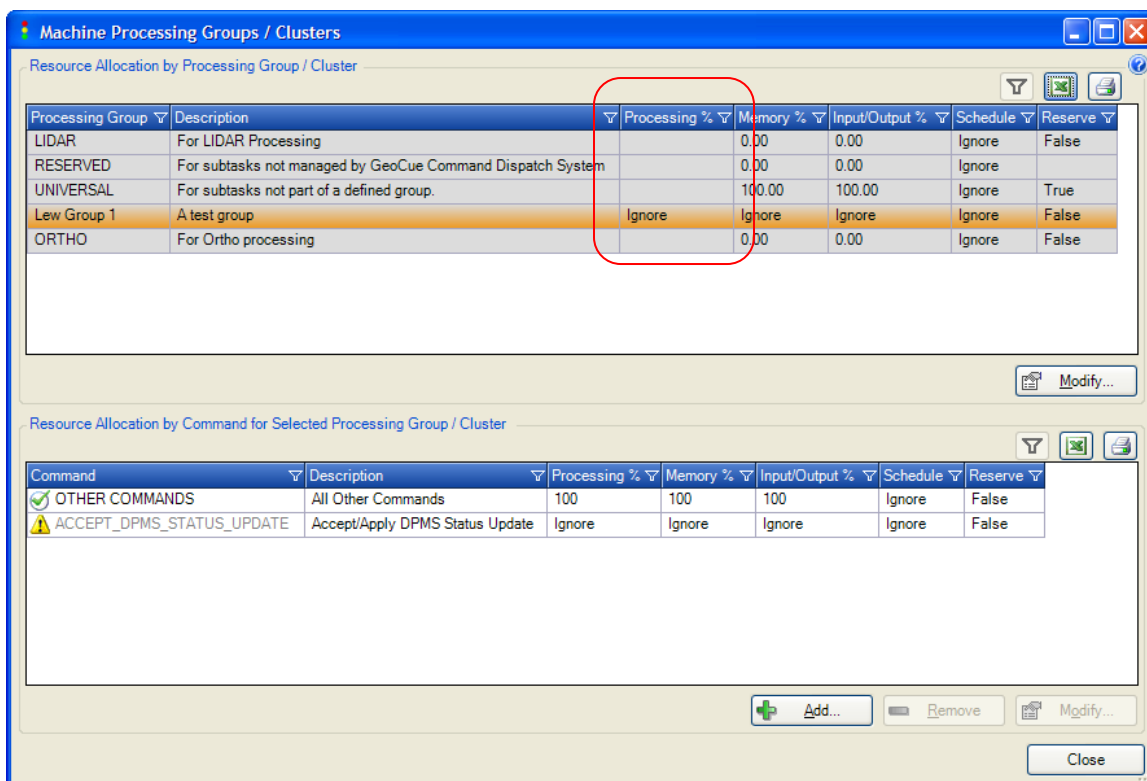


Figure 7-6: Machine Processing Groups dialog

Note that in our example, three processing groups have been defined (recall that Groups are defined from the Groups/Clouds tab of Dispatch Manager and that UNIVERSAL and RESERVED are system defined groups that always are present). Also note that the behavior of fields in GeoCue dialogs for multi-selected items is the same as in Microsoft Office products. Thus several rows in the “Processing %” column of Figure 7-6 are not the same for all three selected machines and hence these fields are blank. The “Lew Group 1” row of the “Processing %” column is indicating a value of “ignore” and hence this is the setting of this value for the three selected machines.

Processing Group values for the selected group are modified by pressing the **Modify...** button at the lower right of the upper pane of the dialog of Figure 7-6. This invokes the Modify Processing Group / Cloud Resource Allocations dialog of Figure 7-7. Note that this dialog will affect the *Machines* that were originally selected (three, in our example) and the *Processing Group* that was selected. You cannot modify more than one processing group at a time using this dialog.

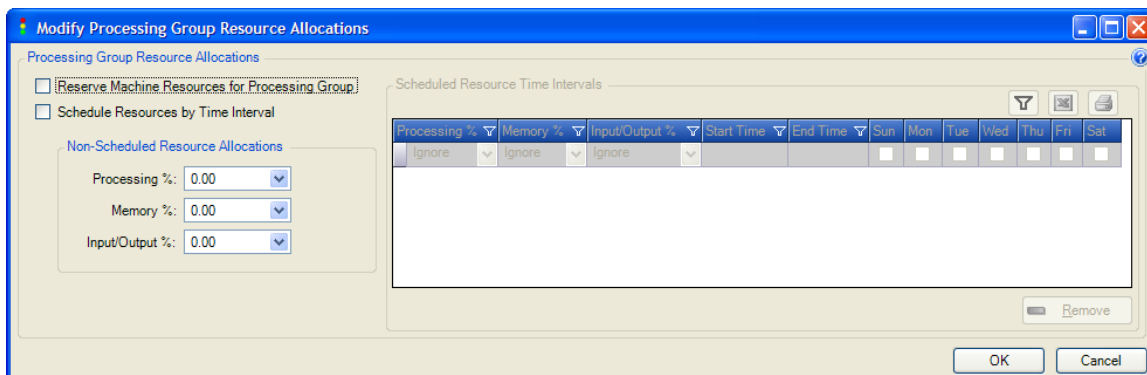


Figure 7-7: Modify Processing Group Resource Allocations dialog

### 7.3.1 Setting Resource Allocations (non-scheduled)

To set the percentage of resources allocated to the selected Resource Group on the selected machines, set a percentage in the Non-Scheduled Resource Allocations section of the dialog (see Figure 7-8).

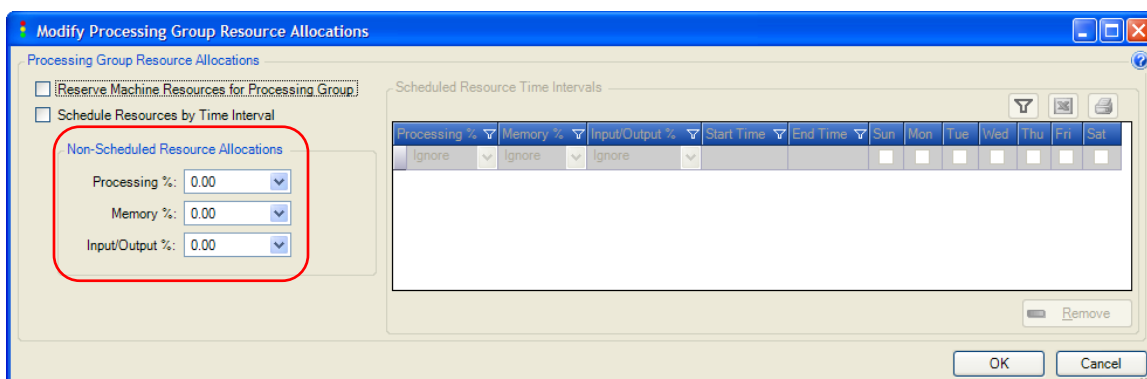


Figure 7-8: Setting the Resource Allocation percentages



## 7.3.2 Reserving Resources

There is an option on the Modify Processing Group Resource Allocations dialog (Figure 7-9) entitled “Reserve Machine Resources for Processing Group/Cloud.” This option has a profound effect on how the CDS treats the resources of a machine. Key a percentage in these fields (e.g. 55.5 = 55.5%) or set the field to “ignore” via the drop-down arrow selector. Setting the field to “ignore” means that the specific parameter will not be considered for these machines for this group.

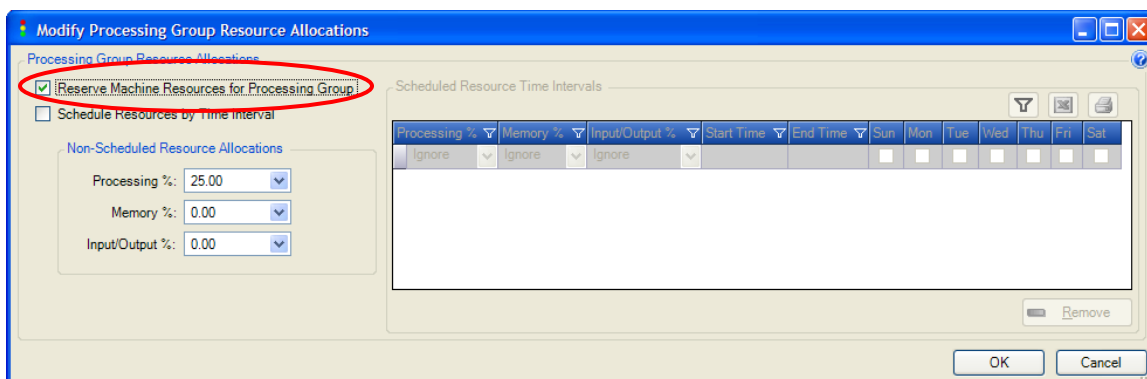


Figure 7-9: Modify Processing Group Resource Allocations

If this option is not set (not checked), then the particular group(s) that the dialog is modifying will not reserve machine resources if the group(s) do not have any subtasks *pending* in the dispatch queue. However, as soon as a subtask enters the dispatch queue that is a member of the group(s) in question, the reserve will enable. However, the Command Dispatch System *never* stops a subtask once it starts<sup>7</sup>. Thus a newly entering subtask of the Process Group will not be able to get its allocated share of resources until that percentage of currently processing subtasks that are not members of the group(s) under discussion complete.

<sup>7</sup> There is one exception to this rule. If a Auto-kill limit has been set on a command, the particular OS on the machine can support the Kill command and the subtask exceeds the auto-kill threshold, then the Command Dispatch System will abort the command currently processing the subtask. However, this is an abnormal situation.

If it is imperative that members of the group(s) immediately have access to their allocated share of resources, set the “Reserve Machine Resources for Processing Group” for the group(s).

There is a potentially very deleterious effect of using this option; the resources will remain reserved for the allocated group even if no subtasks are currently pending for that group. Thus this option should only be used if you fully understand these side effects. There is one common and very useful exception to this guideline discussed in the next subsection.

This option was originally designed for environments in which certain commands could take very long times to process (e.g. 24 hours). In these cases it was not acceptable for subtasks from one group to dominate the machine for these extended periods.

### 7.3.3 Scheduling Resources

You can set schedules for Processing Groups. Like all scheduled times in the Command Dispatch System, you can set any number of schedules for a period of a week. These schedules then repeat on a weekly basis. This facility allows you to do things such as enable the ORTHO group to use 25% of a machine's resources from 7 AM until 5 PM and then allow the ORTHO group to have 100% of the machine's resources from 5 PM until 7AM. You could further refine this by allowing the ORTHO group to have 100% of the machines resources from 5 PM Friday until 7 AM Monday.

## 7.4 The *RESERVED* Group

The Command Dispatch System includes a factory defined group called RESERVED. This group is used for allocating resources to non-CDS operations. Thus, for example, if you set the RESERVED values for Cyclops to 40% then 40% of Cyclops' resources will be excluded from the CDS' launch algorithms. As with other Processing Groups, the RESERVED group can be time scheduled.

Thus, for example, you could set a machine to have only 25% of its resources available during normal working hours and 100% available at other times. This sort of scheduling is very useful for powerful workstations that can give up some percentage of their processing resources during working hours without a big impact on the interactive user. The RESERVED group is discussed in detail in a separate chapter of this document.

## ***7.5 Processing Groups Considerations***

It is very important to understand the side effects of using processing groups. In our example above, projects whose processing group is anything other than ORTHO or LIDAR (projects are set to the default group, UNIVERSAL) will not be able to dispatch to machine Cyclops at all since we allocated all of Cyclops' resources to the LIDAR and ORTHO groups. A second consideration is that if no LIDAR tasks are being dispatched to Cyclops at a particular instant in time, still only 50% of the resources of Cyclops will be available for projects of the LIDAR group (and none for any other projects).

Note that the Command Dispatch System has no control over programs that execute on a node outside the context of the CDS. In our example above, other users can run non-GeoCue commands on Cyclops regardless of the current loading.

## 8 Reserving Machine Resources for non-CDS Use

The Command Dispatch System includes a factory defined group called RESERVED. This group is used for allocating resources to non-CDS operations. Thus, for example, if you set the RESERVED values for the machine "Cyclops" to 40% then 40% of Cyclops' resources will be excluded from the CDS' launch algorithms.

Management of the RESERVED group is the same as managing any other group in the CDS. We have broken this out as a separate chapter since the RESERVED group is most often used to manage non-CDS resources, the most common use of groups.

The RESERVED group is most commonly used to configure machines where you would like the CDS to use varying amounts of resources based on a schedule. For example, you do not want to load down production workstations that are used during normal hours for interactive processing with dispatched CDS tasks. However, you would like to use these resources during non-interactive times. While you can accomplish a simple "available, non-available" schedule for machines using the "Off-Line" facility of the CDS, you cannot set degrees of availability. The RESERVED group allows you to fractionalize the available resources.

### 8.1 Reserving Machine Resources

Setting the reserved resources of a machine is similar to setting any other processing group in the CDS (see the Groups chapter of this document). The exception is that the RESERVED group is always "Reserved" (meaning no other group, including UNIVERSAL, can borrow from the RESERVED group).

You can set the RESERVED resources from either the Machines tab or the Groups/Clouds tab of Dispatch Manager. You can multi-select machines in either tab to accomplish simultaneous settings.

Recall that you can quickly examine the membership of machines in various groups by selecting the machine(s) in the upper pane of the Machine tab of Dispatch Manager and examining the group percentage allocations in the lower pane. Conversely, you can assess the members of groups by selecting the desired Group in the upper pane of the Groups/Cloud tab of Dispatch Manager and examining the allocations in the lower pane.

## **8.2 Scheduling the RESERVED Group**

As with other Processing Groups, the RESERVED group can be time scheduled. Thus, for example, you could set a machine to have only 25% of its resources available during normal working hours and 100% available at other times. This sort of scheduling is very useful for powerful workstations that can give up some percentage of their processing resources during working hours without a big impact on the interactive user.

NOTE: The Off-line schedules for a machine can overlap the RESERVED schedules without conflict. These schedules are independent with the off-line schedule taking precedence.

## 8.3 An Example

Suppose LEW\_DEV\_PC is a dual processor, 4 core machine that is in heavy interactive use during normal working hours (say 6 AM to 6 PM) but is virtually unused in the off-hours. Further suppose that this machine is in moderate use on Saturday mornings from 6 AM to 11 AM.

We will set this machine up as having 8 SPUs of “processor power”, 4 GB of memory and 80 MB/sec of I/O capacity<sup>8</sup>. We will allow the CDS to have access to 25% of the resources of this machine during week day working hours and 50% of the machine on Saturday morning. We will allow the CDS full access at other times. Additionally, this machine is backed up from midnight Saturday night until 3 AM Sunday morning. Thus we will set an off-line schedule for this backup.

We first select LEW\_DEV\_PC in the Machine tab, Available Dispatch Machines pane of Dispatch Manager. Press the **Modify...** button to bring up the configuration dialog.

HINT: You can double click a row in any pane of the Machine tab or Groups/Clouds tab of Dispatch Manager to immediately invoke the corresponding Modify dialog.

Set the values in this dialog according to the above described plan. The populated dialog is depicted in Figure 8-1. Note that we have checked the “Offline Time Intervals” option and set the backup schedule.

---

<sup>8</sup> All commands that are shipped with GeoCue are set to ignore Memory and I/O resources. You can modify these, if desired, using the Command tab of Environment Builder.

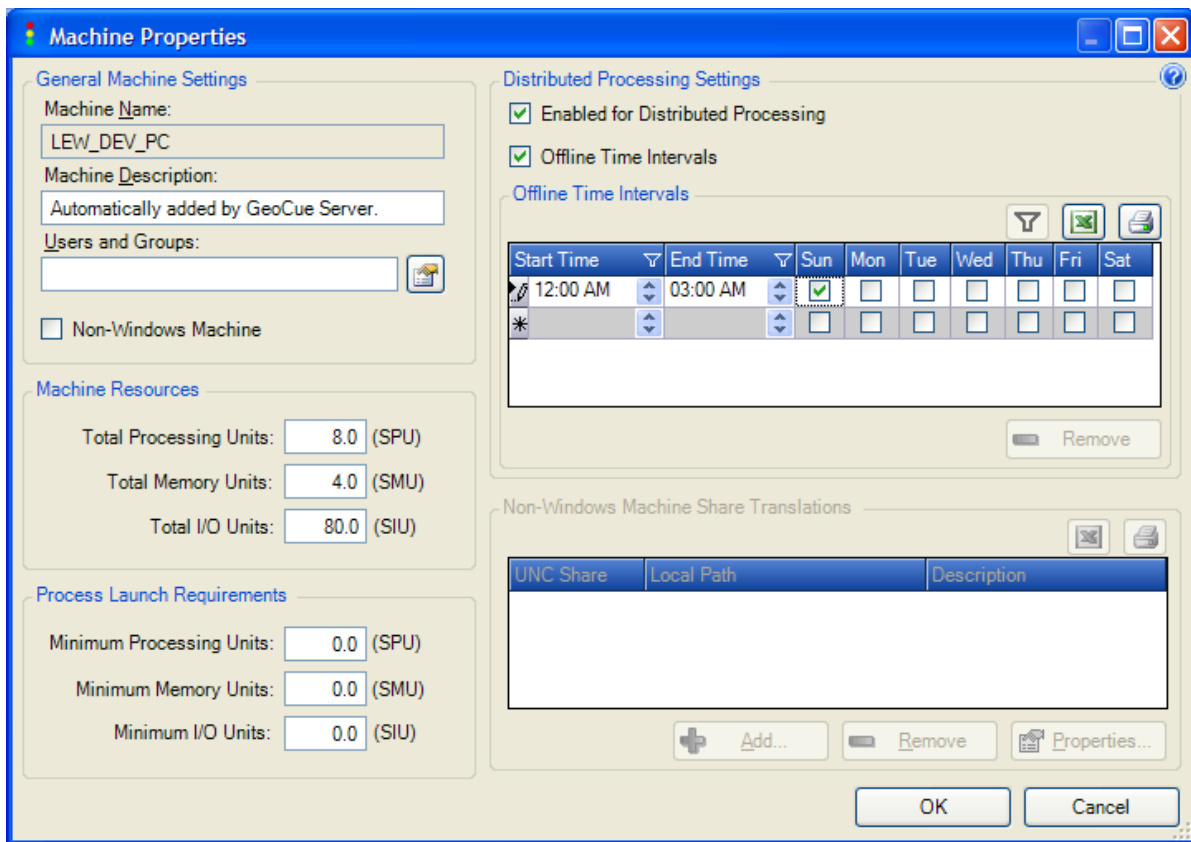


Figure 8-1: Setting Machine configurations for LEW\_DEV\_PC

OK out of this dialog and, still using the Machine tab of Dispatch Manager, select LEW\_DEV\_PC in the “Available Dispatch Machines” pane and the RESERVED group in the “Available Processing Groups/Clouds” pane. Press the **Manage Groups/Clouds...** button in the lower right of the dialog. This will invoke the Machine Processing Groups/Cloud dialog of Figure 8-2 with the RESERVED group selected. Press the **Modify...** button at the lower right of this pane to invoke the “Modify Processing Group/Cloud Resource Allocations” dialog (Figure 8-3).

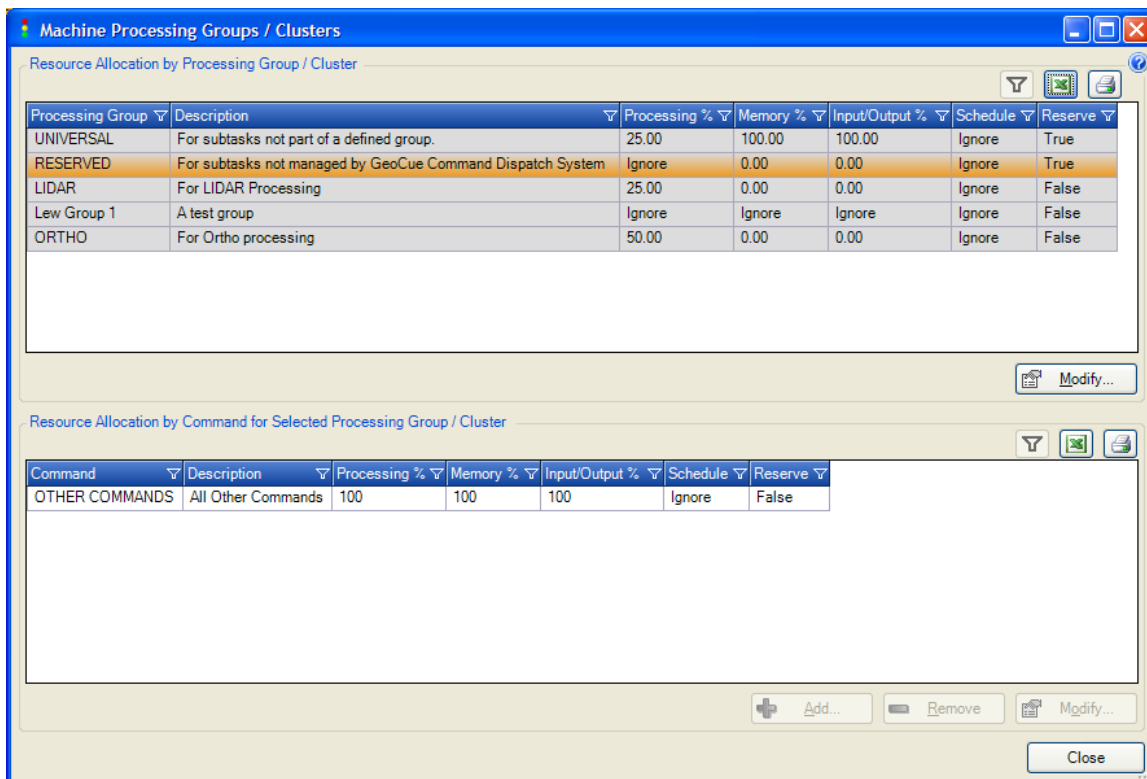


Figure 8-2: The Machine Processing Groups/Clusters dialog

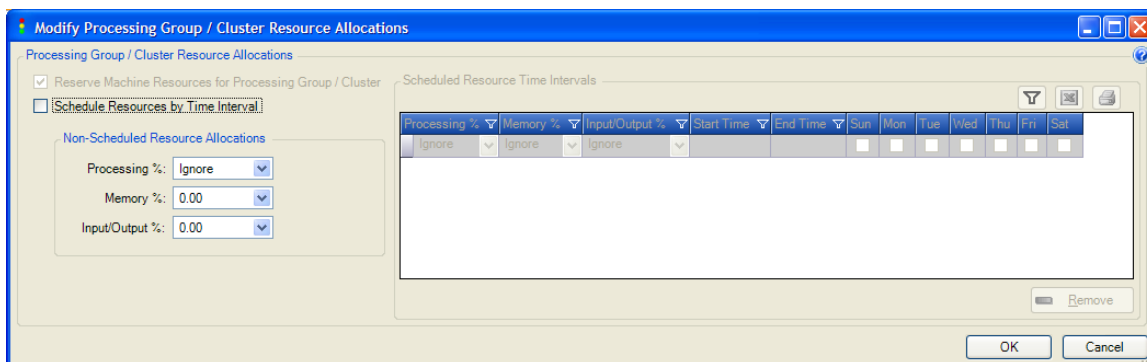


Figure 8-3: Modify Processing Group/Cloud Resource Allocations dialog



Check the Schedule Resources by Time Interval option and set the schedules for workdays and Saturday morning. **OK/Close** out of the dialogs to save your changes.

You have now configured LEW\_DEV\_PC to make a variable percentage of resources available to the CDS on a repeating weekly schedule as well as set an offline time for system backups.

## 9 Configuring Clouds

The Command Dispatch System supports the concept of Machine Clouds or Machine Clustering (or simply Clouds). This capability allows you to collect machines into named configurations and then dispatch Tasks to these named clouds rather than to individual named machines.

The Cloud system uses Process Groups to accomplish this with a few variations from the processing group distribution rules. Thus a view of Clouds looks the same as Processing Groups. Thus Figure 9-1 looks the same as the Processing Groups diagram of Figure 7-1. Processing Groups and Clouds can be freely mixed in the same GeoCue Command Dispatch System constellation.

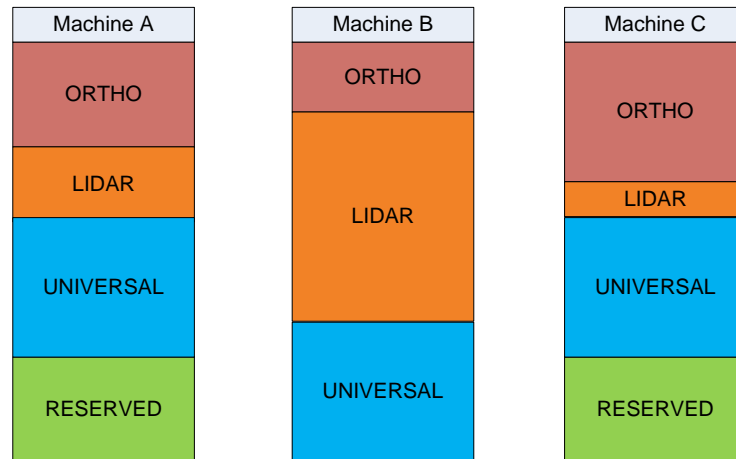


Figure 9-1: Clouds

Clouds are defined in the Groups/Clouds tab of Dispatch Manager (Figure 9-2).

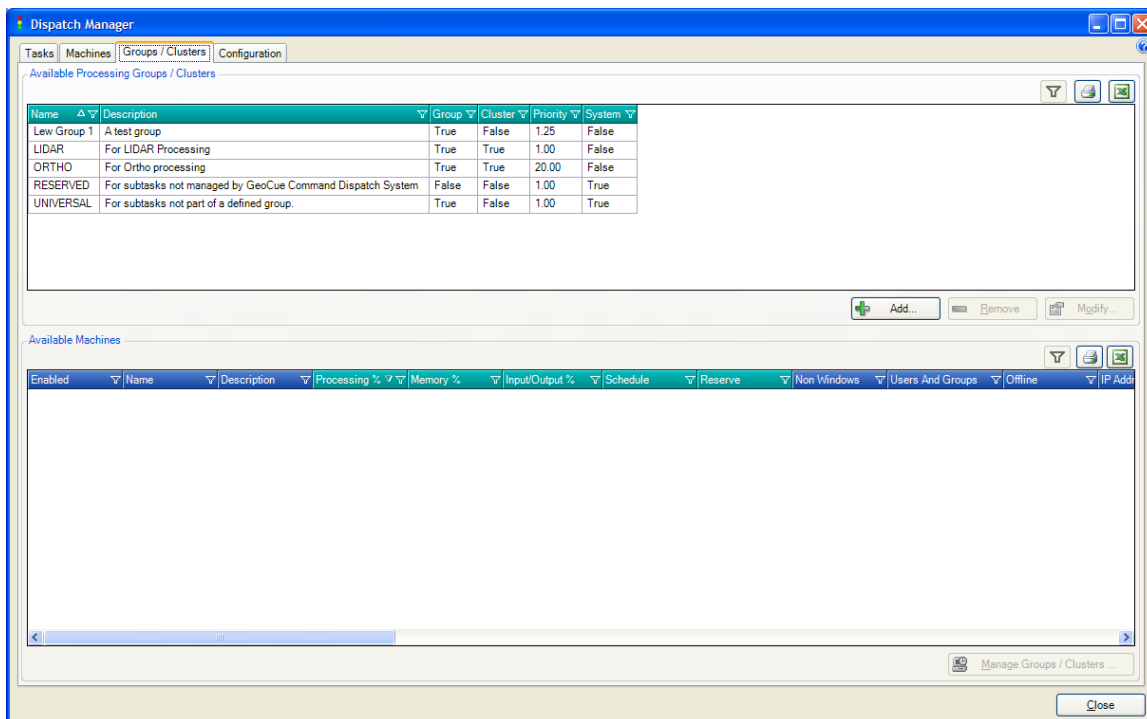


Figure 9-2: The Group/Clouds tab of Dispatch Manager

Press the **Add...** button to define a new Group/Cloud. You can provide any alphanumeric name that is unique to the Group/Cloud list except the reserved names “RESERVED” and “UNIVERSAL.” To enable the definition as a cloud, check the “Processing Cloud” checkbox. Note that an entry can serve as both a Cloud and a Processing Group.

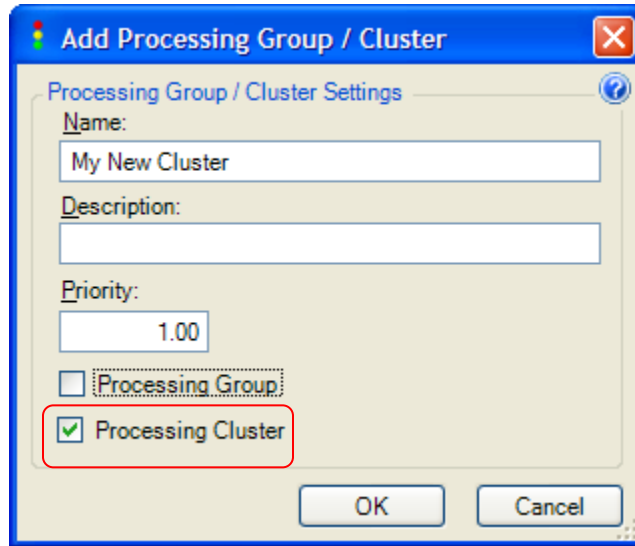


Figure 9-3: Adding a new Cloud definition

You must explicitly enable Clouded dispatch mode in the Configuration tab, Options sub-tab of Dispatch Manager (Figure 9-4). If Cloud mode is enabled, you will see a mode selector in the Dispatch dialog of GeoCue (see the Dispatch chapter of this document). Toggling this to Cloud will cause the dispatch machine list to change to the named clouds. Unlike when in machine mode, only a single cloud can be selected as the processing target. This is because, of course, the cloud comprises a group of 1 or more machines.

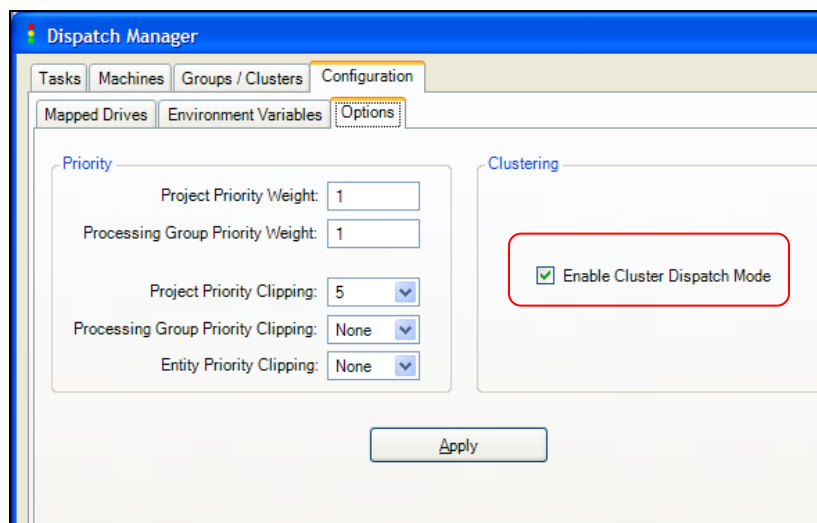


Figure 9-4: Enabling Cloud Mode

NOTE: You can enable the UNIVERSAL group as a Cloud. You cannot enable the RESERVED group.

## 9.1 Configuring Clouds

Clouds are configured using the same techniques as setting Processing Groups (see the Processing Groups chapter). Recall that you must set the checkbox for Cloud mode in Add/Modify Processing Group/Cloud dialog (Figure 9-3).

## 9.2 Processing Groups and Cloud

As you recall from the Processing Groups chapter, groups are used to allocate specific machine resources to *projects*. Processing Groups are assigned at the project level. This means for example, if processing groups are in use<sup>9</sup>, all Tasks dispatched from project A will be members of the group assigned to Project A.

Cloud overrides the assigned Processing Group for the Task being dispatched in Cloud mode. Thus if your project is a member of the LIDAR Processing Group (assigned via Project Properties) and you dispatch a task in Cloud mode to a cloud named "CYCLOPS", all subtasks that result from this dispatched task will be reassigned to the CYCLOPS processing group.

In many circumstances, Tasks that belong to a Processing Group can *borrow* resources from machines that are not members of the group. This borrowing *never* occurs with Cloud mode processing. This ensures that when you define a Cloud for a particular class of processing (e.g. ORTHO), you are assured that your dispatched task will only run on the machines that are members of the Cloud.

## 9.3 The Universal Group/Cloud

Machine resources that are not explicitly allocated to a named processing group / cloud (including the RESERVED group) default to the UNIVERSAL group. You can designate the UNIVERSAL group as a Cloud. This can be useful for circumstances where you are normally performing Group processing (meaning that you have assigned a processing group to your project) but wish to have an override capability. You can accomplish this by dispatching in Cloud mode and selecting the UNIVERSAL Cloud.

---

<sup>9</sup> Technically, Processing Groups are *always* in use since projects default to the UNIVERSAL group. GeoCue's default configuration for the Command Dispatch System makes this behave as if groups were not involved.

## **9.4 Summary**

Clouded processing is a very powerful capability that can be used to insulate users from needing to know the specific machines to which commands should be dispatched. It also provides a mechanism by which IT can schedule backup times for servers without totally disabling the Command Dispatch System. This can be accomplished by devoting two or more machines to a Cloud and staggering their off-line times (using the Machine tab of Dispatch Manager, off-line schedule).

## 10 Configuring Commands within Processing Groups/Clouds

Individual executables can be configured with respect to processing groups and/or clouds. This mode of operation would typically be employed when groups are being used to control the priority of collections of projects (as opposed to controlling the resources of machines). Using groups with commands allows you to prioritize projects based on their membership in a particular group and then allocated resources within a group based on the command itself.

Managing machine resources based on the command that will be dispatched is typically used when you are controlling machines that have a lot of processing resources. For example, these algorithms were designed for use on a 128 processor supercomputer. Most advanced users of the Command Dispatch System will not need to employ this capability.

You can allocate machine resources within a processing group at the command level. For example, if you had two commands, say Rectify and Mosaic, that you wanted to control within the ORTHO group, you can accomplish this by assigning this commands percentages of a the ORTHO groups allocated resources.

For example, you could set the ORTHO group to receive 50% of the resources of a machine (as described above) and then set Rectify to receive 40% of that allocation and Mosaic to receive the other 60%. The net effect of this is that Rectify would receive 20% of the machine's resources (40% allocated to Rectify on a machine to which the group has 50% of the total resources) and Mosaic would receive 30% of the machine's resources.

**NOTE:** If the Group for which you are configuring a command is designated as a Cloud, the behavior is identical to Group behavior. Thus, for example, you could use this feature to reserve a percentage of a Cloud's resources for Rectify commands.



## 10.1 Allocating Resources at the Command Level

Allocating resources at the command level is performed via the **Manage Groups/ Clouds...** button available from either the Machine tab or the Groups/Clouds tab of Dispatch Manager. This invokes the Machine Processing Groups/Clouds dialog of Figure 10-1. The lower pane of the dialog is used for command level resource allocation. To add a command, press the **Add...** button below the "Resource Allocation by Command for Selected Processing Group/Cloud" pane. Alternatively, you can modify an existing command by selecting the desired command and pressing the **Modify...** button. This invokes the Modify Command Resource Allocations dialog of Figure 10-2.

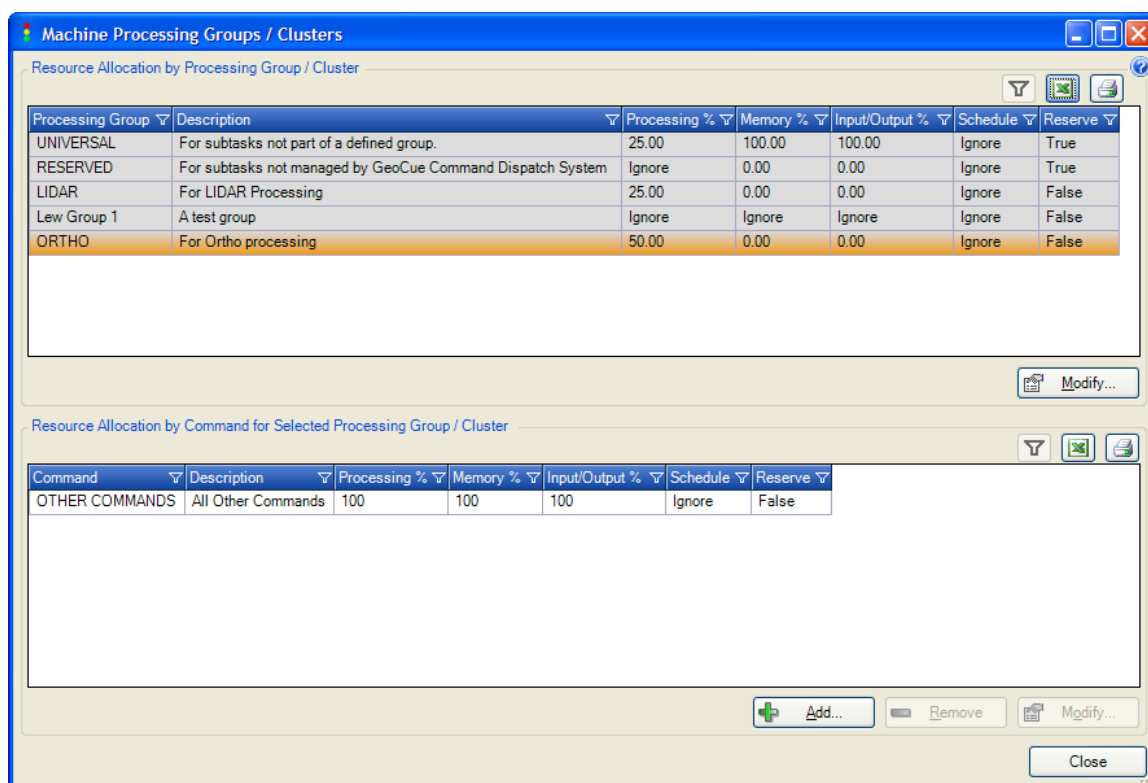
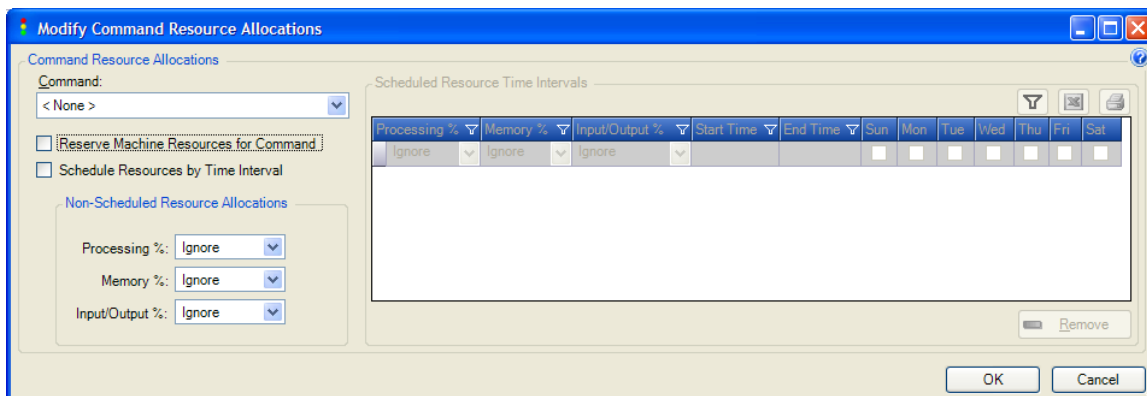


Figure 10-1: The Machine Processing Groups/Clouds dialog



**Figure 10-2: Modify Command Resource Allocations**

Select the desired Command from the drop-down list of available commands (these are containing in the Commands table of Environment Builder).

Resources are allocated in the usual manner using the Non-Scheduled Resource Allocations section of the dialog (scheduling resources is discussed in the next section). The percentages represent percentages of resources assigned to this processing group. Thus if you have assigned 40% of the machine CYCLOPS' resources to the ORTHO group and then assign an SPU value of 50% to the Rectify command, the command will be allocated 20% of CYCLOPS' processor resources (50% of the full machine allocation of 40%)..

## 10.2 Reserving Command Resources

The option "Reserve Machine Resources for Command" will, if checked, prevent other commands from "borrowing" resources reserved for this command. For example, assume your processing group is ORTHO and the command is Rectify. If you have allocated 40% of CYCLOPS (an 8 SPU machine) to the ORTHO group and then assign the Processing % as 50%, then the ORTHO group will receive 3.20 SPU (40% of the total 8 SPU of CYCLOPS) and the Rectify command will receive 50% of this allocation or 1.60 SPU. If you do not check the Reserve Machine Resources for Command option, other commands running in the ORHTO group can "borrow" resources from the Rectify command allocation so long as Rectify subtasks are not pending for CYCLOPS. If you set the option, other commands cannot borrow from the Rectify command.

HINT: You should check the Reserve option if you expect other commands to execute within the process Group that could take a “long” time to execute since these commands can effectively block the command that has allowed resource borrowing.

Use the Reserve options sparingly since this allocation of machine resource will remain idle when commands associated with the reserved resources are not pending.

### **10.3 Scheduling Command Resources**

You can set a schedule for the allocation of a group's resources to a command in a manner similar to setting schedules for the overall resources of a group. This would allow you to, for example, set Rectify to receive 30% of the Processing Group's machine allocation from 7 AM until 5 PM and 70% during other times. Simply check the “Schedule Resources by Time Interval” option and set the schedule in the usual manner.

## 11 Troubleshooting

In this section of the guide we list some common issues and errors with suggestions of how these can be addressed.

### 11.1 Tasks Remain in a Pending State

This is the condition under which you have dispatched a task but the task remains in the pending state (as viewed in the Task pane of Dispatch Manager). The common causes of this are listed in Table 11-1.

**Table 11-1: Stuck in Pending**

<b>Cause</b>	<b>Repair Action</b>	<b>Notes</b>
Target Machine(s) is "off-line"	<ol style="list-style-type: none"> <li>1. Wait for the machine to come back online</li> <li>2. Terminate the Task and re-launch to a different machine(s)</li> </ol>	You can quickly determine if a machine is off-line by viewing the Machines tab of Dispatch Manager and hover your cursor over the off-line column.
Target Machine(s) does not have sufficient resources to execute the command	<ol style="list-style-type: none"> <li>1. Increase the machine resources</li> <li>2. Terminate the Task and re-launch to a different machine(s)</li> </ol>	

<b>Cause</b>	<b>Repair Action</b>	<b>Notes</b>
Machine(s) have RESERVED resources	<ol style="list-style-type: none"> <li>1. Change the RESERVED resources for the machine(s)</li> <li>2. Terminate the Task and re-launch to a different machine(s)</li> </ol>	Be sure to examine the RESERVED schedule for the machine. You can quickly check to see if a machine has RESERVED resources by clicking the machine row in the Machines tab of Dispatch Manager and examining the values in the Available Processing Groups/Clouds pane.
Machine(s) do not have adequate resources for the Processing Group to which your project belongs	<ol style="list-style-type: none"> <li>3. Change the Group resources for the machine(s)</li> <li>4. Change the Project Group</li> </ol>	

## 12 Command Dispatch System Internals

The Command Dispatch System in GeoCue is a Client-Server architecture. The Server portion of the Command Dispatch System performs preliminary processing of subtasks and manages a central dispatch queue. The CDS Queue Processor organizes the subtasks within the dispatch queue and vectors subtasks to processing *nodes* based on a dispatch algorithm and the available of the remote nodes. Each Remote Node is managed by a *Dispatch Service*<sup>10</sup>. The dispatch service accepts subtasks from the Queue Processor, performs preparatory work, executes the command and returns status to the Queue Processor. Following completion (or failure) of the dispatched subtask, the Queue Processor performs finalization work such as setting the exit lock status of objects, recording the processing history and so forth. If the command failed, the Queue Processor may schedule a retasking in which case the subtask is returned to the processing queue.

### 12.1 Processing Granularity

GeoCue manages remote processing using the notion of Tasks and Subtasks. All remote dispatch sessions begin as a Task. A Task is a collection of objects (one or more) that are (transparent to a use) placed in a request queue and sent to the Command Dispatch System. If the objects contained in the Task can be processed via Distributed processing (that is, the objects can be sent, independently to different machines for processing), then the Command Dispatch System will split the Task into the number of Subtasks contained in the original dispatch request. If the dispatch objects must all be processed on the same machine, then a single Subtask will exist for the Task. The settings for Dispatch and Distributable are on the Command tab, Command sub-tab of Environment Builder. An example of a task that cannot be distributed (at least as of Version 5.0 of GeoCue) is Populate LIDAR Working Segments. If, for example, 50 working segments are selected for population and dispatched to machine X, all 50 will be sent to Machine X as a single group. Thus there will be a single subtask associated with the task. On the other hand, the Generate LIDAR Ortho command in GeoCue is *distributable*. Thus, for example, if 40 LIDAR orthos are selected for generation, there will be 1 Task and 40 subtasks. Note that the number of subtasks is completely independent of the number of machines available to run a particular task.

---

<sup>10</sup> This is an NT Service on a Windows machine and a daemon on a UNIX machine.

The CDS always operates on Entities within GeoCue. Thus for a command that can be dispatched but not distributed (spilt up among multiple machines), the subtask will comprise the entire collection of entities. In our above example for Populate Working Segments, the subtask comprises 50 working segments. When a task can be distributed, there is a one-to-one relationship between subtask and GeoCue entities. In the Ortho generation example above, there will be 40 subtasks.

The *Entities* that are associated with a subtask may not always be graphical objects in GeoCue. GeoCue supports the notion of “non-graphical” entities as well as *transient* (temporary) entities that cannot be viewed in a GeoCue Client Map View. These types are used under circumstances in which it is desirable to dispatch a task but physical entities do not exist. An example of this is the GeoCue Import Files command. While the imported file may ultimately have a graphical representation in GeoCue (for example, the Minimum Bounding Rectangle, MBR, of a Shape file), the entity cannot exist prior to the file import. In this particular example, temporary entities are created for the CDS.

Ultimately, the processing granularity of the Command Dispatch System is at the “subtask” level.

## **12.2 Machine Assignment**

Tasks (and hence subtasks) in GeoCue always have a one-to-one relationship with a Command as defined in Environment Builder. The machines (nodes) on which a Command can run can be specified (again, in the Command sub-tab of the Command tab in Environment Builder).

The Command Dispatch System supports two dispatch modes – machine or cloud. This mode is selected on the *Dispatch* dialog at the time a Task is dispatched. Machine dispatch mode is the default mode.

NOTE: By default, Cloud Dispatch Mode is disabled. You must enable this mode by selecting "Enable Cloud Dispatch Mode" on the Options sub-tab of the Configuration tab of Dispatch Manager.

## 12.2.1 Machine Dispatch Mode

When a command is interactively dispatched, GeoCue first displays the Dispatch dialog (Figure 12-1). The list of machines that can be selected for the task is determined from the command machine assignment in Environment Builder.

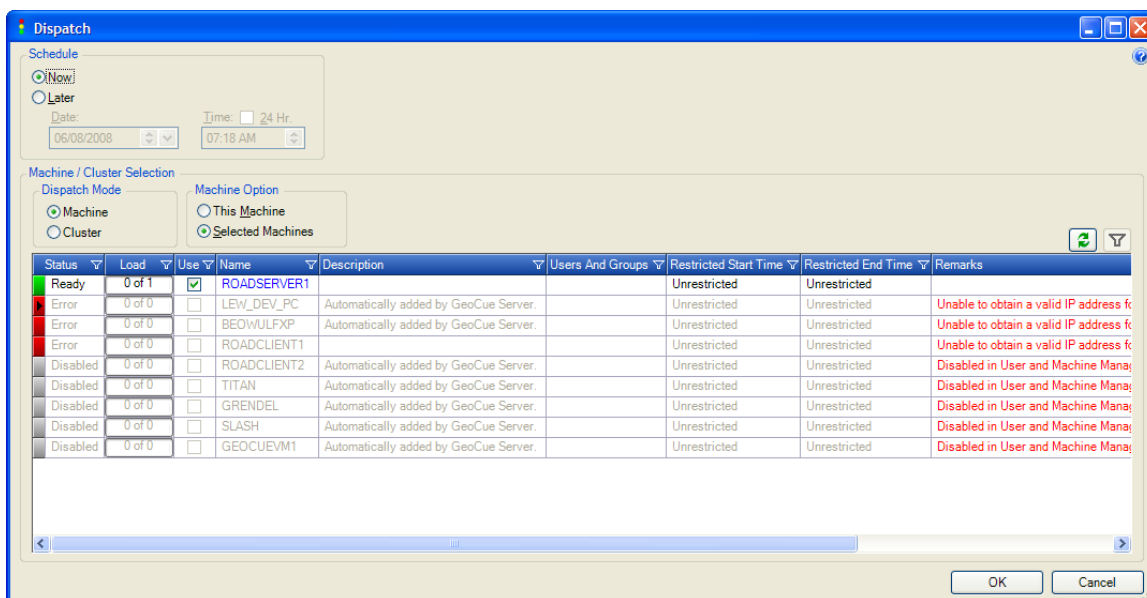


Figure 12-1: The Dispatch dialog



## 12.2.2 Cloud Dispatch Mode

If Cloud Dispatch Mode is enabled (Options sub-tab of the Configuration tab of Dispatch Manager), Cloud Dispatch is selected by toggling this option in the Dispatch dialog (see Figure 12-2). This will result in a list of available clouds as opposed to machines. Simply select the desired Cloud and press OK.

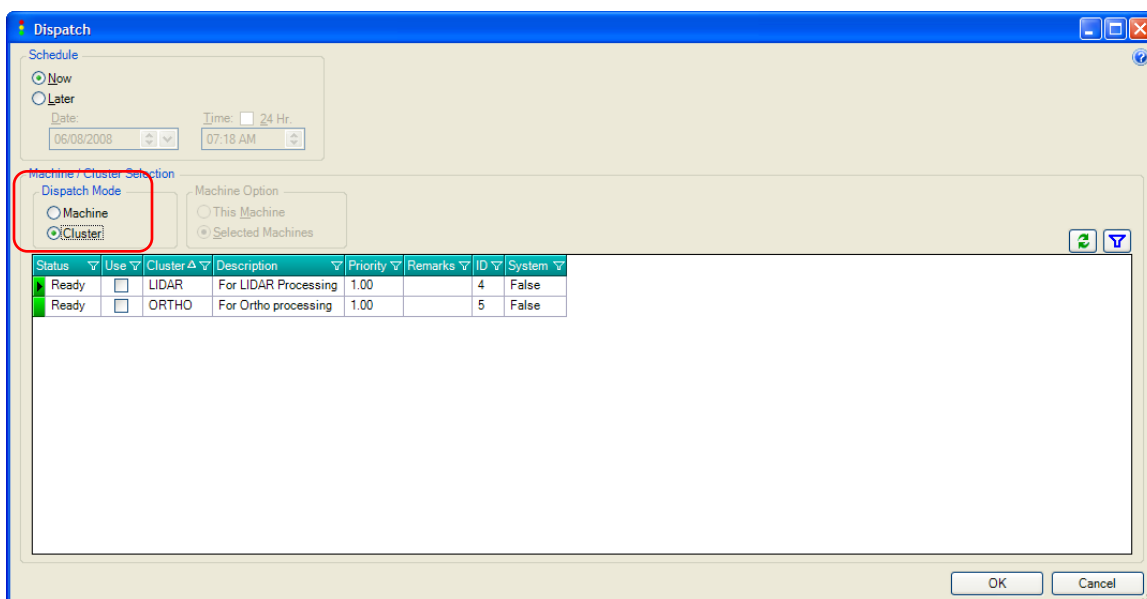


Figure 12-2: Setting Cloud Dispatch mode

## 12.3 Subtask Queuing

The Command Dispatch System maintains a Dispatch Queue in which all unprocessed subtasks are waiting, for various reasons, for processing. Subtasks waiting to be processed are said to be in the *pending* state. Subtasks are generally in a pending state due to one of the following reasons:

1. They are *scheduled* to execute at a later time
2. The subtask has sufficient priority to be dispatched but there is no machine currently available for processing<sup>11</sup>
3. The subtask does not have sufficient priority to move past other pending subtasks

Each subtask in the Dispatch Queue has a *priority*. Priorities are positive (or zero) floating point numbers. The lowest priority in the system is 0.0 and the highest is MAX\_FLOAT. Subtasks are dispatched from the dispatch queue in priority order, highest first. Two subtasks of equal priority are dispatched in the order in which they entered the dispatch queue; oldest first.

## 12.4 Machine (Node) Selection

Subtasks are vectored from the Dispatch Queue to a processing node. Since the subtask is the processing unit in the CDS, a subtask will ultimately be sent to a single, unique machine.

Several parameters are used in the algorithm that the Dispatch Processor used to select the machine to which a subtask will be vectored. The first component is the machine list provided by the Command setting in Environment Builder. If this list was not explicitly set, this implies that this subtask could be vectored to any machine in the Command Dispatch System constellation. Otherwise, the candidate machines will be restricted to those specified in the Environment Builder command table.

The Dispatch Processor next pars down the list of candidate machines to those that are currently available. A machine may not be available because the machine has been *disabled* for dispatch by its Scheduler (set on the Machines tab of Dispatch Manager) or it has been manually disabled by a user (using either the desktop disable tool or via Dispatch Manager).

---

<sup>11</sup> There are a number of reasons why a machine may not be available such as being fully loaded, the subtask not being allocated resources, the machine being off-line and so forth.

The Dispatch Processor next *pings* the remaining candidate machine list to ensure that they are capable of receiving commands. A machine that cannot be reached may have failed, be shut down or have a network communication problem. If all candidate machines fail the *ping* test, the CDS declares an error, sets the associated entity status fields to *error* and removes the subtask from the Dispatch Queue.

The Dispatch Processor next sequentially inspects each machine in the candidate machine list in terms of available resources. If Processor Grouping (Cloud) is not in use, this becomes a check to ensure that the currently available Standard Processing Units (SPU), Standard Memory Units (SMU) and/or Standard IO Units (SIU) are sufficient on the target machine to match the requirements of the current command<sup>12</sup>. The currently available resources of a machine are its total resources (as set in the Machines tab of Dispatch Manager less its current processing load less any reservations made by the Processor Group system). The first machine with sufficient resources to host the command will be the dispatch target for the subtask.

If no machines are capable of hosting the command, the subtask will remain in the pending state and will be retested with the next pass through the Dispatch Queue.

NOTE: Version 5.0 of the CDS does not check to see if a machine will never have sufficient resources to host a command. This means that if you dispatch a subtask that is associated with the command for which no machines have sufficient resources, the command will remain infinitely pending in the Dispatch Queue. For example, if you set up a command called Rectify in Environment Builder that requires 4.0 SPU but you set the maximum SPUs available on all of your dispatch nodes to 2.0, this command can never be matched to a machine with sufficient resources for processing. This is not considered a defect since it is the CDS System Administrator who sets the resources available on machines as well as the resources required by commands.

---

<sup>12</sup> The "load" of each command in terms of SPU, SMU and SIU is set on the Command tab, Command sub-tab in Environment Builder. The default is 1.0 SPU and ignore for SMU and SIU.

If Processing Groups are in effect (discussed in a separate chapter of this document), the above algorithm is simply modified to reduce the available resources as determined by the group allocation (as well as consideration of scheduling).

## **12.5 Subtask Execution**

Once a machine is selected for a subtask, that subtask is vectored to the machine for processing. Each processing machine has a GeoCue Dispatch Service - literally a Windows Service on a Windows machine and a Daemon on a UNIX machine that receives the subtask information (via .NET Remoting on a Windows target and TCP/IP on a UNIX target). The remote service invokes the command and waits for the exit code. The Dispatch Processor, in turn, monitors the remote service for a completion event. The Dispatch Processor tracks the accumulated time used by the command and provides this information to the user via Dispatch Manager, Task pane. If the remote node is equipped with a GeoCue DRIM, then Dispatch Manager provides detailed information on the progress of the command such as accumulated CPU time, memory usage by the executing command and so forth.

## **12.6 How the CDS Dispatches With Processing Groups**

The Command Dispatch System maintains a queue that we will call the Dispatch Queue<sup>13</sup>. When a Task is dispatched, the task is decomposed into subtasks. If the Task can be distributed then it will decompose into N subtasks where N is the number of entities that were placed in the working set when the command was launched via the Dispatch dialog. If the task cannot be distributed, it will decompose into a single subtask. This decomposition process is independent of the number of machines that might be configured in the system.

---

<sup>13</sup> This is a gross simplification of the actual machinery but will suffice for this discussion.

## 12.6.1 Dispatch Queue Entry

The subtasks that result from the above described task decomposition are assigned dispatch parameters according to various rules and algorithms in the CDS. The Processing Group is assigned based on the Project group assignment<sup>14</sup> (UNIVERSAL by default). The priority of each subtask is computed using the Priority Computation Algorithm. The machines on which this particular subtask is permitted to execute are extracted from tables that were set in Environment Builder. The required machine resources (SPU, SMU, SIU) are also obtained from the Environment Builder configuration tables.

After the above (and other) parameters are assigned the subtasks, they are entered into the dispatch queue. All subtasks that result from a Task will enter the dispatch queue with the exact same set of processing parameters (machine lists, priority group membership and so forth). These subtasks can be visualized in the Task tab of Dispatch Manager. Note that once the subtasks enter the dispatch queue, they are independent of one another.

## 12.6.2 Processing Groups and Dispatch

Processing Groups are used for machine allocation based on group membership. As discussed late in this chapter, the resources of the machines in your Dispatch Cloud can be partitioned among your named groups. Subtasks are assigned to group membership at dispatch time. If you are not dispatching in Cloud mode (described in a later chapter), the group assignment is made from the group membership of the project from which the subtask originated.

When a group is assigned to a subtask<sup>15</sup>, the Dispatcher will attempt to match it to a machine that has resources allocated to that same group (those resources must be free, of course).

---

<sup>14</sup> Note that this group assignment is overridden if the Task was dispatched in Cloud mode, discussed in a later chapter.

<sup>15</sup> In reality, subtasks always belong to a group. If a subtask has not been assigned to one of your named groups, it will be a member of the UNIVERSAL group.

The resources assigned to a group for a particular machine can allow “borrow” or forbid borrow (the default is to allow borrowing). Borrowing means that a subtask from group B can use resources allocated to Group A as long as no group A subtasks are trying to acquire those resources. This type of configuration optimizes the use of allocated resources.

### 12.6.3 Dispatch Pass One

In the CDS, subtasks are considered in priority order (highest to lowest). If two subtasks have the same priority, the “oldest” subtask (that is, the one that entered the dispatch queue first) will be considered first. If two subtasks have the same priority and time, the one with the higher Entity ID is considered first. Two entities cannot have the same ID so the tie is guaranteed to be broken at this point.

The selected subtask is examined to obtain a list of the machines on which it can run. The machines are then considered in the order in which they appear in the Machine Table of Dispatch Manager. In order to receive the subtask, the machine must meet the following conditions:

1. The machine must be Enabled (in Dispatch Manager)
2. The machine must be available in terms of being reachable by the Dispatcher (e.g. it is not shut down or not reachable via the network)
3. The machine must be available according to the machine schedule (set in Dispatch Manager)
4. It must have resources available that are assigned to the same group as the subtask
5. The resources in this group must be sufficient to meet the processing specifications of the command associated with the subtask (SPU, SMU, SIU, minimum remaining resources, etc.)
6. The resources of the previous step must be available in terms of the Resource Schedule (set in the Manage groups function of Dispatch Manager)

If the above conditions are met, the subtask is dispatched to the candidate machine.

If the above conditions are not met, the next machine is considered and the tests above are repeated. If the subtask is not dispatched during pass one, it will be considered in pass two after all subtasks in the dispatch queue are considered for pass one.

## 12.6.4 Dispatch Pass Two

If subtasks are considered against all machines in Pass One (see previous section) and cannot be dispatched, then a second pass is made through the machines with the remaining subtasks, in “borrow resources’ mode.

The machines are considered in the same order as pass one. The test this time is:

1. The machine must be Enabled (in Dispatch Manager)
2. The machine must be available in terms of being reachable by the Dispatcher (e.g. it is not shut down or not reachable via the network)
3. The machine must be available according to the machine schedule (set in Dispatch Manager)
4. Resource Groups on the machine are tested in the reverse order of the Priority of the Groups (as set in Dispatch Manager). For each group, if the group allows ‘borrow’ (the default mode) then
  - The resources in this group must be sufficient to meet the processing specifications of the command associated with the subtask (SPU, SMU, SIU, minimum remaining resources, etc.)
  - The resources of the previous step must be available in terms of the Resource Schedule (set in the Manage groups function of Dispatch Manager)
5. If the subtask was not dispatched, consider the next group
6. If the subtask was not dispatched, consider the next machine.

If the subtask has not dispatched after two passes, then it will not be considered for dispatch again until a running subtask completes or a schedule event occurs that adds machines and/or resources to the dispatch list.

## 12.6.5 Starvation Prevention

The astute observer will note that a starvation condition could occur if a subtask with a large resource requirement is in the queue with a number of subtask with more modest resource requirements.

We prevent this by applying the following dispatch rules:

1. If a subtask that belongs to a group (say A) is pending in the dispatch queue, no subtask that belongs to the same group (A) and is of lower priority can be dispatched.
2. If a subtask that belongs to a group (say A) is pending in the dispatch queue, no subtask of lower priority, belonging to a different group, can borrow resources from group A.

It is not necessary that a user understand the logic behind the above rules. We only point them out in case you noticed the starvation condition that could occur.

## 12.7 Subtask Completion Processing

If the command fails via abnormal termination (e.g. the executable crashes), this state is conveyed back to the Dispatch Processor. The Dispatch Processor logs this information in the processing status fields and sets the state of the associated entity(s) to ERROR. If the entity(s) are graphic, the GeoCue Client will display them with ERROR symbology.

When the command completes, the GeoCue Dispatch Service conveys the Exit Code back to the Dispatch Processor. The Dispatch Processor then processes the return code (exit code) via the standard GeoCue Command Return Code machinery (as set on the Return Code sub-tab and Machine sub-tab of the Machine tab of Environment Builder).



## 13 Concluding Remarks

We hope that you have enjoyed working with the GeoCue product family. Hopefully you have not discovered too many software defects (bugs).

The primary idea that we would like for you to recognize in working with GeoCue is that it is a general purpose production management system that, when encoded with the appropriate *Environment*, is suitable to a very wide variety of process management problems. As time moves forward, we and third party companies will be releasing a number of different CuePacs for different production workflows and disciplines. Our ultimate goal is that you employ GeoCue on every production workstation in your company. We have a singular focus on improving your bottom-line profitability through enterprise process management.

## 14 Acronyms

<b>Acronym</b>	<b>Definition</b>
CDS	Command Dispatch System
DSL	Dispatched Subtask License
DRIM	Dynamic Resource Interrogation Module
EPC	Entity Priority Clip
EPW	Entity Priority Weight
ERP	Entity Relative Priority
GPC	Group Priority Clip
GPW	Group Priority Weight
GRP	Group Relative Priority
IP	Internet Protocol
PCA	Priority Computation Algorithm
PPC	Project Priority Clip
PPW	Project Priority Weight
PRP	Project Relative Priority

SDK	Software Developer's Kit
SIU	Standard I/O Units (SIU)
SMU	Standard Memory Units (SMU)
SPU	Standard Processor Units
WS	Working Set
WSQ	Working Set Queue (synonymous with Working Set)