ReST Installer User's Guide

An introduction to the ReST Installer.

Jeff M Larkin, Innovative Computing Laboratory, UT
<larkin@cs.utk.edu>
Eric T Meek, Innovative Computing Laboratory, UT
<meek@cs.utk.edu>

Revision History

Revision 1.0	3/11/2005
	Initial Document Release.

Abstract

This document is written as a basic introduction to using the ReST Installer. No prior knowledge of the ReST Installer is needed.

Table of Contents

ReST Installer Introduction 1
Basic Usage 2
Setup Locations 2
Setup Logical Groups 2
Setup Actions 4
Setup Authentication 4
The Installation 6
Advanced Topics 6
Setup File Substitutions 6
Configuring Command Options 6
Saving an Install Script 6
Examples 7

ReST Installer Introduction

For consumer software installation, Installation "Wizards" have become the de-facto standard; but in a distributed computing environment, such a luxury is almost nonexistent. Software installation in a distributed environment often require a user to login to numerous machines and perform roughly the same tasks on each. Further compounding the problem, most source distributed software installation requires both a lengthy and confusing configuration, compilation, and installation process. The ReST Installer provides a familiar, wizard-like interface to source and binary distributed software that can be used to install software in a heterogeneous environment easily and automatically. At the core of the ReST Installer is the ReST package model. By defining the structure of the ReSI software packages, the ReST package model provides software maintainers a way to define their software's requirements and options so that they can be presented to a user in a simplified manner, shifting the burden of software installation from the user to the software provider. Generic instructions on installing software with the ReST Installer are contained in the sections below. Developers interested in using ReST to distribute their software should instead read the *ReST Package Maker's Guide*.

Basic Usage

Setup Locations

The Setup Locations screen allows for choosing a one or more locations on which to install the package. The user must have a valid login for each machine listed and the machine must be reachable via secure shell (SSH). Locations can be inserted using a regular expressions to define multiple locations with similar names. For example, entering a location as someone@somewhere[0-2].example.com would create three locations of someone@somewhere0.example.com, someone@somewhere1.example.com, someone@somewhere2.example.com. If a group of machines will be used for the installation of several packages, the user may wish to write a script file containing the list of machines. An example of such a file exists later in this document.

Figure 1. Setup Locations

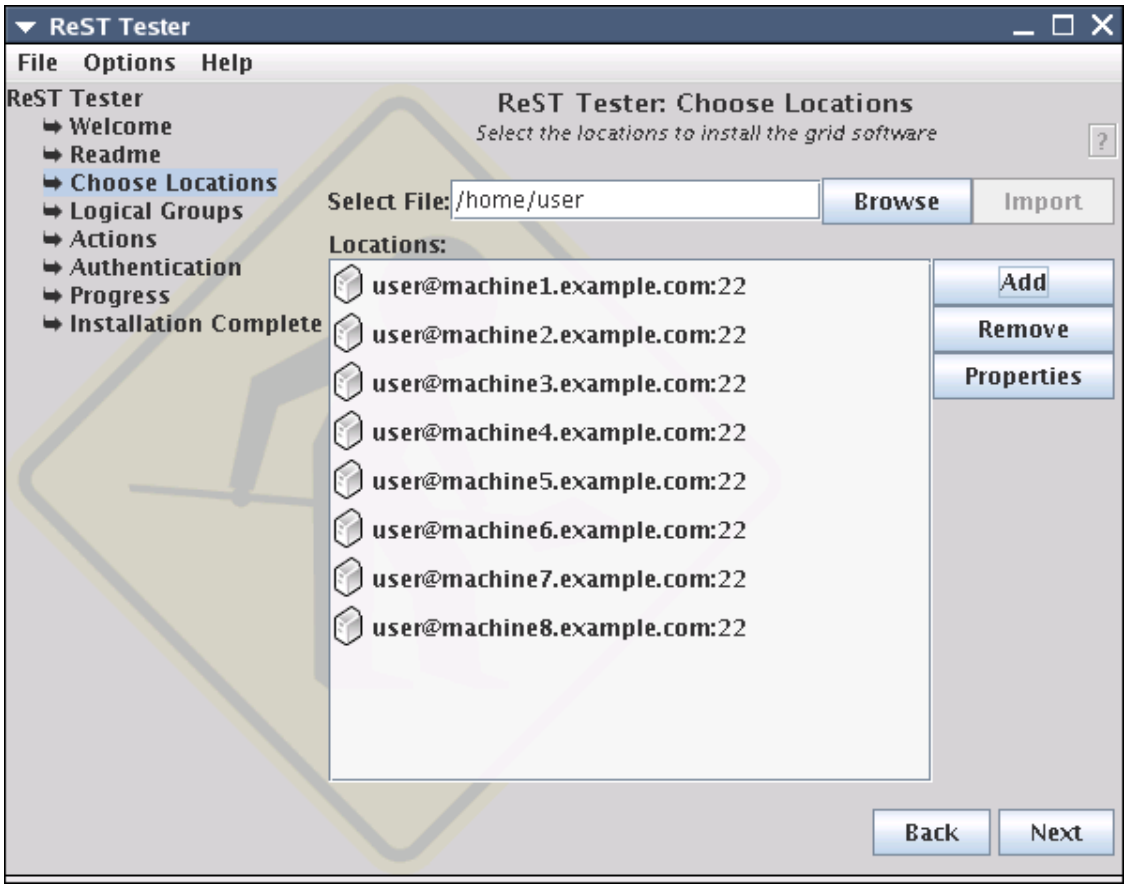

This screenshot shows the setup locations dialog of the ReST Installer. This user has setup 8 locations on which the software should be installed.

Setup Logical Groups

A logical group should consist of locations grouped by filesystems and binary compatibility. In other words, if five machines share a common filesystem and programs compiled on one machine may be run on all, then these machines should be placed in a logical group. By creating logical groups, the Rest In-

staller is able to reduce the network traffic, file space and time needed to install the package. Two important entities exist among logical groups: build masters and file masters. When installing a package, some parts of the package may only need to be executed on one machine for every machine in the group to benefit from that action. In this case the command will only be run on the build master of the group. Every group must have a build master. If no build master is explicitly selected by the user, one will be automatically chosen from the group. Once a build master is selected, a single asterisk will appear by its name. To explicitly choose a build master, select one machine in the group and press the *Set Buildmaster* button below the group list.

Multiple logical groups may share a common filesystem, making it possible to transfer files to one machine and have them available to machines within multiple logical groups. For this reason file master entities exist. A file master is an extension of a build master, so if a machine is a file master, it is also the build master of its logical group. To select a filemaster, choose a machine and press the *Set Filemaster* button below the groups list. This will cause a window to appear, which will include all of the defined groups. To select multiple groups within this window, hold the *ctrl* key while clicking on the group name. When you have selected the desired groups, press *Ok* . Selecting a file master will only affect the build master of the group to which the new file master belongs. All other groups will remain unchanged. When a machine has been selected as a file master, two asterisks will appear by its name.

By default a logical group is created for all machines unselected in the groups panel. This group is known as RSI_DEFAULT and may be treated as a normal logical group or treated as separate machines, each in their own logical group. By default the RSI_DEFAULT group is treated as a logical group. To treat each machine as an independent group, simply uncheck the *RSI_DEFAULT grouped* checkbox on the groups panel.
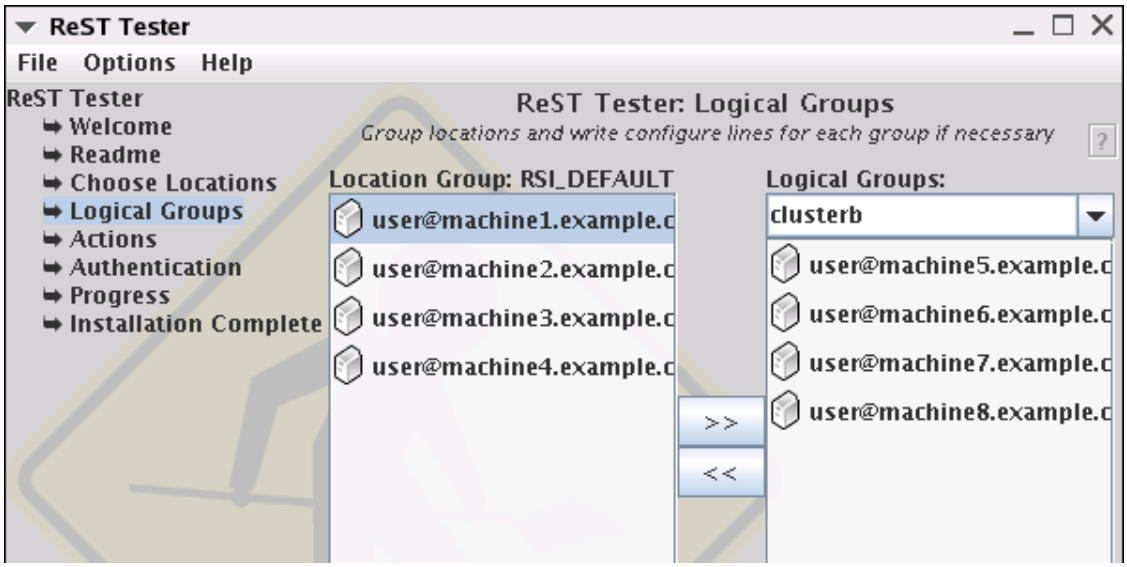
Figure 2. Setup Logical Groups