

RStudio Package Manager: Admin Guide

Version 1.0.6-6

Abstract

This guide will help an administrator install and configure RStudio Package Manager on a managed server. You will learn how to install the product on different operating systems, configure authentication, and monitor system resources.

Contents

1	Introduction	1
2	Get Started: Installation	1
2.1	Need Help?	1
2.2	Prerequisites	2
2.3	Installation	3
2.4	Initial Configuration	3
2.5	Licensing	3
2.6	Restart RStudio Package Manager	4
2.7	Visit the web UI	4
2.8	Try out the CLI	4
2.9	Next Steps	4
3	Get Started: Configuration	4
3.1	Serving CRAN Packages	5
3.2	Distributing Local Packages	5
3.3	Serving Local Packages from Git	6
3.4	Distributing Local Packages along with CRAN Packages	6
3.5	Supplementing CRAN with Bleeding Edge Packages from GitHub	7
3.6	Serving a Subset of Approved CRAN Packages	7
3.7	Serving a Subset of Approved CRAN Packages and Local Packages	9
4	Frequently Asked Questions	9
4.1	Setting Up RStudio Package Manager	9
4.2	Configuring RStudio Package Manager	10
4.3	Using RStudio Package Manager from R or RStudio	11
4.4	Controlling Access to Packages	11
4.5	Package Updates	11
4.6	Managing Change Control	11
4.7	RStudio Package Manager and Docker	12
5	Server Management	12
5.1	Stopping and Starting	12
5.2	Upgrading	14
5.3	Purging RStudio Package Manager	14
5.4	Backups	15
5.5	Privileged Ports	15
6	Licensing	15
6.1	Proxy Servers	16
6.2	Offline Activation	16
6.3	Licensing errors	17

6.4	Repository Count	17
6.5	Floating Licenses	17
7	Files and Directories	21
7.1	Changing Ownership	21
7.2	Program Files	21
7.3	Configuration	21
7.4	Server Log	21
7.5	Access Logs	21
7.6	Variable Data	22
7.7	Variable Data Classes	22
7.8	Destinations	23
8	AWS S3 Configuration	24
8.1	Destinations	24
8.2	AWS Documentation	24
8.3	Configuration Examples	24
9	Database	25
9.1	SQLite	26
9.2	PostgreSQL	26
9.3	Usage Data	27
10	Outbound Proxy	27
11	Running with a Proxy	28
11.1	Using a Proxy for TLS/SSL	28
11.2	Nginx Configuration	28
11.3	Apache Configuration	30
12	Security & Auditing	30
12.1	Browser Security	30
12.2	Package Security	32
13	High Availability and Load Balancing	33
13.1	HA Checklist	33
13.2	HA Time Synchronization Requirements	33
13.3	HA Limitations	33
13.4	Updating HA Nodes	34
13.5	Downgrading	34
14	Admin CLI	35
14.1	Privileges	35
14.2	Location	35
14.3	Getting Help	35
14.4	Examples	35
15	Repositories and Sources	35
15.1	Repository Structure	36
15.2	Repository Versioning	36
15.3	Sources	36
15.4	The CRAN Source	37
15.5	Curated CRAN Sources	39
15.6	Git Sources	40

16 Building R Packages	43
16.1 Job Launcher	43
16.2 R Configuration	43
16.3 R Installation	43
16.4 Process Management	44
17 Process Management	44
17.1 Quick Start	44
17.2 Sandboxing	45
17.3 Temporary Directory	46
17.4 User Namespace Support	47
17.5 User Account for Processes	47
18 Configuring RStudio Server	48
18.1 A Single Repository	48
18.2 Internal Packages and CRAN Packages	48
18.3 Allow Users to Optionally Add Additional Repos	48
18.4 Precedence of Settings	49
19 Managing Change Control	49
19.1 Approach 1: Client Side Management	49
19.2 Approach 2: Repository Versioning	49
19.3 Approach 3: Locked Down Repositories	50
19.4 What about versions of R?	50
19.5 What about Docker?	50
Appendix	50
A Configuration Options	50
A.1 Server	53
A.2 API	54
A.3 HTTP	55
A.4 HTTPS	55
A.5 HttpRedirect	56
A.6 Git	56
A.7 Launcher	57
A.8 Licensing	58
A.9 CRAN	58
A.10 Database	59
A.11 SQLite	59
A.12 Postgres	60
A.13 Storage	60
A.14 FileStorage	61
A.15 S3Storage	61
B Package Ecosystem	62
B.1 Packages	62
B.2 Repositories	63
B.3 Git(hub)	63
B.4 Libraries	63
C Changing RunAs User	64
C.1 Default Account	64
C.2 Changing The RunAs Account (Service)	64
C.3 Using the New RunAs Account (CLI)	66

D	Air-Gapped RStudio Package Manager	66
D.1	Initial Setup	66
D.2	Regular Updates	68
D.3	Upgrading RStudio Package Manager	68
E	Manual Installation for Minimal Root Use	69
E.1	Discussion	69
E.2	Extracting Files	69
E.3	Create Directories	69
E.4	Licensing	70
E.5	Edit config file	70
E.6	Start the RStudio Package Manager Service	70
E.7	Use the CLI to Manage RStudio Package Manager	70

1 Introduction

RStudio Package Manager organizes and centralizes R packages across your team, department, or entire organization.

Traditionally, R packages entered the organization from a variety of sources including CRAN, Bioconductor, Github, and even internally developed package sources. RStudio Package Manager empowers R users to access packages and reproduce environments while giving IT control and visibility into package use.

2 Get Started: Installation

This chapter outlines the minimum steps to install RStudio Package Manager. More information on each step is available in subsequent chapters. Your environment will likely require one or more additional steps.

2.1 Need Help?

RStudio is happy to provide support for anyone using the product.

To focus on building the best products, RStudio does not provide installation or professional services. If you require such assistance, please refer to our certified partners.

RStudio support is limited to our professional products. If your question is related to R, R code, or specific R packages, please visit <https://community.rstudio.com>.

If your question is related to RStudio Package Manager, you can open a support ticket by:

1. Running the command `/opt/rstudio-pm/bin/run-diagnostics`. *Note, if you have not been able to download the software, skip this step. If you unpacked the software in a non-default location, navigate to that directory and use `./bin/run-diagnostics`.*
2. If the cause of the problem is not evident to you, send an email to support@rstudio.com with:
 - Your name and company
 - Your goal
 - What you have tried
 - The current error
 - The diagnostics output from step 1

RStudio support is via email, see our SLA for details.

2.2 Prerequisites

Please review the necessary system requirements, account privileges, and internet access requirements prior to installing RStudio Package Manager.

2.2.1 System Requirements

- Red Hat Enterprise Linux/CentOS Linux 6.0+
- Red Hat Enterprise Linux/CentOS Linux 7.0+
- SUSE Enterprise Linux 12 SP3+
- Ubuntu 14.04 (**not recommended**; end of life scheduled for April 2019)
- Ubuntu 16.04
- Ubuntu 18.04

RStudio Package Manager should run on a server with a minimum of 2GB of RAM. Package sources can be lazily cached from CRAN or eagerly downloaded. RStudio Package Manager recommends 50-200GB of disk storage.

2.2.2 Internet Access Requirements

RStudio Package Manager acts as a “bridge” between offline servers running R and upstream package sources like CRAN. RStudio Package Manager should be installed on a server in the network with outbound access to:

`https://rspm-sync.rstudio.com`

See Repository Syncing for more details on what information is downloaded from the RStudio CRAN service.

RStudio Package Manager can be configured to use a HTTP proxy to access the internet, see Outbound Proxy for more details.

Additional steps are necessary if your internal servers can not access an online server through your internal network. See the appendix on running RStudio Package Manager in an air-gapped environment.

2.2.3 Root Requirements

RStudio Package Manager does not run as root, see 2.2.4. By default, root privileges are required to:

- Install RStudio Package Manager
- Start and Stop RStudio Package Manager via the service daemons
- Activate the RStudio Package Manager license

RStudio Package Manager can be installed and licensed without requiring root privileges. Instructions are available in the appendix.

2.2.4 RunAs User

RStudio Package Manager starts and runs under an unprivileged account. It never requires root.

The installer creates a user account and group named `rstudio-pm` and runs the RStudio Package Manager service under this account. If you wish to change the account under which the service runs, please see C.

2.2.5 R Installation

If you plan to use git sources, RStudio Package Manager will require a valid installation of R version 3.0.0 or greater. See 16 for more information.

2.3 Installation

Download and install RStudio Package Manager for your platform. These commands will install the product into `/opt/rstudio-pm`. See E to install into a non-default location or to perform a non-root install.

```
# for ubuntu
sudo apt-get install gdebi-core
wget https://s3.amazonaws.com/rstudio-package-manager/ubuntu/amd64/rstudio-pm_1.0.6-6_amd64.deb
sudo gdebi rstudio-pm_1.0.6-6_amd64.deb

# for redhat and centos
wget https://s3.amazonaws.com/rstudio-package-manager/centos/x86_64/rstudio-pm-1.0.6-6.x86_64.rpm
sudo yum install --nogpgcheck rstudio-pm-1.0.6-6.x86_64.rpm

# for suse
wget https://s3.amazonaws.com/rstudio-package-manager/opensuse/x86_64/rstudio-pm-1.0.6-6.x86_64.rpm
sudo zypper --no-gpg-checks install rstudio-pm-1.0.6-6.x86_64.rpm
```

2.4 Initial Configuration

RStudio Package Manager settings are controlled by a configuration file located at `/etc/rstudio-pm/rstudio-pm.gcfg` by default. At a minimum, complete the **Address** property within the **Server** section, by specifying the URL used to access RStudio Package Manager.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
[Server]
Address = http://r-packages.example.com
```

Replace `http://r-packages.example.com` with the value appropriate to your organization.

There are many other default options you may want to configure now or at a later point, such as:

- The on-disk location where data will be stored, default: `/var/lib/rstudio-pm`
- The database where essential metadata should be stored, default: built-in SQLite.
- Important security settings settings such as HTTPS certificates and ports, default: HTTP on port 4242.

For multi-node or highly available (HA) setups, please refer to the REQUIRED config settings listed at 13.

2.5 Licensing

Verify the status of your license using:

```
/opt/rstudio-pm/bin/license-manager status
```

If you have a license key, activate the license using

```
/opt/rstudio-pm/bin/license-manager activate LICENSE_KEY
```

Replace `LICENSE_KEY` with the value of your license key. If your server is offline or you need more information, consult 6.

2.6 Restart RStudio Package Manager

After changing the configuration file or license you must restart the RStudio Package Manager service. Follow the instructions for your operating system:

```
# for Ubuntu 16.04, Ubuntu 18.04, RHEL/CentOS 7, SUSE 12  
sudo systemctl stop rstudio-pm  
sudo systemctl start rstudio-pm
```

```
# for Ubuntu 14.04 and RHEL/CentOS 6  
sudo stop rstudio-pm  
sudo start rstudio-pm
```

2.7 Visit the web UI

You should be able to visit RStudio Package Manager in a web browser. **The default port is 4242**, so navigate to: <https://r-packages.example.com:4242>. RStudio Package Manager should present a welcome page, though no repositories are yet available.

2.8 Try out the CLI

You should be able to invoke the admin CLI by running: `/opt/rstudio-pm/bin/rsfm list`. If you are not able to use the CLI, ensure your user is a member of the `rstudio-pm` group:

1. Run: `sudo usermod -aG rstudio-pm USER`, replacing `USER` with your username. This command will add `USER` to the `rstudio-pm` group.
2. **Close the terminal or SSH session and begin a new session.**
3. In the **new** session, run the command `groups` to enumerate the groups the current user is a member of. Confirm that `rstudio-pm` is listed.
4. Re-run the command: `/opt/rstudio-pm/bin/rsfm list`.

Note, you can elect a different group to grant access to the CLI, see C.

2.9 Next Steps

Now that RStudio Package Manager is installed, you can complete the second part of the setup by configuring RStudio Package Manager for your desired use case. Follow the instructions in [Get Started: Configuration](#).

3 Get Started: Configuration

These guides are designed to help administrators configure RStudio Package Manager for common scenarios. Before beginning, follow the [Getting Started with Installation](#) steps to install and license RStudio Package Manager.

The commands in the quick start guides need to be run as an appropriately privileged user. By default, the user should be a member of the `rstudio-pm` group. In addition, the instructions assume you've setup an alias: `alias rsfm='/opt/rstudio-pm/bin/rsfm'` or have added the binary to your path.

Once you are done setting up RStudio Package Manager, share the URL to RStudio Package Manager with users. Users will be able to setup R or RStudio to use RStudio Package Manager by following the instructions included in each repository's [Setup](#) page.

If you prefer, you can configure RStudio Server (Pro) to use RStudio Package Manager without requiring user setup, see the configuration instructions.

Pick the quick start guide(s) most suited to your use case(s):

1. Serving CRAN Packages 3.1
2. Serving Local Packages 3.2
3. Serving Local Packages from Git 3.3
4. Serving CRAN and Local Packages 3.4
5. Supplementing CRAN with Bleeding Edge Packages from GitHub 3.5
6. Serving an Approved Subset of CRAN 3.6
7. Serving an Approved Subset of CRAN and Local Packages 3.7

3.1 Serving CRAN Packages

A common use case for RStudio Package Manager is making CRAN packages available in environments with restricted internet access. To do so, start by ensuring that RStudio Package Manager has the appropriate metadata using the `sync` command. RStudio Package Manager pulls packages and metadata from the RStudio CRAN service, it is *not* necessary to configure an upstream CRAN URL. Then, create a repository and subscribe it to the built-in source named “cran”.

```
# Initiate a sync:
rspm sync --wait

# Create a repository:
rspm create repo --name=prod-cran --description='Access CRAN packages'

# Subscribe the repository to the cran source:
rspm subscribe --repo=prod-cran --source=cran
```

The behavior of the repository will depend on whether your server is set up for lazy or eager downloading. This choice is made in the RStudio Package Manager config file. Likewise, future updates will occur on a schedule dictated in the configuration file. See 15.4.3 for more information.

After completing these steps, the `prod-cran` repository will be available in the web interface.

3.2 Distributing Local Packages

Many teams have a handful of internally built packages. If your internal packages are tracked in Git, see the next section. Otherwise, create the **bundled** version of each package and copy the resulting tar files to the RStudio Package Manager server. If you are unfamiliar with building the bundled version of a package, reach out to the R developer maintaining the package.

```
# Create a local source:
rspm create source --name=prod-internal-src

# Add each local package tar file to the source:
# The tar file must be rwx by the user running the CLI and the account running
# RSPM (rstudio-pm by default)
rspm add --source=prod-internal-src --path='/path/to/package_1.0.tar.gz'

# Create a repository:
rspm create repo --name=prod-internal --description='Stable releases of our internal packages'
```



```
# Subscribe the repository to the source:
rspm subscribe --repo=prod-internal --source=prod-internal-src
```

RStudio Package Manager automatically supports multiple versions of each package. When the R developers are ready for the next release of a package, simply run:

```
rspm add --source=prod-internal-src --path='/path/to/package_2.0.tar.gz'
```

RStudio Package Manager will ensure that version 2.0 is the default for new installations, but will keep version 1.0 available in the repository's archive for users who wish to use the older version.

Most internal packages will depend on packages from CRAN. In this case, the easiest option is to create a repository that includes the local packages and their dependencies.

See the quickstart sections for serving local packages and all of CRAN or serving local packages and a curated subset of CRAN.

3.3 Serving Local Packages from Git

The previous configuration using a local source requires manual steps to add and update packages. If your organization uses Git to store internal R packages, you can automate this process using a git source.

git sources require a valid R installation.

```
# Create a git source:
rspm create source --type=git --name=prod-internal-src

# Add a Git endpoint, configured to surface tagged commits:
rspm add --git-url=https://bitbucket.example.com/r-pkg.git --source=prod-internal-src --git-build-trigger=

# Create a repository and subscribe it to the source:
rspm create repo --name=prod --description='Stable releases of our internal packages'
rspm subscribe --source=prod-internal-src --repo=prod
```

Packages can be built using Git endpoints accessed via HTTP(s) or SSH URLs - `https://github.com/user/repo.git` vs `git@github.com:user/repo.git`. If the Git URL uses SSH it will require an SSH key for authentication. In this case, first import the key and then use the add command.

SSH keys must NOT use a passphrase.

```
# Import the SSH key:
rspm import --name=read-r-pkg --path=/path/to/ssh/key

# Optionally, remove the key from disk:
# rm /path/to/ssh/key

# Add the package:
rspm add --git-url=user@bitbucket.example.com/r-pkg.git --source=prod-internal-src --git-build-trigger=
```

For more information on Git sources and SSH key security, see chapters 15.6 and 15.6.1.1.

3.4 Distributing Local Packages along with CRAN Packages

Often, for convenience, organizations opt to distribute their local packages along with CRAN packages in a single repository. This setup gives a single URL for all the organization's R packages. To do so, follow the quick start guides above to create a local source or a git source along with a cran source, then:

```

# Confirm sources exist:
rspm list sources; # should have: prod-internal-src, cran

# Create repository:
rspm create repo --name=prod --description='Production R packages from CRAN and our local packages'

# Subscribe the repository to the sources:
rspm subscribe --repo=prod --source=prod-internal-src; rspm subscribe --repo=prod --source=cran

```

In the final step, the order of subscriptions is important. Packages from the local source will be favored if there are conflicts. See 15.3.2 for details on how conflicts are resolved.

3.5 Supplementing CRAN with Bleeding Edge Packages from GitHub

In addition to a production repository, some organizations allow advanced R users to access “bleeding-edge” versions of packages available on GitHub.

```

# Create a git source:
rspm create source --type=git --name=github

# Add the GitHub version of ggplot2:
rspm add --source=github --git-url=https://github.com/tidyverse/ggplot2.git --git-build-trigger=commits

# Sync CRAN:
rspm sync --wait

# Create a repository, and subscribe it to the GitHub source and CRAN:
rspm create repo --name=bleeding-edge --description='CRAN supplemented with bleeding edge packages from CRAN and GitHub'
rspm subscribe --repo=bleeding-edge --source=github
rspm subscribe --repo=bleeding-edge --source=cran

```

The result of these steps is a repository that contains the GitHub version of the package in addition to CRAN packages. The GitHub package is installed from RStudio Package Manager using `install.packages` NOT `devtools`.

In the example above, the ordering of the source subscriptions is important. Imagine the desired GitHub package is `ggplot2`. `ggplot2` will be available from the `cran` source and the `git` source. When `ggplot2` is requested, RStudio Package Manager looks in the order of the subscriptions. First it will look for `ggplot2` in the `github` source. Since `ggplot2` is in the set of bleeding edge packages, RStudio Package Manager will find `ggplot2` in the `github` source and serve the bleeding edge version to users. Any `ggplot2` dependencies that are not in the `github` source will be pulled from `cran`. RStudio Package Manager displays whether the package came from CRAN or an alternative source in the package page for the repository.

3.6 Serving a Subset of Approved CRAN Packages

Some organizations only want to give access to an approved list of CRAN packages. A curated CRAN source enables administrators to serve the approved list of packages as well as any dependencies, while enabling admins to preview changes, add new packages, and run updates.

Create a file containing one package name per line. For example, `/tmp/packages.csv`:

```

plumber
shiny
ISLR

```

Start by creating a `curated-cran` source. Then use the `add` command to preview the changes needed to add the packages.

```
# Ensure you have CRAN metadata:
rspm sync --wait

# Create the curated-cran source:
rspm create source --name=subset --type=curated-cran

# Dry run to see proposed packages:
rspm add --file-in='/tmp/packages.csv' --source=subset --dryrun
```

The result will contain information on all the packages that will be added. The proposal can be saved to a csv file using the `csv-out` flag. **The required dependencies for the named packages are automatically discovered and included.** Optionally use the `--include-suggests` flag to also discover and add suggested packages.

This action will add the following packages:

Name	Version	Path	License	Needs	Compilation	Dependency
BH	1.66.0-1		BSL-1.0	false		true
crayon	1.3.4		MIT + file LICENSE	false		true
digest	0.6.15		GPL (>= 2)	false		true
htmltools	0.3.6		GPL (>= 2)	false		true
httpuv	1.4.3		GPL (>= 2) file LICENSE	false		true
ISLR	1.2		GPL-2	false		false
jsonlite	1.5		MIT + file LICENSE	false		true
later	0.7.3		GPL (>= 2)	false		true
...

To complete this operation, execute this command without the `--dryrun` flag. You will need to include the `--transaction-id=281` flag.

If the proposal is acceptable, run the command again, but use the `transaction-id` indicated in the dry run output. Finally, create a repository and subscribe the repository to the source.

```
# Commit the changes:
rspm add --file-in='/tmp/packages.csv' --source=subset --transaction-id=281

# Create a repository:
rspm create repo --name=approved-cran --description='Approved packages from CRAN'

# Subscribe the repository to the source:
rspm subscribe --repo=approved-cran --source=subset
```

When the source is created, RStudio Package Manager automatically pins the “subset” source to a frozen point in time on CRAN. New packages can be added from this frozen snapshot of CRAN by repeating the process.

To update the entire set of packages to the latest data available from CRAN, use the `update` command. (Note that the latest data will depend on the server’s sync schedule).

```
rspm update --source=subset --dryrun

# Like the add command, a preview of the changes is printed out
# along with a transaction-id. Use this id to commit the changes
```

```
rspm update --source=subset --transaction-id=281
```

Please review the full description of curated CRAN sources to understand further which points in time `add` and `update` will use.

3.7 Serving a Subset of Approved CRAN Packages and Local Packages

An extension of the previous use case is serving a subset of CRAN and a set of internal packages from within a single repository. Start by following the steps to create a source with the subset of approved packages and a local source with the desired internal packages.

Next, for each internal package, obtain the list of the package's dependencies from the package developer. Create a csv file with one line per package name, e.g. `internal_deps.csv`:

```
xml2  
jsonlite
```

Use the `add` command to add these additional dependencies to the curated CRAN source.

```
# List sources, should include  
# prod-internal-source and subset:  
rspm list sources  
  
# Add the local packages' dependencies:  
rspm add --source=subset --file-in='internal_deps.csv' --dryrun  
  
# Proposed changes will be printed out, to commit the changes  
# run the command again with the transaction-id:  
rspm add --source=subset --file-in='internal_deps.csv' --transaction-id=283
```

Finally, create a repository and subscribe the repository to both sources. The order of the subscription commands is important, see 15.3.2 for details.

```
# Create a repo:  
rspm add repo=prod-pkgs --description='Stable release of our internal packages and approved CRAN packages'  
  
# Subscribe the repo to both sources:  
rspm subscribe --repo=prod-pkgs --source=prod-internal-source  
rspm subscribe --repo=prod-pkgs --source=subset
```

4 Frequently Asked Questions

4.1 Setting Up RStudio Package Manager

Where should I install RStudio Package Manager? What server specifications do you recommend?

RStudio Package Manager can be installed on a physical server, cloud instance, virtual machine, or a long-living Docker container with persistent storage. RStudio Package Manager is supported on most Linux distributions, specifically listed in the installation section of the admin guide.

RStudio Package Manager requires 2 GB of RAM. Normally, RStudio Package Manager fetches data lazily and requires 50-100 GB of disk storage. For air-gapped networks, all CRAN data must be downloaded up front which requires 200 GB of disk storage.

Does RStudio Package Manager require internet access?

RStudio Package Manager can be used without internet access, but most organizations will want RStudio Package Manager to have outbound internet access to RStudio's CRAN service.

R clients using RStudio Package Manager do not need internet access, just access to RStudio Package Manager. RStudio Package Manager is designed to help organizations with restricted internet access in their R environment.

Can RStudio Package Manager be setup for High Availability?

Yes. RStudio Package Manager can be setup on 2 or more servers and placed behind a load balancer. This setup requires a PostgreSQL database and shared storage; see the admin guide for details.

Can I use RStudio Package Manager in an offline, air-gapped environment?

Yes, there are a number of options for air-gapped environments; see the instructions in the getting started guide.

What permissions are necessary to install RStudio Package Manager?

RStudio Package Manager's default installation requires root privileges in order to configure the service daemon for RStudio Package Manager, setup log files, unpack RStudio Package Manager, and activate the license. It is possible to manually install RStudio Package Manager without root; see the associated appendix.

What permissions are needed to manage repositories in RStudio Package Manager?

Repositories and packages are managed with a command line interface located at `/opt/rstudio-pm/bin/rspm`. By default, users must be a member of the `rstudio-pm` Unix group to use the CLI. The group that determines access can be customized.

How do I configure RStudio Package Manager?

RStudio Package Manager's primary configuration file is located at `/etc/rstudio-pm/rstudio-pm.gcfg` by default.

Do I need a database for RStudio Package Manager?

RStudio Package Manager uses a database to store meta-data and serve packages. Administrators do not need to interact with the database. By default, RStudio Package Manager will use an included SQLite database. Optionally, RStudio Package Manager can use an external PostgreSQL database and in a cluster RStudio Package Manager requires PostgreSQL.

4.2 Configuring RStudio Package Manager

Can I use RStudio Package Manager to access CRAN?

Yes, see the CRAN quickstart guide. RStudio Package Manager can be configured to serve CRAN packages to servers behind your firewall. Behind the scenes, RStudio Package Manager maintains metadata on all CRAN packages, but only downloads content when necessary.

How do I use RStudio Package Manager to share local packages?

See the quickstart guide for local packages. RStudio Package Manager makes it easy to distribute internal packages and automatically handles archiving older versions of the package when new releases are added.

Can I setup multiple repositories in RStudio Package Manager?

Yes, RStudio Package Manager allows administrators to create multiple repositories with different package and source compositions. For example, you may have a repo with CRAN packages and a repo with local packages, or you may create a single repository that contains local and CRAN packages. You could create a production repo with stable releases and a development repo with the bleeding edge versions of packages. RStudio Package Manager optimizes storage to prevent needless duplication.

Where does RStudio Package Manager get packages?

RStudio Package Manager currently supports packages from CRAN, GitHub, and local packages stored in Git or available as tar files. CRAN packages are managed by an RStudio service and can be downloaded eagerly or lazily on-demand.

Does RStudio Package Manager support BioConductor?

Today, RStudio Package Manager does not have direct support for BioConductor. However, BioConductor package tar files can be added manually or admins can configure RStudio Package Manager to track specific BioConductor Git endpoints. Please reach out to sol-eng@rstudio.com for more information.

4.3 Using RStudio Package Manager from R or RStudio

How do I tell R and RStudio to use RStudio Package Manager?

Each repository in RStudio Package Manager will contain a Setup web page with instructions for configuring R to use RStudio Package Manager. Once configured, R users can access packages using standard tools including `install.packages`, `packrat`, and `devtools`.

Can I configure RStudio Server to use RStudio Package Manager?

Yes, an administrator can configure RStudio Server (Pro) to use RStudio Package Manager by default, see the configuration instructions.

4.4 Controlling Access to Packages

How do I specify which packages should be available in RStudio Package Manager?

Administrators configure the available packages, sources, and repositories using an admin command line interface.

4.5 Package Updates

How do I update packages in RStudio Package Manager?

RStudio Package Manager can automatically update packages from CRAN on an admin-defined schedule. Local packages can be updated at any time by adding the tar file for the new version of the package. Packages from Git or GitHub will be updated automatically. Updates can track new commits or new tags.

RStudio Package Manager automatically archives older versions of packages during updates, ensuring older versions are available to users. Older versions are listed on the web page for each package.

4.6 Managing Change Control

Does RStudio Package Manager help users manage the versions of packages installed?

RStudio Package Manager was built to support all the common strategies organizations use to manage different versions of R packages. See 19 for more information.

Is RStudio Package Manager compatible with packrat?

Yes, RStudio Package Manager works with packrat for CRAN and local packages.

4.7 RStudio Package Manager and Docker

Can RStudio Package Manager run in Docker?

Yes, RStudio Package Manager can run in a Docker container that has persistent, mounted storage.

If I use Docker to manage package dependencies, should I use Package Manager?

Docker images can be used to create an environment for a certain project or analysis. However, the use of Docker is orthogonal to RStudio Package Manager. For example, inside a Dockerfile for an R project you will normally see a line that installs R packages, e.g.

```
RUN Rscript -e 'install.packages(...)'
```

This installation step benefits from RStudio Package Manager in the same way that non-Docker users benefit from RStudio Package Manager.

5 Server Management

This section describes common administrative tasks for RStudio Package Manager.

5.1 Stopping and Starting

Occasionally it is necessary to start and stop the RStudio Package Manager service. Stopping and starting is handled by `systemd` or Upstart.

The specific stop/start commands depend on the service daemon. Commands for `systemd` and Upstart are listed below.

After a restart, any scheduled syncs that were missed during downtime will automatically begin. See 15.4 for details.

5.1.1 systemd

Distributions using `systemd` include Red Hat/CentOS 7, SUSE 12, Ubuntu 16.04, and Ubuntu 18.04.

`systemd` is a management and configuration platform for Linux. The newest versions of most major Linux distributions have adopted `systemd` as their default init system.

The RStudio Package Manager installer installs a `systemd` service called `rstudio-pm`, which causes the RStudio Package Manager to be started and stopped automatically when the machine boots up and shuts down. The `rstudio-pm` service is also automatically launched during installation.

Use the following commands to manually start and stop the server:

```
sudo systemctl start rstudio-pm
```

```
sudo systemctl stop rstudio-pm
```

You can restart the server with:

```
sudo systemctl restart rstudio-pm
```

If you wish to keep the server running without interruption, but reload the configuration, you can use the `systemctl` command to send a HUP signal:

```
sudo systemctl kill -s HUP --kill-who=main rstudio-pm
```

The HUP signal causes the server to re-initialize but does not interrupt the current processes or any of the open connections to the server.

Use a HUP signal when your configuration changes are limited to properties marked as reloadable. See Appendix A to learn which settings may be reloaded via HUP. Perform a full restart of RStudio Package Manager when changing other properties.

You can check the status of the `rstudio-pm` service using:

```
sudo systemctl status rstudio-pm
```

And finally, you can use the `enable/disable` commands to control whether RStudio Package Manager should be run automatically at boot time:

```
sudo systemctl enable rstudio-pm
```

```
sudo systemctl disable rstudio-pm
```

5.1.2 Upstart (Ubuntu 14.04, Red Hat 6)

Distributions using Upstart include Red Hat/CentOS 6 and Ubuntu 14.04. We do **not recommend** Ubuntu 14.04 since its end of life is scheduled for April 2019.

Upstart is a system used to automatically start, stop and manage services. The installer writes an Upstart configuration file to `/etc/init/rstudio-pm.conf`. This instructs the Upstart to initialize RStudio Package Manager as soon as the network is activated on the machine and stop when the machine is being shut down.

The Upstart configuration also ensures that the `rstudio-pm` process is respawned if the process unexpectedly terminates. However, in the event that there is an issue which consistently prevents RStudio Package Manager from being able to start (such as a bad configuration file), Upstart will give up on restarting the service after approximately 5 failed attempts within a few seconds. For this reason, you may see multiple repetitions of a bad RStudio Package Manager startup attempt before it transitions to the “stopped” state.

To start or stop the server, run the following commands, respectively.

```
sudo start rstudio-pm
```

```
sudo stop rstudio-pm
```

To restart the server you can run:

```
sudo stop rstudio-pm
```

```
sudo start rstudio-pm
```

The `restart` command re-initializes the server.

We recommend `stop` and `start` over `restart` because some configuration changes are not incorporated into a restart. In particular, `restart` *does not re-read the Upstart definition at `/etc/init/rstudio-pm.conf`*. Changes to this file need `astopandstart` to take effect.

If you wish to reload the configuration and keep the server and all R processes running without interruption, you can use the `reload` command:

```
sudo reload rstudio-pm
```

This command causes the server to re-initialize but does not interrupt the current processes or any of the open connections to the server.

Use a HUP signal when your configuration changes are limited to properties marked as reloadable. See Appendix A to learn which settings may be reloaded via HUP. Perform a full restart of RStudio Package Manager when changing other properties.

To check the status or retrieve the process ID associated with `rstudio-pm`, run the following:

```
sudo status rstudio-pm
```

5.2 Upgrading

Upgrading RStudio Package Manager requires limited downtime. During an upgrade users will not be able to install packages. We recommend upgrading during a period of downtime.

The latest version is available on the download page along with release notes. The current version is available by running:

```
cat /opt/rstudio-pm/VERSION
```

To upgrade:

1. Download the latest `.rpm` or `.deb` file
2. Run the install command:

Ubuntu:

```
sudo gdebi <rstudio-pm-version.deb>
```

Red Hat/CentOS:

```
sudo yum install --nogpgcheck <rstudio-pm-version.rpm>
```

SUSE:

```
sudo zypper --no-gpg-checks install <rstudio-pm-version.rpm>
```

The new version of RStudio Package Manager will install on top of an earlier installation. Existing configuration settings are respected. During installation the RStudio Package Manager service is restarted.

5.3 Purging RStudio Package Manager

You can fully remove RStudio Package Manager and all its data from your server using the following steps:

1. Stop the RStudio Package Manager service. (See 5.1 for details)
2. Uninstall the RStudio Package Manager package from your system.

Ubuntu:

```
sudo apt-get purge rstudio-pm
```

Red Hat/CentOS:

```
sudo yum remove rstudio-pm
```

SUSE:

```
sudo zypper remove rstudio-pm
```

3. Remove `/opt/rstudio-pm` if it still exists.
4. Remove logs from `/var/log/rstudio-pm*`

5. Purge the databases
 - When using SQLite, remove the `SQLite.Dir` directory. This has a default location of `/var/lib/rstudio-pm/db`.
 - When using PostgreSQL, drop the databases used by RStudio Package Manager. You may also wish to remove the PostgreSQL user associated with RStudio Package Manager.
6. Remove the `Server.DataDir` directory. By default, this is `/var/lib/rstudio-pm`.
7. Remove configuration files from `/etc/rstudio-pm` if they still exist.
8. Remove any variable data directories outside the `Server.DataDir` directory that are defined by `FileStorage` configuration sections. See section 7.7 for more information.
9. Remove any data stored on S3. See section 7.8.

5.4 Backups

We recommend including the RStudio Package Manager configuration file in `/etc/rstudio-pm` as well as the variable data directory which defaults to `/var/lib/rstudio-pm` in your system backups. If you have configured the databases to be stored outside the data directory, ensure that it is also included in the backup. If you have customized the storage locations for any variable data classes (see section 7.7), each customized directory should also be included in the backup.

If you use S3 to store variable data, consider whether you need additional backups for your S3 data. See section 7.8 for more information about storing variable data on S3.

A running RStudio Package Manager server may be writing into the data directory. You should stop the RStudio Package Manager server before taking a backup.

```
sudo stop rstudio-pm
# Run appropriate backup steps here.
sudo start rstudio-pm
```

Your platform may need alternate commands to restart RStudio Package Manager. Please see Section 5.1 for instructions specific to your operating system version.

5.5 Privileged Ports

RStudio Package Manager listens on HTTP port 4242 by default. When you modify the `HTTP.Listen` or `HTTPS.Listen` configuration properties to use a privileged port under 1024, the service will fail to start, and you will see an error like the following in `/var/log/rstudio-pm.log`:

```
2017/11/28 13:41:59 Error: Could not initialize the HTTP listener: listen tcp :80: bind: permission denied
```

If you wish to listen with HTTP or HTTPS on a privileged port (< 1024), you can grant the RStudio Package Manager binary permission to do so by issuing the following command as root:

```
sudo setcap 'cap_net_bind_service=+ep' /opt/rstudio-pm/bin/rstudio-pm
```

After issuing the above command, restart the service (see 5.1).

6 Licensing

When RStudio Package Manager is first installed on a system it operates in an evaluation mode for a period of time and then subsequently requires activation for continued use.

To determine the current license status of your system, you can use the following command:

```
sudo /opt/rstudio-pm/bin/license-manager status
```

After purchasing a license to RStudio Package Manager, you will receive a product key that is used to activate the license on a given system.

You can activate your license key with the command:

```
sudo /opt/rstudio-pm/bin/license-manager activate <product-key>
```

After activation, restart the RStudio Package Manager server.

```
sudo stop rstudio-pm
sudo start rstudio-pm
```

Your platform may need alternate commands to restart RStudio Package Manager. Please see Section 5.1 for instructions specific to your operating system version.

If you want to move your license of RStudio Package Manager to another system, you should first deactivate it on the old system.

```
sudo /opt/rstudio-pm/bin/license-manager deactivate
```

6.1 Proxy Servers

If your server is behind an internet proxy, you may need to add an additional command line flag indicating the address and credentials required to communicate through the proxy. This may not be necessary if either the `http_proxy` or `all_proxy` environment variable is defined (these are read and used by the license manager when available).

If you do need to specify a proxy server explicitly you can do so using the `--proxy` command line parameter. For example:

```
sudo /opt/rstudio-pm/bin/license-manager \
  --proxy=http://127.0.0.1/ activate <product-key>
```

Proxy settings can include a host-name, port, and username/password if necessary. The following are all valid proxy configurations:

```
http://127.0.0.1/
http://127.0.0.1:8080/
http://user:pass@127.0.0.1:8080/
```

If the port is not specified, the license manager will default to using port 1080.

6.2 Offline Activation

If your system has no connection to the internet it's also possible to perform an offline activation. To do this, we recommend using our offline activation app which will walk you through the process: RStudio Offline Activation

You first generate an offline activation request as follows:

```
sudo /opt/rstudio-pm/bin/license-manager activate-offline-request <product-key>
```

Executing this command will print an offline activation request to the terminal which you should copy and paste into an email to RStudio customer support (support@rstudio.com). You will receive a reply with a file attachment that can be used to activate offline as follows:

```
sudo /opt/rstudio-pm/bin/license-manager activate-offline <activation-file>
```

After activation, restart the RStudio Package Manager server. Please see Section 5.1 for instructions specific to your operating system version.

If you want to renew your license of RStudio Package Manager or move it to another system you can also perform license deactivation offline. You can do this as follows:

```
sudo /opt/rstudio-pm/bin/license-manager deactivate-offline
```

Executing this command will print an offline deactivation request to the terminal which you should send to support@rstudio.com.

You can also perform an offline check of your current license status using the following command:

```
sudo /opt/rstudio-pm/bin/license-manager status-offline
```

6.3 Licensing errors

RStudio Package Manager uses the `license-manager` to determine if a valid license is available. Should an error occur when interacting with the license manager, RStudio Package Manager indicates that problem in the `/var/log/rstudio-pm.log` log. The license manager sends details about the error to the system messages (syslog). You should consult both locations should RStudio Package Manager be unable to obtain license status.

6.4 Repository Count

RStudio Package Manager is licensed, in part, by the number of available repositories. In some situations the web interface may display the message:

```
The number of repos exceeds the licensed repo count. Please see a system administrator.
```

This error may occur if you activate your license after the evaluation and have created more repositories than your license allows. The error may also occur if you downgrade your license. A similar error will also occur if you try to create a new repository after you reach your repository limit.

This error indicates that you are *in violation of your license* and steps should be taken *immediately*:

- If you require all of your repositories, or if a mistake has been made with the license, email support@rstudio.com or your customer success rep.
- Consider restructuring your sources and repository subscriptions. Multiple sources can be combined into a single repository.
- Delete unused repositories with the `delete` command.

RStudio Package Manager will continue to serve packages from all defined repositories to provide a grace period while you apply changes. However, you are obligated to bring your configuration back into compliance with your license agreement.

6.5 Floating Licenses

Traditionally, an admin must activate a license on each RStudio Package Manager instance by running the `activate` and `deactivate` commands. If you start and stop RStudio Package Manager frequently, for instance because you are running inside a virtual machine or docker container, it is preferable to use floating licensing. Floating licensing allows you to run a standalone licensing server and then have each RStudio Package Manager instance communicate with this licensing server to lease a license. This configuration can

be useful if RStudio Package Manager is ephemeral. In this case, you can use the RStudio Package Manager License Server (RSPM License Server).

6.5.1 Overview

As an administrator you'll need to follow these steps:

1. Install the RSPM License Server on a persistent server with network access to the RStudio Package Manager instance(s).
2. Activate your license key on the RSPM License Server
3. Configure RStudio Package Manager to use the RSPM License Server
4. (Optional) Change the RSPM License Server defaults

6.5.2 RSPM License Server

Follow these instructions in a persistent server that will run the RSPM License Server.

Download the RSPM License Server and then install the server and activate the license:

Ubuntu/Debian

```
sudo dpkg -i <rspm-license-server-version.deb>
sudo rspm-license-server activate <license-key>
sudo rspm-license-server start
```

RedHat/Fedora

```
sudo yum install --nogpgcheck <rspm-license-server-version.rpm>
sudo rspm-license-server activate <license-key>
sudo rspm-license-server start
```

OpenSUSE/SUSE

```
sudo zypper --no-gpg-checks install <rspm-license-server-version.rpm>
sudo rspm-license-server activate <license-key>
sudo rspm-license-server start
```

A specific license key is necessary for floating licensing. Please reach out to support@rstudio.com to swap the default license key for a key that supports floating licensing.

6.5.2.1 Offline Activation

If you are running the RStudio Package Manager Licensing Server in an offline environment you will need to perform an additional offline activation. First, generate an offline activation request:

```
sudo rspm-license-server activate-offline-request <license-key>
```

Executing this command will print an offline activation request. Copy this request and send it to support@rstudio.com. You will receive a reply with a file attachment, use the attached file to complete the activation:

```
sudo rspm-license-server activate-offline <activation-file>
sudo rspm-license-server restart
```

6.5.3 Configuring RStudio Package Manager to Use the RSPM License Server

Once the RSPM License Server is installed and activated, you need to configure RStudio Package Manager to use the License Server.

Follow these instructions in RStudio Package Manager or script them into your configuration management tools that will be building, starting, and stopping RStudio Package Manager.

1. Edit the RStudio Package Manager configuration file:

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Licensing]
LicenseType = remote
```

2. In the server running RStudio Package Manager, execute:

```
sudo /opt/rstudio-pm/bin/license-manager license-server <rspm-license-server-host-or-ip:port>
```

By default, the RSPM License Server runs on port 8969. The `license-server` command only needs to be run once. If the same RStudio Package Manager server subsequently starts and stops it will remember the RSPM License Server address. However, if you rebuild the server or image running RStudio Package Manager you will need to re-run the `license-server` command.

3. Restart RStudio Package Manager using the commands specific to your operating system:

```
sudo stop rstudio-pm; sudo start rstudio-pm;
```

```
sudo systemctl stop rstudio-pm; sudo systemctl start rstudio-pm
```

6.5.4 Changing the RSPM License Server Defaults

These instructions refer to configuration settings on the persistent RSPM License Server.

6.5.4.1 Port

The port used by the RSPM License Server can be changed:

```
<!-- /etc/rspm-license-server.conf-->
<bind port="8969"/>
```

After modifying `/etc/rspm-license-server.conf`, restart the RSPM License Server:

```
sudo rspm-license-server restart
```

6.5.4.2 License Leases

When using floating licenses, you can optionally determine how long the license leases last by setting the `lease length` value on the licensing server. This value is in seconds, so for instance to make license leases last 30 minutes you would use the following syntax:

```
<!--/etc/rspm-license-server.conf-->
<lease length="1800"/>
```

The lease length controls how frequently the RStudio Package Manager instances need to contact the licensing server to renew their license leases in order for the lease to remain valid.

A *shorter* lease length will increase tolerance to failures on RStudio Package Manager instances by making leases available for reuse more quickly. RStudio Package Manager will release its lease immediately if shut down normally, but if abnormally terminated, the lease will not be released until it expires.

A *longer* lease length will increase tolerance to transient failures on the network and the RStudio Package Manager License Server. Any such issues that can be resolved before the lease is due for renewal won't interrupt use of RStudio Package Manager.

We generally recommend using a longer lease length. Use a short lease length only if your environment routinely encounters abnormal terminations of the server or the container/instance on which it runs.

6.5.5 Changing the RStudio Package Manager Lease Expiration and Renewal Default

This configuration applies to RStudio Package Manager and can be scripted into the build process for ephemeral instances.

Under normal conditions RStudio Package Manager will automatically renew its license lease in a configurable interval as described above. However, there are situations in which it will be unable to do so, such as a network problem, or an issue on the host running the license server.

When RStudio Package Manager cannot obtain a license lease, either because there are no leases currently available or because it can't reach the licensing server, it will begin automatically attempting to acquire a lease every 10 seconds. This interval is configurable; for instance, to retry every 30 seconds instead you would set the following value:

```
; /etc/rstudio-pm/rstudio-pm.gcfg  
  
[Licensing]  
RemoteRetryFrequency = 30s
```

If you don't want RStudio Package Manager to attempt to reestablish a license lease automatically, set the value to 0 to disable retries. In this case you will need to manually restart RStudio Package Manager in order to reestablish the lease. This can be useful if you often run more instances than you have keys for, and wish to have more control over which RStudio Package Manager instances receive license leases from the limited pool on the license server.

6.5.6 Troubleshooting Floating Licensing

Run these troubleshooting steps on the persistent RSPM License Server.

To validate that the RSPM License Server has been successfully activated, run the `activation-status` command. This will report the version of the server as well as the license key and the number of available slots.

```
sudo rspm-license-server activation-status
```

If your server is activated but you're still having trouble with floating licensing, you can tell the RSPM License Server to emit more detailed logs. Change the log level to notification:

```
<!--/etc/rspm-license-server.conf-->  
  
<log file="/var/log/rstudio-licensing.log" level="notification"/>
```

Then, restart the License Server, tail the licensing log, and start your RStudio Package Manager instances.

```
sudo rspm-license-server restart  
tail -f /var/log/rstudio-licensing.log
```

At the notification level, the licensing log will tell you the total number of licenses associated with your key, and how many are currently in use. It will also notify you when RStudio Package Manager instances acquire leases, and when those leases are released, renewed, or expired. No rotation is done for this log, so it's recommended to use the warning level in production.

7 Files and Directories

7.1 Changing Ownership

Many of the files and directories mentioned in this chapter are, by default, owned by the `rstudio-pm` user. If you change the RunAs user for the RStudio Package Manager service, you will need to change ownership of these files and directories. See C for details on changing the RStudio Package Manager service RunAs user.

7.2 Program Files

The RStudio Package Manager installers place all program files into the `/opt/rstudio-pm` directory.

You should not need to change any files in the `/opt/rstudio-pm` hierarchy. Any alterations will be overwritten by subsequent re-installs or upgrades of RStudio Package Manager.

7.3 Configuration

The RStudio Package Manager configuration file is `/etc/rstudio-pm/rstudio-pm.gcfg`. This file is initially owned by `rstudio-pm` with permissions `0640`. You will edit this file to properly configure RStudio Package Manager for your organization.

A configuration management tool like Puppet or Chef can be used to maintain the `rstudio-pm.gcfg` file. We recommend that it remain owned by `rstudio-pm` and have permissions `0640`, as your configuration may need to contain passwords and other sensitive information.

RStudio Package Manager upgrades will not overwrite customizations to the `rstudio-pm.gcfg` file.

7.4 Server Log

The RStudio Package Manager server log is located at `/var/log/rstudio-pm.log`. This file is owned by `rstudio-pm` with permissions `0600`.

If `logrotate` is available when RStudio Package Manager is installed, a `logrotate` configuration will be installed. The default configuration is to rotate the logfile daily. The old log file will be stored alongside the original with a numeric extension, `.1`, `.2`, etc. The rotated log files are compressed after one day. The `.1` log file is retained uncompressed, but older logs are compressed. Most systems use `gzip` for compression, giving log files with extensions like `.2.gz`, `.3.gz`. Logs will be maintained for 30 days.

The manual for `logrotate` has more information.

7.5 Access Logs

The RStudio Package Manager HTTP access logs are located at `/var/log/rstudio-pm.access.log`. This file is owned by `rstudio-pm` with permissions `0600`. Log files are stored in Apache Combined Log Format. See <http://httpd.apache.org/docs/2.2/logs.html#combined> for a description of this format.

If `logrotate` is available when RStudio Package Manager is installed, a `logrotate` configuration will be installed. The default configuration is to rotate the logfile daily. The old logfile will be compressed and stored alongside the original log file with a `.1.gz` extension (then `.2.gz`, etc.). Logs will maintained for 30 days.

7.6 Variable Data

RStudio Package Manager manages R packages and repositories. All package source bundles are stored in the server's data directory. The RStudio Package Manager handles incoming requests for packages across repositories. Only a single copy of each package source is stored, even if the package is referenced in multiple repositories.

The RStudio Package Manager data directory also contains information used by the server to manage repositories including the RStudio Package Manager SQLite databases and encryption key if SQLite is used.

The default location for the RStudio Package Manager data directory is `/var/lib/rstudio-pm`. This can be customized by specifying an alternate `DataDir` in the `Server` section of your configuration file.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Server]
```

```
DataDir = /mnt/rstudio-pm
```

If you customize the RStudio Package Manager data directory, make sure that the `rstudio-pm` user has permission to read, write, and create directories in the data directory.

The RStudio Package Manager SQLite databases **must** exist on local storage. If the location for `DataDir` is not local storage but a networked location over NFS, configure the `Dir` setting in the `SQLite` section of your server configuration file.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Server]
```

```
DataDir = /mnt/rstudio-pm
```

```
[SQLite]
```

```
Dir = /var/lib/rstudio-pm/db
```

7.6.1 Permissions

`/var/lib/rstudio-pm` is owned by `rstudio-pm` with permissions `0700`.

7.7 Variable Data Classes

All variable data storage locations default to subdirectories of the `Server.DataDir` setting. There are four classes of variable data, listed below

- **Cache** - Stores data to increase performance for computationally intensive operations. Certain operations, such as Git package building, also temporarily cache data here. Defaults to `<DataDir>/cache`.
- **Launcher** - Stores data for Job Launcher operations. This location currently stores the stdout and stderr data associated with each Git package builder operation. Defaults to `<DataDir>/launcher`.
- **Metrics** - This directory contains aggregated metrics data to improve Usage Stats performance. Defaults to `<DataDir>/metrics`.
- **Packages** - Package tarballs and README files are stored here. Defaults to `<DataDir>/packages`. This includes:
 - Package tarballs and README files for local packages,
 - Package tarballs and README files git packages,
 - Package tarballs for CRAN packages that have been downloaded lazily or eagerly, and
 - README files for CRAN packages (downloaded either lazily or eagerly).

You can customize the storage directory for each storage class. For example:

```
;/etc/rstudio-pm/rstudio-pm.gcfg

[FileStorage "cache"]
Location = /mnt/rstudio-pm-cache

[FileStorage "launcher"]
Location = /mnt/rstudio-pm-launcher

[FileStorage "metrics"]
Location = /mnt/rstudio-pm-metrics

[FileStorage "packages"]
Location = /mnt/rstudio-pm-packages
```

Again, if you customize any of the RStudio Package Manager storage directories, make sure that the `rstudio-pm` user has permission to read, write, and create directories in each data directory.

7.8 Destinations

IMPORTANT: AWS S3 support is in beta. Please do not use S3 for production data at this time.

The four variable storage classes (see section 7.7, above) default to storing data on disk. Each storage class can optionally be configured to store data on S3. For example, to configure all four variable data storage classes for S3, use the following configuration:

```
;/etc/rstudio-pm/rstudio-pm.gcfg

[Storage]
Cache = s3
Launcher = s3
Metrics = s3
Packages = s3

;/ Default S3 settings. This is the minimum-required setting for using S3.
[S3Storage]
Bucket = your-s3-bucket

;/ Override default S3 settings for the "packages" class. This demonstrates
;/ all the available S3 configuration settings.
[S3Storage "packages"]
Bucket = another-s3-bucket
Prefix = rspm-packages
Profile = dev-rspm
Region = us-west-1
EnableSharedConfig = true
```

RStudio Package Manager's AWS S3 support utilizes the AWS S3 SDK, which documents configuration and credential standards for interacting with S3 services.

See chapter 8 for information on configuring your system to use AWS S3.

8 AWS S3 Configuration

IMPORTANT: AWS S3 support is in beta. Please do not use S3 for production data at this time.

8.1 Destinations

You can configure variable data storage classes to use AWS S3 for storage. See section 7.8 for more information on configuring a variable data storage class to use AWS S3.

8.2 AWS Documentation

RStudio Package Manager's AWS S3 support utilizes the AWS S3 SDK, which documents configuration and credential standards for interacting with S3 services.

See the AWS CLI Configuration for detailed documentation on configuring your environment for interaction with AWS.

Below are some configuration examples.

8.3 Configuration Examples

To interact with S3, you will typically configure either environment variables or a credentials file.

8.3.1 Environment Variables

Refer to AWS Environment Variables for more information on environment variables. Here are two examples:

8.3.1.1 Ubuntu 14 and Red Hat/CentOS 6

Edit the file `/etc/init/rstudio-pm.override`

```
;/etc/init/rstudio-pm.override  
  
env AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE  
env AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY  
env AWS_DEFAULT_REGION=us-west-2
```

Restart the RStudio Package Manager service with `sudo restart rstudio-pm`.

8.3.1.2 Ubuntu 16+, Red Hat/CentOS 7, SUSE 12

Create and edit a new file at `/etc/systemd/system/rstudio-pm.service.d/aws.conf`

```
;/etc/systemd/system/rstudio-pm.service.d/aws.conf  
  
Environment="AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE"  
Environment="AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY"  
Environment="AWS_DEFAULT_REGION=us-west-2"
```

Reload the `systemd` process and restart the RStudio Package Manager service with:

```
sudo systemctl daemon-reload  
sudo systemctl start rstudio-pm
```

8.3.2 Credentials and Config Files

Refer to AWS Configuration and Credential Files for more information on configuration files.

Alternatively, you can configure RStudio Package Manager to use AWS S3 services by creating credential and configuration files in the service account's home directory. Since the `rstudio-pm` service account does not have a home directory by default, you will first need to create it.

```
sudo mkdir /home/rstudio-pm
sudo chown rstudio-pm:rstudio-pm /home/rstudio-pm
```

Next, create a `credentials` file to contain your AWS credentials

```
; /home/rstudio-pm/.aws/credentials

[default]
aws_access_key_id=AKIAIOSFODNN7EXAMPLE
aws_secret_access_key=wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

If you wish to configure your AWS region without an environment variable, you can use a config file. By default, RStudio Package Manager will only read the `~/.aws/credentials` file. To also read the `~/.aws/config` file, you must configure one of the following:

- Set an environment variable: `AWS_SDK_LOAD_CONFIG=1`, or
- Set the `S3Storage.EnableSharedConfig = true` configuration value.

When the `AWS_SDK_LOAD_CONFIG` environment variable is true or the `S3Storage.EnableSharedConfig` setting is true, both the shared config file (`~/.aws/config`) and the shared credentials file (`~/.aws/credentials`) will be loaded. Values in the shared credentials file take priority when there are duplicated values.

Create a shared config file:

```
; /home/rstudio-pm/.aws/config

[default]
region=us-west-2
```

After creating the `credentials` and/or `config` files, restart the RStudio Package Manager service.

8.3.3 RStudio Package Manager Configuration

See section 7.8 for examples of configuring RStudio Package Manager to use S3.

See section A.15 for more details on the S3 storage configuration settings available to RStudio Package Manager.

See section A.13 for details on configuring a variable data storage class to use S3.

9 Database

RStudio Package Manager supports multiple database options. Currently, the supported databases are SQLite and PostgreSQL.

Customize the `Database.Provider` property with a database scheme appropriate for your organization. See Section A.10 for details.

Here is a partial configuration which chooses to use SQLite.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Database]  
Provider = sqlite
```

9.1 SQLite

SQLite is the default database provider.

RStudio Package Manager will use SQLite database if the `Database.Provider` setting has a value of `sqlite` or if `Provider` is not present in the configuration file.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Database]  
Provider = sqlite
```

You can also specify the directory to store the SQLite file on your file system. This can be done by specifying `SQLite.Dir` in the configuration file.

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[SQLite]  
Dir = /mnt/rstudio-pm/sqlite
```

If this field is not specified, it will default to `{Server.DataDir}/db`. This location **must** exist on local storage.

If the location for `Server.DataDir` is not local storage but a networked location over NFS, configure the `SQLite.Dir` setting so it still resides on some local volume.

9.2 PostgreSQL

PostgreSQL is an available database provider which is more powerful and performant than SQLite.

You must provide your own Postgres server which will likely be a separate server from your RStudio Package Manager server (but not required). We currently support any 9.x version greater than or equal to 9.2. Your Postgres server does not have to be dedicated to RStudio Package Manager, but it must have its own dedicated database.

To use Postgres, select it as your provider with `Database.Provider = postgres`. You will also need to provide a fully qualified Postgres URL in `Postgres.URL`. The user credentials supplied in this URL must have read/write permissions to the database referenced at the end of url. Please ensure that you have already created a blank database with the name given at the end of your URL.

Using a separate database for usage data:

```
; /etc/rstudio-pm/rstudio.gcfg
```

```
[Database]  
Provider = postgres
```

```
[Postgres]  
URL = "postgres://username:password@db.seed.co/rstudio-pm"  
UsageDataURL = "postgres://username:password@db.seed.co/rstudio-pm-usage"
```

Using a single database with a separate schema for usage data:

```
; /etc/rstudio-pm/rstudio-pm.gcfg

[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/rstudio-pm"
UsageDataURL = "postgres://username:password@db.seed.co/rstudio-pm?search_path=metrics"
```

9.3 Usage Data

RStudio Package Manager relies on two databases by default. The primary database stores information needed to run the service including the arrangement of repositories, sources, and packages. Another database is used to record usage data like the number of times a package was downloaded.

If using SQLite, these two databases will be created automatically in the configured directory. If using PostgreSQL, you will need to define two different databases or schemas: `Postgres.URL` and `Postgres.UsageDataURL`.

If you do not wish to track usage data, you can disable this feature by setting `Server.UsageDataEnabled = false`. If disabled, usage data will not be tracked or displayed. You can use the `Server.UsageDataRetention` setting to alter the amount of usage data you wish to retain (the default is 365 days). Increasing this value will consume more disk space for the usage data database and may negatively impact performance slightly over time for busy servers.

10 Outbound Proxy

If you are syncing RStudio Package Manager to create your own CRAN repository (see 15.4), the server will need access to the internet to download manifest and package data. If you need to use an outbound proxy server, RStudio Package Manager uses the `HTTP_PROXY` and `HTTPS_PROXY` environment variables to configure HTTP or HTTPS proxy behavior.

Typically, it is best to define these environment variables in the startup script used to run RStudio Package Manager. This varies depending on the service daemon. Instructions for `systemd` and `Upstart` are listed below.

10.0.1 systemd (Red Hat/CentOS 7, SUSE 12, Ubuntu 16.04+)

1. Stop the RStudio Package Manager service. See 5.1.

```
sudo systemctl stop rstudio-pm
```

2. Define an `HTTP_PROXY` or `HTTPS_PROXY` environment variable in the service configuration file:

```
; /etc/systemd/system/rstudio-pm.service.d/user.conf

[Service]
Environment="HTTP_PROXY=http://username:password@proxy.example.com:8080"
Environment="HTTPS_PROXY=http://username:password@proxy.example.com:8080"
```

3. Reload the `systemd` configuration.

```
sudo systemctl daemon-reload
```

4. Start the RStudio Package Manager service. See 5.1.

```
sudo systemctl start rstudio-pm
```

10.0.2 Upstart (Ubuntu 14.04, Red Hat 6)

1. Stop the RStudio Package Manager service. See 5.1.

```
sudo stop rstudio-pm
```

2. Define an HTTP_PROXY or HTTPS_PROXY environment variable in the service configuration file:

```
# /etc/init/rstudio-pm.override  
  
env HTTP_PROXY=http://username:password@proxy.example.com:8080  
env HTTPS_PROXY=http://username:password@proxy.example.com:8080
```

3. Start the RStudio Package Manager service. See 5.1.

```
sudo start rstudio-pm
```

11 Running with a Proxy

If you are running RStudio Package Manager behind a proxy server, you need to be sure to configure the proxy server so that it correctly handles all traffic to and from RStudio Package Manager. This section describes how to correctly configure a reverse proxy with Nginx or Apache HTTPD.

When RStudio Package Manager is behind a proxy, it is important to send the original request URL information to RStudio Package Manager so that it can generate fully qualified URLs and return them the requester. For this reason, when proxying to RStudio Package Manager, we recommend adding a header, `X-RSPM-Request`, to the request. This header value should be the absolute URL of the original request made by the user or browser (i.e. `https://rspm.company.com/some/path`)

Some proxies (like Amazon Web Services Elastic Load Balancer), do not make it possible to add custom headers. Because of this, if this header is not supplied, “best efforts” are made utilizing the standard headers `X-Forwarded-Proto`, `X-Forwarded-Host`, and `X-Forwarded-Port` to parse the original request URL. If your proxy removes a server prefix from the path, `X-Forwarded` headers will not work for your use case, and you should use `X-RSPM-Request`. If both `X-RSPM-Request` and `X-Forwarded` headers are supplied, `X-RSPM-Request` takes precedence.

11.1 Using a Proxy for TLS/SSL

See 12.1.3 for information on configuring a proxy to handle HTTPS requests.

11.2 Nginx Configuration

On Ubuntu, a version of Nginx that supports reverse-proxying can be installed using the following command:

```
apt-get install nginx
```

On Red Hat/CentOS, you can install Nginx using the following command:

```
yum install nginx
```

On SUSE, you can install Nginx using the following command:

```
sudo zypper install nginx
```

To enable an instance of Nginx running on the same server to act as a front-end proxy to RStudio Package Manager you would use a configuration like the following in your `nginx.conf` file. This configuration assumes RStudio Package Manager is running on the same host as Nginx and listening for HTTP requests on the `:4242` port. If you are proxying to RStudio Package Manager on a different machine or port, replace the `localhost:4242` references with the correct address of the server where RStudio Package Manager is hosted.

```
http {
    map $http_upgrade $connection_upgrade {
        default upgrade;
        ''      close;
    }
    server {
        listen 80;
        location / {
            proxy_set_header    X-RSPM-Request $scheme://$host:$server_port$request_uri;
            proxy_pass http://localhost:4242;
        }
    }
}
```

If you want to serve RStudio Package Manager from a custom path (e.g. `/rspm`) you would edit your `nginx.conf` file as shown below:

```
http {
    map $http_upgrade $connection_upgrade {
        default upgrade;
        ''      close;
    }
    server {
        listen 80;
        location /rspm/ {
            rewrite ^/rspm/(.*)$ /$1 break;
            proxy_set_header    X-RSPM-Request $scheme://$host:$server_port$request_uri;
            proxy_pass http://localhost:4242;
            proxy_redirect / /rspm/;
        }
    }
}
```

After adding these entries you'll then need to reload Nginx so that the proxy settings take effect.

On `systemd` systems (Red Hat/CentOS 7, SUSE 12, Ubuntu 16.04+):

```
systemctl restart nginx
```

On Upstart systems (Ubuntu 14.04, Red Hat 6):

```
restart nginx
```


11.3 Apache Configuration

The Apache HTTPD server can act as a front-end proxy to RStudio Package Manager by first enabling three modules:

```
a2enmod rewrite
a2enmod headers
a2enmod proxy_http
```

The following configuration will permit proxying to RStudio Package Manager from the `:3737` port. Depending on the layout of your Apache installation, you may need the `Listen` and `VirtualHost` directives in different files.

```
Listen 3737
```

```
<VirtualHost *:3737>
    RewriteEngine on
    RequestHeader set X-RSPM-Request "%{REQUEST_SCHEME}s://%{HTTP_HOST}s%{REQUEST_URI}s"
    ProxyPass / http://172.17.0.1:4242/
    ProxyPassReverse / http://172.17.0.1:4242/
</VirtualHost>
```

You can serve RStudio Package Manager from a custom path (e.g. `/rspm`) with a configuration like the following:

```
Listen 3737
```

```
<VirtualHost *:3737>
    RewriteEngine on
    RedirectMatch ^/rspm$ /rspm/
    RequestHeader set X-RSPM-Request "%{REQUEST_SCHEME}s://%{HTTP_HOST}s%{REQUEST_URI}s"
    ProxyPass /rspm/ http://172.17.0.1:4242/
    ProxyPassReverse /rspm/ http://172.17.0.1:4242/
    Header edit Location ^/ /rspm/
</VirtualHost>
```

12 Security & Auditing

12.1 Browser Security

There are a variety of security settings that can be configured in RStudio Package Manager. Some of these settings are enabled by default but can be customized while others are opt-in. Below are some of the security features worth considering.

12.1.1 Guaranteeing HTTPS

If you can guarantee that your server should only ever be accessed over a TLS/SSL connection (HTTPS), then you can consider enabling the `HTTPS.Permanent` setting. This elevates the security of your server by requiring that future interactions between your users and this server must be encrypted.

Enabling this setting may keep users from being able to access your RStudio Package Manager instance if you later disable HTTPS or if your certificate expires. Use this setting only if you will permanently provide a valid TLS/SSL certificate on this server.

Behind the scenes, this makes two changes:

1. Introduces HTTP Strict Transport Security (HSTS) by adding a `Strict-Transport-Security` HTTP header with a `max-age` set to 30 days. HSTS ensures that your users' browsers will not trust a service hosted at this location unless it is protected with a trusted TLS/SSL certificate.
2. Enforces the `Secure` flag on cookies that are set. This prohibits your users' browsers from sending their RStudio Package Manager cookies to a server without an HTTPS-secured connection.

12.1.2 Strong HTTPS

Even with HTTPS, you may wish for more constraints on transport settings. Two settings are available to configure HTTPS differently:

1. `HTTPS.MinimumTLS`: This setting sets the minimum TLS version. By default, the minimum TLS version supported by RStudio Package Manager is TLS1.1. You can choose to set this configuration setting to "1.2" or back to "1.0"; before doing so, you should check the SSL Labs User Agent List to ensure your browsers will be compatible with the version you select.
2. `HTTPS.ExcludedCiphers`: This setting allows you to remove ciphers from the list of ciphers available. This may be useful if your organization has a security policy that disallows certain ciphers to be used.

An example exclusion list could be:

```
[HTTPS]
Listen = 443
MinimumTLS = 1.2
ExcludedCiphers = TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
ExcludedCiphers = TLS_RSA_WITH_3DES_EDE_CBC_SHA
ExcludedCiphers = TLS_RSA_WITH_AES_128_GCM_SHA256
ExcludedCiphers = TLS_RSA_WITH_AES_256_GCM_SHA384
ExcludedCiphers = TLS_RSA_WITH_AES_128_CBC_SHA256
ExcludedCiphers = TLS_RSA_WITH_AES_128_CBC_SHA
ExcludedCiphers = TLS_RSA_WITH_AES_256_CBC_SHA
ExcludedCiphers = TLS_RSA_WITH_3DES_EDE_CBC_SHA
```

The strings for the ciphers may be found in the IANA TLS Parameters List. For a list of supported ciphers on your platform that may be excluded, start up RStudio Package Manager and look for the log line beginning `Enabled HTTPS Ciphers:`.

12.1.3 Using a Secure Proxy

If you would prefer that the RStudio Package Manager process not have access to your TLS/SSL certificates, you may wish to configure a proxy to handle HTTPS requests. To accomplish this:

- Ensure that your `Server.Address` is set to the proxy address and uses the `https` scheme.
- Set `HTTP.ForceSecure` to `true`, which will set the `Secure` flag on all cookies.
- Set `HTTP.NoWarning` to `true` to suppress the warning regarding running RStudio Package Manager over an unsecured connection, because the connection between the client and the proxy will be secured.
- If necessary, enable the `HTTPRedirect.Listen` option to redirect proxied plain HTTP connections to HTTPS.

NOTE: Because the connection between the proxy and RStudio Package Manager is not secured in this case, please ensure that the proxy and RStudio Package Manager are connecting on a trusted network where an adversary would not be able to capture plain text credentials. For example, many cloud providers allow isolating servers from the internet while permitting load balancers to access them. Please see your cloud provider's documentation for more details.

12.1.4 Content Sniffing

The `Server.ContentTypeSniffing` setting can be used to configure the `X-Content-Type-Options` HTTP header. This protects your users from a certain class of malicious uploads and is enabled by default.

When disabled (the default), the `X-Content-Type-Options` HTTP header will be set to a value of `nosniff` to tell browsers not to sniff the content type. If enabled, no such header will be provided.

12.1.5 Content Embedding

The `X-Frame-Options` HTTP header is used to control what content can be embedded inside other content in a web browser. The relevant attack is commonly referred to as a “clickjack attack” and involves having your users interact with a sensitive service without their knowledge.

Some advertised values for this header are not supported across all browsers. RStudio Package Manager does not restrict the values of these headers.

12.1.6 Custom Headers

If you need to include additional HTTP headers that are not covered by any of the above features, you can include your own custom headers on all responses from RStudio Package Manager using the `Server.CustomHeader` setting.

This feature can be used to accommodate various other security practices that are not explicitly available as options elsewhere in RStudio Package Manager. For instance, X-XSS-Protection, Content Security Policy (CSP), HTTP Public Key Pinning (HPKP), and Cross-origin Resource Sharing (CORS) could all be configured using custom headers.

Custom headers are added to the HTTP response early during request processing. Values may later be overwritten or modified by other header settings. This includes both the security preferences described earlier in this chapter and other headers used internally by RStudio Package Manager. You should not depend on a custom header that conflicts with a header already in use by RStudio Package Manager.

The `Server.CustomHeader` takes a value of the header name and its value separated by a colon. Whitespace surrounding the header name and its value are trimmed. You can use this setting multiple times as in the following example:

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[Server]
```

```
CustomHeader = "HeaderA: some value"
```

```
CustomHeader = "HeaderB: another value"
```

12.2 Package Security

RStudio Package Manager allows administrators to decide what packages are allowed into an organization. For CRAN packages, RStudio Package Manager relies on the RStudio Package Service. This service is responsible for creating metadata about the daily changes on CRAN and providing package tar files from CRAN. The metadata and package files are reviewed and tested for consistency and accuracy. The service **does not** run any security checks. CRAN itself, prior to accepting new packages or packages updates, runs a series of checks that include package installation and compatibility tests.

In the event that a malicious CRAN package was identified, RStudio maintains a package security blog and RSS feed. This blog would be updated with remediation instructions for all RStudio Package Manager users.

13 High Availability and Load Balancing

Multiple instances of RStudio Package Manager can share the same data in highly available (HA) and load-balanced configurations. In this document, we refer to these configurations as “HA” for brevity.

13.1 HA Checklist

Follow the checklist below to configure multiple RStudio Package Manager instances for HA:

1. Ensure that all node clocks are synchronized 13.2
2. Ensure that all server configurations (i.e. contents of the `/etc/rstudio-pm` directory) are identical.
3. Install and Configure the same version of RStudio Package Manager on each node - 2
4. HA requires using a PostgreSQL database. All nodes in the cluster must use the same PostgreSQL database.
5. When using NFS for shared storage, configure each server’s `Server.DataDir` (7.6) to point to the same shared location. Be sure to read 13.3.3 for additional information on the recommended settings for the shared directory. For more granular control of data directories, see 7.7 for information on customizing the locations of each storage class.
6. When using S3 for shared storage, each server’s `Server.EncryptionKeyPath` must point to a file that contains the same encryption key.

13.2 HA Time Synchronization Requirements

The clocks on all nodes in an HA configuration must be synchronized. We recommend configuring NTP for clock synchronization.

13.3 HA Limitations

13.3.1 Node Management

RStudio Package Manager nodes in an HA configuration are not self-aware of HA. The load-balancing responsibility is fully assumed by your load balancer, and the load balancer is responsible for directing requests to specific nodes and checking whether nodes are available to accept request

13.3.2 Database Requirements

RStudio Package Manager only supports HA when using a PostgreSQL database.

13.3.3 Shared Data Directory Requirements

RStudio Package Manager manages repository content within the server’s data and variable data directories. These directories must be at shared locations, and each node must be configured to point to the same shared locations. See section 7.6 for more information on the server’s data directories.

RStudio Package Manager supports using either NFS or AWS S3 storage for shared data directories. You can also use a combination of both NFS and AWS S3 for different variable data classes.

13.3.3.1 NFS

We recommend and support NFS version 3 or 4 for file sharing.

RStudio Package Manager relies on being able to efficiently detect new files inside of the NFS-shared `DataDir`. By default, NFS clients are configured to cache responses for up to 60 seconds, which means that it can take up to a minute before an RStudio Package Manager service is able to respond to certain requests. For most deployments, this is an unacceptably long delay.

Therefore, we strongly recommend that you modify your NFS client settings for the mount on which you'll be hosting your `DataDir`. Typically, the best way to accomplish this is to set `lookupcache=pos` for your NFS mount, which will allow existing files to be cached but will contact the NFS server directly to check for the existence of new files. If this setting is not acceptable for your mount, you could alternatively consider shortening `acdirmax` or `actimeo` so that your client becomes aware of new files within, say, 5 seconds, instead of the default of 60.

13.3.3.2 S3

IMPORTANT: AWS S3 support is in beta. Please do not use S3 for production data at this time.

When using S3 for shared storage, each server's `Server.EncryptionKeyPath` must point to a file that contains the same encryption key. See also A.1. The easiest way to ensure a consistent encryption key on all nodes is to start RStudio Package Manager on one of the nodes and then copy the key file created at `/var/lib/rstudio-pm/rstudio-pm.key` to the same location on the other nodes. Set each key's file mode to 0600.

Please refer to Data Destinations for information on configuring RStudio Package Manager to store variable data on S3. For help configuring your server with the credentials and settings you need to interact with S3, see S3 Configuration

13.4 Updating HA Nodes

When applying updates to the RStudio Package Manager nodes in your HA configuration, you should follow these steps to avoid errors due to an inconsistent database schema:

1. Stop all RStudio Package Manager nodes in your cluster.
2. Upgrade one RStudio Package Manager node. The first update will upgrade the database schema (if necessary) and start RStudio Package Manager on that instance - 5.2.
3. Upgrade the remaining nodes.

If you forget to stop any RStudio Package Manager nodes while upgrading another node, these nodes will be using a binary that expects an earlier schema version, and will be subject to unexpected and potentially serious errors. These nodes will detect an out-of-date database schema within 30 seconds and shut down automatically.

13.5 Downgrading

If you wish to move from an HA environment to a single-node environment, please follow these steps:

1. Stop all RStudio Package Manager services on all nodes
2. Reconfigure your network to route traffic directly to one of the nodes, unless you wish to continue using a load balancer.
3. If you wish to move all shared file data to the node, then
 1. Configure the server's `Server.DataDir` to point to a location on the node, and copy all the data from the NFS share to this location - 7.6

4. If you wish to move the databases to this node, install PostgreSQL on the node and copy the data. Moving the PostgreSQL databases from one server to another is beyond the scope of this guide. Please note that we do not support migrating from PostgreSQL to SQLite.
5. Start the RStudio Package Manager process 5.1

14 Admin CLI

RStudio Package Manager is administered through a command-line interface (CLI). Administrators can use the CLI to create repositories and sources, add local packages to sources, and setup sync schedules for CRAN sources.

The CLI is installed at `/opt/rstudio-pm/bin/rspm`. The CLI uses the configuration defined in `/etc/rstudio-pm/rstudio-pm.gcfg` unless you specify an alternate configuration file with the `--config` flag.

14.1 Privileges

Users must be a member of the `rstudio-pm` group in order to use the RStudio Package Manager CLI. See C for instructions on changing the required group.

The RStudio Package Manager CLI uses a Unix domain socket for communicating with the RStudio Package Manager server. By default, the domain socket file is located at `/var/run/rstudio-pm/rstudio-pm.sock`. You can customize the location by configuring the `Server.SockFileDir` setting.

Any user invoking the RStudio Package Manager CLI must have read/write access to the Unix domain socket in order to communicate with the RStudio Package Manager server.

14.2 Location

The CLI is located at `/opt/rstudio-pm/bin/rspm`. We recommend that users make an alias, add this location to their PATH, or navigate to this directory if they will be running multiple commands in one session, for example:

```
alias rspm='/opt/rstudio-pm/bin/rspm'
```

14.3 Getting Help

Simply run `rspm` without any arguments to display the top-level help for the RStudio Package Manager CLI. For help with a specific command, you can use the `help` command. For example, to display help about the `create` command, you can use `rspm help create`.

14.4 Examples

For examples using the CLI, please refer to the Configuration guide.

15 Repositories and Sources

R repositories contain package tar files and are the primary vehicle for organizing and distributing R packages. For more information on packages and repositories see B.

In RStudio Package Manager, repositories are created from one or more sources. The documentation in this chapter outlines repositories as well as the types and structure of sources.

15.1 Repository Structure

R package repositories have a specific structure that enables client commands like `install.packages` to query the repository's contents and download packages.

A regular CRAN repository is just a set of files served from disk. RStudio Package Manager *does not create repositories on disk*. Instead, RStudio Package Manager maintains a single copy of each package source and uses a database and specialized web server to handle HTTP requests from R.

Some example requests that can be served by the RStudio Package Manager:

PACKAGES file

`http://pkg-manager.example.com/repo/latest/src/contrib/PACKAGES`

This serves a PACKAGES file. The PACKAGES file for a repository is human-readable and contains information on each package available in the repository. RStudio Package Manager can also serve requests for PACKAGES.gz and PACKAGES.rds.

Package Source

`http://pkg-manager.example.com/repo/latest/src/contrib/package_2.1.0.tar.gz`

This request downloads the package source to the client.

Archived Package Source

`http://pkg-manager.example.com/repo/latest/src/contrib/archive/package/package_1.1.0.tar.gz`

This request downloads the tar file for an older, archived version of the package.

Most importantly, a RStudio Package Manager repository **is a CRAN-like repository** which means users can access and install packages using their regular R functions: `install.packages`, `available.packages`, `packrat`, and `devtools::install`.

15.2 Repository Versioning

RStudio Package Manager tracks every change to a repository (or source) and associates each change with a checkpoint. Together, these checkpoints create a full versioned history of each repository. If a user wants to install packages from a prior point in the repository's history they can do so by replacing the `/latest` component of the request URL with a checkpoint id. Checkpoint ids can be obtained in a repository's "Activity" log. The current checkpoint id is available in the "Setup" page. **For projects that require strict reproducibility, we recommend configuring R to use a repository URL with a checkpoint id.** Versioning is available for all repository and source types.

15.3 Sources

15.3.1 About Sources

RStudio Package Manager repositories are composed of one or more **sources**. There are currently four types of sources:

1. *cran* source - A single cran source is automatically created. This source contains metadata and packages from RStudio's CRAN service. The source can be used directly in a repository to give users access to all CRAN packages, or it can be used indirectly by curated-cran sources.

While the cran source is created automatically, an administrator must use the CLI before any metadata or packages are downloaded to RStudio Package Manager. See the CLI section for more information on making CRAN available through RStudio Package Manager.

2. *curated-cran* source - A curated CRAN source allows administrators to specify specific sets of approved CRAN packages. Administrators can add or remove packages from the set, and they can also update the set. See 15.5 for more information.
3. *local* source - A local source is used as a mechanism to distribute locally developed packages or other packages without native support in RStudio Package Manager. Administrators add packages to local sources by specifying a path to a package’s tar file.
4. *git* source - A git source allows RStudio Package Manager to automatically make packages in Git available to R users through `install.packages` (without requiring `devtools`). Git sources work for internal packages as well as external sites such as GitHub. Packages can be automatically updated on each commit or when a new Git tag is pushed.

15.3.2 Repositories with Multiple Sources

A repository can have more than one source. If you wish to serve both local packages and CRAN packages from a single repository, you can create a single repository that subscribes to multiple sources. For example:

- public (a repository)
 - internal (local source)
 - cran (CRAN source)

The “public” repository above gives users access to both local and CRAN packages, and its `PACKAGES` list could be accessed, for example, at `http://pkg-manager.example.com/public/latest/src/contrib/PACKAGES`. A repository subscribes to sources, which means that changes to a source will be reflected in the repository. For example, if an admin adds a new package to the `internal` source, users will automatically be able to access the new package via the `public` repository.

15.3.3 Package Conflicts Between Sources

If a repository has multiple sources and a package with the same name exists in both sources, RStudio Package Manager eliminates duplicates, giving preference *in the order the sources are subscribed*. In the example repository above, if a package named “plumber” exists in both the “cran” and “internal” sources, the “plumber” package from the “internal” source would be served and listed since it is the first source for the repository. The same conflict resolution occurs as sources change. For example, in the sample above, even if a new package is added to CRAN with the same name as an internal package, the internal package will continue to be served. The precedence is also maintained during updates. In the example above, the internal version of plumber will continue to be served even if the CRAN version of plumber is updated. The order of sources within a repository can be re-arranged using the `reorder` command.

15.4 The CRAN Source

A primary use case for RStudio Package Manager is making packages in public repositories, like CRAN, available to enterprise users. Administrators can elect to make all of CRAN available, or to make only curated subsets of CRAN available.

15.4.1 What is RStudio’s Package Service?

RStudio Package Manager doesn’t download packages directly from CRAN. Instead, RStudio maintains a curated s3 bucket that contains metadata about CRAN and package tar files. The metadata is used to track

CRAN's day-to-day changes.

See D if your environment does not have access to the RStudio Package Service.

During a sync, the metadata is downloaded to RStudio Package Manager. The metadata is compared against the RStudio Package Manager database to determine what changes need to be applied. Package tarballs are then downloaded to the cache either eagerly or lazily depending on the sync mode.

Lazy sync mode is recommended.

Chapter 12.2 details the security measures in place for the RStudio Package Service.

15.4.2 Eager vs Lazy

The sync mode is configured using the `CRAN.SyncMode` property. The property defaults to “lazy”, but can be set to “eager” for eager package fetching. See A.1

15.4.2.1 Eager

For eager fetching, RSPM attempts to download any package sources that could be served to users. This means that RSPM will download packages under a number of scenarios:

1. When a repository subscribes to the CRAN source.
2. If a repository is subscribed to the CRAN source, then packages will be downloaded when a sync occurs, initiated with the CLI `sync` command or initiated based on the configured schedule.
3. When a package is added to a curated CRAN source with the `add` command (and at least one repository has subscribed to the curated CRAN source).
4. When a curated CRAN source is updated with the `update` command (and at least one repository has subscribed to the curated CRAN source).

All four of these actions happen automatically. Additionally, an admin can force RSPM to download packages using the CLI `fetch` command.

If a repository subscribes to the CRAN source, then all CRAN packages will be downloaded. Downloading all of CRAN during an initial sync can take a significant amount of time, bandwidth, and disk space. However, downloading packages **does not block end users** from accessing repositories that subscribe to the source. End users can request any package as soon as the metadata is sync'd. If a user requests a package that is not already downloaded, the package will be immediately downloaded and served to the client.

15.4.2.2 Lazy

If RStudio Package Manager is set up for lazy fetching, it downloads packages as the packages are requested by end users. Package Manager will still download the metadata from CRAN on the sync schedule to keep the RStudio Package Manager database updated. The database serves as the source of truth for package availability. The benefit of lazy fetching is a smaller footprint in terms of network bandwidth and disk space.

15.4.2.3 Package Caching

In either mode, each version of a package is only downloaded once. RStudio Package Manager always checks the local cache to see if the required tar file is already available.

15.4.2.4 Changing Modes

If you change the sync mode, the following occurs:

Lazy to Eager - Eager fetching is applied to future syncs. When you next synchronize, all packages that have not yet been downloaded will be downloaded.

Eager to Lazy - Lazy fetching is applied to all future syncs. A sync with in-progress downloads will be completed.

15.4.3 Updates from CRAN

The cran source is updated according to a schedule set using the `CRAN.SyncSchedule` property in the RStudio Package Manager configuration file. This property accepts a string in crontab format. See A.1.

By default, the configuration file includes a crontab that will cause RStudio Package Manager to sync once a day at midnight (in the server's timezone), **if any repository subscribes to “cran”, if a “curated-cran” source is used by any repository, or if a manual sync has been run with the sync command.** A sync schedule will **not** be applied if any of those 3 conditions are not met. **If you only want manual syncs**, change the configuration file to have a blank value for `CRAN.SyncSchedule`:

```
;/etc/rstudio-pm/rstudio-pm.gcfg
[CRAN]
SyncSchedule = ""
```

The `SyncSchedule` property does not necessarily determine when a repository will make updated packages available to users. If the repository subscribes directly to the cran source, users will see updates according to the sync schedule. In contrast, if the repository subscribes to a curated CRAN source, an administrator must explicitly update the source in order for updates to become available.

In addition, updating the repository does not automatically push updated packages to R clients. **A repository specifies what packages are available, but the R user is in control of when and how to update the packages used by a project.**

See the section on managing change control for more information.

RStudio Package Manager keeps track of old versions of packages as well. Old versions of packages are available in the repository's archive, and are listed in the RStudio Package Manager web UI. This allows users to roll back updates if necessary or install packages as they existed at a prior time.

15.5 Curated CRAN Sources

Curated CRAN sources allow administrators to create and update approved subsets of CRAN. The behavior is best explained in an example.

Assume that RStudio Package Manager has been configured to sync CRAN updates daily.

January 1st - An administrator creates a curated CRAN source and is given a list of desired packages.

January 2nd - The administrator can use the `add` command with the `dryrun` flag, supplying the list of desired packages. RStudio Package Manager will identify all of the required dependencies and create a proposal. The proposal includes the set of packages to be added as well as information about each package, such as license type. This information can be used to facilitate an external review process.

January 15th - The proposal is approved. The administrator returns to RStudio Package Manager and runs the `add` command again, replacing the `dryrun` flag with a transaction ID included in the proposal. **The set of packages is added from CRAN as they existed on January 1st, the date the source was created.**

January 20th - The administrator receives a request to add a new package to the set of approved packages. The admin uses the `add` command with the `dryrun` flag, supplying the new package as an argument. **RStudio Package Manager will create a proposal using the version of CRAN as it existed on January 1st.** In order to ensure compatibility between the packages added to the source, RStudio Package Manager will add to the set of packages by pulling from CRAN as it existed the day the source was created. As before, if the proposal is accepted, the admin can commit the changes.

January 30th - Now the administrator gets a request to update the approved packages. In order to keep all packages consistent, the entire set is updated at once using the `update` command. Like the `add` command, the `update` command supports a `dryrun` flag that will enumerate all the changes needed to update the set of packages from January 1st to January 30th.

February 1st - The proposal is approved and the administrator completes the `update` command by using the transaction ID included in the dry run update. **The set of packages is now tied to CRAN on January 30th.** Future `add` commands will use this pinned date, until another update sequence occurs.

To summarize, curated CRAN sources allow admins to create a subset of CRAN at a point in time. Administrators can add packages to the subset from the same frozen point in time. Administrators can also update the subset to a newer point in time. Each change supports a dry run that creates a proposal and a confirmation run that applies the proposal.

Given a list of desired packages, RStudio Package Manager automatically determines the full set of dependencies and also tracks those dependencies over time. Admins can elect to include suggested dependencies or only required dependencies by using the `include-suggests` flag. During each update, older versions of packages are archived, ensuring that tools like packrat and RStudio Connect work seamlessly with the curated CRAN subset.

The `update` command will be impacted by the sync schedule defined on the server. If the server only syncs every few weeks, `update` will only reference the latest data from CRAN **available on the server**.

15.6 Git Sources

Git sources allow RStudio Package Manager to automatically expose R packages tracked in Git. Git sources work with internal packages as well as external sites such as GitHub.

Git sources require a configured R installation.

An admin follows these steps:

1. Create a git source
2. Add a Git endpoint to the source, specifying whether to watch for commits to a branch or tags. The endpoint can be HTTP or SSH (see below). See the `add` command for full details, e.g. how to track a specific branch.
3. Based on the selection specified with the `add` command, RStudio Package Manager clones the Git endpoint and runs an R job to transform the Git clone into a package bundle. The package bundle is made available to any repositories subscribing to the source.
4. RStudio Package Manager polls the Git endpoint to watch for either new commits or new tags (based on the selection specified with the `add` command). If an update is available, RStudio Package Manager automatically pulls the new changes and launches an R job. The R job creates a package bundle from the updated Git clone and updates the package available in the git source. Previous versions are archived.
5. Users install the package from the repository via `install.packages` NOT `devtools`.

See the quickstart guide for a specific example.

15.6.1 Access restricted Git endpoints using SSH keys

Many internal Git endpoints require authentication. RStudio Package Manager can be configured to use SSH keys to authenticate against the endpoint.

Begin by creating an SSH key and granting the SSH key access to the Git endpoint. The specific steps will depend on your Git provider. Once you have the path to the SSH key, use the `import` command to securely name and store the SSH key for later use by RStudio Package Manager. If desired, you can now remove the SSH key file. Multiple keys can be imported.

To use the newly imported SSH key, use the SSH identifier for the Git endpoint in the `add` command and reference the name of the key with the `--git-ssh-key` flag.

15.6.1.1 SSH key Security

RStudio Package Manager encrypts and stores imported SSH keys in the metadata database. Any person (by default, members of the `rstudio-pm` unix group) with access to the admin CLI can:

- Associate an imported key with a Git endpoint using the `add` command
- List the names of available SSH keys

Users **cannot** access the contents of the key, nor is the key available for arbitrary actions. We recommend granting SSH keys imported to RStudio Package Manager limited read-only access to only the endpoints you wish to expose as R packages.

When imported, the keys are encrypted at rest, during Git operations which require SSH, the keys are added to an ssh-agent and thus never written to the filesystem or written to STDIN.

Although RStudio Package Manager allows the use of SSH keys with no passphrase, it is still recommended to use a strong SSH key with a passphrase.

15.6.2 Commits vs Tags

A package based on a Git endpoint can be configured to watch one of two types of changes:

1. *Commits* - RStudio Package Manager will update the package any time new commits are discovered in a branch. In this mode, RStudio Package Manager *automatically modifies the package's version*, assigning a unique version number to each build. The version number is created based on the commit time-stamp and is designed to avoid conflicts with the version scheme used by the package author. For example, if the Description file for a package indicates a version of `1.1-3`, the automatic version number would be: `1.1-3.0.0.0.1537204599`. If the author updates the package with a new commit, but keeps the version in the Description file the same, the new automatic version number would reflect the new commit time-stamp, e.g. `1.1-3.0.0.0.1537218677`. This process ensures that users of the package always get the correct behavior from `install.packages`, with newer commits being associated with a semantically higher version number.
2. *Tags* - RStudio Package Manager will update the package any time a new Git tag is discovered. In this mode, RStudio Package Manager retains the version specified in the package's Description file. This mode is designed to work when a Git tag is used to indicate a package release. *Note: The name of the tag must match the version in the Description file.* For example, if your package's Description file has `Version: 5.4.2` your tag must be either `5.4.2` or `v5.4.2`. If two tags reference the same version, preference is given to the newer tag. If a newer tag references an older version than a prior tag, the new tag is built as an archived package. If a tag is removed from a Git endpoint, the package is deleted.

Commit mode is recommended for bleeding edge repositories, whereas tag mode is suitable for exposing stable releases of packages.

A git source can support different packages with different modes. However, a given package can only have one mode in a source. If you would like to surface the same package in both commit and tag mode, you must create two git sources.

15.6.3 Git directories

By default, packages will be built from the git root directory. If the R package exists in a different location, it can be specified using the `--git-sub-dir` flag when adding a git package.

15.6.4 Managing Packages from Git

RStudio Package Manager automatically handles updating and archiving packages in git sources as the Git endpoints change. Additionally, the package artifacts themselves can be manually removed using the `remove` command.

Deleting a git source with the `delete` command will remove all the packages generated by the git source and will remove all the metadata about the Git endpoints.

Finally, it is possible to keep the package artifacts already created but stop RStudio Package Manager from tracking the Git endpoint. To do so, use:

```
rspm delete git-builder --name=[name of package] --source=[name of source]
```

To view information about the current Git endpoints that are being tracked, use:

```
rspm list git-builders
```

A `git-builder` is automatically created when a package is added to a Git source using the `add` command. Git builders do not need to be created manually.

15.6.5 Combining packages from Git(Hub) with other package sources

Local packages cannot be added manually to a git source, but a repository can surface packages from a git source alongside local packages and CRAN packages by subscribing to multiple sources. Take care when managing a repository's subscriptions as order is important, see 15.3.2.

15.6.6 Polling Frequency

You can control how frequently RStudio Package Manager checks for updates using the `Git.PollInterval` configuration field. If multiple commits occur between checks, RStudio Package Manager will create a single version representing all of the changes. If multiple tags are created or removed between checks, RStudio Package Manager will build each tag individually, automatically archiving tags representing older versions of the package.

Repository Versioning is identical in all source types, including git sources.

15.6.7 Tracking Changes and Errors

If a repository subscribes to a git source, you can view the git source's history in the Activity Log. The Activity Log will identify each change to a package including the new version, and a message will indicate the associated Git tag or commit as appropriate. If an error is encountered attempting to clone, poll, or bundle a package, the Activity Log will record the attempt and include a message with the CLI command to be run to view a full error log.

RStudio Package Manager automatically tries to build updates from a Git source 3 times. If the build fails more than 3 times, the update causing the failure is ignored. New updates are still discovered and built.

To force RStudio Package Manager to retry building an update, use the `rerun` command:

```
rspm rerun git-builder \
  --name=[package name] \
  --source=[source name] \
  --tag=[tag to rebuild, only required if the build trigger is tags]
```

16 Building R Packages

R packages can come in multiple formats:

- Source: A collection of directories and files containing source code.
- Bundle: A specially created tar file containing bundled source code. The result of `R CMD build`.
- Binary: A binary file specific to an operating system and architecture, containing compiled source code. Not an executable. The result of `R CMD INSTALL`.

In some configurations, RStudio Package Manager will need to run R in order to transform a package from one state to another. RStudio Package Manager uses the RStudio Job Launcher to perform this task.

Currently, the Job Launcher and R installation are only relevant if you use git sources.

16.1 Job Launcher

The RStudio Job Launcher is a service responsible for running jobs in support of RStudio Package Manager. The Job Launcher is automatically installed when RStudio Package Manager is installed. Settings for the Job Launcher can be configured directly in the RStudio Package Manager configuration file, `/etc/rstudio-pm/rstudio-pm.gcfg`, see A.

RStudio Package Manager can only use its own instance of the Job Launcher and runs all tasks locally. Other RStudio products may have separate Job Launcher instances managed independently from RStudio Package Manager. The RStudio Package Manager's Job Launcher cannot be shared by other products, nor can RStudio Package Manager use a Job Launcher instance set up for another product.

16.2 R Configuration

The Job Launcher requires a valid installation of R. Any version of R greater than 3.0.0 is sufficient. Specify the installation of R in the RStudio Package Manager configuration file:

```
; /etc/rstudio-pm/rstudio-pm.gcfg
[Server]
RVersion = /opt/R/R-3.4.1
```

Replace `/opt/R/R-3.4.1` with the path to your R installation. Use the path to the R installation directory, not the path to the binary (do not include `/bin/R`). While multiple versions of R may be installed on the server, **only one version of R** may be specified for use by RStudio Package Manager.

The `RVersion` configuration option and an installation of R *is not required*, however if the `RVersion` field is included it must be valid and must only appear once in the configuration file. Check the server log after starting and stopping for messages relevant to the R configuration.

16.3 R Installation

We highly recommend installing R from source. Installing R from source makes it is easy to manage multiple versions of R, prevents accidental upgrades of R, and ensures R is built with your architecture. The RStudio Server documentation contains instructions for building and installing R from source.

This RStudio Support article also includes R installation instructions.

16.4 Process Management

See 17 for information on how RStudio Package Manager securely runs R processes when building R packages for git sources.

17 Process Management

RStudio Package Manager executes R to build packages for Git sources. Use the `Server.RVersion` setting to specify the R installation. See Chapter 16 for details on installing and configuring R.

Since R processes can execute arbitrary code and commands, RStudio Package Manager prefers to run R in a sandbox.

17.1 Quick Start

This section walks you through the scenarios you are likely to encounter and helps you quickly decide how to configure RStudio Package Manager for building R packages from Git sources.

17.1.1 No Git Sources

If you are not using Git sources, you do not need to worry about sandboxing.

17.1.2 For Git Sources

If you are using Git sources, then:

17.1.2.1 Built-in OS Support

If your operating system is configured for user namespace sandboxing, you will see a message like this in `/var/log/rstudio-pm.log` at startup:

```
Attempting to detect R version in a user namespace sandbox.  
Detected R version <version> at /path/to/R
```

If you see the above messages, no further action is required.

17.1.2.2 No built-in OS Support

If your operating system does not support user namespaces, you will see a message like this in `/var/log/rstudio-pm.log` at startup:

```
Warning: Unable to use R on this system. Building Git packages will be disabled.
```

If you see the above message, choose one of the following:

- If you are using **CentOS / RedHat 7**, enable user name spaces by following the instructions at 17.4.3,
- Change RStudio Package Manager to run as `root` by following the instructions at C, *or*
- Use the configuration setting `Git.AllowUnsandboxedGitBuilds = true` to run R processes without a sandbox.

If you choose to enable user namespaces on CentOS / RedHat 7, you should see this message at startup:

```
Attempting to detect R version in a user namespace sandbox.  
Detected R version <version> at /path/to/R
```

If you choose to run RStudio Package Manager as `root` or use the `Git.AllowUnsandboxedGitBuilds` configuration setting, then you should see this message at startup:

```
Attempting to detect R version in a call to rsandbox
Detected R version <version> at /path/to/R
```

17.2 Sandboxing

RStudio Package Manager prefers to run R in a sandbox. RStudio Package Manager can run R processes in three different environments:

1. **User Namespace Sandbox** - When RStudio Package Manager is running under an unprivileged service account (by default, the `rstudio-pm` user), it attempts to run R in a user namespace.
2. **Unshare Sandbox** - When RStudio Package Manager is running as the `root` user, it runs R processes as an unprivileged user and uses `unshare` to create sandboxes for the R processes.
3. **No Sandbox** - If neither user namespace or `unshare` sandboxing is available, you can configure `Git.AllowUnsandboxedGitBuilds = true` to allow R processes to run without a sandbox. This is insecure, should only be used when you implicitly trust the sources of all R packages that will be built.

The three environments above are described in more detail below.

17.2.1 User Namespace Sandbox (the default)

User namespaces are supported as follows. See 17.4 for details.

- **Fully Supported** on Ubuntu, SUSE 12
- **Supported with System Configuration** on CentOS/Red Hat 7
- **Unsupported** on CentOS/Red Hat 6
- **Varies by Kernel** in Docker containers

The RStudio Package Manager process runs as the `rstudio-pm` user and runs R securely in a new user namespace. When using a user namespace, parts of the file system are hidden from the R process, and certain OS features are isolated in new namespaces. For details on the file system isolation, see 17.2.4.

You can learn more about cloning namespaces here, and you can learn about the Linux user namespace here. Your local man pages will document their behavior specific to your system.

See 17.2.4 for information on the bind mounts that are used.

17.2.2 Unshare Sandbox

When the RStudio Package Manager process runs as the `root` user, R processes are securely invoked as an unprivileged user, `rstudio-pm` by default.

The “unshare” environment created for R execution involves first establishing a number of bind mounts and then switching to the target unprivileged user. The RStudio Job Launcher uses `unshare` to alter the execution context available to R processes. Within this newly established environment, a number of `mount` calls are made in order to hide or isolate parts of the filesystem.

To configure the RStudio Package Manager process to run as `root`, please see C.

See 17.2.4 for information on the bind mounts that are used.

You can learn more about `unshare` here. The `mount` call is detailed here. Your local man pages will document their behavior specific to your system.

17.2.3 No Sandbox

If your operating system does not support user namespaces (see 17.4) and you cannot run the RStudio Package Manager service as `root`, then you cannot isolate R processes from the rest of the system. However, you can configure the `Git.AllowUnsandboxedGitBuilds = true` setting to allow R processes to run without a sandbox. This can be a security risk, and you should only use the `Git.AllowUnsandboxedGitBuilds = true` setting if you trust the sources of the R packages that will be tracked from Git. The setting does not impact other sources such as CRAN packages.

17.2.4 Bind Mounts

When sandboxing is supported, the following locations are masked during R execution:

- The `Server.DataDir` directory containing all variable data used by RStudio Package Manager.
- The `SQLite.Dir` directory, which can optionally be placed outside the data directory.
- The `Git.BuilderDir` directory, which can optionally be placed outside the data directory
- Configuration directories, including `/etc/rstudio-pm`.
- The `Server.SockFileDir`, which contains the domain socket used by the `rspm` CLI to communicate with the RStudio Package Manager service. This directory defaults to `/var/run/rstudio-pm`.
- The `bin` directory (usually `/opt/rstudio-pm/bin`) that contains the RStudio Package Manager binaries.
- Any persistent storage directories that are defined outside of the `Server.DataDir` directory. These locations are configured with `FileStorage.Location` configuration settings, and are not used by default.
- The `Server.TempDir/rspm-workspaces` directory, which contains temporary directories, one per process.

The following information is exposed during R execution:

- The Git repository source directory when building an R package from a Git source.
- The R package build destination directory when building an R package for a Git source.
- A per-process temporary directory specified in the `TMPDIR` environment variable of the process. This temporary directory is created under `Server.TempDir/rspm-workspaces`.

Running R processes have write access to the directory containing the unpackaged R code. Data written here will be visible to all processes running under the RStudio Package Manager service account, but are not visible to other R processes started by RStudio Package Manager when these processes run in a sandbox.

17.3 Temporary Directory

Each R process started by RStudio Package Manager is given its own unique temporary directory. These directories are created under `Server.TempDir/rspm-workspaces`.

`Server.TempDir`'s default value is obtained by first checking the `TMPDIR` environment variable for a path and falls back to `/tmp` otherwise.

You may wish to override `Server.TempDir` if the default temporary directory has too little space or is mounted with the `noexec` option.

Note: If you do override `Server.TempDir` please ensure the location can be reached by, read from, and written to by any user on the system. On most systems, temporary directories typically have permissions of `1777`.

You can learn more about the `noexec` option here.

17.4 User Namespace Support

17.4.1 Ubuntu, SUSE 12

User namespace sandboxing is fully supported by Ubuntu and SUSE 12.

17.4.2 CentOS 6, Red Hat 6

CentOS/Red Hat 6 does not support user namespaces. If you wish to securely run R processes on CentOS/Red Hat 6, consider running RStudio Package Manager as `root`. See 17.2.2 for details.

17.4.3 CentOS 7, Red Hat 7

CentOS/Red Hat 7 supports user namespaces, but they must be enabled. To see if user namespace support is enabled, run `cat /proc/sys/user/max_user_namespaces`. If a value greater than 0 (zero) is returned, user namespace support is enabled.

To enable user namespace support permanently on CentOS/Red Hat 7, edit `/etc/sysctl.d/files.conf` with a change similar to the following. This change will be effective when the server is rebooted.

```
# /etc/sysctl.d/files.conf
user.max_user_namespaces=15000
```

You can temporarily enable user namespace support by simply echoing a new value to `/proc/sys/user/max_user_namespaces`. For example:

```
echo 1000 > /proc/sys/user/max_user_namespaces
```

17.4.4 Docker

User namespace support is a feature of Linux kernel version 3.8 and greater. If you are running RStudio Package Manager in a Docker environment, user namespace support will depend upon the kernel of the host operating system.

Some of the sandbox operations, like bind mounting, require capabilities that are only available in privileged Docker processes. You must start the container with the `--privileged` option if you wish to use R process sandboxing in a Docker environment.

If you prefer to run without the `--privileged` flag, you can do so by setting `Git.AllowUnsandboxedGitBuilds = true`.

17.5 User Account for Processes

The RStudio Package Manager installation creates a local `rstudio-pm` user account. This account runs all R processes even if you have reconfigured RStudio Package Manager to run as `root`. If you would like a different user to run R processes, customize the `Launcher.ServerUser` property. If you reconfigure RStudio Package Manager to run under a different unprivileged account, then R processes will automatically run under that account, as well.

18 Configuring RStudio Server

Administrators can configure RStudio Server (Pro) to automatically use RStudio Package Manager when users install packages.

The necessary configuration files are documented in the RStudio Server admin guide, example configurations are provided here for common scenarios.

18.1 A Single Repository

In the most common scenario, users will install all packages from a single RStudio Package Manager repository that may contain CRAN packages and internal packages. An admin can discourage users from changing this repository setting. In this scenario, configure RStudio Server:

1. In `/etc/rstudio/rsession.conf`, set:

```
r-cran-repos=https://<your-package-manager-server-address>/<repo-name>/latest
```

The exact URL to use is available in the RStudio Package Manager Setup page for the repository.

2. (RStudio Server Pro only) To disable the repository option menu and discourage users from changing the repository setting, also include in `/etc/rstudio/rsession.conf`:

```
allow-r-cran-repos-edit=0
```

18.2 Internal Packages and CRAN Packages

Another common scenario is to setup two repositories in RStudio Package Manager, one to serve CRAN packages and one to serve internal packages. (This setup also applies if you're only using RStudio Package Manager for internal packages and want to allow users to access a public CRAN mirror).

1. In `/etc/rstudio/rsession.conf`, remove any existing `r-cran-repos` configuration.
2. Create a file called `/etc/rstudio/repos.conf`. The file should contain:

```
CRAN=https://<your-package-manager-server-address>/<cran-repo-name>/latest  
Internal=https://<your-package-manager-server-address>/<internal-repo-name>/latest
```

The name `Internal` can be replaced with whatever name makes sense for your organization, but the repo containing CRAN packages should be indexed with the keyword `CRAN`.

18.3 Allow Users to Optionally Add Additional Repos

In addition to the two scenarios described above, some organizations may have additional repositories for certain groups of users, e.g. “bleeding edge” packages available in a dedicated “dev” repository. Administrators can configure RStudio’s `Global Options` menu to automatically offer RStudio Package Manager repositories as optional, secondary repos that a user can opt into using.

1. In `/etc/rstudio/rsession.conf`, set:

```
r-cran-repos-url=https://<your-package-manager-server-address>/__api__/repos  
allow-r-cran-repos-edit=1
```

18.4 Precedence of Settings

If you have an `Rprofile.site` file that modifies the repository setting, that file **will take precedence** over these settings.

Some Ubuntu and Debian R binaries (installed via `apt-get install`) include an `Rprofile.site` file that modifies the repository option. Unless you intended to have this file, we recommend removing `/usr/lib/R/etc/Rprofile.site`.

19 Managing Change Control

An important concern in any organization adopting R is code stability. Tooling is required to meet the needs of R users which include access to the latest packages and reproducible environments. In regulated environments, an additional level of control may be required to ensure certain teams or projects only use groups of packages that have undergone additional vetting. Administrators also need to worry about the availability of different versions of R and system dependencies that R packages may rely on.

There is not a single solution that meets the needs of every organization. RStudio Package Manager is designed to support a number of strategies. Three common strategies are outlined below.

19.1 Approach 1: Client Side Management

The most common approach to managing change control is for the user to manage package dependencies for a project. An easy way R users can do this is by creating project-specific libraries of packages. In this model, administrators specify what packages are available in the repository. Users create a specific library for each project and are responsible for installing the packages required for their project. Over time, users can upgrade packages or add new packages.

For reproducibility, users must maintain a record of which package versions are included in their libraries. This record allows the user to recreate the environment from a clean slate by requesting the specific version of each package from the repository. There are a number of R packages designed to help users accomplish this task, including the `packrat` package. RStudio Connect automates this approach for users when they deploy content.

This model relies on the ability of the repository to handle multiple versions of each package. R repositories handle this task using a specific structure called an archive. RStudio Package Manager automatically handles archiving packages and knows how to respond to requests for older packages from tools like `packrat` and `devtools`.

19.2 Approach 2: Repository Versioning

Another common approach is for users to pin a project to a specific version of the package repository. In contrast to the first approach, this approach does not rely on the repository managing an archive of package versions. Instead, the entire repository is versioned. Like the first approach, this strategy relies on the user creating a specific package library for each project. However, instead of recording the versions of each package in use, the user simply records the version of the repository and the names of the necessary packages. To recreate the package environment, the user simply has to install the requisite packages from the recorded version of the repository. The most popular example of this approach is the `checkpoint` package which relies on Microsoft's online copy of CRAN.

RStudio Package Manager supports this strategy by automatically associating every action in a repository with a checkpoint id. See 15.2 for details. Users can get started by configuring the Repository URL in the "Setup" page to always use the current packages.

19.3 Approach 3: Locked Down Repositories

Some organizations prefer to manage which versions of packages are used across the organization instead of on a project-by-project basis. A common strategy entails:

1. Administrators test a set of desired packages and then freeze that set. Users are able to use only those versions of the packages.
2. A few times per year, administrators test and update the set of packages. They might also approve new packages.

RStudio Package Manager enables this strategy through curated CRAN sources. RStudio Package Manager also enhances this strategy by adding the option for a step between 1 and 2. When an administrator creates a curated CRAN source and defines a set of packages (step 1), RStudio Package Manager also records the state of all of CRAN at that moment. This enables an administrator to later add packages to the approved subset without worrying about updating the entire set.

Curated CRAN sources can be used with any of these strategies if organizations want to apply additional governance policies such as limiting packages to those with approved licenses.

19.4 What about versions of R?

In order to capture and recreate the environment for an R project, organizations must account for managing the version of R in addition to managing R packages.

In general, a library of installed packages is only compatible with a single version of R. For a new version of R, users or administrators must re-install the desired set of R packages. Luckily, managing a repository of packages facilitates re-installing packages into different libraries.

19.5 What about Docker?

Docker can play an important role in creating reproducible environments. Docker specifies the steps for creating an environment in a Dockerfile. A Docker image is created by running each of those steps.

Normally, a Dockerfile for an R project will include one or more steps that install R packages:

```
RUN Rscript -e 'install.packages(...)'
```

Using Docker in this way facilitates reproducibility, but **Docker alone is not sufficient to guarantee reproducibility**. The reason Docker is insufficient is because each time the Docker image is recreated the R command to install packages is re-run. Just like a Dockerless environment, this command can return different results over time.

Luckily, the same approaches outlined above also work with Docker. Simply replace the `install.packages` command with a variation that uses a frozen repo or restores a specific environment using a tool like `packrat`.

A Configuration Options

This appendix documents the RStudio Package Manager configuration file format and enumerates the user-configurable options.

The RStudio Package Manager configuration file is located at `/etc/rstudio-pm/rstudio-pm.gcfg`. This configuration is read at startup and controls the operation of the service.

The RStudio Package Manager configuration file uses the **gcfg** (Go Config) format, which is derived from the Git Config format.

Here is an example of that format showing the different property types:

```
; Comment
[BooleanExamples]
property1 = true
property2 = off
property3 = 1

[IntegerExamples]
Property1 = 42
Property2 = -123

[DecimalExamples]
Property1 = 3.14
Property2 = 7.
Property3 = 2
Property4 = .217

[StringExamples]
Property1 = simple
Property2 = "quoted string"
Property3 = "escaped \"quote\" string"

[MultiStringExamples]
ListProperty = black
ListProperty = blue
ListProperty = green

[DurationExamples]
Property1 = 1000000000
Property2 = 500ms
Property3 = 1m15s ; comment with a property
```

Comments always start with a semi-colon (;) and continue to the end of the line. Comments can be on lines by themselves or on a line with a property or section definition.

Configuration sections always begin with the name of the section bounded by square brackets. A section may appear multiple times and are additive with the last value for any property being retained. The following two configuration examples are equivalent.

```
[Example]
A = aligator
B = 2
```

```
[Example]
A = aardvark
C = shining
```

```
[Example]
A = aardvark
B = 2
C = shining
```

Each configuration property must be included in its appropriate section. Property and section names are interpreted case-insensitively.

Property definitions always have the form:

```
Name = value
```

The equals sign (=) is mandatory.

If a property happens to be given more than once, only the last value is retained. The “multi” properties are an exception to this rule; multiple entries are aggregated into a list.

```
[MultiExample]
Color = black
Color = blue

[NonMulti]
Animal = cat
Animal = dog
```

If `Color` is a multi-string property, both the “black” and “blue” values are used. If `Animal` is a normal string property, only the value “dog” is retained.

Configuration properties all have one of the following types:

string A sequence of characters. The value is taken as all characters from the first non-whitespace character after equal sign to the last non-whitespace character before the end-of-line or start of a comment. Double-quotes (") are supported, but usually unnecessary. A literal double-quote MUST be escaped and quoted itself like `QuotedValue = "J.R. \"Bob\" Dobbs"`.

multi-string A property that takes multiple string values. The property name is listed with each individual input value. For example, providing `Color = black` and `Color = blue` results in two separate values.

boolean A truth value. The values `true`, `yes`, `on`, and `1` are interpreted as true. The values `false`, `no`, `off`, and `0` are interpreted as false.

integer An integral value.

decimal A numeric value with an optional fractional component. Values with and without a decimal point are allowed.

duration A value specifying a length of time. When provided as a raw number, the value is interpreted as nanoseconds. Duration values can also be specified as a sequence of decimal numbers, each with optional fraction and unit suffix, such as `300ms`, `1.5h`, or `1m30s`.

Valid time units are `ns` (nanoseconds), `us` (microseconds), `ms` (milliseconds), `s` (seconds), `m` (minutes), and `h` (hours).

version A string representing a version. A version may have one to four numeric components, separated by periods or hyphens. Examples include `2`, `2.5`, `2.5.6`, `2.5.6.1`, and `2.5-6-11`.

bytesize A string representing a size in bytes. Valid examples include `10MB`, `1 tb`, `3 terabytes`, and `2 GB`.

Each configuration property documented in this appendix includes its description, data type, and default value.

Some properties are marked as “reloadable”. Sending a HUP signal to the Package Manager process causes the on-disk configuration to be reread. The server is reconfigured with the latest values of these reloadable properties. See 5.1 for details about sending a HUP signal to your Package Manager process.

Use a HUP signal when your configuration changes are limited to properties marked as reloadable. Perform a full restart of RStudio Package Manager when changing other properties.

A.1 Server

The **Server** section contains configuration properties which apply across the whole of RStudio Package Manager and are not appropriate for the other sections, which are generally narrower.

The properties which follow all must appear after `[Server]` in the configuration file.

DataDir The directory where RStudio Package Manager will store its variable data.

Type: string

Default: `/var/lib/rstudio-pm`

TempDir The directory that will contain all temporary directories needed by R processes. May not be a child path of `Server.DataDir`. If `TMPDIR` environment variable is not defined then `/tmp`.

Type: string

Default: `$TMPDIR`

CacheTimeout The amount of time we wait for files to appear in the cache directory after the database reports they are available. The default of 65 seconds should accommodate the NFS default directory caching limit of 60 seconds

Type: duration

Default: `65s`

RVersion Path to an R installation root.

Type: multi-string

Default: *unspecified*

Address A public URL for this RStudio Package Manager server.

Type: string

Default: *<empty-string>*

ContentTypeSniffing If disabled, sets the `X-Content-Type-Options` HTTP header to `nosniff`. When enabled, removes that header, allowing browsers to mime-sniff responses.

Type: boolean

Default: `false`

ServerName By default, Package Manager sets the `Server` HTTP header to something like `RStudio Package Manager v1.2.3`. This setting allows you to override that value.

Type: string

Default: *<empty-string>*

AccessLog Path to the file that RStudio Package Manager will use for its access logs. Disabled when empty.

Type: string

Default: `/var/log/rstudio-pm.access.log`

CustomHeader Custom HTTP header that should be added to responses from Package Manager in the format of `key: value`. The left side of the first colon in the string will become the header name; everything after the first colon will be the header value. Both will be trimmed of leading/trailing whitespace. This will always add a new header with the specified value; it will never override a header that Package Manager would otherwise have set. Multiple definitions can be used to provide multiple custom headers.

Type: multi-string

Default: *unspecified*

FrameOptionsUI The value for the `X-Frame-Options` HTTP header for the Package Manager UI and all other Package Manager pages. If empty, no header will be added.

Type: string

Default: DENY

HideVersion When `true`, the `/version` API returns a blank string and the product version is not displayed in the UI.

Type: boolean

Default: `false`

Reloadable: true

SockFileDir RStudio Package Manager will use this directory to create a socket file for admin connections.

Type: string

Default: `/var/run/rstudio-pm`

UsageDataEnabled If `true`, will enable tracking of package downloads.

Type: boolean

Default: `true`

UsageDataRetention The number of days for which usage data will be retained. If 0, data will be preserved forever.

Type: integer

Default: 365

UsageDataSchedule The schedule for usage data aggregation. Must be a valid crontab string.

Type: string

Default: `0 0 * * *`

Reloadable: true

EncryptionKeyPath The filepath of the encryption key used to store sensitive data at rest. Must be 512 bits (64 bytes) and file mode must be 0600. For cluster environments every server must have a copy of the same key file.

Type: string

Default: `{Server.DataDir}/rstudio-pm.key`

A.2 API

The API section contains configuration properties which apply to the RStudio Package Manager API.

The properties which follow all must appear after `[API]` in the configuration file.

MaxApiResults Limits the maximum number of results that can be returned from API endpoints that support a `_limit` parameter. When zero, there is no limit. Please consider performance implications before raising this limit.

Type: integer

Default: 1000

DefaultApiResultsLimit The default limit on the number of results returned from API endpoints that support a `_limit` parameter. When zero, there is no limit.

Type: integer

Default: 100

A.3 HTTP

The HTTP section contains configuration properties which control the ability of RStudio Package Manager to listen for HTTP requests. RStudio Package Manager must be configured to listen for either HTTP or HTTPS requests (allowing both is acceptable).

These properties must appear after [HTTP] in the configuration file.

Listen RStudio Package Manager will listen on this network address for HTTP connections. The network address can be of the form `:80` or `192.168.0.1:80`. Either `HTTP.Listen` or `HTTPS.Listen` is required.

Type: string

Default: *<empty-string>*

NoWarning Disables warnings about insecure (HTTP) connections.

Type: boolean

Default: `false`

ForceSecure RStudio Package Manager will mark cookies secure on its unsecured connection. This option should be used when RStudio Package Manager is behind a secure proxy.

Type: boolean

Default: `false`

A.4 HTTPS

The HTTPS section contains configuration properties which control the ability of RStudio Package Manager to listen for HTTPS requests. RStudio Package Manager must be configured to listen for either HTTP or HTTPS requests (allowing both is acceptable).

These properties must appear after [HTTPS] in the configuration file.

Listen RStudio Package Manager will listen on this network address for HTTPS connections. The network address can be of the form `:443` or `192.168.0.1:443`. Either `HTTP.Listen` or `HTTPS.Listen` is required.

Type: string

Default: *<empty-string>*

Key Path to a private key file corresponding to the certificate specified with `HTTPS.Certificate`. Required when `HTTPS.Certificate` is specified.

Type: string

Default: *<empty-string>*

Certificate Path to a TLS certificate file. If the certificate is signed by a certificate authority, the certificate file should be the concatenation of the server's certificate followed by the CA's certificate. Must be paired with `HTTPS.Key`.

Type: string

Default: *<empty-string>*

Permanent Advertises to all visitors that this server should only ever be hosted securely via HTTPS. WARNING: if this is set to true – even temporarily – visitors may be permanently denied access to your server over an unsecured (non-HTTPS) protocol. This sets the `secure` flag on all session cookies and adds a `Strict-Transport-Security` HTTP header with a value of 30 days.

Type: boolean

Default: `false`

MinimumTLS Minimum TLS version supported. Valid values are 1.0, 1.1, and 1.2.

Type: string

Default: 1.1

ExcludedCiphers List of IANA strings for ciphers to exclude. See <https://www.iana.org/assignments/tls-parameters/tls-parameters.xml> for a list of all known cipher strings. Multiple definitions can be used to exclude multiple ciphers.

Type: multi-string

Default: *unspecified*

A.5 HttpRedirect

The `HttpRedirect` section contains configuration properties which control the ability of RStudio Package Manager to listen for HTTP requests and then redirect all traffic to some alternate location. This is useful when paired with an `HTTPS.Listen` configuration.

These properties must appear after `[HttpRedirect]` in the configuration file.

Listen RStudio Package Manager will listen on this network address for HTTP connection and redirect to either the `HttpRedirect.Target` or `Server.Address` target location. The network address can be of the form `:8080` or `192.168.0.1:8080`. Useful when you wish all requests to be served over HTTPS and send users to that location should they accidentally visit via an HTTP URL. Must be paired with either `HttpRedirect.Target` or `Server.Address`.

Type: string

Default: *<empty-string>*

Target The target for redirects when users visit the `HttpRedirect.Listen` HTTP service. `Server.Address` is used as a redirect target if this property is not specified.

Type: string

Default: *<empty-string>*

A.6 Git

The `Git` section contains configuration properties which RStudio Package Manager uses to build R packages from Git repositories.

These properties must appear after `[Git]` in the configuration file.

BuilderDir The directory where RStudio Package Manager will store git repos.

Type: string

Default: `{Server.DataDir}/git`

Reloadable: true

PollInterval The amount of time to wait between polling git repos to look for package changes.

Type: duration

Default: 5m

Reloadable: true

TagMatching The regular expression used to determine whether a git tag is a valid version.

Type: string

Default: `<empty-string>`

BuildRetries The maximum number of times to attempt building a git package when the build fails.

Type: integer

Default: 3

AllowUnsandboxedGitBuilds When the Linux user namespace is unavailable, allow unsafe Git package builds without a sandbox.

Type: boolean

Default: false

A.7 Launcher

The **Launcher** section contains configuration properties which control the RStudio Job Launcher. The RStudio Job Launcher provides the capability to launch executables, allowing RStudio Package Manager to launch R processes for tasks like building R packages.

These properties must appear after `[Launcher]` in the configuration file.

Disabled Disables the job launcher

Type: boolean

Default: `false`

Port Port number (0-65535)

Type: integer

Default: 5561

ServerUser User that runs the job launcher executable.

Type: string

Default: `rstudio-pm`

AdminGroup Name of the admin group of users that are able to see/control all jobs in the system of other users.

Type: string

Default: `rstudio-pm`

DebugLogging Enables verbose debug logging for the job launcher.

Type: boolean

Default: `false`

A.8 Licensing

The **Licensing** section contains configuration properties which control how RStudio Package Manager interacts with its licensing system.

These properties must appear after `[Licensing]` in the configuration file.

LicenseType Enable remote or local validation. `local` is traditional activation, whereas `remote` uses floating licensing.

Type: string

Default: `local`

RemoteRetryFrequency When Package Manager loses its lease, it will begin automatically attempting to acquire a lease by `RemoteRetryFrequency`. Use a value of `0` to disable retries.

Type: duration

Default: `10`

A.9 CRAN

The **CRAN** section contains configuration properties related to the CRAN source in RStudio Package Manager.

ManifestURL The URL to the CRAN manifest service. Do not use this setting unless you want to get CRAN data from your own server.

Type: string

Default: `https://rspm-sync.rstudio.com`

SyncSchedule The schedule for CRAN synchronization. Must be a valid crontab string. Use a blank value for manual syncing only.

Type: string

Default: `0 0 * * *`

Reloadable: `true`

SyncMode The sync mode for CRAN synchronization. Must be either `eager` or `lazy`

Type: string

Default: `lazy`

Reloadable: `true`

CRANTimeout The timeout for downloading checkpoints or changes (JSON) from the CRAN manifest

Type: duration

Default: 30m

FetchTimeout The timeout for downloading a CRAN package tarball.

Type: duration

Default: 10m

A.10 Database

The **Database** section contains configuration properties which control the location of and how RStudio Package Manager interacts with its databases.

These properties must appear after **[Database]** in the configuration file.

Provider The type of database to use

Type: string

Default: sqlite3

MaxIdleConnections The maximum number of database connections that should be retained after they become idle. If this value is less-than or equal-to zero, no idle connections are retained.

Type: integer

Default: 0

MaxOpenConnections The maximum number of open connections to the database. If this value is less-than or equal-to zero, then there is no limit to the number of open connections.

Type: integer

Default: 0

ConnectionMaxLifetime The maximum amount of time a connection to the database may be reused. If this value is less-than or equal-to zero, then connections are reused forever.

Type: duration

Default: 0

A.11 SQLite

The **SQLite** section contains configuration properties which control the location of and how RStudio Package Manager interacts with the SQLite database.

These properties must appear after **[SQLite]** in the configuration file.

Dir The directory containing the RStudio Package Manager database. Must reference a directory on the local filesystem and not on a networked volume like NFS.

Type: string

Default: `{Server.DataDir}/db`

A.12 Postgres

The **Postgres** section contains configuration properties which control the location of and how RStudio Package Manager interacts with its postgres.

These properties must appear after `[Postgres]` in the configuration file.

URL The fully qualified URL to connect to a Postgres database

Type: string

Default: `<empty-string>`

UsageDataURL The fully qualified URL to connect to a Postgres database where usage information will be written.

Type: string

Default: `<empty-string>`

A.13 Storage

The **Storage** configuration section contains properties that control whether variable data should be stored on a local filesystem or S3. You can configure the storage destination for four classes of data. By default, each storage class defaults to `file`, but you can configure any storage class for S3.

IMPORTANT: AWS S3 support is in beta. Please do not use S3 for production data at this time.

```
; /etc/rstudio-pm/rstudio-pm.gcfg

[Storage]
; Each class defaults to `file`
; Cache = file
; Launcher = file
;
; You can override the destination to use S3
Metrics = s3
Packages = s3
```

Cache Storage location for the cache. Must be either `file` or `S3`.

Type: string

Default: `file`

Packages Storage location for packages. Must be either `file` or `S3`.

Type: string

Default: `file`

Metrics Storage location for usage metrics aggregation. Must be either `file` or `S3`.

Type: string

Default: `file`

Launcher Storage location for job launcher logs. Must be either `file` or `S3`.

Type: string

Default: file

A.14 FileStorage

The `FileStorage` section contains configuration properties which control where files are stored. You can configure a `FileStorage` section for any of four classes of data:

- `cache`
- `launcher`
- `metrics`
- `packages`

If you wish to change the base storage location for all RStudio Package Manager data, please use the `Server.DataDir` property instead. Use a `FileStorage` configuration section only when you wish to change the storage location for a particular class of data.

For example, to configure a custom storage location for the cache, you would create the following configuration section:

```
;/etc/rstudio-pm/rstudio-pm.gcfg  
  
[FileStorage "cache"]  
Location = /path/to/cache
```

Location The directory to use for this storage class.

Type: string

Default: `{Server.DataDir}/{StorageClass}`

A.15 S3Storage

IMPORTANT: AWS S3 support is in beta. Please do not use S3 for production data at this time.

The `S3Storage` section contains configuration properties which control where files are stored on S3. You can configure an `S3Storage` section for any of four classes of data:

- `cache`
- `launcher`
- `metrics`
- `packages`

If you wish to use S3 for all variable data, please use a generic `S3Storage` configuration section. You can then override the generic settings for any storage class as necessary. For example:

```
;/etc/rstudio-pm/rstudio-pm.gcfg  
  
[Storage]  
Cache = s3  
Launcher = s3  
Metrics = s3  
Packages = s3
```



```

; Default S3 settings. This is the minimum-required setting for using S3.
[S3Storage]
Bucket = your-s3-bucket

; Override default S3 settings for the "packages" class. This demonstrates
; all the available S3 configuration settings.
[S3Storage "packages"]
Bucket = another-s3-bucket
Prefix = rspm-packages
Profile = dev-rspm
Region = us-west-1
EnableSharedConfig = true

```

Bucket S3 bucket to use. Required.

Type: string

Default: *<empty-string>*

Prefix Prefix to prepend to item address. This is useful for organizing a shared S3 bucket. Optional.

Type: string

Default: {StorageClass}

Profile AWS Profile to use. AWS profiles can be defined in *~/.aws/config* or *~/.aws/credentials*. Optional.

Type: string

Default: *<empty-string>*

Region AWS Region to use. Optional

Type: string

Default: *<empty-string>*

EnableSharedConfig Overrides the `AWS_SDK_LOAD_CONFIG` env var value and enables the shared config functionality. Optional

Type: boolean

Default: false

B Package Ecosystem

The R package ecosystem has a few key components.

B.1 Packages

Packages are the primary extension mechanism for R. They can be used to share functions, datasets, and documentation. An R package can exist in a few states:

B.1.1 Source

An R package is composed of a series of directories and files. The source of an R package is just a top-level directory containing the components of the package. Package authors work with source packages during development. Git(hub) repositories store source packages.

B.1.2 Bundle

A bundled package is a package that has been compressed into a single file. By convention, package bundles in R use the extension `.tar.gz`.

B.1.3 Binary

A binary package is the result of building a source package for a specific operating system. Binary packages are single files that are ready for installation on their specific operating systems.

B.1.4 Installed

An installed package is a binary package that has been decompressed into a package library and is ready for use by R.

B.2 Repositories

Repositories organize R packages for distribution to end users. Repositories contain package bundles and binaries that are organized in a specific way so that users can install packages from the repository using R's `install.packages` command. CRAN and Bioconductor are examples of R repositories.

B.3 Git(hub)

Many R package sources are stored in version controlled directories. A popular versioning tool is Git. Github, as an extension of Git, houses many package sources. The `devtools` R package includes convenience functions for installing packages from the package source contained on a Git repository, including Github. Used in this manner, git repositories and Github are one way to distribute R packages, but Github and Git repositories *are not* R package repositories.

B.4 Libraries

End users of R typically interact with installed packages that live in libraries. Package libraries are just directories containing installed packages. When a package is requested by R, R searches the different library directories to find the installed package.

R libraries are very flexible. In the past, R users have set up libraries for specific projects or set up a system-wide library used across multiple projects. In multi-tenant servers it has been common to have both a system library shared by all users and user-specific libraries.

A best practice is to set up per-project libraries alongside a package cache.

C Changing RunAs User

C.1 Default Account

The installer creates a user account and group named `rstudio-pm` and runs the RStudio Package Manager service under this account. See 2.2.4 for more information.

C.2 Changing The RunAs Account (Service)

You can configure RStudio Package Manager to run under another account. The steps below serve as a guide for reconfiguring RStudio Package Manager to run under an account named `thor` with a primary group of `heroes` instead of the default `rstudio-pm:rstudio-pm`.

1. Stop the RStudio Package Manager service. See 5.1.

```
# Ubuntu 14 and Red Hat/CentOS 6
sudo stop rstudio-pm

# Ubuntu 16+, Red Hat/CentOS 7, SUSE 12
sudo systemctl stop rstudio-pm
```

2. Create a new group and user account

```
sudo groupadd heroes
sudo useradd -r -g heroes -M -s /sbin/nologin thor
```

In order to use the CLI tool, a user must be a member of the primary group of the user that starts the RStudio Package Manager service.

In this example, RStudio Package Manager is started by the user `thor`. The primary group of the `thor` account is `heroes`, so users must be members of the `heroes` group to use the CLI.

3. Edit the service configuration

Ubuntu 14 and Red Hat/CentOS 6

```
sudo vi /etc/init/rstudio-pm.override
```

Change these lines:

```
env RSTUDIO_PM_USER=thor
env RSTUDIO_PM_GROUP=heroes
```

Ubuntu 16+, Red Hat/CentOS 7, SUSE 12

```
sudo vi /etc/systemd/system/rstudio-pm.service.d/user.conf
```

Change these lines:

```
[Service]
User=thor
Group=heroes
```

4. Change ownership of files and directories

```
# Configuration file
sudo chown thor:heroes /etc/rstudio-pm/rstudio-pm.gcfg

# Log files
sudo chown thor:heroes /var/log/rstudio-pm.*
```

```

# Data directory (or `Server.DataDir`, if configured for a custom location)
sudo chown -R thor:heroes /var/lib/rstudio-pm

# Run directory
sudo chown -R thor:heroes /var/run/rstudio-pm

# If you have a custom `Sqlite.Dir` (e.g., `Sqlite.Dir = /database/directory`)
sudo chown -R thor:heroes /database/directory

# Temporary workspace directory for Git package builds
# This directory is usually found under `/tmp`. If you use a custom
# `Server.TempDir`, then it will be found at `Server.TempDir`/rspm-workspaces.
sudo chown -R thor:heroes /tmp/rspm-workspaces

# If you have any custom FileStorage configuration settings
#
# For example:
#
# [FileStorage "cache"]
# Location = /path/to/cache
#
# [FileStorage "launcher"]
# Location = /path/to/launcher
#
# [FileStorage "metrics"]
# Location = /path/to/metrics
#
# [FileStorage "packages"]
# Location = /path/to/packages
#
sudo chown -R thor:heroes /path/to/cache
sudo chown -R thor:heroes /path/to/launcher
sudo chown -R thor:heroes /path/to/metrics
sudo chown -R thor:heroes /path/to/packages

```

5. Edit the Job Launcher configuration (unless running as root)

If you wish to build Git packages, you may need to update the Job Launcher configuration to build packages under the correct account and group.

When Running Unprivileged

When RStudio Package Manager is running unprivileged (the default), the Job Launcher automatically runs under the same user and group as the RStudio Package Manager service. Skip this step. Using a different account for the `Launcher.ServerUser` configuration setting is prohibited.

When Running Privileged

If you are changing the `RunAs` user to `root`, the Job Launcher will automatically use the `rstudio-pm` user and group to build Git packages, and you may skip this step. However, you may customize the Job Launcher user and admin group by using the `Launcher.ServerUser` and `Launcher.AdminGroup` configuration settings.

```
sudo vi /etc/rstudio-pm/rstudio-pm.gcfg
```

Add or change these lines:

```
[Launcher]
ServerUser = thor
AdminGroup = heroes
```

6. Remove remaining domain socket file (if any)

```
sudo rm /var/run/rstudio-pm/rstudio-pm.sock
```

7. Start the RStudio Package Manager service. See 5.1.

```
# Ubuntu 14 and Red Hat/CentOS 6
```

```
sudo start rstudio-pm
```

```
# Ubuntu 16+, Red Hat/CentOS 7, SUSE 12
```

```
sudo systemctl daemon-reload # Reload the systemd process
```

```
sudo systemctl start rstudio-pm
```

8. Verify that the `rstudio-pm` service is running under the `thor` account.

```
ps -axj | grep `id -u thor`
```

9. Check `/var/log/rstudio-pm.log` to verify that the server started up with no errors.

C.3 Using the New RunAs Account (CLI)

After changing the service RunAs user, your CLI users must be members of the `heroes` group. For example:

```
sudo useradd -g heroes hulk
sudo passwd hulk
su hulk
/opt/rstudio-pm/bin/rspm <command>
```

D Air-Gapped RStudio Package Manager

RStudio Package Manager communicates with a RStudio CRAN service to access CRAN packages and metadata. In offline environments, it is possible to directly download the necessary data from the RStudio CRAN service and then copy it to an offline RStudio Package Manager server.

This guide includes 3 sections:

- Instructions for an initial air-gapped setup (D.1)
- Instructions for performing regular updates (D.2)
- Instructions for upgrading RStudio Package Manager in an air-gapped environment (D.3)

D.1 Initial Setup

1. Install the AWS CLI tools in an online machine. Confirm by running `aws help`.
2. Run the command `rspm air-gap` in the offline RStudio Package Manager server. See 14 for more information about the `rspm` commands.

The `air-gap` command will print information and output a command similar to:

```
MAJOR=`curl https://rspm-sync-staging.rstudio.com/v2/version.txt`; \  
echo "aws s3 sync --no-sign-request --exclude=* --include=v2/version.txt \  
--include=v2/${MAJOR}/* s3://rstudio-pm-sync/ ./"
```

3. In the online machine, create a new directory, e.g `cran-data`. Copy the command from step 2 and run it from within the directory on the online machine. A new command will be printed starting with `aws s3 sync`. Copy the entire output and run it in the the same location. The sequence of input and output looks like:

```
# Input:  
MAJOR=`curl https://rspm-sync-staging.rstudio.com/v2/version.txt`; \  
echo "aws s3 sync --no-sign-request --exclude=* --include=v2/version.txt \  
--include=v2/${MAJOR}/* s3://rstudio-pm-sync/ ./"  
  
# Output:  
aws s3 sync --no-sign-request --exclude=* --include=v2/version.txt \  
--include=v2/2/* s3://rstudio-pm-sync/ ./  
  
# Input:  
aws s3 sync --no-sign-request --exclude=* --include=v2/version.txt \  
--include=v2/2/* s3://rstudio-pm-sync/ ./"
```

These steps will begin the download which includes 50+ GB of data and may take some time to complete.

The result will be a directory with files prefixed by a version indicator, e.g.

```
cran-data/v2/version.txt  
cran-data/v2/2  
cran-data/v2/2/...
```

4. Create a directory in the offline RStudio Package Manager server. If you have a cluster of nodes, use shared storage for this directory.

```
mkdir /var/lib/rstudio-pm/cran
```

Copy the data downloaded in step 3 from the online machine to the directory on the offline RStudio Package Manager server. For completely isolated servers, you may need to copy the data to a physical drive in order to move it to the offline environment.

Confirm by checking that the offline directory has the same number of files as the original directory on the online machine.

Finally, modify the permissions on the directory in the offline RStudio Package Manager server, changing ownership to the unix account running RStudio Package Manager, `rstudio-pm` by default:

```
sudo chown -R rstudio-pm /var/lib/rstudio-pm/cran
```

5. Next, configure the offline RStudio Package Manager server to use the downloaded data. To do so, modify the RStudio Package Manager Configuration file to include the following properties in the `CRAN` configuration section:

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[CRAN]
```

```
ManifestURL = A URL in the form, `file:///<the directory you created in step 4>`
```

For example, if your CRAN data directory is at `/var/lib/rstudio-pm/cran`, the file `/etc/rstudio-pm/rstudio-pm.gcfg` should contain:

```
; /etc/rstudio-pm/rstudio-pm.gcfg
```

```
[CRAN]
```

```
ManifestURL = file:///var/lib/rstudio-pm/cran
```

Once the file is updated, restart the RStudio Package Manager Service.

6. Follow the remainder of the instructions in the admin guide for setting up sources and repositories using the `rspm` commands in the offline RStudio Package Manager server. The `rspm sync` command will use the downloaded data to populate the necessary CRAN data and packages.

D.2 Regular Updates

It is important to regularly update data available on the offline server. Updating this data does not automatically make new packages available to end users. The following process pushes updates from the RStudio CRAN service to RStudio Package Manager's metadata. Follow the instructions in the admin guide to make updates available to end users. See chapter 14 for more information.

1. In the offline RStudio Package Manager server, run the command `rspm air-gap`. This command will output information along with a command similar to:

```
MAJOR=`curl https://rspm-sync.rstudio.com/v2/version.txt`; \  
echo "aws s3 sync --no-sign-request --exclude=* \  
--include=v2/version.txt --include=v2/${MAJOR}/* s3://rstudio-pm-sync/ ./"
```

2. In the online machine, navigate to the directory created during the initial setup, e.g. `cran-data`, and paste the command from step 1. A second command will be printed. Copy and run the second command in the same location. Any newly available data will be downloaded, but existing data will not be re-downloaded.
3. Copy the directory from the online machine to the folder created in the offline RStudio Package Manager during the initial setup, e.g. `/var/lib/rstudio-pm/cran`. Ensure that the directory is still owned by the unix account running RStudio Package Manager, `rstudio-pm` by default.

D.3 Upgrading RStudio Package Manager

A new version of RStudio Package Manager may require data from a new version of the RStudio CRAN service. To ensure a smooth upgrade with limited downtime, we recommend the following steps:

1. You will need a staging environment that mirrors your offline production server. After creating this environment, begin by upgrading the offline staging server to the latest RStudio Package Manager release.
2. Follow the instructions for the *Initial Setup* of an Air-Gapped server in D.1, using the offline staging server.
3. After you have validated that everything works as expected, copy the the CRAN data, e.g. `/var/lib/rstudio-pm/cran`, from the offline staging server to the offline production server.
4. Upgrade the offline production server to the new version of RStudio Package Manager.
5. (Optional) After an upgrade, navigate to the directory storing CRAN data, e.g. `/var/lib/rstudio-pm/cran`. This directory will contain versioned folders, e.g.

```
/v2
```

```
/v3
```

Run the command `rspm air-gap` in the offline production server. The output will indicate the version of the RStudio CRAN service required by the current version of RStudio Package Manager, e.g.

Your RStudio Package Manager uses a CRAN Manifest Schema Version of “v3”.

You can safely remove all of the other folders in the directory. In this example, only `/var/lib/rstudio-pm/v3` is necessary.

E Manual Installation for Minimal Root Use

E.1 Discussion

The RStudio Package Manager installer requires root privileges and installs the RStudio Package Manager to `/opt/rstudio-pm`. Additionally, it configures services for Upstart or systemd (depending on your OS), adds log files to `/var/log`, adds a configuration file under `/etc/rstudio-pm`, creates an `rstudio-pm` user account and group, and places a logentries configuration.

It is possible to extract the RStudio Package Manager files manually and configure the service to use configuration and data files at a custom location. Most of these steps can be performed without root privileges.

E.2 Extracting Files

First, download the latest version of RStudio Package Manager from the link provided by RStudio. Then, extract the files from the installer by following the instructions for your operating system, below.

E.2.1 Extracting Files (Ubuntu)

Note you may need to install `binutils` with `sudo apt install binutils` to use the `ar` utility.

```
mkdir rspm
ar x rstudio-pm-<version>_amd64.deb
tar -xzvf data.tar.gz -C rspm/
cd rspm
```

E.2.2 Extracting Files (Red Hat/CentOS 7, SUSE 12)

```
mkdir rspm
cd rspm
rpm2cpio ../rstudio-pm-<version>.x86_64.rpm | cpio -div
```

E.3 Create Directories

Create the directories you need to run RStudio Package Manager.

```
mkdir log
touch log/rstudio-pm.log
touch log/rstudio-pm.access.log
mkdir run
mkdir data
mkdir temp
```


E.4 Licensing

The license manager used by RStudio Package Manager supports userspace license activation. If you wish to activate a system-wide license, root privileges are required. See 6 for more details on system-wide licensing.

```
# Initialize a trial in userspace
./opt/rstudio-pm/bin/license-manager initialize --userspace

# Activate a userspace license key, if you have a license key
./opt/rstudio-pm/bin/license-manager activate --userspace <key>
```

E.5 Edit config file

Next, edit the RStudio Package Manager configuration file to point to the directories and access log you created. Add the following lines to the [Server] configuration section.

```
; ./etc/rstudio-pm/rstudio-pm.gcfg

[Server]
DataDir = <path-to-unbundled-rspm-dir>/data
SockFileDir = <path-to-unbundled-rspm-dir>/run
AccessLog = <path-to-unbundled-rspm-dir>/log/rstudio-pm.access.log
TempDir = <path-to-unbundled-rspm-dir>/temp
```

E.6 Start the RStudio Package Manager Service

When RStudio Package Manager is installed without root, the system daemons are not used to automatically start RStudio Package Manager. Instead, the user is required to start RStudio Package Manager manually and ensure the process continues to run. The command to start RStudio Package Manager is:

```
./opt/rstudio-pm/bin/rstudio-pm --config \
./etc/rstudio-pm/rstudio-pm.gcfg >> ./log/rstudio-pm.log 2>&1
```

As an example, `nohup` can be used to launch this command without blocking or depending on the terminal using:

```
nohup ./opt/rstudio-pm/bin/rstudio-pm --config \
./etc/rstudio-pm/rstudio-pm.gcfg >> ./log/rstudio-pm.log 2>&1 &
```

E.7 Use the CLI to Manage RStudio Package Manager

```
./opt/rstudio-pm/bin/rspm --config ./etc/rstudio-pm/rstudio-pm.gcfg help
```

In this mode, every CLI command will require the config flag that points at the configuration file.