

# **Wide-field Infrared Survey Explorer (WISE)**

## **Science Operations and Data Analysis Software Management Plan**

**Version 0.1**

**20-Mar-06**

**Prepared by: Tim Conrow**



**Infrared Processing and Analysis Center  
California Institute of Technology**

**WSDC D-XXXX**

**Approved By:**

---

Ned Wright, WISE Principal Investigator

---

Ingolf Heinrichsen, WISE Mission Operations Center Manager

---

Roc Cutri, WISE Science Data Center Manager

---

Tim Conrow, WISE Science Data Center Lead Engineer

**Revision History**

<b>Date</b>	<b>Version</b>	<b>Author</b>	<b>Description</b>
March 20, 2006	0.1	Tim Conrow	Initial Draft

## Table of Contents

<b>1</b>	<b>INTRODUCTION.....</b>	<b>6</b>
1.1	Document Scope.....	6
<b>2</b>	<b>PROBLEM TRACKING .....</b>	<b>6</b>
2.1	Introduction .....	6
2.2	Reference PR System; Bugzilla .....	7
2.3	PR Resolution.....	7
2.4	Which Problems Should Be Tracked? .....	7
2.5	PR Creation Procedure .....	8
<b>3</b>	<b>REVISION CONTROL .....</b>	<b>8</b>
3.1	Introduction .....	8
3.2	Reference Revision Control Tool: Subversion.....	8
3.3	Revision Number and Release Tags .....	8
3.4	Branches.....	9
3.5	Repository Browsing.....	9
3.6	What Is Revision Controlled? .....	9
3.7	Revision Control Instructions .....	10
<b>4</b>	<b>BUILD TOOLS.....</b>	<b>10</b>
4.1	Introduction .....	10
4.2	Reference Build Tool: make .....	10
4.3	Use of the Build Tool.....	11
4.4	Makefile Templates.....	11
<b>5</b>	<b>DELIVERY .....</b>	<b>12</b>

<b>5.1</b>	<b>Introduction .....</b>	<b>12</b>
<b>5.2</b>	<b>Who Delivers? .....</b>	<b>12</b>
<b>5.3</b>	<b>Delivery Target Isolation.....</b>	<b>12</b>
5.3.1	The Ops Target.....	13
5.3.2	The Int Target.....	13
5.3.3	The Dev Target.....	13
5.3.4	Other Targets.....	13
<b>6</b>	<b>OPERATIONAL SYSTEM CONFIGURATION CONTROL.....</b>	<b>13</b>
<b>6.1</b>	<b>Introduction .....</b>	<b>13</b>
<b>6.2</b>	<b>The CCB.....</b>	<b>14</b>
<b>6.3</b>	<b>Configuration Validation.....</b>	<b>14</b>
<b>6.4</b>	<b>WSDC Operational System Build Procedure .....</b>	<b>15</b>
<b>7</b>	<b>SOFTWARE BACKUP .....</b>	<b>15</b>
<b>7.1</b>	<b>Introduction .....</b>	<b>15</b>
<b>7.2</b>	<b>Daily Copying of the SVN Repository.....</b>	<b>16</b>
<b>7.3</b>	<b>Daily Partial And Weekly Full Backups To Tape .....</b>	<b>16</b>

## **1 INTRODUCTION**

### **1.1 Document Scope**

This document is about software management, i.e. the handling and disposition of source code and any binary files derived from them. It will cover the following topics:

- report tracking;
- software and parameter revision control;
- software delivery;
- operational system configuration control, including testing;
- software backup.

### **1.2 What is Not Covered**

These topics, though important, are beyond the scope of this document:

- computing hardware and infrastructure configuration control;
- document revision control;
- reference data (e.g. calibration data) version control.

### **1.3 Applicable Documents**

This plan conforms to the specifications in the following project documents:

- TBD

### **1.4 Acronyms**

The acronyms used in this document are normally defined with their first use.

## **2 PROBLEM TRACKING**

### **2.1 Introduction**

In software management, a problem report (PR) tracking system provides traceability and documentation of the full life cycle of a problem report, from initial report to resolution. It is used for reporting and documenting the characteristics and resolution of:

- bugs – software doesn't perform as documented.
- deficiencies – the software as designed is inadequate.
- change requests – alter designed behavior.
- feature requests.

## 2.2 Reference PR System; Bugzilla

As a reference PR system, we will assume the WISE WSDC uses the open source Bugzilla suite of tools (<http://www.bugzilla.org/>) for problem tracking. Whatever tool is selected will have the key attributes discussed below. Bugzilla is used as a reference system – i.e. an existence proof – to show that systems meeting the WSDC's needs are readily available without cost. The WSDC may actually use Bugzilla, or may select another compliant system.

Bugzilla is a collection of tools that provide access to a problem report (PR) database. These tools are most conveniently utilized through the Bugzilla web interface web. An example of the web interface can be found at <http://landfill.bugzilla.org/bugzilla-2.20-branch/>.

Via the web interface, users can

- be authenticated,
- create PRs,
- update previous PRs,
- query the PR database,
- create summaries of all or a cross-section of PRs, and
- trace the history of all action taken on a PR.

When submitting a PR, a submitter categorizes the PR by software category and supplies detailed information, possibly including sample code, data files, environment, etc. Once a PR is completed a predetermined stuckee for each category is automatically notified by email. Either the submitter or the stuckee may subsequently modify the PR.

The PR database may be searched using any of a dizzying array of conditions and filters.

## 2.3 PR Resolution

All PR's should be resolved, which may mean the bug was fixed, or that software was used incorrectly (user error), or a feature request was implemented or denied, etc. After resolution, the PR should be formally closed in Bugzilla (i.e. the PR state is changed) and the stuckee should document how it was resolved and why.

## 2.4 Which Problems Should Be Tracked?

Any problem with operational software should be entered into Bugzilla if the resolution will take longer than a few days, or if the user explicitly wishes to create a problem record, e.g. to document a series of small but ongoing problems with a design.

During development, either prior to the existence of an operational system or in parallel with it, users and developers may enter PRs whenever tracking an issue would be useful. The only limitation is to avoid making PRs so trivial that it is merely a waste of time dealing with them (e.g. It's probably excessive to enter document typos as a PR).

## **2.5 PR Creation Procedure**

Here are the steps to creating a PR:

- i. If desired, consult with cognizant parties to determine if a PR is appropriate, and, if so, what information to supply.
- ii. Browse to the WSDC PR web site.
- iii. Fill in the form. Be expansive in supplying information. Submit the PR.
- iv. Make any updates to the PR as new information becomes available.
- v. Make noise if a resolution is not forthcoming.
- vi. Concur or demur with the proposed resolution.

## **3 REVISION CONTROL**

### **3.1 Introduction**

Revision control refers to the ability to label, document, track and recall prior states of a file or related group of files. Any file can be revision controlled, but only files containing source code or parameters are considered in this document.

### **3.2 Reference Revision Control Tool: Subversion**

For reference, we will assume the WISE WSDC uses Subversion (<http://subversion.tigris.org>) – SVN for short -- for revision control. SVN maintains a central repository into which users check files in and out. For each check-in SVN will record all changes from the previous state of the file, as well as meta-data about the check-in, such as the date and time and any comments from the user. Whatever tool is selected will have the key attributes discussed below. SVN is used as a reference system – i.e. an existence proof – to show that systems meeting the WSDC's needs are readily available without cost. The WSDC may actually use SVNa, or may select another compliant system.

### **3.3 Revision Number and Release Tags**

SVN maintains a revision number that uniquely identifies the checked-in state of each file. This number is retrievable from the repository as part of a file's check-in meta-data, and it can be embedded in any plaintext file by including a special text string that SVN will replace with the current revision number upon check-in.



The revision number is strictly file-local; it does not by itself say anything about the developmental state of a file, e.g. whether it's ready for release. Information about developmental state is explicitly supplied as part of the configuration control process (see below) using release tags.

Tags give a user-supplied name to the state of the entire repository or any part of it and can be applied by a user any time a particular state of (part of) the repository is of interest, such as prior to a release. Thus, a given tag will allow a user to retrieve a particular point in the revision history of any group of files so tagged. Each tagged file will have its own revision number since it has its own, separate history. Thus a tag essentially marks a cross-section of revision numbers within the repository. As an example, prior to a release, the text for a tag will be composed (e.g. "wsdc-release-1\_0\_2") and applied to the appropriate source code within the repository. Afterwards, all code associated with this release can be retrieved and examined at any time in the future.

### **3.4 Branches**

Changes to source code which are too speculative, broad, or otherwise risky to be on the main development track (called the trunk in SVN) can be split out onto a side track (called a branch) and developed in parallel. If/when a side track reaches sufficient maturity (as decided by the CCB) it can be merged back onto the trunk. As an example, after a particular release has become operational, further development for the next release will occur on the trunk. But if bugs are subsequently discovered in the ops system, a branch can be split off to work on the bug fix by creating a branch at the release tag for the prior release. That fix can then be applied to the operational system and separately merged into the trunk so it is accounted for in future development.

### **3.5 Repository Browsing**

The SVN repository can be browsed or searched via a web CGI program. Users may access the complete revision history (any state of any file which is revision controlled), including meta-data, about any checked-in file. Repository browsing will be a convenient way to access all WSDC software and parameters without having to navigate the WSDC filesystem.

### **3.6 What Is Revision Controlled?**

All locally developed source code necessary for operations or intended to be used by anyone other than the author should be revision controlled. Also, a user may place any plaintext file desired under revision control for any reason.

Source code developed elsewhere should not be revision controlled as long as it is not modified locally. Externally developed source code that is locally modified might need

to be revision controlled. This is an issue, which the CCB (see below) should address on a case-by-case basis.

### **3.7 Revision Control Instructions**

Here are some instructions a user should follow in maintaining a file under revision control:

- i. Maintain a repository-friendly directory hierarchy. Details elsewhere.
- ii. Put the revision ID replacement tag in all files so the revision ID string is present for review in all code and binary objects.
- iii. Add new software to the SVN repository.
- iv. Check changes in at least once per week as well as when necessary to prepare for a new delivery (see configuration control below).
- v. Periodically do a checkout of all code to test that the repository and a user's work areas are in sync.
- vi. Add tags whenever they might be useful for future reference (e.g. “before-big-bug-fix” etc.), but always prior to deliveries and releases in accordance with configuration control rules.

## **4 BUILD TOOLS**

### **4.1 Introduction**

Build tools are the command line tools used to take source code and produce an executable application or a library and to deliver it to a specified delivery target. All developers at the WSDC will use a common set of build tools, making the build and delivery process uniform.

### **4.2 Reference Build Tool: make**

For reference, we will assume building and delivery will be done at the WSDC using the common command line build tool “make.” Specifically, the implementation provided by GNU make is assumed. Other make implementations may be incompatible. Whatever tool is selected will have the key attributes discussed below. Make is used as a reference system – i.e. an existence proof – to show that systems meeting the WSDC's needs are readily available without cost. The WSDC may actually use make, or may select another compliant system.

### **4.3 Use of the Build Tool**

The operation of the make tool is controlled by a “Makefile” which contains a specification of targets (things to be built), the dependencies (prerequisites) for that target, and rules used to resolve (provide) the dependencies. Each code unit, or closely related group of code units, will have its own Makefile occupying the same directory as the code.

Makefile can refer to other Makefiles and can thus cascade or chain, potentially building an entire system.

See below for how the build tool will be employed in the delivery process.

### **4.4 Makefile Templates**

Delivery target isolation (see below) and a uniform build and installation process are enforced through the mandatory use of Makefile templates by all developers. These templates are a small set of files of rules, macros and targets to be used by the make utility. All users include these templates in their Makefiles. These Makefiles, in turn, form the complete set of rules by which code is built and delivered when the make utility is invoked.

The make utility employed at the WSDC is GNU make. Making utilities with a different origin may or may not work.

Here are the main areas controlled by the Makefiles:

- which source files depend on which other source files, and the order in which they must be compiled;
- what external standard libraries are used and their locations;
- which internal libraries are linked to and their locations (usually the same delivery target);
- the set of compiler flags employed for compilation;
- what is delivered, and where.

Users may override some of the standard behavior, but such overrides are subject to review.

## **5 DELIVERY**

### **5.1 Introduction**

Delivery is the process of integrating developer – local copies of executable programs and associated files (libraries, include files, etc.) derived from configuration controlled source code, into a common area for use by testers or users. The common area may be an operational release, a public development area, or a private area for integrated testing (see Delivery Target Isolation below).

Delivery proceeds as follows:

- i. Stabilize – halt code development and document changes
- ii. Build – compile source to binary, if necessary
- iii. Unit test – test at the unit (component) level
- iv. Check-in – update the source code’s revision number
- v. Install – move the program to the delivery target
- vi. Validate – test the integrated code

Both the build and install steps are carried out through the use of standard tools based on the “make” (or the equivalent) utility and makefile templates (see Makefile Templates above). Unit testing is testing on small units of coding prior to delivery – e.g. checking if a change to a particular subroutine worked or not, as opposed to testing that subroutine as part of the fully integrated system. Validation is discussed below.

### **5.2 Who Delivers?**

Software authors are responsible for checking-in their files to the SVN (or equivalent) repository) prior to deliveries. Building and installation to test or development areas may be carried out by a designated build specialist, but can also be done by any developer. Building and delivery to integration and operational areas will be done by a build/delivery specialist in consultation with the CCB and all developers.

### **5.3 Delivery Target Isolation**

In order to test new, developing configurations, developers must be able to integrate their code with code from other developers to allow extensive use of all its features and interfaces without affecting an operational release. This calls for a delivery process that controls dependencies such that delivered programs depend only on libraries, data, parameters, and other programs that are part of the same delivery. The WSDC delivery process ensures this dependency isolation (see Makefile Templates below).

All deliveries specify a delivery target within which all dependencies will be resolved. Any number of separate targets may coexist. Here are the most commonly employed delivery targets and their uses:

### **5.3.1 The Ops Target**

Ops deliveries produce the operational system as used for all WSDC activities other than testing and development. Greatest care is taken in the validation step for ops deliveries and the CCB is always consulted.

There are actually two parallel ops deliveries; one accessible from development machines, and a separate one used by the operations servers. These deliveries are identical.

### **5.3.2 The Int Target**

Deliveries to the int target are made prior to a delivery to ops, for integrated testing of the next ops configuration. The int target thus has only a short lifespan. Use of the int target is controlled by the CCB.

### **5.3.3 The Dev Target**

The dev target is used for ongoing development of the next operational release. Only minimal CCB involvement is required for delivery to the dev area. It is, therefore, an unstable area usually containing the most up-to-date versions of all code.

### **5.3.4 Other Targets**

Other delivery targets can be created at will by any user for special purposes, such as testing new, highly experimental, features or large changes (refactoring, etc.)

## **6 OPERATIONAL SYSTEM CONFIGURATION CONTROL**

### **6.1 Introduction**

For the purposes of this document, the “WSDC system configuration” is the state of all source code (and any derived binary files) which implements the capabilities required at the WSDC. In this context, configuration control is the ability to precisely maintain or restore the behavior of the system until change is deliberately called for, and carefully controlling the change process.

Exerting this control requires these elements:

- a Change Control Board (CCB), which decides when the configuration may change;
- revision control (see above), to provide traceability of changes and thereby a means of uniquely identifying and, if necessary, recalling a particular state of the configuration;
- a delivery system, to orchestrate the changeover to a new configuration.

Revision control and delivery have already been addressed. The CCB and build validation are discussed below.

## **6.2 The CCB**

The CCB is responsible for deciding when the WSDC system configuration may change and establishing when the change should occur. The criteria for when it is appropriate to alter an existing configuration will be situational, but will include:

- importance (does it add or fix something important);
- scope (how much of the system is affected by the change);
- risk (what might break);
- nearness of critical mission events;
- quality (has the change been carefully considered and tested)
- cost (utilization of project resources).

In addition to controlling the configuration of the operational system, the CCB will also modify the change procedure (see below) as necessary, and decide if development of a proposed change should proceed on the main development track, or a separate branch (see revision control above).

## **6.3 Configuration Validation**

Newly installed configurations are validated through the use of a regression test baseline (RTB) – i.e. running a series of standardized tests on the new configuration and comparing the results with those of previous configurations. For ops deliveries, this comparison is considered by the CCB (see below) as part of the approval process.

The RTB is not a deep and thorough test of all aspects of a particular configuration, but confirms that a set of basic capabilities are present to confirm that the delivery was complete – i.e. it's the delivery process which is being tested, not the code itself. Thorough code testing is assumed to have occurred during the development process.

## 6.4 WSDC Operational System Build Procedure

Here are the steps to replacing an operational build with a new one.

- i. Halt new development.
- ii. Convene CCB to review changes from current operational state.
- iii. Make any approved hardware and infrastructure changes to the ops machines. Any such changes will already have been made and validated on development machines.
- iv. Deliver to integration area on ops machine.
- v. Validate int delivery (run RTBs)
- vi. Tag SVN repository with release candidate tag.
- vii. Get CCB final approval.
- viii. Suspend operations
- ix. Deliver to operational area from a fresh checkout using the release candidate tag. Deliver to both the ops and dev servers so the machines are kept in sync.
- x. Check dependencies. Check all scripts, binary programs, and libraries for illegal references outside the ops delivery area.
- xi. Validate ops delivery (run RTBs).
- xii. Tag SVN repository with final release tag.
- xiii. Restart operations.

## 7 SOFTWARE BACKUP

### 7.1 Introduction

The goal of the WSDC software backup process is to ensure the ability to reform an operational system and continue operations and development within TBD days of an ordinary outage (e.g. a disk failure) and within TBD weeks of a disaster (e.g. loss of the WSDC facility). Details of the recovery process are beyond the scope of this document.

## **7.2 Daily Copying of the SVN Repository**

The repository will be copied over the network to backup locations on disk both within the WSDC, and, via sftp, to a remote location. These network backups will occur automatically overnight.

Since all internal software necessary for operations is kept in the repository, this alone should be sufficient to rebuild all necessary executable programs. Calibration data and other reference data will need to be recovered from on- or off-line WSDC backup media.

## **7.3 Daily Partial And Weekly Full Backups To Tape**

As part of the WSDC facility backup system, all code and programs will be backed up along with all other files on a nightly basis. Off-site copies of tapes will be maintained.