
Dart: Testing, Reports and Dashboards

Daniel J. Blezek
James V. Miller

September 27, 2005

Acknowledgments

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

This product includes software developed by the Visigoth Software Society (<http://www.visigoths.org/>).

This work funded in part by the National Library of Medicine of the National Institutes of Health as a component of the Insight Segmentation and Registration Toolkit (<http://www.itk.org>), contract number N01-LM-9-3531.

This work is part of the National Alliance for Medical Image Computing (NA-MIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>

License

Copyright (c) 2004-2005, The Insight Consortium
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Insight Consortium nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Introduction

1.1 Dart Statement of Purpose

Dart shall aggregate data across many independent distributed build and test hosts, summarizing the software quality aspects of the project in a concise and informative fashion cross-sectionally and longitudinally.

Tutorial

2.1 Quick Start

If you are building Dart from source, please refer to Section 4.1 to build `DartServer.jar`, then return to this Section. If you've downloaded the jar file, you may begin here. The distributions are build using JDK 1.4. For users upgrading from a pre 0.4 release, please see Section 2.8.

1. Create the Server directory and configuration

```
java -jar DartServer.jar --createserver TestServer
```

The `--createserver` flag creates a new Server directory and an default configuration file `Server.xml`.

2. Initialize the server

```
java -jar DartServer.jar --initializeserver TestServer
```

Initializes the Server database.

3. Create the project

```
java -jar DartServer.jar --create TestProject
```

The `--create` flag requires a directory name argument. Dart will create this directory for the Project (`TestProject` in this example).

4. Configure the Project By default, a reasonable settings file is found in `TestProject/Dart.xml`. If desired, this file may be edited to change the projects settings (see Section 2.5). Now is a good time to do this.
5. Start the project

```
java -jar DartServer.jar \  
    --initialize --refresh --refreshServer TestServer TestProject
```

The `--initialize` flag instructs the server to create the database tables that Dart requires for the project, while `--refresh` copies the project resources into the `TestProject` directory. `TestServer` is the name of the Dart Server, while any projects to be started can be configured in `TestServer/Server.xml` and are overridden by the commandline arguments. Note that `\` indicates line continuation; that is, the code above should be typed on one line. The `--refreshServer` puts some server specific files in the correct locations in the `Server` directory.

6. View the dashboard. Point your browser to `http://localhost:8081/TestProject/Dashboard/` to view the (empty) Dashboard.

The project `TestProject` is now up and running accepting XML-RPC submissions on port 8080 and serving HTML pages on port 8081. The ports and other configurations are covered in Section 2.4.

2.2 Submission

Dart ships with a utility called `DartClient` to submit results to the server. The basic use is:

```
# java -jar DartClient.jar TestProject Results.xml
```

This submits `Results.xml` to the `TestProject` Project on the Server running on `localhost`. Submission is only a copy, and does not provide feedback on the XML validity.

`DartClient` also provides other options:

```
# java -jar DartClient.jar --help  
0 [main] INFO dart.DartClient - Starting DartClient  
usage: DartClient [options] Project <foox.xml> <foo2.xml> ... <fooN.xml>  
-g,--getstatus    Get Server status  
-h,--help         Print help message  
-p,--port         XML-RPC Port to connect to, 8080 is default  
-q,--shutdown     Shutdown the Server  
-r,--refresh      Refresh Project resources  
-s,--server       Server to connect to, localhost is default
```

To connect through a proxy or firewall use:

```
java -Dhttp.proxyPort=8080 -Dhttp.proxyHost=proxyhost.mydomain.org \  
    -jar DartClient.jar --help
```

with `http.proxyPort` and `http.proxyHost` replaced by your proxy port and server.

2.3 Software Installation

To be completed.

2.4 Server Setup

The Server has several command line options.

```
# java -jar DartServer.jar
usage: DartServer [options] Server.xml <Project0.xml> <Project1.xml> ...
        <ProjectN.xml>
-R,--refreshServer      Refresh server resources
-a,--archive            Archive the project
-c,--create            Create a new project in the directory specified
-d,--database          At project creation time, configure the
                        Schema.sql file for generic, Postgres, Derby
-h,--help              Print help message
-i,--initialize        Initialize the database from the Schema.sql file
                        in the project directory
-j,--initializeserver  Initialize the database from the
                        ServerSchema.sql file in the dart server directory
-k,--createserver      Create a new server in the directory specified
-l,--logconfiguration  File to configure log4j from, defaults are used
                        if not present
-q,--createdqed        Create a new QED in the directory specified
-r,--refresh           Refresh project resources
-t,--projecttemplate   Create a new Project using the specified default
                        template: dart/Resources/Server/DartDefault.xml in the jar file is the
                        default
```

The `--archive` flag dumps all the Submissions in each project in Dart XML format into the Project/Archive directory. This is the best way to archive a Project all at once.

The `--projecttemplate` flag is used in conjunction with the `--create` argument at project creation time. If not specified, the file `dart/Resources/Server/DartDefault.xml` is used as the Project template. To extract this file, use the command:

```
jar fxv DartServer.jar dart/Resources/Server/DartDefault.xml
```

The `DartDefault.xml` file may be edited to suit the site specific needs. Take care to preserve the FreeMarker tags in the original file, as they are specific to some parts of the Project creation process.

2.5 Server Configuration

In the `TestServer` directory, there is a file named `Server.xml`. This contains the default settings for a dart server. The sections of the Server configuration are as follows.

2.5.1 Server Info

```
<?xml version="1.0" encoding="utf-8"?>
```

```
<Server>
  <Title>TestServer</Title>
  <BaseDirectory>f:\Source\Dart\TestServer</BaseDirectory>
```

This is the XML preamble followed by the Server information: Title and BaseDirectory.

2.5.2 Ports

```
<HTTPPort>8081</HTTPPort>
<XMLRPCPort>8080</XMLRPCPort>
```

These are the HTTP and XML-RPC port specifications. They are required to be different.

2.5.3 Scheduler

```
<Scheduler>
  <ThreadPoolSize>10</ThreadPoolSize>
</Scheduler>
```

The Scheduler has a default ThreadPoolSize of 10, indicating that 10 jobs may be executed concurrent by all the Projects managed by this Server instance.

2.5.4 Database

```
<!-- Configure the database parameters derby-->
<Database>
  <!-- Derby database -->
  <Driver>org.apache.derby.jdbc.EmbeddedDriver</Driver>
  <URL>jdbc:derby:f:\Source\Dart\TestServer/Database/TestServer;create=true</URL>
  <ShutdownURL>jdbc:derby:f:\Source\Dart\TestServer/Database/TestServer;shutdown=true</ShutdownURL>
  <Username/>
  <Password/>
  <!-- Maximum active / idle connections, -1 is infinite -->
  <MaxActive>10</MaxActive>
  <MaxIdle>3</MaxIdle>
</Database>
```

This section specifies the connection to the Server's database. In this example, the database is Derby. The Driver tag specifies the class implementing the JDBC connection. URL is the connection string, ShutdownURL is used to shutdown the Derby database cleanly, it may be safely left blank for other JDBC packages. The Username and Password tags specify the connection parameters. Dart uses a connection pooling mechanism for the database with two parameters: MaxActive specifies the maximum number of active connections, and MaxIdle specifies the maximum number of idle threads. If a connection is needed when MaxActive threads are already active, the connection will hang until a connection is return to the pool.

2.5.5 Servlet Manager

```
<!-- Servlet configuration -->
<ServletManager>
  <Servlet>
    <Class>dart.server.servlet.Server</Class>
    <Context>/Dart/*</Context>
    <Properties/>
  </Servlet>
</ServletManager>
```

The Servlet manager is responsible for configuring Jetty. Different Servlets can respond to different URLs. In this case, the `dart.server.servlet.Server` class is configured to respond to requests starting with `/Dart/*`.

2.6 Project Configuration

After following the directions in Section 2.1, in the `TestProject` directory, there is a file named `Project.xml`. This is a preconfigured Project configuration file containing all the settings required to run a basic Dart Project. The contents of the file are not presented in their entirety, but will be discussed section by section.

2.6.1 Project Info

```
<?xml version="1.0" encoding="utf-8"?>
<Project>
  <Title>TestProject</Title>
  <BaseDirectory>/projects/Dart/TestProject</BaseDirectory>
```

The first line is the `xml` preamble, and is required for all `xml` files. The `<Project>` tag indicates the start of the Project configuration. `<Title>` is the project title, and `<BaseDirectory>` is the absolute path name to the Project directory. If the Project is moved to a new location on the file system, `<BaseDirectory>` must be changed to reflect the new location.

2.6.2 Project Properties

Certain aspects of a Dart project can be customized by properties assigned to the project.

```
<Properties>
  <Property name="MaxTestsPerSubmission">1500</Property>
  <Property name="RepositoryURL">http://www.itk.org/cgi-bin/viewcvs.cgi/</Property>
  <Property name="RepositoryURL.Type">viewcvs</Property>
  <Property name="RepositoryURL.Repository">Insight</Property>
  <Property name="Menu">
    <![CDATA[
      <ul>
```

```

<li><a href="#">Repository &gt;</a>
  <ul>
    <li><a href="http://www.itk.org/cgi-bin/viewcvs.cgi?cvsroot=Insight">
      Insight</a>
    </li>
    <li>
      <a href="http://www.itk.org/cgi-bin/viewcvs.cgi?cvsroot=InsightApplications">
        Insight Applications
      </a>
    </li>
  </ul>
</li>
<li><a href="http://www.itk.org/Bug">Bugs</a></li>
<li><a href="#">Documentation &gt;</a>
  <ul>
    <li><a href="http://www.itk.org/Insight/Doxygen/html/">Doxygen (API)</a></li>
    <li><a href="http://www.itk.org/">www.itk.org</a></li>
    <li><a href="http://www.itk.org/Wiki/ITK">ITK Wiki</a></li>
    <li><a href="http://www.insightsoftwareconsortium.org/">
      Insight Software Consortium</a>
    </li>
  </ul>
</li>
<li><a href="http://www.itk.org/">Home</a></li>
</ul>
]]>
</Property>
</Properties>

```

MaxTestsPerSubmission Defines a throttle for the number of tests that can be included in a submission.

RepositoryURL URL for accessing a project's software repository via the web.

RepositoryURL.Type Identifies the type of the web portal for the repository (viewcvs, cvsweb, websvn).

RepositoryURL.Repository Identifies the repository to reference at the specified RepositoryURL.

Menu Definition of a menu to display on web pages. This menu can include links to navigate within Dart or to navigate to external sites such as documentation, bug trackers, etc. Menu is defined as an HTML list (encapsulated in a CData section), with submenus being sublists. Each menu item is modelled with an HREF within an LI. If a menu item is the title of a submenu, then the HREF should link to "" and the item label should indicate a drop down menu is available (for instance, suffixing the name with a > sign).

2.6.3 Database Configuration

```

<!-- Derby database -->
<Database>

```

```

<Driver>org.apache.derby.jdbc.EmbeddedDriver</Driver>
<URL>jdbc:derby:/projects/Dart/TestProject/Database/TestProject;create=true</URL>
<ShutdownURL>jdbc:derby:/projects/Dart/TestProject/Database/TestProject
;shutdown=true</ShutdownURL>

<Username/>
<Password/>
</Database>

```

This section configures the Database connection that the Project will use. The `<Driver>` tag indicates the JDBC Java class for the particular type of relational database management system (RDBMS). In the example, `org.apache.derby.jdbc.EmbeddedDriver` is the driver for the Derby Open Source embedded RDBMS system. Note: the `<ShutdownURL>` was broken across two lines for display purposes, and should not be broken in an actual configuration file.

The `<URL>` tag specifies the connection URL for the RDBMS. This is a RDBMS specific string. In the above example, the `create=true` property indicates that the driver should create the database if it does not exist. Please consult your RDBMS documentation for the proper setting for the `<URL>` tag. Because Derby is an embedded RDBMS, it must be properly shutdown to leave the database in a consistent state. This is specified in the `<ShutdownURL>` tag. If the RDBMS does not require special shutdown processing, leave this tag empty and it will be ignored.

`<Username/>` and `<Password>` tags specify the authentication settings for the RDBMS. In the case of Derby, no Username/Password is required.

2.6.4 CommandManager Configuration

```

<CommandManager>
  <Command>
    <Name>Submit</Name>
    <Class>dart.server.command.Submit</Class>
    <Properties>
      <Property name="DeleteWhenDigested">true</Property>
    </Properties>
  </Command>
</CommandManager>

```

The Dart Server provides an XML-RPC server for results to be submitted to a Project. This server operates through a Servlet configured in the ServletManager (see Section 2.6.5 below). For the CommandManager to operate, a `dart.server.servlet.CommandServlet` object must be added to the ServletManager. In addition, the `CommandServer` can be configured to respond to any query using specialized Commands. The `<CommandServer>` section specifies the settings for the Project specific settings.

In the instance above, the `<Command>` tag specifies an object that the Project will use to respond to XML-RPC calls. Commands must implement the `dart.server.command.Command` interface. `<Name>` is the object name, `<Class>` is the name of the class that the CommandManager instantiates and any Properties for the object are specified using the `<Properties>` tags. Any public methods of the object are exposed to XML-RPC calls.

2.6.5 ServletManager Configuration

```
<!-- Servlet configuration -->
<ServletManager>
  <Servlet>
    <Class>dart.server.servlet.Dashboard</Class>
    <Context>/Dashboard/*</Context>
    <Properties>
    </Properties>
  </Servlet>
  <Servlet>
    <Class>dart.server.servlet.CommandServlet</Class>
    <Context>/Command/</Context>
    <Properties/>
  </Servlet>
  <Servlet>
    <Class>dart.server.servlet.ChartServlet</Class>
    <Context>/Chart</Context>
    <Properties/>
  </Servlet>
</ServletManager>
```

To generate Dashboard pages, the Server uses the Jetty Servlet engine in conjunction with the FreeMarker template engine. Stock Project Servlets are automatically configured at project creation time. User defined Servlets may be added if desired. The `<Class>` tag indicates the class of the Servlet, `<Context>` tag indicates how the Servlet is found by Jetty. By default, the Project title is stored in the Servlet's initial parameters as ```project``` and may be accessed as `getInitParameter (``project``)` within the Servlet. Parameters in the `<Properties>` section are also put in the initial parameters map.

The second Servlet in the stock configuration is `dart.server.servlet.CommandServlet`. `CommandServlet` accepts XML-RPC calls and delegates them to the appropriate handler object as configured in the `CommandManager`. The last Servlet is the `ChartServlet` used to generate charts for the dashboard.

To call an XML-RPC method, the URL needed is determined by the root project URL, *i.e.* `http://localhost:8081/ProjectName/Command/Command.Method`. For example, the URL to submit some results to the Dart project `TestProject` running on the local system is: `http://localhost:8081/TestProject/Command/` and the method is `Submit.put`.

2.6.5.1 Adding a new servlet to Dart

If you are building Dart from source, you can add additional servlets to the `DartServer.jar` by adding the servlet's source code to `Dart/Source/dart/server/servlet` directory and rebuilding. The servlet can be activated for a project as in the above example.

If you are running from a pre-built `DartServer.jar`, you can still add additional servlets to Dart. This is done by adding the `classpath` property to the servlet definition in the `Project.xml` file.

```
<Servlet>
  <Class>user.servlet.MyDashboard</Class>
```



```

<Context>/MyDashboard/*</Context>
<Properties>
  <Property name="Foo">10</Property>
  <Property name="classpath">foo.jar</Property>
</Properties>
</Servlet>

```

This example adds the class `user.servlet.MyDashboard` from the `foo.jar` archive to the project and assigns it the url space of `/MyDashboard/*`.

2.6.6 Task Configuration

```

<!-- Scheduled tasks. The Schedule tag is in cron format. -->
<Task>
  <Type>dart.server.task.QueueManager</Type>
  <Schedule>0/10 * * * * ?</Schedule>
  <Properties>
    <Property name="MaxTasks">10</Property>
  </Properties>
</Task>

```

Tasks configured in the `Dart.xml` file are periodically scheduled. Tasks must implement the `dart.server.task.Task` interface. In the above example, the `dart.server.task.QueueManager` is scheduled to run every 10 seconds. The `QueueManager` class processes other Tasks that have been placed in the `TaskQueue`. The `Properties` tag specifies settings that are passed into the Task when it executes. For `QueueManager`, the “MaxTasks” property indicates how many queued tasks will be processed at during each execution, providing a “throttling” mechanism. The format of the `<Schedule>` tag is detailed at <http://quartz.sourceforge.net/javadoc/org/quartz/CronTrigger.html> and is reproduced here for clarity.

2.6.6.1 Cron Expressions

For those unfamiliar with “cron”, this means being able to create a firing schedule such as: “At 8:00am every Monday through Friday” or “At 1:30am every last Friday of the month”.

A “Cron-Expression” is a string comprised of 6 or 7 fields separated by white space. The 6 mandatory and 1 optional fields are as follows:

Field Name	Allowed Values	Allowed Special Characters
Seconds	0-59	, - * /
Minutes	0-59	, - * /
Hours	0-23	, - * /
Day-of-month	1-31	, - * ? / L W C
Month	1-12 or JAN-DEC	, - * /
Day-of-Week	1-7 or SUN-SAT	, - * ? / L C #
Year (Optional)	empty, 1970-2099	, - * /

The ‘*’ character is used to specify all values. For example, “*” in the minute field means “every minute”.

The '?' character is allowed for the day-of-month and day-of-week fields. It is used to specify 'no specific value'. This is useful when you need to specify something in one of the two fields, but not the other. See the examples below for clarification.

The '-' character is used to specify ranges. For example "10-12" in the hour field means "the hours 10, 11 and 12".

The ',' character is used to specify additional values. For example "MON,WED,FRI" in the day-of-week field means "the days Monday, Wednesday, and Friday".

The '/' character is used to specify increments. For example "0/15" in the seconds field means "the seconds 0, 15, 30, and 45". And "5/15" in the seconds field means "the seconds 5, 20, 35, and 50". You can also specify '/' after the '*' character - in this case '*' is equivalent to having '0' before the '/'.

The 'L' character is allowed for the day-of-month and day-of-week fields. This character is short-hand for "last", but it has different meaning in each of the two fields. For example, the value "L" in the day-of-month field means "the last day of the month" - day 31 for January, day 28 for February on non-leap years. If used in the day-of-week field by itself, it simply means "7" or "SAT". But if used in the day-of-week field after another value, it means "the last xxx day of the month" - for example "6L" means "the last friday of the month". When using the 'L' option, it is important not to specify lists, or ranges of values, as you'll get confusing results.

The 'W' character is allowed for the day-of-month field. This character is used to specify the weekday (Monday-Friday) nearest the given day. As an example, if you were to specify "15W" as the value for the day-of-month field, the meaning is: "the nearest weekday to the 15th of the month". So if the 15th is a Saturday, the trigger will fire on Friday the 14th. If the 15th is a Sunday, the trigger will fire on Monday the 16th. If the 15th is a Tuesday, then it will fire on Tuesday the 15th. However if you specify "1W" as the value for day-of-month, and the 1st is a Saturday, the trigger will fire on Monday the 3rd, as it will not 'jump' over the boundary of a month's days. The 'W' character can only be specified when the day-of-month is a single day, not a range or list of days.

The 'L' and 'W' characters can also be combined for the day-of-month expression to yield 'LW', which translates to "last weekday of the month".

The '#' character is allowed for the day-of-week field. This character is used to specify "the nth" XXX day of the month. For example, the value of "6#3" in the day-of-week field means the third Friday of the month (day 6 = Friday and "#3" = the 3rd one in the month). Other examples: "2#1" = the first Monday of the month and "4#5" = the fifth Wednesday of the month. Note that if you specify "#5" and there is not 5 of the given day-of-week in the month, then no firing will occur that month.

The 'C' character is allowed for the day-of-month and day-of-week fields. This character is short-hand for "calendar". This means values are calculated against the associated calendar, if any. If no calendar is associated, then it is equivalent to having an all-inclusive calendar. A value of "5C" in the day-of-month field means "the first day included by the calendar on or after the 5th". A value of "1C" in the day-of-week field means "the first day included by the calendar on or after sunday".

The legal characters and the names of months and days of the week are not case sensitive.

Here are some full examples:

Expression	Meaning
"0 0 12 * * ?"	Fire at 12pm (noon) every day
"0 15 10 ? * *"	Fire at 10:15am every day
"0 15 10 * * ?"	Fire at 10:15am every day
"0 15 10 * * ? *"	Fire at 10:15am every day
"0 15 10 * * ? 2005"	Fire at 10:15am every day during the year 2005
"0 * 14 * * ?"	Fire every minute starting at 2pm and ending at 2:59pm, every day
"0 0/5 14 * * ?"	Fire every 5 minutes starting at 2pm and ending at 2:55pm, every day
"0 0/5 14,18 * * ?"	Fire every 5 minutes starting at 2pm and ending at 2:55pm, AND fire every 5 minutes starting at 6pm and ending at 6:55pm, every day
"0 0-5 14 * * ?"	Fire every minute starting at 2pm and ending at 2:05pm, every day
"0 10,44 14 ? 3 WED"	Fire at 2:10pm and at 2:44pm every Wednesday in the month of March.
"0 15 10 ? * MON-FRI"	Fire at 10:15am every Monday, Tuesday, Wednesday, Thursday and Friday
"0 15 10 15 * ?"	Fire at 10:15am on the 15th day of every month
"0 15 10 L * ?"	Fire at 10:15am on the last day of every month
"0 15 10 ? * 6L"	Fire at 10:15am on the last Friday of every month
"0 15 10 ? * 6L"	Fire at 10:15am on the last Friday of every month
"0 15 10 ? * 6L 2002-2005"	Fire at 10:15am on every last friday of every month during the years 2002, 2003, 2004 and 2005
"0 15 10 ? * 6#3"	Fire at 10:15am on the third Friday of every month

Pay attention to the effects of '?' and '*' in the day-of-week and day-of-month fields!

NOTES:

- Support for the features described for the 'C' character is not complete.
- Support for specifying both a day-of-week and a day-of-month value is not complete (you'll need to use the '?' character in one of these fields).
- Be careful when setting fire times between mid-night and 1:00 AM - "daylight savings" can cause a skip or a repeat depending on whether the time moves back or jumps forward.

Another example Task is the `SaveStatistics` Task which writes the Project statistics to the file system on a regular bases. The Task does not have any parameters, and saves the statistics in a file named `Statistics.txt` in the Project directory.

```
<Task>
  <Type>dart.server.task.SaveStatistics</Type>
  <Schedule>0 * * * * ?</Schedule>
  <Properties>
  </Properties>
</Task>
```

An extremely important Task is the `ArchiveTask`. This Task is detailed in Section 2.7.

2.7 Archiving a Dart Server

Though the `ArchiveTask` strictly belongs in the Task section of the configuration file, it's options are quite complex and deserve a section of it's own. The basic job of the `ArchiveTask` is to copy data from the Dart

database to the file system in a easy to restore format and, optionally, delete the data from the database.

A project can have 0 or more "Archivers". Each Archiver selects a set of Submissions that matches some criteria. The current selection criteria are:

- AgeInDays A submission must be older than this to be considered (may be BuildStamp or Created-TimeStamp[when the submission happened])
- MatchTrack a comma seperated list of regular expressions to match against the track of each Submission
- MatchSite a comma seperated list of regular expressions to match against the Site
- MatchBuildName a comma seperated list of regular expressions to match against the BuildName
- MatchTest a comma seperated list of regular expressions to match against each test in the submission

Once a submission is matched, it is archived, if not already archive, to a user specified directory (default is ProjectDirectory/Archive). When the "Working" directory reaches a specific size (MaxDirectorySizeMB default is 700), it is renamed using the current date time stamp and a new Working directory is created. In this way, a set of 700MB directories is created and may be safely moved to CDROM or other archival media and deleted.

Each Archiver has specified levels of archive, specified by the ArchiveLevel:

- 0 default, all data is still in the Database
- 1 remove least amount of data. All bulk data (images, logs, etc...) are to be removed
- 2 remove all leaf tests and data, leaving only intermediate levels in the Test hierarchy
- 3 remove all non-root Tests, leaving rollout info at the root level of the Test hierarchy
- 4 remove all data and the Submission itself. This is as if the submission had never existed

This is an example of an ArchiveTask definition:

```
<Task>
  <!-- Archive task, runs at 3am each morning -->
  <Type>dart.server.task.ArchiveTask</Type>
  <Schedule>*/30 * * * * ?</Schedule>
  <Properties>
    <Property name="ArchiverList">Nightly</Property>

    <Property name="Archiver.Nightly.WriteArchive">true</Property>
    <Property name="Archiver.Nightly.FileNamePattern">Archive-%L-%P-%S-%B-%T-%D.xml.gz</Property>
    <Property name="Archiver.Nightly.Template">ArchiveSubmission.xml</Property>
    <Property name="Archiver.Nightly.MatchTrack">.*</Property>
    <Property name="Archiver.Nightly.AgeInDays">0.010</Property>
    <Property name="Archiver.Nightly.ArchiveLevel">3</Property>
    <Property name="Archiver.Nightly.MatchTest">.*</Property>
```

```

    <Property name="Archiver.Nightly.MatchSite">.*</Property>
    <Property name="Archiver.Nightly.ArchiveBy">CreatedTimeStamp</Property>
    <Property name="Archiver.Nightly.MaxDirectorySizeMB">700</Property>
    <Property name="Archiver.Nightly.ArchiveDirectory">/tmp/DartArchive</Property>

  </Properties>
</Task>

```

This tag defines the list of Archivers (comma separated). So we only have one called "Nightly"

```

  <Property name="ArchiverList">Nightly</Property>

```

Each Archiver is run in the order of the ArchiverList Property.

All properties for this Archiver are indicated by "Archiver.Foo.PropertyName" where Foo is the name of the property. For example: "Archiver.Nightly.AgeInDays". Below, when I refer to a Property, it is assumed there is a "Archiver.Nightly." prefix to it. So ArchiveLevel really refers to "Archiver.Nightly.ArchiveLevel".

The first check is to see match Submissions by AgeInDays. In this case, 0.010 is about 10 minutes. The Database column to match is the CreatedTimeStamp. This is when the submission hit the Dart Server. The other option for ArchiveBy is TimeStamp, which is the datetime reported by the Client. The last criteria is that the Submission must have an ArchiveLevel less than the one specified for this Archiver (this is explained below). Thus to match, a Submission must be older than 0.01 days by CreatedTimeStamp, and have an archive level of ≤ 3 .

```

  <Property name="Archiver.Nightly.AgeInDays">0.010</Property>
  <Property name="Archiver.Nightly.ArchiveBy">CreatedTimeStamp</Property>
  <Property name="Archiver.Nightly.ArchiveLevel">3</Property>

```

From the list of Submissions return, each is examined in turn. The first test is by Site and BuildStamp:

```

  <Property name="Archiver.Nightly.MatchSite">.*</Property>
  <Property name="Archiver.Nightly.MatchBuildName">.*</Property>

```

MatchSite is a comma separated list of regular expressions. ".*" matches all sites. If the Submission does not match any of the expressions, it is ignored. The same matching is applied for MatchBuildName.

The next test is to Match the Track:

```

  <Property name="Archiver.Nightly.MatchTrack">.*</Property>

```

Again, a comma separated list. Here we match all tracks.

If a Submission matches by Site, Track and BuildName, the Submission's ArchiveDateStamp and ArchiveLevel are examined. If ArchiveDateStamp is null, then go ahead and archive the Submission to disk as outlined below. If the CreatedTimeStamp is later than the ARchiveDateStamp, some new data has hit the database since the last archive, so go ahead and archive the Submission to disk.

To archive to Disk, the Archiver writes all the Submission data to one big XML file using the new Dart v2 XML format. This is specified as a FreeMarker template by this Property.

```
<Property name="Archiver.Nightly.Template">ArchiveSubmission.xml</Property>
```

ArchiveSubmission.xml is the default. This file is suitable for re-submitting to the Dart Server. However, if the WriteArchive property is false, the data will not be written to disk. This is useful for Experimental and continuous builds that you may not care about. "true" is the default.

```
<Property name="Archiver.Nightly.WriteArchive">true</Property>
```

Once the file has been generated, it is written to disk. Each Archiver has a location to write to. The default is the "ProjectDirectory/Archive" directory that is created when the project starts up. In our case we write to /tmp/DartArchive:

```
<Property name="Archiver.Nightly.ArchiveDirectory">/tmp/DartArchive</Property>
```

Inside /tmp/DartArchive two directories will be created: Working and Temporary. Once the XML is generated in Temporary, it is gzipped and moved to Working. The filename is generated by this Property:

```
<Property name="Archiver.Nightly.FileNamePattern">Archive-%L-%P-%S-%B-%T-%D.xml.gz</Property>
```

Where:

- %L is replaced by ArchiveLevel
- %P is replaced by the project name
- %S is replaced by the Site name
- %B is replaced by the BuildName
- %T is replaced by the TrackName
- %D is replaced by the DateTimeStamp
- %N is replaced by the current time "now" in UTC

So you can have as much or as little info encoded in the filename as you like.

If moving the file from Temporary would push the Working Directory over the limit set by:

```
<Property name="Archiver.Nightly.MaxDirectorySizeMB">700</Property>
```

The Working directory is renamed by the current time in UTC and a new Working directory is started. We could specify that the Working directory is rolled over on a regular time schedule rather than a file size criteria if that would be useful.

Finally, the Archiver decides which data from the Submission to remove (since it's all been archived by now). This is decided by ArchiveLevel

```
<Property name="Archiver.Nightly.ArchiveLevel">3</Property>
```

ArchiveLevel indicates how much data to remove

- 1 remove least amount of data. All bulk data (images, logs, etc...) are to be removed
- 2 remove all leaf tests and data, leaving only intermediate levels in the Test hierarchy
- 3 remove all non-root Tests, leaving rollup info at the root level of the Test hierarchy

- 4 remove all data and the Submission itself

Each Test is matched using the MatchTest Property:

```
<Property name="Archiver.Nightly.MatchTest">.*</Property>
```

Except at level 4, where everything is removed. If a test does not match, it is ignored and no data is deleted from it.

Finally, the Submission's ArchiveLevel is updated, as is the ArchivedTimeStamp. Now this Submission will not match the initial select, as it's ArchiveLevel is equal to the ArchiveLevel for this Archiver.

The Archiver then processes the next Submission until no more are found.

2.7.1 Default Archivers

By default, several Archivers are defined in the Project, however the Archivers are not active. They are:

- Nightly
 - Level 1 after 2 months (NightlyLevel1)
 - Level 2 after 4 months (NightlyLevel2)
 - Level 3 after 6 months (NightlyLevel3)
- Continuous
 - Level 2 after 1 week (ContinuousLevel2)
 - Level 4 after 1 month (ContinuousLevel4)
- Experimental
 - Level 4 after 1 week - No files saved (ExperimentalLevel4)

To enable the default Archivers, uncomment this line(broken for clarity):

```
<!--
  <Property name="ArchiverList">NightlyLevel1,NightlyLevel2,NightlyLevel3,
    ContinuousLevel2,ContinuousLevel4,ExperimentalLevel4</Property>
-->
```

The definition of the default Archivers are:

```
<!-- Nightly archivers -->
<Property name="Archiver.NightlyLevel1.AgeInDays">60</Property>
<Property name="Archiver.NightlyLevel1.ArchiveLevel">1</Property>
<Property name="Archiver.NightlyLevel1.MatchTrack">Nightly</Property>

<Property name="Archiver.NightlyLevel2.AgeInDays">120</Property>
<Property name="Archiver.NightlyLevel2.ArchiveLevel">2</Property>
```

```

<Property name="Archiver.NightlyLevel2.MatchTrack">Nightly</Property>

<Property name="Archiver.NightlyLevel3.AgeInDays">180</Property>
<Property name="Archiver.NightlyLevel3.ArchiveLevel">3</Property>
<Property name="Archiver.NightlyLevel3.MatchTrack">Nightly</Property>

<!-- Continuous archivers -->
<Property name="Archiver.ContinuousLevel2.AgeInDays">7</Property>
<Property name="Archiver.ContinuousLevel2.ArchiveLevel">2</Property>
<Property name="Archiver.ContinuousLevel2.MatchTrack">Continuous</Property>

<Property name="Archiver.ContinuousLevel4.AgeInDays">30</Property>
<Property name="Archiver.ContinuousLevel4.ArchiveLevel">4</Property>
<Property name="Archiver.ContinuousLevel4.MatchTrack">Continuous</Property>

<!-- Experimental archiver -->
<Property name="Archiver.ExperimentalLevel4.AgeInDays">7</Property>
<Property name="Archiver.ExperimentalLevel4.ArchiveLevel">4</Property>
<Property name="Archiver.ExperimentalLevel4.MatchTrack">Experimental</Property>
<Property name="Archiver.ExperimentalLevel4.WriteArchive">false</Property>

```

2.8 Upgrade from pre-0.4

If you have been building Dart from source and would like to migrate to the first stable release (0.5), you must archive your Projects and re-create them. The basic structure of the Dart database has changed significantly, requiring a complete Archive and reloading of you data. The basic steps are:

1. Download the 0.4 and 0.5 releases of Dart from <http://na-mic.org/Wiki/index.php/Dart2Summary>.
2. Stop your Dart server

```
java -jar DartClient.jar --shutdown TestProject
```

3. Extract the 0.4 jar file, and Archive your Projects

```
java -jar DartServer.jar --archive DartServer TestProject
```

Note: Replace DartServer and TestProject with you installation specific directories.

4. Check the status of the Archive project by looking at the Project dashboard. When the Archive is complete, no data should remain on the Dashboard. The Server will not start up the HTTP server until all the data is Archived. In addition, log messages should be generated indicating the status of the Archive process.
5. Shutdown the Server.

```
java -jar DartClient.jar --shutdown TestProject
```


6. Rename Project directories to OldTestProject. Inside each Project directory is a sub-directory called Archive that contains all of the data from the Project that must be re-submitted to the 0.5 version.
7. Extract the 0.5 jar file.
8. Recreate the Projects. DartServers do not need to be recreated. Be sure to use the 0.5 jar.
9. Restart your Server
10. Resubmit old data for each Archived project.

```
java -jar DartClient.jar TestProject OldTestProject/Archive/Working/*
```

After all the Archived Submissions are digested you are finished!

2.9 Using Apache to proxy requests to Dart

While Dart ships with its own internal HTTP server and servlet container (thanks to jetty, <http://jetty.mortbay.org/>), Dart can also be run from behind any web server that will proxy and rewrite urls. Running Dart from behind another web server allows a site to provide a single access point to their website and web applications. It also allows for a limited type load balancing, where the Dart server can be a different machine than the public web server.

In this section we describe how to configure Apache to proxy requests to a Dart server. Apache must be configured to proxy requests and to rewrite urls. Several Apache modules must be enable in the Apache web server if they are not already built in. Add these lines to the Apache configuration file (http.conf):

```
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_connect_module modules/mod_proxy_connect.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule rewrite_module modules/mod_rewrite.so
```

These modules may need to be activated by added the lines

```
AddModule mod_rewrite.c
AddModule mod_proxy.c
```

Next, you need to establish the proxy. In this example, we assume the Dart server is on the same machine as the Apache server and we will map the Dart server to Apache's url space under /Dart/.

```
ProxyPass /Dart/ http://localhost:8081/
ProxyPassReverse /Dart/ http://localhost:8081/
```

Finally, we need to rewrite some urls. Each Dart project has a set of resources (icons, images, and style sheets) that are referenced by the dynamically generated web pages. The urls to these resources need to be redirected to the Dart server. This requires the activation of the rewrite engine

```
RewriteEngine on
```

and the establishing a rewrite rule for *each* Dart project. If your Dart project is named `TestProject`, then the rewrite rule will be

```
RewriteRule ^/TestProject/(.*)$ http://localhost:8081/TestProject/$1 [P]
```

If the Dart server is hosted on a different machine than the Apache web server, then you'll need to replace the use of `localhost` with the url to your Dart server.

Tool Integration

3.1 Cruise Control

Cruise Control (<http://cruisecontrol.sourceforge.net/>) is a build/test control project aimed to integrate with Java projects using Ant. Built into Dart are the capabilities to parse Cruise Control XML log files. Dart itself is tested using Cruise Control. This section is written for Cruise Control version 2.3.

3.1.1 Dart Integration

For Java, Cruise Control functions as a Dart client. To enable a Cruise Control project to submit to a Dart Server, first set up the Dart Server as described in Section 2.1. The instructions below assume you have installed Cruise Control, have configured Cruise Control according to the quick start guide at <http://cruisecontrol.sourceforge.net/gettingstarted.html>, and are running a Dart Project named MyProject running on a server called MyServer. This tutorial will cover Continuous builds and Nightly builds.

There are only a few extra requirements when submitting a standard Cruise Control log to a Dart Server.

3.1.2 Configure the Cruise Control Project

In this example, we will have two Cruise Control Projects: a Nightly and Continuous build. It is best to work through the Cruise Control quick start guide to ensure the builds function correctly. The Nightly build configuration should look something like this:

```
<!-- Nightly build at 10pm local time, see the schedule/ant tag, time attribute -->
<!-- Should update code from the last nightly build -->
<project name="Nightly">
  <dateformat format="yyyy-MM-dd'T'HH:mm:ss.SZ"/>
  <schedule interval="300">
    <ant buildfile="build-MyProject.xml"
      target="build"
      uselogger="true"
      usedebug="false"
      time="2200"
    />
  </schedule>
```

This Project will perform a build at 10:00pm local time each night, capturing the code changes from the previous 24 hours.

For the Continuous Project, the `config.xml` file should contain tags similar to this:

```
<project name="Continuous" buildafterfailed="false">
  <dateformat format="yyyy-MM-dd'T'HH:mm:ss.SZ"/>
```

The `<dateformat>` tag specifies that Cruise Control format dates according to UTC, a uniform time format. This allows Dart to easily parse dates. The “buildafterfailed” attribute instructs the Continuous project to build only if files have been updated.

3.1.2.1 Describe the Submission

Dart requires three items of information to properly categorize a Submission, the BuildName, Site and Track. Edit a file called `BuildInfoNightly.xml`. This file should look like this:

```
<?xml version="1.0" encoding="utf-8"?>
<BuildInfo>
  <BuildName>Linux-JDK-1.5</BuildName>
  <Site>MyClient</Site>
  <Track>Nightly</Track>
</BuildInfo>
```

The tags in `BuildInfoNightly.xml` instruct the Dart Server that this Submission is for the Nightly track, from a Site called MyClient and is a Linux build using JDK 1.5. For the continuous build, create a file called `BuildInfoContinuous.xml` which contains almost the same information:

```
<?xml version="1.0" encoding="utf-8"?>
<BuildInfo>
  <BuildName>Linux-JDK-1.5</BuildName>
  <Site>MyClient</Site>
  <Track>Continuous</Track>
</BuildInfo>
```

In this case, the Submission will go in the Continuous Track. It is also possible to have Ant configure this file automatically. Create a template XML file called `BuildInfoTemplate.xml` with these contents:

```
<?xml version="1.0" encoding="utf-8"?>
<BuildInfo>
  <BuildName>@Dart.BuildName@</BuildName>
  <Site>@Dart.Site@</Site>
  <Track>@Dart.Track@</Track>
</BuildInfo>
```

Ant can configure this file during a copy operation using the following Task definition and Properties:

```

<!-- Access environment variables -->
<property environment="env"/>

<!-- Get the HOSTNAME in an os independant way -->
<!-- Under Windows, COMPUTERNAME is set, for Unix-like systems
      HOSTNAME is used -->
<property name="env.HOSTNAME" value="\${env.COMPUTERNAME}"/>

<!-- BuildName is composed of a number of Ant and Java properties -->
<property name="Dart.BuildName"
      value="\${os.name}-\${os.arch}-\${os.version}-JDK-\${java.version}"/>
<property name="Dart.Site" value="\${env.HOSTNAME}"/>

<target name="configure.dart">
  <!-- Create the BuildName info, if necessary -->
  <filter token="Dart.BuildName" value="\${Dart.BuildName}"/>
  <filter token="Dart.Site" value="\${Dart.Site}"/>

  <!-- Filter tokens while copying the file -->
  <filter token="Dart.Track" value="Continuous"/>
  <copy file="BuildNameTemplate.xml" tofile="BuildNameContinuous.xml" filtering="true"/>

  <!-- Now do the Nightly version -->
  <filter token="Dart.Track" value="Nightly"/>
  <copy file="BuildNameTemplate.xml" tofile="BuildNameNightly.xml" filtering="true"/>
</target>

```

In this manner, each Client submitting to the Dart Dashboard will be self-configuring. The Ant task should produce the `BuildNameNightly.xml` and `BuildNameContinuous.xml` exactly as above. These two files will be merged into the Cruise Control log automatically.

3.1.2.2 Merge Build Name information

The Cruise Control getting started directions show how to merge JUnit logs into the Cruise Control log. The same mechanism is used to merge the Build Name information from the previous section into the final log. In your Cruise Control `config.xml` file, add or modify this section:

```

<!-- directory to write build logs to -->
<log logdir="logs/MyProject">
  <merge dir="checkout/MyProject/build/junit-reports/" />
  <merge file="BuildNameContinuous.xml" />
</log>

```

This will merge the `BuildNameContinuous.xml` file into the resultant log. If this build loop is for a Nightly build, substitute `BuildNameNightly.xml` instead.

3.1.2.3 Submit the log

To submit the Cruise Control log, we use AntPublisher. The getting started page recommends that you write a short Ant script to drive the build process. To submit the log after a build, edit this script (assume it's named `build-MyProject.xml`) adding a new target:

```
<target name="publish">
  <!-- Optional call to set the HTTP proxy if behind a firewall -->
  <setproxy proxyhost="proxy.host.com" proxyport="8080"/>

  <java classname="dart.DartClient">
    <classpath>
      <pathelement location="DartClient.jar"/>
    </classpath>
    <arg value="--server"/>
    <arg value="MyServer"/>
    <arg value="MyProject"/>
    <arg value="\${logdir}/${logfile}"/>
  </java>
</target>
```

`DartClient.jar` is distributed as part of Dart, and contains a minimal Client suitable for submitting Cruise Control logs. Once this Ant target is established, edit `config.xml` to instruct Cruise Control to publish the log via the publish target:

```
<publishers>
  <antpublisher buildfile="build-MyProject.xml" target="publish">
  </antpublisher>
</publishers>
```

Now Cruise Control will submit the log as part of the Publish step in the build loop.

3.1.3 Testing Dart with Cruise Control

The steps required to test Dart using Cruise Control are:

- Install Cruise Control from <http://cruisecontrol.sourceforge.net/>
- Check out Dart (be sure svn is in your path):

```
svn co http://svn.na-mic.org:8000/svn/Dart
```

- Run Cruise Control

```
/path/cruisecontrol-2.3.0.1/main/bin/cruisecontrol.sh
```

The default Cruise Control configuration (see config.xml) will run a Continuous build every 5 minutes, only building if changes have occurred and a nightly build each night at 10:00pm local time (22:00 on the 24 hour clock). Though the default settings should work in most cases, they may be over ridden using a properties file "build.properties". The properties honored in this file are:

- proxyhost : Name of the proxy to use for XML/HTTP
- proxyport : Port number on the proxy
- Dart.BuildName : Name of this build
- Dart.Site : Build site

Dart.BuildName is taken from the environment variable `COMPUTERNAME` on Windows and `HOSTNAME` on Unix-based systems. If an HTTP proxy is required, that setting is taken from the `HTTP_PROXY` and `HTTP_PROXY_PORT` on both Unix-based and Windows systems. The properties file takes precedence over the environment variables.

3.1.4 ctest Integration

The latest CVS version of ctest (<http://www.cmake.org/>) can be used to submit to a Dart Server. Instructions may be found at <http://na-mic.org/Wiki/index.php/User:Andy>. The settings needed in the project's CMake file are:

```
ENABLE_TESTING()
INCLUDE(Dart)

SET (DROP_METHOD "xmlrpc")
SET (DROP_SITE "http://myserver.org:8081")
SET (DROP_LOCATION "TestProject")
SET (COMPRESS_SUBMISSION ON)

ADD_TEST(name executable arg1 arg2 ...) # see CMake documentation
```

In DartConfiguration.tcl, this translates to:

```
DropSite: http://myserver.org:8081/
DropLocation: TestProject
DropMethod: xmlrpc
```


Development

4.1 Requirements

To work on Dart, you will need:

- Subversion (<http://subversion.tigris.org/>). Dart source code is maintained in a Subversion repository.
- Java SDK (<http://java.sun.com>). Version 1.4.2 or later is needed.
- Apache Ant (<http://ant.apache.org/>), version 1.6.2 or greater. This is a build system, similar in concept to Unix Makefiles.
- JUnit (<http://www.junit.org/>). Java unit testing framework. This is used to define and run regression tests on the Dart source. The JUnit jar file is included in the checkout. Drop the `junit.jar` file in `ant/lib` directory to enable JUnit to run as an ant task.
- The Dart source (see below).

The other packages required by Dart, such as Quartz and Jaxor, are available as part of the Dart source. You do not need to obtain these separately.

4.2 Obtaining the source

Obtain a copy of the source code by checking it out of the repository:

```
cd MySrc
svn co http://svn.na-mic.org:8000/svn/Dart
```

This will create a directory `MySrc/Dart` containing the current Dart source.

If you have a HTTP proxy server, you will need to specify the variables `http-proxy-exceptions`, `http-proxy-host` and `http-proxyt-port` in your `/.subversion/servers` (Unix) or `c:/Documents and Settings/User/Application Data/Subversion/servers` (Windows) file. Refer to the Subversion documentation for more details.

4.3 Build the source

The most straight forward method of building is

```
cd MySrc/Dart
ant all
```

basic steps are

```
cd MySrc/Dart
ant wrap
ant compile
ant jar
ant test
```

Each of “wrap”, “compile”, “jar” and “test” are compile targets, similar to Makefile targets. The full list is:

wrap Generate the Jaxor wrapping code. This generates Java objects to wrap the SQL queries defined in `Source/Wrap`. The wrapping process can be time consuming, and so is not run automatically for every compile. Wrap must be run when any of the Jaxor sources changes.

compile Compile the `.java` files to `.class` files. This is the default target.

jar Generate `DartServer.jar` containing the compiled Dart code.

test Run regression tests, with summary output.

testverbose Run regression tests with verbose output.

clean Clean the `.class` files.

fullclean Clean the `.class` files and the `.java` files generated by “wrap” above.

doc Runs JavaDoc to generate the API documentation into `Documentation/api`.

all Does a clean compile of Dart, runs the tests and builds the jar file.

4.4 Troubleshooting

- ‘Unexpected element “setproxy” ’
 - ▷ You need a newer version of Apache Ant.
- ‘package dart.server.wrap does not exist’, while compiling.
 - ▷ You didn’t run “wrap” before “compile”.
- ‘org.apache.velocity.runtime.exception.ReferenceException: reference : template = FinderImpl.vm [line 45,column 36] : \${primaryKeyQuery.getMethodName()} is not a valid reference’, while wrapping.
 - ▷ The wrapping process did not execute correctly. This could be due to clock skew on NFS mounted file systems, which incorrectly causes some rules to not fire.

Dart Requirements and Design

This chapter describes the requirements and design criteria for the next version of Dart.

5.1 Dart Statement of Purpose

Dart shall aggregate data across many independent distributed build and test hosts, summarizing the software quality aspects of the project in a concise and informative fashion cross-sectionally and longitudinally.

5.2 User Requirements

1. A single server instance shall process multiple projects, with simple, flexible configuration and management.
2. Presentation of results shall be configurable, allowing results to persist on the dashboard for different periods. For instance, coverage information is time consuming to produce but slowly changing and ought to persist for more than one day.
3. Dashboards may be aggregated into Meta-Dashboards. For instance, Slicer depends on VTK, ITK, gsl and Tcl/Tk. The Slicer Meta-Dashboard shall present summary information from these dependencies.
4. Dart shall support submission authentication and selectively reject or expire unauthenticated submissions.
5. Dart shall provide resource management tools for disk space, bandwidth and processing time allowing both Clients and Servers to efficiently manage resources.

5.3 Design Requirements

Basic

1. The server shall contain all components required and shall not require any external packages, nor operating system applications. The server shall run as a daemon and shall include these components:
 - (a) Scheduler: Dart shall include an internal scheduling system for routine systems tasks, *etc.* . .

- (b) RDBMS: Dart shall include an embedded database to handle small Projects.
 - (c) Web Server: Dart shall include an embedded web server to publish dashboard pages.
 - (d) Web Services: Dart shall communicate using an established protocol for web services, allowing Results submissions and query of Project status from remote, homogeneous clients.
2. The server shall be extensible with user supplied components, including:
 - (a) RDBMS: Dart shall use JDBC compliant drivers for all DB access allowing different database systems such as MySQL, Postgres, Oracle, *etc.*
 - (b) Web Server: Apache and other web servers shall be capable of serving Dart generated pages.
 - (c) Web Services: Dart shall allow the ability to communicate using external web servers such as Apache, Tomcat, *etc.*
 - (d) Portal Server: If desired, a Portal server such as Jetspeed may be used to interface with Dart results database. This capability is currently unspecified.

Resource Management (Section 5.3.1)

1. Dart shall, as an option, maintain compressed XML files, using on-the-fly decompression. This will result in approximately 10:1 spacing savings for the XML.
2. Dart shall provide a policy mechanism to selectively delete or archive unnecessary Builds. An archived Build shall consume less than 10K of disk space by retaining only summary information.

Storage, Processing and Presentation Engines

1. Dart shall comprise three engines: Storage, Processing and Presentation.
2. The Storage engine shall accept submissions from clients parse the input and store results in a generic format with large data items stored in the file system, *i.e.* images, with numeric and shorter text information stored in a database system.
3. The Processing engine shall process and summarize the results organized by the Storage engine at regular intervals and upon user-defined event triggering processing actions.
4. The Presentation engine shall provide a customizable view of data: both “raw” data from the Storage engine, and from the Processing engine. In the first instance, the Presentation engine shall simply be HTML, potentially migrating to a Portal based server.

Customization

1. Dart shall provide an easy to modify template engine for summarizing results.
2. Dart shall provide a server side plug in mechanism allowing custom data aggregation and flexible reporting.
3. Dart shall make provide mechanisms for simple localization and internationalization, where appropriate.

5.3.1 Resource Management

5.3.1.1 Disk Space

The Insight toolkit is the largest Dart project to date. Currently, with compressed HTML files, one day consumes over 650M of disk space. This includes (from November 4, 2004):

- One Doxygen run (400K for XML, 13K for HTMLZ)
- One Master Update (13K for XML, 4K for HTMLZ)
- One Dashboard (12M for XML, 9K for HTMLZ)
- One BuildOverview (250K for XML, 8K for HTMLZ)
- One TestOverview (25M for XML, 2 x 29K for HTMLZ)
- One Coverage build (24M for XML, 21M for HTMLZ)
- Builds (average of 9M for XML and HTMLZ)
 - 52 Nightly Builds
 - 19 Continuous Builds
 - 22 Experimental Builds

Breaking down an example day, we have:

File	Size	Notes
Build.xml	140K	111 Warnings, average of 1.2K per warning
Configure.xml	0.8K	
Test.xml	4.3M	859 Tests, average of 5K per test
Update.xml	1.4K	1 Update, 1.4K per updated file
TestSummary.xml	215K	Total for 8 HTMLZ files
All.htmlz	86K	

The largest generator of data is test output. Errors/Warnings and Update information are rather verbose, capturing context information. In general, XML is verbose with low entropy. A 4.3M Test.xml file is 522K compressed with gzip.

5.3.1.2 Dashboard Generation Time

5.3.1.3 Bandwidth

5.3.2 Historical Data

Dart currently preserves data from previous days, it is not linked across temporal Builds on the same system. While simple, this restriction increases the difficulty of monitoring the quality of a project. To overcome this limitation, Dart shall link data in a temporal fashion.

5.3.3 Hierarchical Data

5.3.3.1 Dashboards

5.3.3.2 Tests

5.3.3.3 Builds

5.3.4 Persistence of Builds/Results

5.3.4.1 Stream Concept

5.3.5 Documentation

5.3.6 Submissions

5.3.6.1 Incremental Submission

5.3.6.2 Mechanisms

5.3.6.3 Authentication

5.3.7 Configuration

5.3.7.1 Initial Setup

5.3.7.2 Options

5.3.8 Customization

5.3.8.1 Dashboard presentation

5.3.9 Extensibility

Design

6.1 Server

The Dart server is implemented in Java. It is composed of several different services, outlined below.

6.2 DartServer

The DartServer is responsible for starting up the other services. Projects are created, configured, loaded and started by the DartServer class.

6.2.1 Command Manager

Commands to Dart are passed to the DartServer via XML-RPC. The DartServer starts up a the Apache XML-RPC server on the same port as the HTTP Server by default.

6.2.2 Scheduler

The Quartz enterprise Scheduler is initialized and passed to each Project. In turn, each Project adds Tasks to the Scheduler to be executed as needed.

6.2.3 HTTP Server

Jetty is used in an embedded mode to serve static content, and generate dynamic content.

6.3 Project

Each Project hosted on a DartServer is created by loading the Dart.xml file in the Project directory. A Project is composed of several components.

6.3.1 Database

The Database object coordinates all access to the underlying RDBMS. The Database provides Connections to other Project components as needed.

6.3.2 ResultServer

The ResultServer object is responsible for handling XML-RPC requests. During startup, a service is added to the DartServer SubmissionServer.

6.3.3 ServletManager

The ServletManager is responsible for creating the Project specific Servlets and adding them to the Server's HTTP Server. User Servlets may be added to the Project's Plugins directory.

Implementation Ideas

This section captures some implementation ideas.

7.1 Server

Language Of all cross-platform languages, Java provides the most robust set of libraries suitable for Dart. Java also allows simple distribution of compiled libraries, *i.e.* jar files, as plug-ins. Potentially, a client could query the server for a list of available plug-ins downloading and installing as needed.

RDBMS There are several embeddable Java RDBMS available, two of the more interesting projects are Cloudscape, recently released from IBM, and renamed Derby on the Apache site and Hypersonic SQL (HSQLDB) project hosted on SourceForge. Dart is envisioned to have a RDBMS holding summary data; embedding a database into the server should help to make it transparent and invisible to the casual Dart user. For more scalability, the backing store could be any RDBMS with a JDBC driver. MySQL and Postgres come to mind.

Transport Though over-designed and complex, SOAP has the elements need to transmit XML files to the server from the client. Specifically, SOAP with attachments could deliver chunks of compressed XML to the server via HTTP, since most (all?) firewalls allow HTTP traffic. SOAP could also be used for Dashboard to Dashboard (D2D?) communication and remote management and monitoring of Dart servers. XML-RPC is a much simpler API, and identically suitable. XML-RPC will be considered at the same level as SOAP. Another possible use is dissemination of plug-ins for clients. The Java Messaging Service (JMS) is another possibility. JMS gives great flexibility to transport mechanisms and can operate asynchronously.

Scheduling Quartz is an open source enterprise strength scheduling system for Java. Quartz will drive scheduled events such as Dashboard roll ups, DB tasks, and archiving/deletion of old results. Quartz will replace cron.

Template Engine There are several competing Template engines for Java. Velocity is an Apache sponsored project and has some great features including close integration with other Apache projects. FreeMarker is another engine that is more sophisticated than Velocity, but not as integrated. The Template engine will be the driver to produce HTML and other reports replacing XSLT.

Jakarta The Apache Jakarta project provides several packages of immediate use.

- Digester builds objects from XML, greatly simplifying configuration from XML files. Each object is constructed as needed and automatically configured.
- CLI should provide a great command line parsing interface.
- Commons eMail provides a simple java email client.
- ORO and RegExp, two regular expression packages.

Portal Though the current Dart HTML pages serve the purpose well, adding a portal on top would allow custom portlets to be developed for specific purposes. For instance, one portlet could be configured to show a particular build over the last several days, or perhaps graph the performance of a Test or Result through time across several architectures. Dynamic generation of all the Dart results places undo burden on the server, where a Portal could dynamically generate limited data in an efficient manner. One Portal project that is interesting is Jetspeed 2, an Apache sponsored project.

Portals do add administrative overhead. It is preferable to have the ability to use Dart without a Portal, but easily being able to add the increased utility if desired.

7.2 Client

External Packages

8.1 Packages

Dart is built upon many Open Source packages. Each of these packages has different licenses. To comply with the licenses of each of these, we have listed the packages, their licenses and copyrights in this chapter.

Apache v1.1 Apache XML-RPC, Apache CLI

Apache Version 2.0 Bean Utilities, Derby, Collections, DBCP, Digester, Pool, VFS, Jetty

BSD License Jaxor

Common Public License, v1.0 JUnit

BSD-Like license Quartz

Freemarker License Freemarker

GNU Lesser General Public License JFreeChart (<http://www.jfree.org/jfreechart/>)

8.2 Apache License, Version 1.1

The Apache Software License, Version 1.1

Copyright (c) 2001 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions

are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment:
 "This product includes software developed by the
 Apache Software Foundation (<http://www.apache.org/>)."
Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.
4. The names "XML-RPC" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.
5. Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

8.3 Apache License, Version 2.0

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain

separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the

origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.

8.4 BSD License

Copyright (c) <YEAR>, <OWNER>
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:

- * Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.

- * Redistributions in binary form must reproduce the above
* copyright notice, this list of conditions and the following
* disclaimer in the documentation and/or other materials provided
* with the distribution.

- * Neither the name of the <ORGANIZATION> nor the names of its
* contributors may be used to endorse or promote products derived
* from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8.5 Freemarker License

Copyright (c) 2003 The Visigoth Software Society. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. The end-user documentation included with the redistribution, if any, must include the following acknowledgment:
"This product includes software developed by the
Visigoth Software Society (<http://www.visigoths.org/>)."
Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.
3. Neither the name "FreeMarker", "Visigoth", nor any of the names of the project contributors may be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact visigoths@visigoths.org.
4. Products derived from this software may not be called "FreeMarker" or "Visigoth" nor may "FreeMarker" or "Visigoth" appear in their names without prior written permission of the Visigoth Software Society.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE VISIGOTH SOFTWARE SOCIETY OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8.6 Quartz License

Copyright James House (c) 2001-2004

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8.7 GNU Lesser General Public License

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether

this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or

other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.

b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the

ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the

Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy

from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are

prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add

an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH

DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or  
modify it under the terms of the GNU Lesser General Public  
License as published by the Free Software Foundation; either  
version 2.1 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU  
Lesser General Public License for more details.
```

```
You should have received a copy of the GNU Lesser General Public  
License along with this library; if not, write to the Free Software  
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the  
library 'Frob' (a library for tweaking knobs) written by James Random Hacker.
```

```
<signature of Ty Coon>, 1 April 1990  
Ty Coon, President of Vice
```

That's all there is to it!

