

RStudio Connect: Admin Guide

Version 1.4.4.1-16

Abstract

This guide will help an administrator install and configure RStudio Connect 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	3
1.1	System Requirements	3
2	Getting Started	4
2.1	Installation	4
2.2	Initial Configuration	6
3	Licensing & Activation	8
3.1	Proxy Servers	8
3.2	Offline Activation	9
4	Files & Directories	9
4.1	Program Files	9
4.2	Configuration	10
4.3	Server Log	10
4.4	Access Logs	10
4.5	Application Logs	10
4.6	Variable Data	10
4.7	Backups	12
5	Server Management	12
5.1	Stopping and Starting	12
5.2	System Messages	14
5.3	Upgrading	14
5.4	Purging RStudio Connect	15
6	Running with a Proxy	15
6.1	Overview	15
7	Security & Auditing	17
7.1	Browser Security	17
7.2	Audit Logs	18
8	Authentication	19
8.1	Migrating to a different Authentication Provider	19
8.2	Password	19
8.3	LDAP and Active Directory	20
8.4	OAuth2 (Google)	25
8.5	PAM	26
8.6	Proxied Authentication	28

9	User Management	29
9.1	User Roles	29
9.2	User Permissions	30
9.3	Locked Accounts	30
9.4	User Renaming	31
9.5	Command-Line Interface	31
10	Process Management	32
10.1	Sandboxing	32
10.2	Shiny Applications	34
10.3	User Account for R Processes	34
10.4	Current user execution	34
10.5	PAM sessions	35
10.6	Path Rewriting	35
10.7	Program Supervisors	36
11	Content Management	37
11.1	Vanity Paths	37
11.2	Bundle Management	38
12	R	38
12.1	Installing R	38
12.2	R Versions	38
12.3	R Version Matching	40
13	Package Management	40
13.1	Package Installation	40
13.2	Private Repositories	41
13.3	Private Packages	42
14	Historical Metrics	42
14.1	Historical Metrics Settings	42
14.2	Historical Metrics Process Management	43
14.3	Historical Metrics Process Logging	43
15	Appendix: Configuration Options	43
15.1	Server	45
15.2	Http	47
15.3	Https	47
15.4	HttpRedirect	48
15.5	Database	49
15.6	Authentication	49
15.7	OAuth2	50
15.8	LDAP	50
15.9	PAM	52
15.10	Proxied Authentication	52
15.11	Authorization	52
15.12	Applications	53
15.13	Performance/Scheduler	54
15.14	Jobs	55
15.15	Historical Metrics	56
16	Appendix: LDAP/AD Configuration Examples	56
17	Appendix: RStudio Connect Deployment Guide	58

17.1 Overview	58
17.2 Programmatic Deployment	58
17.3 Step 1: Building the Bundle	59
17.4 Step 2: Push Bundle to Connect	62
17.5 Step 3: Bundle is deployed on Connect	62
17.6 Other Frequently Asked Questions	63
18 Using Continuous Integration to Deploy Content	64
18.1 Overview	64
18.2 Prerequisites	64
18.3 Configuring a CI Server to Deploy Content to Connect	64
18.4 Warning and Security Information	65
19 Appendix: Programmatic Deployment with rsconnect	66
19.1 Overview	66
19.2 Use Case: A Shiny Application	66
19.3 Warning and Security Information	67
19.4 Example Shiny Application	68

1 Introduction

RStudio Connect allows users to share and collaborate on the results they produce with R such as R Markdown documents, Shiny applications, and plots. Source code or rendered artifacts can be deployed into RStudio Connect and selectively shared with other viewers and collaborators within the organization. Some content can even be scheduled to be re-executed and emailed on a given schedule.

RStudio Connect can also help simplify the role of the system administrator tasked with supporting R by offering:

- Detailed metrics for the server and the associated R processes
- Logs for all R processes spawned by Connect
- Secure deployments and interactions with artifacts using SSL/TLS
- Scale a Shiny application beyond a single R process to support additional visitor load

1.1 System Requirements

RStudio Connect is supported on the following distributions of the Linux operating system:

- Red Hat Enterprise Linux/CentOS Linux 6.0+
- Red Hat Enterprise Linux/CentOS Linux 7.0+
- Ubuntu 12.04
- Ubuntu 14.04
- Ubuntu 16.04

We currently only offer installers for the x86-64 architecture and require root privileges both to install and run Connect.

RStudio Connect is supported against the latest versions of the following browsers:

- Chrome
- Safari
- Firefox
- Internet Explorer 10
- Internet Explorer 11
- Microsoft Edge

2 Getting Started

This chapter helps you install RStudio Connect on Ubuntu or Red Hat Enterprise Linux/CentOS Linux, learn to manage the server, and perform some initial configuration.

We built this checklist to guide you through that process.

1. Install R - Ubuntu 2.1.1, Red Hat/CentOS 2.1.2
2. Download RStudio Connect installer
3. Install RStudio Connect - Ubuntu 2.1.1, Red Hat/CentOS 2.1.2
4. Set `Server.SenderEmail` - 2.2.1
5. Set `Server.Address` - 2.2.1
6. Configure Authentication - 2.2.2, 8
7. Restart RStudio Connect - 5.1
8. Sign into RStudio Connect - 2.2.3
9. Configure email sending - 2.2.4

2.1 Installation

This section explains how to install R using the public package repositories for your Ubuntu or Red Hat/CentOS server. Chapter 12 explains how to configure RStudio Connect to access multiple versions of R on the same server.

2.1.1 Ubuntu (12.04+)

Connect recommends an installation of R version 3.0 or higher. To install the latest version of R you should first add the CRAN repository to your system as described here:

<http://cran.rstudio.com/bin/linux/ubuntu/README.html>

You can then install R using the following command:

```
$ sudo apt-get install r-base
```

Note: If you do not add the CRAN Ubuntu repository as described above this command will install the version of R corresponding to your current system version. This version of R may be a year or two old. It is strongly recommended that you add the CRAN repositories so you can run the most up-to-date version of R.

RStudio Connect can be configured to use versions of R other than the default system R. Please see Section 12 for details.

You will use `gdebi` to install Connect and its dependencies. It is installed via the `gdebi-core` package.

```
$ sudo apt-get install gdebi-core
```

You should have been provided with a `.deb` installer for RStudio Connect. If you only have a link to this file, you can use `wget` to download the file to the current directory.

```
$ wget https://download-url/rstudio-connect-1.4.4.1-16.deb
```

Once the `.deb` file is available locally, run the following command to install RStudio Connect.

```
$ sudo gdebi rstudio-connect-1.4.4.1-16.deb
```

This will install Connect into `/opt/rstudio-connect/`, and create a new `rstudio-connect` user.

At this point you should have Connect installed on your server. You can now configure the server following the instructions in Section 2.2.1 However, we recommend that you consider installing some additional system

dependencies that common R packages require. Without these system dependencies, your users may not be able to use the R packages they require on the server.

Recommended Packages

The following system dependencies are required by many common R packages and nearly all deployments will need to provide these. These package names may vary slightly between different versions of Ubuntu.

```
build-essential
libcurl4-gnutls-dev
openjdk-7-* # may require also executing `R CMD javareconf`
libxml2-dev
libssl-dev
texlive-full # very large dependency, but needed to render PDF documents from R Markdown
```

Supplemental Packages

There are additional system dependencies that may be required for some R packages depending on the types of R packages your users are leveraging. You could consider providing these packages for your users now, or wait until they are requested.

An up-to-date resource to help understand the system dependencies of various R packages in Ubuntu is available at <https://github.com/rstudio/shinyapps-package-dependencies>.

```
libgmp10-dev
libgs10-dev
libnetcdf6
libnetcdf-dev
netcdf-bin
libdigest-hmac-perl
libgmp-dev
libgmp3-dev
libgl1-mesa-dev
libglu1-mesa-dev
libglpk-dev
tdsodbc
freetds-bin
freetds-common
freetds-dev
odbc-postgresql
libtiff-dev
libsndfile1
libsndfile1-dev
libtiff-dev
tk8.5
tk8.5-dev
tcl8.5
tcl8.5-dev
libgs10-dev
libv8-dev
```

2.1.2 Red Hat Enterprise Linux/CentOS Linux (6.0+)

Prerequisites

RStudio Connect recommends an installation of R version 3.0 or higher. Connect has several dependencies on packages (including R itself) found in the Extra Packages for Enterprise Linux (EPEL) repository. If you

don't already have this repository available you should add it to your system using the instructions found here: <https://fedoraproject.org/wiki/EPEL>. On some distributions of Red Hat Enterprise Linux/CentOS Linux, the R package references dependencies that are not available by default. In this case, you may need to edit the `/etc/yum.repos.d/redhat.repo` file to enable the `rhel-6-server-optional-rpms` (by setting `enabled = 1`) before you can install the R package.

After enabling EPEL you should then ensure that you have installed the version of R available from EPEL. You can do this using the following command:

```
$ sudo yum install R
```

RStudio Connect can be configured to use versions of R other than the default system R. Please see Section 12 for details.

You can now begin the installation of RStudio Connect. You should have been provided with an RPM file which contains Connect and all of its dependencies (other than R). You can install this rpm file using yum. If you have only a link to the RPM file, you can use `wget` to download the file to the current directory. You can install this RPM file using yum.

```
$ sudo yum install --nogpgcheck rstudio-connect-1.4.4.1-16.rpm
```

This will install Connect into `/opt/rstudio-connect/` and create a new `rstudio-connect` user.

At this point you should have Connect installed on your server. You can now configure the server following the instructions in Section 2.2.1 However, we recommend that you consider installing some additional system dependencies that common R packages require. Without these system dependencies, your users may not be able to use the R packages they require on the server.

Recommended Packages

The following system dependencies are required by many common R packages and nearly all deployments will need to provide these. These package names may vary slightly between different versions of Red Hat Enterprise Linux/CentOS Linux.

```
make
gcc
gcc-c++
libcurl-devel
libxml2-devel
java-1.7.0-openjdk-devel # may require also executing `R CMD javareconf`
openssl-devel
texlive-* # VERY large dependency, but needed to render PDF documents from R Markdown
```

2.2 Initial Configuration

RStudio Connect is installed and running, but probably needs a bit of customization for your organization. This section will help you specify the public URL of your server, configure authentication, and validate that RStudio Connect is able to send email.

2.2.1 Editing the Configuration File

RStudio Connect is controlled by the `/etc/rstudio-connect/rstudio-connect.gcfg` configuration file. You will edit this file to make server-wide configuration changes to the system. See the configuration appendix 15 for details about this file, its syntax, and the available settings.

We strongly recommend setting the `SenderEmail` and `Address` server properties. Both must be specified in the `[Server]` section of your configuration file.

The `Server.SenderEmail` property is the email address from which Connect sends emails. It is important that the sendmail or SMTP configuration RStudio Connect uses be willing and able to send email from this `SenderEmail` address. Otherwise, Connect will not be able to successfully send email. See Section 2.2.4 for more details about mail sending.

The `Server.Address` property is the public URL used to access the server. When accessible over a non-standard port, this URL must specify both hostname and port. If this isn't configured, Connect will not be able to include links in emails that send users to the appropriate location on the server.

The standard HTTP port is 80; the standard HTTPS port is 443.

Whenever RStudio Connect is deployed behind a proxy, you *must* configure the `Server.Address` setting with the proxied location. RStudio Connect normally returns URLs that are in terms of its local address. The `Server.Address` property causes Connect to use an alternate base location when building URLs. Setting `Server.Address` to the location of your proxy will produce URLs in terms of your proxy address instead of the Connect local address.

Using RStudio Connect with a proxy is **experimental**. There are known issues with URL construction, especially in situations where the proxy is performing path rewriting. This limitation will be removed in upcoming releases.

Here is a sample configuration specifying both `SenderEmail` and `Address`.

```
[Server]
SenderEmail = rstudio-connect@company.com
Address = https://rstudio-connect.company.com/
```

Use the instructions in Section 5.1 to restart RStudio Connect after altering the `rstudio-connect.gcfg` configuration file.

2.2.2 Authentication

It is important that you specify the correct style of authentication for your organization. If you are using RStudio Connect on a trial basis, feel free to use the default password authentication. RStudio Connect also supports a number of external authentication integrations, which are detailed in Section 8.

You must establish the correct form of authentication before using RStudio Connect. Changing the style of authentication later is not supported.

2.2.3 Sign In!

Use a web browser to visit the RStudio Connect dashboard. This has a default location of `http://your-connect-server:3939/`. Click the “Sign In” link. If you are using an external authentication provider, specify your login credentials. If you are using password authentication, follow the “Create a new account” link and configure your account.

The first account will be marked as an RStudio Connect administrator. Please use this account to configure mail sending. Connect needs to send email for account confirmation messages (when using the default password auth provider) and when distributing content.

2.2.4 Email Sending

Visit the RStudio Connect dashboard and sign in as an administrator. Visit the Admin>Settings screen and configure mail sending for your organization.

RStudio Connect supports two options for sending mail:

- Sendmail - The `sendmail` command is used to send messages locally on your server. This relies on a working sendmail configuration or some equivalent replacement.
- SMTP - Mail is sent using an SMTP endpoint and supports SSL and authentication.

Please contact your system administrator if you have questions about which of these options are appropriate.

Be sure to verify your settings by sending a test message!

3 Licensing & Activation

When RStudio Connect 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-connect/bin/license-manager status
```

After purchasing a license to RStudio Connect, you will receive a product key that is used to activate the license on a given system. Each product key given limits usage of RStudio Connect in the following ways...

- Number of user accounts that have signed into RStudio Connect. Once this limit is reached, additional users will not be permitted to sign into RStudio Connect. This limit is enforced the first time each user logs in. Locked users are not counted against this quota.
- Number of users that can access Shiny applications at one moment in time. If this number is exceeded, new anonymous users will be unable to view the Shiny application requested. This limitation does not affect logged in users.

How many are allowed of each metric depends on the license purchased from RStudio.

You can activate your license key with the command:

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

After activation, we recommend restarting the RStudio Connect server. A change in license status will eventually be detected by the product; a forced restart ensures that change is seen immediately.

```
$ sudo stop rstudio-connect  
$ sudo start rstudio-connect
```

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

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

```
$ sudo /opt/rstudio-connect/bin/license-manager deactivate
```

3.1 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-connect/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.

3.2 Offline Activation

If your system has no connection to the internet it's also possible to perform an offline activation. To do this, you first generate an offline activation request as follows:

```
$ sudo /opt/rstudio-connect/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 and then send 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-connect/bin/license-manager activate-offline <activation-file>
```

After activation, we recommend restarting the RStudio Connect server. A change in license status will eventually be detected by the product; a forced restart ensures that change is seen immediately.

```
$ sudo stop rstudio-connect  
$ sudo start rstudio-connect
```

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

If you want to move your license of RStudio Connect to another system you can also perform license deactivation offline. You can do this as follows:

```
$ sudo /opt/rstudio-connect/bin/license-manager deactivate-offline
```

Executing this command will print an offline deactivation request to the terminal which you should copy and paste and then send to RStudio customer support (support@rstudio.com).

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

```
$ sudo /opt/rstudio-connect/bin/license-manager status-offline
```

4 Files & Directories

4.1 Program Files

The RStudio Connect installers place all program files into the `/opt/rstudio-connect` directory.

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

4.2 Configuration

The RStudio Connect configuration file is `/etc/rstudio-connect/rstudio-connect.gcfg`. You will edit this file to properly configure RStudio Connect for your organization.

A configuration management tool like Puppet or Chef can be used to maintain the `rstudio-connect.gcfg` file. We recommend that it be owned by `root` and have permissions `0600`, as your configuration may need to contain passwords and other sensitive information.

RStudio Connect upgrades will not overwrite customizations to the `rstudio-connect.gcfg` file.

4.3 Server Log

The RStudio Connect server log is located at `/var/log/rstudio-connect.log`. This file is owned by `root` with permissions `0600`.

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

4.4 Access Logs

The RStudio Connect HTTP access logs are located at `/var/log/rstudio-connect.access.log`. This file is owned by `root` 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 Connect 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.

4.5 Application Logs

Each R process launched by RStudio Connect produces output that is retained within the `jobs` subdirectory of the RStudio Connect data directory (see Section 4.6 for details). These directories and files are managed by the server. They are retained for 30 days and subsequently removed from the system.

Application logs are available in the RStudio Connect dashboard. The dashboard settings page for deployed content contains a **Logs** section containing execution details for each launched R process. Standard output and standard error are captured and available.

4.6 Variable Data

RStudio Connect manages uploaded Shiny applications, R Markdown documents, and plots. All of the variable data associated with this content is stored within the server's data directory. This includes:

- Deployment bundles as uploaded by the user.
- Directories containing unpacked bundles, including R source code.
- R packages, as demanded by the deployed code.
- Rendered R Markdown documents.

The RStudio Connect data directory also contains information used by the server in managing your deployed content. This includes:

- The RStudio Connect `sqlite` database and encryption key.

- R process execution information including logged output.
- Parameter overrides for R Markdown documents.

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

```
[Server]
DataDir = /mnt/rstudio-connect
```

The RStudio Connect database **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 `Database` section of your server configuration file.

```
[Server]
DataDir = /mnt/rstudio-connect

[Database]
Dir = /var/lib/rstudio-connect/db
```

4.6.1 Permissions

Data directory permissions are established by RStudio Connect as files are created. This section documents the general ownership patterns you will find under the RStudio Connect data directory.

Directories directly accessed from R applications will usually be owned by the `Applications.RunAs` user. This setting defaults to use an `rstudio-connect` account created during RStudio Connect installation. The `rstudio-connect` account has a default primary group also named `rstudio-connect`. We use the account and group name `rstudio-connect` throughout this section instead of referencing the property name.

Directories used during metrics collection are owned by the `rstudio-connect` user (customizable via the `Metrics.User` setting).

Learn more about customizing metrics collection in Section 14.1.

Directories not accessed by R applications or by the monitoring system will be owned by `root`.

`/var/lib/rstudio-connect` is owned by `root` with permissions `0701`.

The `R` subdirectory contains R packages used when content is deployed. The entire `R` directory hierarchy needs to be owned by `rstudio-connect`. Files must have `0600` permissions and directories need `0700` permissions.

The `packrat` subdirectory contains R packages installed on behalf of deployed content. These packages are installed when content is deployed and subsequently used when an application or report executes. The entire `packrat` directory hierarchy needs to be owned by the `rstudio-connect` and the `rstudio-connect` group. Files must have `0640` permissions while directories need `0750` permissions.

The `reports` subdirectory is owned by `root` with `0711` permissions. This contains generated output for report content deployed with source. The nested directories are written to by R processes and are owned by `rstudio-connect` with `0700` permissions. Files contained in this hierarchy will have `0600` permissions.

The `apps` directory contains directories for each deployment. The top-level directory is owned by `root` with `0711` permissions. The first level of the `apps` hierarchy is a directory for each content deployment. These `apps/A_ID` directories are owned by `rstudio-connect` with `0700` permissions.

Beneath each `apps/A_ID` directory is a set of directories for each deployed bundle. The ownership and permissions for this hierarchy depend on whether or not the content is configured with a custom `RunAs` setting. Without a custom `RunAs` setting, permissions are simple: owned by `rstudio-connect` with directories having `0700` and files having `0600` permissions.

Learn more about using a custom `RunAs` in Section 10.3.

RStudio Connect needs a more complicated permission structure when content is configured with a custom `RunAs` setting. This is because the `rstudio-connect` user (`Applications.RunAs`) is used to install the necessary packages while the content-specific custom `RunAs` is used when running the deployed R code. The `apps/A_ID/B_ID` directory is owned by the custom `RunAs` with group ownership set to `rstudio-connect`. Permissions on this directory are `0750`. The `packrat` subdirectory is owned by `rstudio-connect` with group ownership of `rstudio-connect`. File permissions on this directory and its sub-directories are `0750` while files have `0640` permissions. Other than the `packrat` directory, all files underneath `apps/A_ID/B_ID` have `0600` permissions and directories are given `0700`.

All other data subdirectories are owned by `root` with `0700` permissions.

4.7 Backups

We recommend including the RStudio Connect configuration file in `/etc/rstudio-connect` as well as the variable data directory which defaults to `/var/lib/rstudio-connect` in your system backups. If you have configured the database to be stored outside the data directory, ensure that it is also included in the backup.

A running RStudio Connect server may be writing into the data directory if there are any active deployments, applications or documents. You should stop the RStudio Connect server before taking a backup.

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

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

5 Server Management

This section describes common administrative tasks for RStudio Connect.

5.1 Stopping and Starting

Occasionally it is necessary to start and stop the RStudio Connect service. Stopping and starting is handled by `systemd` or `Upstart`. On stop/start or restart the following occurs:

Stop:

- The RStudio Connect process is stopped.
- R processes serving Shiny applications are stopped.
- R processes rendering R Markdown documents run through completion.
- In-progress deployments will fail. R processes running as part of the deployment may run to completion.

Start:

- RStudio Connect process is resumed.
- Shiny applications with a minimum number of R processes are started.
- Scheduled R Markdown updates missed during system downtime are run at most once.

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

5.1.1 systemd (Red Hat/CentOS 7, Ubuntu 15.04+)

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

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

```
$ sudo systemctl start rstudio-connect
```

```
$ sudo systemctl stop rstudio-connect
```

You can restart the server with:

```
$ sudo systemctl restart rstudio-connect
```

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

```
$ sudo systemctl kill -s HUP --kill-who=main rstudio-connect
```

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

We do not recommend using `reload` or a `HUP` signal to reconfigure RStudio Connect, as many property changes will not fully propagate through the system. This may adversely affect the behavior of RStudio Connect. A full restart of the service is preferred.

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

```
$ sudo systemctl status rstudio-connect
```

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

```
$ sudo systemctl enable rstudio-connect
```

```
$ sudo systemctl disable rstudio-connect
```

5.1.2 Upstart (Ubuntu 12.04 through 14.10, Red Hat 6)

Upstart is a system used to automatically start, stop and manage services. The installer writes an Upstart configuration file to `/etc/init/rstudio-connect.conf`. This instructs the Upstart to initialize RStudio Connect 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 `connect` process is respawned if the process unexpectedly terminates. However, in the event that there is an issue which consistently prevents RStudio Connect 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 Connect startup attempt before it transitions to the “stopped” state.

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

```
$ sudo start rstudio-connect
```

```
$ sudo stop rstudio-connect
```

To restart the server you can run:

```
$ sudo stop rstudio-connect
$ sudo start rstudio-connect
```

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-connect.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-connect
```

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

Known Bug: Due to a bug in the version of Upstart which comes with Ubuntu 13.04, `reload` does not behave as expected on that platform and should not be used.

We do not recommend using `reload` or a HUP signal to reconfigure RStudio Connect, as many property changes will not fully propagate through the system. This may adversely affect the behavior of RStudio Connect. A full restart of the service is preferred.

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

```
$ sudo status rstudio-connect
```

5.2 System Messages

Administrators can add a message to the RStudio Connect welcome page and content page.

Messages are set in the `/etc/rstudio-connect/rstudio-connect.gcfg` file. `Server.PublicWarning` defines the message for the welcome page. `Server.LoggedInWarning` defines the message for the content page. The messages are supplied as HTML snippets. For example:

```
[Server]
PublicWarning = "<strong>Warning:</strong> Scheduled downtime this weekend."
LoggedInWarning = "Data Science Team Meeting Tomorrow"
```

Messages can be added or modified without restarting the connect service. After adding the message property to the config file, use the reload commands for either `systemd` (Red Hat/CentOS 7, Ubuntu 15.04+):

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

or `Upstart` (Ubuntu 12.04 through 14.10, Red Hat 6):

```
sudo reload rstudio-connect
```

5.3 Upgrading

Upgrading RStudio Connect requires limited downtime. Scheduled R Markdown documents are not interrupted. Connections to running Shiny applications are closed. We recommend upgrading during a period of downtime. Users can be warned ahead of an upgrade with system messages.

The RStudio Connect version number is visible on the lefthand navigation pane. The latest version is available on the download page along with release notes.

To upgrade:

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

Ubuntu:

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

Red Hat/CentOS:

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

The new version of RStudio Connect will install on top of an earlier installation. Existing configuration settings are respected. During installation the RStudio Connect service is restarted. Total downtime is less than 10 minutes.

5.4 Purging RStudio Connect

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

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

Ubuntu:

```
sudo apt-get purge rstudio-connect
```

Red Hat/CentOS:

```
sudo yum remove rstudio-connect
```

3. Remove `/opt/rstudio-connect` if it still exists.
4. Remove logs from `/var/log/rstudio-connect*`
5. Remove the `Database.Dir` directory. This has a default location of `/var/lib/rstudio-connect/db`.
6. Remove the `Server.DataDir` directory. By default, this is `/var/lib/rstudio-connect`.
7. Remove configuration files from `/etc/rstudio-connect` if they still exist.

6 Running with a Proxy

6.1 Overview

If you are running RStudio Connect 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 Connect. This section describes how to correctly configure a reverse proxy with Nginx.

When RStudio Connect is behind a proxy, it is important that we send the original request URL information to Connect so that it can generate FQDN URLs and return them the requester. For this reason, when proxying to Connect, we must add a header, `X-RSC-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://connect.company.com/some/path`)

6.1.1 Nginx Configuration

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

```
sudo apt-get install nginx
```

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

```
sudo yum install nginx
```

To enable an instance of Nginx running on the same server to act as a front-end proxy to RStudio Connect you would add commands like the following to your `nginx.conf` file. This configuration assumes RStudio Connect is running on the same host as Nginx and listening for HTTP requests on the `:3939` port. If you are proxying to RStudio Connect on a different machine or port, replace the `localhost:3939` references with the correct address of the server where RStudio Connect is hosted.

```
http {
    map $http_upgrade $connection_upgrade {
        default upgrade;
        ''      close;
    }
    server {
        listen 80;

        location / {
            proxy_set_header    X-RSC-Request $scheme://$host:$server_port$request_uri;
            proxy_pass http://localhost:3939;
            proxy_set_header Upgrade $http_upgrade;
            proxy_set_header Connection $connection_upgrade;
            proxy_http_version 1.1;
        }
    }
}
```

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

```
http {
    map $http_upgrade $connection_upgrade {
        default upgrade;
        ''      close;
    }
    server {
        listen 80;

        location /rstudio/ {
            rewrite ^/rstudio/(.*)$ /$1 break;
            proxy_set_header    X-RSC-Request $scheme://$host:$server_port$request_uri;
            proxy_pass http://localhost:3939;
            proxy_redirect / /rstudio/;
            proxy_set_header Upgrade $http_upgrade;
            proxy_set_header Connection $connection_upgrade;
            proxy_http_version 1.1;
        }
    }
}
```

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

On `systemd` systems (Red Hat/CentOS 7, Ubuntu 15.04+):


```
sudo systemctl restart nginx
```

On `upstart` systems (Ubuntu 12.04 through 14.10, Red Hat 6):

```
sudo restart nginx
```

7 Security & Auditing

7.1 Browser Security

There are a variety of security settings that can be configured in RStudio Connect. 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.

7.1.1 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 Connect 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 Connect cookies to a server without an HTTPS-secured connection.

7.1.2 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.

7.1.3 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.

For the purposes of the `X-Frame-Options` header, RStudio Connect distinguishes between “dashboard” and “user” content. Dashboard content are any of the internal services or assets that are shipped with RStudio Connect. User content is anything uploaded by a user (reports, Shiny applications, etc.)

`Server.FrameOptionsContent` configures the `X-Frame-Options` header value for user-uploaded content. By default it is empty, meaning that the header will not be set. This allows user-provided content to be embedded in iframes from any location. If you do not intend for others to embed user content on their sites, you can set this to a value of `SAMEORIGIN` to ensure that only sites on the same server will be able to embed your users'

content. The RStudio Connect dashboard itself uses iframes to present user content in the dashboard, so it is not recommended to set this option to `DENY`.

`Server.FrameOptionsDashboard` configures the `X-Frame-Options` header value for internal services and assets provided with RStudio Connect and defaults to a value of `DENY`. This means that other sites will not be able to embed the RStudio Connect dashboard. This setting is more secure in that it protects against clickjacking attacks against the dashboard, but if you plan to embed the dashboard elsewhere you may need to tune this setting.

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

7.1.4 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 Connect 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 Connect. 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 Connect or by Shiny. You should not depend on a custom header that conflicts with a header already in use by RStudio Connect.

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:

```
[Server]
CustomHeader = "HeaderA: some value"
CustomHeader = "HeaderB: another value"
```

7.2 Audit Logs

The following events are logged by the auditing system:

Event	Description
<code>add_user</code>	Create a user
<code>edit_user</code>	Change an existing user
<code>update_lock_user</code>	Set or remove a lock for an existing user
<code>add_application</code>	Add new content
<code>upload_bundle</code>	Upload a bundle for a content
<code>deploy_application</code>	Deploy content to the server. Note that content still needs to be published after deployment.
<code>edit_application</code>	Change content settings
<code>remove_application</code>	Delete content
<code>activate_token</code>	Activate a token. Tokens are used by the <code>rsconnect</code> package to authenticate a user.
<code>add_group</code>	Create a group
<code>remove_group</code>	Delete a group

Event	Description
<code>add_group_member</code>	Add a user to a group
<code>remove_group_member</code>	Remove a user from a group
<code>assign_user_app_role</code>	Give a user view or edit access to content
<code>remove_user_app_role</code>	Remove a user from view or edit access list
<code>assign_group_app_role</code>	Give a group view or edit access to content
<code>remove_group_app_role</code>	Remove a group from view or edit access list
<code>clear_app_viewer_acl</code>	Change from a specific list of viewers to “just me”
<code>add_api_key</code>	Added API key
<code>remove_api_key</code>	Removed API key
<code>add_vanity</code>	Add vanity url
<code>update_vanity</code>	Update vanity url
<code>remove_vantiy</code>	Remove vanity url
<code>remove_bundle</code>	Remove a bundle
<code>download_bundle</code>	Download a bundle

8 Authentication

RStudio Connect supports a variety of user authentication options. Without customization, a locally-backed password scheme is used. You can learn more about password authentication in Section 8.2.

When signing into RStudio Connect, a session cookie is used to keep a user logged in for 30 days.

External authentication is available through the following integrations:

- LDAP and Active Directory (Section 8.3)
- OAuth 2.0 using Google Apps accounts (Section 8.4)
- PAM (Section 8.5)
- Proxied Authentication (Section 8.6)

Customize the `Authentication.Provider` property with an authentication scheme appropriate for your organization. See Section 15.6 for details

Here is a partial configuration which chooses to use LDAP.

```
[Authentication]
Provider = ldap
```

8.1 Migrating to a different Authentication Provider

Migrating from one authentication provider to another (for example, password → LDAP) is **NOT SUPPORTED**. If you absolutely need to, you’ll need to completely purge and reinstall RStudio Connect. See Section 5.4 for instructions.

8.2 Password

Password authentication is the default authentication provider used by RStudio Connect. This is a local user account backed by the RStudio Connect database and is not integrated with a third-party service.

Users will be able to create accounts when they first visit the system and will provide profile details at that time. An administrator will also be able to create new accounts.

Password authentication may be appropriate in small organizations without centralized IT systems.

RStudio Connect will use password authentication if the `Authentication.Provider` setting has a value of `password` or if `Provider` is not present in the configuration file.

```
[Authentication]
Provider = password
```

8.3 LDAP and Active Directory

RStudio Connect can integrate with your company's LDAP or Active Directory (AD) infrastructure. User authentication and user search requests will be directed to the LDAP/AD server.

LDAP and Active Directory support in RStudio Connect has the following constraints:

- Your LDAP/AD user objects must contain a user's first name, last name, email address, and username.
- Changes to a user (e.g. their name, email address, or username) will not propagate to RStudio Connect once the user is created internally.
- When using single bind, the DN of a user must contain their username (i.e. must utilize the `UsernameAttribute`). For example, it is not supported if the DN for a user is `cn=SueJacobs,ou=People,dc=company,dc=com` but their actual username is stored in the `uid` or `SAMAccountName` LDAP attribute. You must use double bind when the DN does not contain the username.
- When using a single bind configuration, searches will only include users who have previously logged into RStudio Connect.
- When using a single bind configuration, groups will be unavailable.
- A username or DN containing a forward slash (/) is not supported.

When attempting to troubleshoot a problem relating to LDAP, you can enable more verbose logging by adding `ldap` to `Debug.Log` section in the configuration.

```
[Debug]
Log = ldap
```

8.3.1 Defining an LDAP or AD section

RStudio Connect does support the notion of having multiple LDAP or AD servers. This can be utilized by defining multiple LDAP sections.

To define an LDAP or AD section in the configuration file, add a header like the following:

```
[LDAP "European LDAP Server"]
...
```

An LDAP/AD configuration section header is always bounded by square brackets (`[]`). After the section type `LDAP` is the effective name of the LDAP or AD server ("`European LDAP Server`" in the example). Make sure that this text is unique per LDAP or AD section you configure. The LDAP section name is treated *case sensitively*.

RStudio Connect can support more than one LDAP server through multiple, uniquely named LDAP configuration sections. Other complex LDAP configurations can also be achieved by using multiple LDAP sections.

If multiple LDAP sections have the same name, they will be combined as described in Appendix 15. As this is unlikely your intent, please take care to give unique names to each LDAP configuration section.

Here is an sample configuration using two LDAP sections.

```
[LDAP "European LDAP Server"]
...

[LDAP "Statistics Department LDAP Server"]
...
```

Each of these sections will have a variety of configuration settings, which are explained below.

8.3.2 Complete Configuration Example

Here is a complete LDAP configuration as an example. Here, we are communicating with an OpenLDAP server on the local host; see the documentation for `ServerAddress` to learn how to direct requests elsewhere. The other settings will probably need adjustment for your environment. Talk to your LDAP administrator if you need help with your organization's LDAP hierarchy.

```
[LDAP "Sample OpenLDAP Configuration"]
ServerAddress = 127.0.0.1:389
BindDN = "cn=admin,dc=example-openldap"
BindPassword = "XXXXXXXX"
UserSearchBaseDN = "ou=People,dc=example-openldap"
UsernameAttribute = "uid"
UserObjectClass = "posixAccount"
UserEmailAttribute = mail
UserFirstNameAttribute = givenName
UserLastNameAttribute = sn
```

This sample configuration assumed a very simple OpenLDAP structure; here is a sample user record to show the mapping between LDAP records and RStudio Connect LDAP configuration.

```
dn: uid=john,ou=People,dc=example-openldap
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
uid: john
sn: Doe
givenName: John
cn: John Doe
displayName: John Doe
uidNumber: 10000
gidNumber: 5000
userPassword: johnldap
gecos: John Doe
loginShell: /bin/bash
homeDirectory: /home/john
mail: john@example.com
```

More LDAP configuration scenarios can be found in Appendix 16.

8.3.3 LDAP or AD Configuration Settings

8.3.3.1 ServerAddress

`ServerAddress` (required) is used to define the location of the LDAP/AD server. This should contain an IP address or DNS address, and a port (colon separated). Most LDAP/AD servers operate on port 389 or 636 (for SSL). But you can specify any port that fits your environment.

Examples

```
ServerAddress = 127.0.0.1:389
ServerAddress = ldap.company.com:389
ServerAddress = ldaps.company.com:636
ServerAddress = private.internal.local:7554
```

8.3.3.2 TLS

TLS is a Boolean (true/false) attribute that causes all connections to your LDAP/AD server to use TLS (SSL). The default value for this is `false`. This cannot be enabled if `StartTLS` is `true`.

Examples

```
TLS = true
TLS = false
```

8.3.3.3 StartTLS

`StartTLS` is a Boolean (true/false) attribute that causes connections to your LDAP/AD server to initially use an unencrypted channel but then upgrade to a TLS connection using “Opportunistic TLS”. The default value for this is `false`. This cannot be enabled if `TLS` is `true`.

Examples

```
StartTLS = true
StartTLS = false
```

At present, the error messages associated with `StartTLS` problems can be cryptic. If you’re encountering issues while configuring `StartTLS`, consider adding debug logging for LDAP by including the following line in your configuration file.

```
[Debug]
Log = ldap
```

8.3.3.4 TLSCACertificate

`TLSCACertificate` is a file location that is a certificate authority that is used to connect to an LDAP server securely. This file should be in PEM format.

Examples

```
TLSCACertificate= /etc/ssl/cert/ca.pem
```

8.3.3.5 ServerTLSInsecure

`ServerTLSInsecure` is a Boolean (true/false) attribute that allows insecure TLS connections. This controls whether a client will verify the server’s certificate chain and host name. If this is `true`, RStudio Connect will accept any certificate presented by the server and any host name in that certificate. Setting to `true` is susceptible to man-in-the-middle attacks, but is required if, for example, your server uses a self-signed certificate. The default value is `false`.

Examples

```
ServerTLSInsecure = true
ServerTLSInsecure = false
```

8.3.3.6 BindDN and BindPassword

`BindDN` and `BindPassword` are credentials used to connect to an LDAP/AD server to authenticate, search for users, and other functionality. While it is encouraged to specify these two attributes (a.k.a. “double bind”), it is not required (a.k.a. “single bind”). These credentials should have read-only administrator’s rights, if configured.

If you do not specify these attributes, some functionality of RStudio Connect will not work. For example, searching for users to add as collaborators, or sending email documents will only work partially.

The `BindDN` can be a DN, UPN, or NT-style login.

Examples

```
# Example DN
BindDN = uid=john,ou=People,dc=company,dc=com
BindPassword = johnpassword

# Example UPN
BindDN = admin@company.com
BindPassword = adminpassword

# Example NT-style login
BindDN = COMPANY\\admin # we use double slashes (\\) to character escape the last slash
BindPassword = adminpassword
```

8.3.3.7 AnonymousBind

`AnonymousBind` instructs RStudio Connect to establish an anonymous bind to your LDAP/AD server. For organizations that support anonymous binds, you may use this option instead of `BindDN` and `BindPassword`.

For this to work properly, your LDAP server must allow anonymous binds to search and view all pertinent groups, group memberships, and users.

Examples

```
AnonymousBind = true
```

8.3.3.8 UserSearchBaseDN

`UserSearchBaseDN` (required) is the starting point from which RStudio Connect will search for user entries in your LDAP/AD server.

Examples

```
UserSearchBaseDN = dc=company,dc=com
UserSearchBaseDN = ou=People,dc=company,dc=com
```

8.3.3.9 UserObjectClass

`UserObjectClass` (required) is the `objectClass` that a user in your LDAP/AD structure will have. Common examples of this are `user`, `posixAccount`, `organizationalPerson`, `person`, and `inetOrgPerson`.

Examples

```
UserObjectClass = user
UserObjectClass = posixAccount
```

8.3.3.10 UsernameAttribute

UsernameAttribute (required) is the LDAP entry attribute that contains the username of a user.

Examples

```
UsernameAttribute = uid  
UsernameAttribute = sAMAccountName
```

8.3.3.11 UserFirstNameAttribute

UserFirstNameAttribute (required) is the LDAP entry attribute that contains the first name of a user.

Examples

```
UserFirstNameAttribute = givenName
```

8.3.3.12 UserLastNameAttribute

UserLastNameAttribute (required) is the LDAP entry attribute that contains the last name of a user.

Examples

```
UserLastNameAttribute = sn
```

8.3.3.13 UserEmailAttribute

UserEmailAttribute (required) is the LDAP entry attribute that contains the email address of a user.

Examples

```
UserEmailAttribute = mail
```

8.3.3.14 WhitelistedLoginGroup

WhitelistedLoginGroup defines a group DN that a user must be a member of in order to login into Connect. You can specify this attribute multiple times. Be aware that this feature restricts *only* the ability for users to login. Users not in this group could still be referenced when setting access controls for content or as email recipients. Because the users could not login, they would not be able to access content even if they were added as a viewer or collaborator, but they might still be able to receive emailed versions of reports.

Examples

```
WhitelistedLoginGroup = cn=admins,ou=group,dc=company,dc=com  
WhitelistedLoginGroup = cn=scientists,ou=group,dc=company,dc=com
```

8.3.3.15 GroupObjectClass

GroupUserObjectClass is the objectClass that a group in your LDAP/AD structure will have. Common examples of this are `group`, and `posixGroup`.

Examples

```
GroupObjectClass = group  
GroupObjectClass = posixGroup
```


8.3.3.16 GroupNameAttribute

GroupNameAttribute is the LDAP entry attribute that contains the name of a group.

Examples

```
GroupNameAttribute = cn  
GroupNameAttribute = sAMAccountName
```

8.3.3.17 GroupSearchBaseDN

GroupSearchBaseDN is the starting point from which RStudio Connect will search for group entries in your LDAP/AD server.

Examples

```
GroupSearchBaseDN = dc=company,dc=com  
GroupSearchBaseDN = ou=Groups,dc=company,dc=com
```

8.4 OAuth2 (Google)

OAuth2 authentication is available to authenticate against the Google OAuth2 service.

RStudio Connect will use OAuth2 authentication if the `Authentication.Provider` setting has a value of `oauth2`.

```
[Authentication]  
Provider = oauth2
```

Appendix 15.7 contains information about each OAuth2 configuration option.

In order for RStudio Connect to use Google as an OAuth2 service, you will need a client ID and client secret.

8.4.1 Obtaining a Client ID and Client Secret

These instructions tell you how to obtain an OAuth2 client ID and client secret. We recommend a distinct set of credentials for each application you configure to use the Google OAuth2 service.

1. Visit the Google Developers Console and create a new project. Give it a name of your choosing, such as “rstudio-connect”.
2. Once the project is created, locate and enable the “Google+ API”.
3. In the left navigation window, click on “Credentials”, then goto the “OAuth consent screen” tab, fill in the information requested and click “Save”.
4. Once again, click “Credentials” in the left navigation window. Then click the dropdown button “New credentials”, then “OAuth client ID”.
5. For “Application Type”, select “Web Application”. Then give your client ID a descriptive name. For “Authorized JavaScript origins”, enter your RStudio Server URL (i.e. `https://HOST:PORT`). For “Authorized redirect URIs”, use your RStudio Connect server address with `/__login__/callback` (i.e. `https://HOST:PORT/__login__/callback`).
6. Click “Create”. Your client ID and client secret will be shown to you.

Add the client ID and secret to your configuration file as shown in the example below.

```
[OAuth2]  
DiscoveryEndpoint = https://accounts.google.com/.well-known/openid-configuration  
ClientId = <CLIENT ID>  
ClientSecret = <CLIENT SECRET>
```

With `DiscoveryEndpoint`, `ClientId` and either `ClientSecret` or `ClientSecretFile` configured, you can use your Google Apps account to sign into RStudio Connect!

8.4.2 Restricting Access

The default configuration allows all Google account holders to access RStudio Connect. We recommend that you limit access to specific domains that are used by your organization.

Verify that you can use your Google Apps account to sign into RStudio Connect before attempting to configure access restrictions.

The `OAuth2.AllowedDomains` setting specifies the set of domains that are allowed to access your RStudio Connect server. Multiple domains should be space-separated.

```
[OAuth2]
AllowedDomains = company.com subsidiary.com
```

You may also restrict access by email address if using domain alone is insufficient. The `OAuth2.AllowedEmails` setting specifies the set of email addresses that are allowed to access your RStudio Connect server. Multiple addresses should be space-separated.

```
[OAuth2]
AllowedEmails = jdoe@company.com asmith@subsidiary.com
```

It is important to understand how the `AllowedDomains` and `AllowedEmails` properties interact.

If only `AllowedDomains` is configured, *only* email addresses with a listed domain will be permitted access.

If only `AllowedEmails` is configured, *only* listed email addresses will be permitted access.

When both `AllowedDomains` and `AllowedEmails` are specified, email addresses given in `AllowedEmails` are permitted access *in addition* to email addresses with a domain listed in `AllowedDomains`.

8.4.3 Searches

RStudio Connect allows users to search for collaborators against the user directory associated with your Google Apps account. That search is performed on behalf of the current user. Different accounts may have different visibility within the user directory and therefore will see different results. This is most obvious when you have configured RStudio Connect to allow access to two different domains. Users in `company.com`, for example, will likely not be able to search for colleagues in `subsidiary.com`.

RStudio Connect augments the Google Apps user directory search with a local search across its set of known accounts. Once your colleague has created their own RStudio Connect account, they will become discoverable.

8.5 PAM

RStudio Connect can use PAM for user authentication. PAM authentication is used if the `Authentication.Provider` setting has a value of `pam`.

```
[Authentication]
Provider = pam
```

See Section 10.5 for information about using PAM sessions when launching R processes.

You can change the PAM service name used for authentication by customizing the `PAM.Service` setting. The default PAM service name used for authentication is `rstudio-connect`.

```
[PAM]
```

```
Service = rstudio-connect
```

We assume that RStudio Connect is configured to use the `rstudio-connect` PAM service name in the examples that follow.

8.5.1 Ubuntu

RStudio Connect does not create a PAM service on Ubuntu systems. When RStudio Connect attempts to use the `rstudio-connect` service name for authentication, PAM will recognize that there is no service with that name and fall back to the default `other` service located at `/etc/pam.d/other`.

The default Ubuntu `other` service is configured to inherit from a set of common PAM services:

```
# Ubuntu default "other" PAM service.
@include common-auth
@include common-account
@include common-password
@include common-session
```

If the `other` service is appropriate for your organization, no further configuration is needed.

You need a custom `rstudio-connect` PAM service for RStudio Connect only if the `other` service is not fitting for your users. Create and configure `/etc/pam.d/rstudio-connect` to prevent PAM from falling back to the `other` service. PAM will use this service for subsequent authentication attempts using the `rstudio-connect` service name.

8.5.2 Red Hat/CentOS

Red Hat/CentOS systems deny access to unknown PAM service names by default. This is because the `other` configuration in `/etc/pam.d/other` contains only “deny” rules.

```
##%PAM-1.0
# The Red Hat/CentOS default "other" PAM service.
auth      required      pam_deny.so
account   required      pam_deny.so
password  required      pam_deny.so
session   required      pam_deny.so
```

The RStudio Connect RPM installs an `rstudio-connect` PAM service at `/etc/pam.d/rstudio-connect`. This service is configured to require a user-id greater than 500 and authenticates against local system accounts.

```
##%PAM-1.0
# The RStudio Connect default PAM service.
auth      requisite     pam_succeed_if.so uid >= 500 quiet
auth      required     pam_unix.so nodelay
account   required     pam_unix.so
```

This default PAM service may not reflect the authentication behavior that you want for RStudio Connect. Feel free to customize this service for your organization.

8.5.3 Configuring a PAM service

This section may be helpful if your organization has different requirements from the default behavior of the `rstudio-connect` PAM service name. Please consult with your PAM/systems administrator to be sure that

the RStudio Connect PAM service configuration fits your needs.

If your system already has a PAM service (e.g. `/etc/pam.d/login`) with the desired behavior, it may be enough to simply include that service from within the RStudio Connect service. For example:

```
# RStudio Connect PAM service that defers to the existing login service.  
@include login
```

You could also copy that existing service into the RStudio Connect service, meaning the copy can be changed and evolve independently from the source service.

```
$ sudo cp /etc/pam.d/login /etc/pam.d/rstudio-connect
```

Lastly, you could configure the `PAM.Service` setting to reference that PAM service. This would be appropriate if you have a common `rstudio` service that you use across all the RStudio products, for example.

```
[PAM]  
Service = rstudio
```

If you change the `PAM.Service` setting from its default `rstudio-connect` value, the PAM service defined in `/etc/pam.d/rstudio-connect` will not be used.

8.5.4 Groups

Groups are not supported when using PAM authentication.

8.6 Proxied Authentication

RStudio Connect supports proxied authentication. This allows an external system to intercept requests and handle the authentication of users visiting the Connect dashboard or applications Connect is hosting.

8.6.1 How this Works

A service (like Apache, for example) runs as your customized authentication server. It is responsible for intercepting all requests to RStudio Connect and performing the required authentication and authorization. Requests from authenticated users will have a custom HTTP header added before the request is proxied through to RStudio Connect. That HTTP header contains the username of that visitor. RStudio Connect will take the value from the HTTP header and treat the current user as the username specified in the header.

We have no means of validating that this HTTP header was added by your authentication server and not by the user directly. So it is very important from a security perspective that the RStudio Connect server is properly firewalled off in your network and that all access to the Connect server is proxied through your authentication server.

Important Note

The username HTTP header should never be set by the requester. In all cases, your authentication server should delete that header if it exists before authenticating the user and adding the header itself.

RStudio Connect does not currently support directing users to a login page when using proxied authentication. Therefore, we recommend that your proxy prevent anonymous access to RStudio Connect; only allow authenticated users.

8.6.2 Deployment from the RStudio IDE

Deploying from the RStudio IDE is a unique situation. The IDE uses an R package `rsconnect` to obtain deployment credentials from RStudio Connect. Those credentials are used to sign deployment requests.

Deployment requests are signed with credentials obtained during an earlier, authenticated session, and should pass through your proxy without alteration.

The following three headers when used together identify deployment requests and should pass through your proxy without attempting to authenticate the user:

- `X-Auth-Token`
- `X-Auth-Signature`
- `X-Content-Checksum`

8.6.3 Configuring Proxied Authentication

To configure RStudio Connect to use proxied authentication, set `Authentication.Provider` to `proxy`.

```
[Authentication]
Provider = proxy
```

Proxied authentication requires that you set `Server.Address` to point at your proxy server. If you do not configure `Server.Address`, the browser may not have all its requests routed through your authenticating proxy. See Section 2.2.1 for more information about `Server.Address`.

```
[Server]
Address = https://myproxy.company.com/
```

You can customize the name of the header that your authentication server will send upon a successful authentication. By default, this key name is `X-Auth-Username`.

```
[ProxyAuth]
UsernameHeader = X-Auth-Username
```

8.6.4 Groups

Groups are not supported when using proxied authentication.

9 User Management

9.1 User Roles

Every RStudio Connect user account is configured with a role that controls their default capabilities on the system. Data scientists, analysts and others working in R will most likely want “publisher” accounts. Other users are likely to need only “viewer” accounts.

The `DefaultUserRole` property within the `Authorization` configuration section specifies the role for new accounts and defaults to “publisher”.

Administrator RStudio Connect administrator accounts have permissions which allow them to manage the service. This includes setting the role of an account and configuring email settings. Administrators may or may not be system administrators.

Publisher Accounts with a “publisher” role are allowed to deploy content into RStudio Connect. They can also help manage another user’s content when made a “collaborator” of that content.

Viewer “Viewer” accounts can be added as a viewer to specific content. They can discover that content through the RStudio Connect dashboard and see its settings. Viewers can also email themselves copies of documents they are permitted to see.

Anonymous An anonymous visitor to RStudio Connect who is not authenticated with the system can view content that has been marked as viewable by “Everyone”.

9.2 User Permissions

Administrators and Publishers can be assigned permissions for content published to RStudio Connect.

9.2.1 All Content

Anonymous Visitors Anonymous users can access content listed for **Everyone**. Anonymous viewers access content through direct URLs and will not have any view into Connect.

Viewers “Viewers” can sign into the Connect dashboard and discover and access content listed for **Everyone**, **All logged-in users**, and content for which they are granted access.

Collaborators “Collaborators” can change access controls and add Viewers and other Collaborators.

Administrators “Administrators” have all the permissions of Collaborators. Administrators are not automatically added to content and will not see all content on their homepage. Administrators can proactively add themselves as Collaborators or Viewers to any content. Administrators can set vanity URLs and change the **RunAs** user. Administrators and the original content owner can delete content.

9.2.2 R Markdown Reports

Access controls and user privileges apply to every public version of a report. For example, if the default version of a report is accessible to **Everyone**, all public versions will be accessible to **Everyone**.

Anonymous Visitors Every version of a report has a unique URL (accessible by opening the content with ‘Open Solo’). Reports must be listed for **Everyone** for the URL to be available to anonymous users.

Viewers “Viewers” have the ability to view a report through the Connect dashboard. They can discover and toggle between public versions of a report. They can email themselves the current version of a report. They can not see parameters for different versions of a report. They can see the distribution and schedule for public versions.

Collaborators “Collaborators” have the privileges of Viewers and additionally can: view parameters for public versions, change parameters and run ad hoc reports, create new versions, schedule versions, setup distribution lists, and request reports to be refreshed. Collaborators can also create private versions that are not discoverable or accessible by any other user.

9.2.3 Shiny Applications

Collaborators “Collaborators” can change the performance settings for applications.

9.3 Locked Accounts

You can prohibit a user from accessing RStudio Connect by “locking” their account. This control is available to administrative users when editing user profile information in the RStudio Connect dashboard.

Locked users are prohibited from signing into RStudio Connect, deploying content, and otherwise interacting with the service.

A locked account is not deleted and deployed content continues to be available. A non-personal report configured with scheduling and distribution will continue to execute according to its schedule. A locked user no longer receives scheduled content at their email address.

Content owned by a locked user can be deleted by a collaborator or by an administrative user. Each piece of deployed content must be deleted individually; there is no bulk removal.

A locked user can be subsequently unlocked. All their previously allowed abilities are immediately restored.

9.4 User Renaming

Administrators may alter the usernames of existing users on the system regardless of the current authentication system. Users will still be able to access their deployed content and content that has been shared with them. If they have existing vanity URLs with their username incorporated, none of those will be altered. They will, of course, need to use the new username when logging in.

If the user has authenticated inside of the RStudio IDE, they will still be able to deploy using a previous connection; however, the IDE will continue displaying their old username during deployments. To minimize the risk of future ambiguity, we recommend that the user disconnect and reconnect their IDE to RStudio Connect so that the valid username is displayed.

9.5 Command-Line Interface

Connect includes a `usermanager` command for some basic user management tasks. This utility helps you list users and modify user roles in the event that no one can access a Connect administrative user account.

The user management utility is installed at `/opt/rstudio-connect/bin/usermanager`. It uses the configuration defined in `/etc/rstudio-connect/rstudio-connect.gcfg` unless you specify an alternate configuration file with the `--config` flag.

The `usermanager` utility must be run as `root`.

The `usermanager` utility can only be run when Connect is stopped. See Section 5.1 for information on stopping and restarting Connect.

9.5.1 Commands

The `usermanager` utility supports two commands:

- `list`: Lists users
- `alter`: Changes a user's role

9.5.2 Flags

Configuration for `usermanager`:

- `--config`: The full or relative path to a Connect configuration file (`.gcfg`). Defaults to `/etc/rstudio-connect/rstudio-connect.gcfg`.

Flags for the `list` command:

- `--include-locked`: Includes locked user accounts in the list.

Flags for the `alter` command:

- `--username`: Specifies the user name of the user to alter.
- `--role`: Specifies the role to set for the user. Allowed roles are `viewer`, `publisher`, and `administrator`.
- `--force`: Force demotion of the last remaining administrator.

9.5.3 Examples:

List unlocked users:

```
sudo ./bin/usermanager list
```

List all users (locked and unlocked):

```
sudo ./bin/usermanager list --include-locked
```

Specify a custom configuration file

```
sudo ./bin/usermanager --config /etc/connect/mycustomconfig.gcfg list
```

Promote the user `john` to an administrator role

```
sudo ./bin/usermanager alter --username john --role administrator
```

Demote the last remaining administrator to a non-administrative role

```
sudo ./bin/usermanager alter --username admin --role publisher --force
```

10 Process Management

RStudio Connect launches R to perform a variety of tasks. This includes:

- Installation of R packages
- Rendering of R Markdown documents
- Running Shiny Applications
- Running a Shiny application to customize a parameterized R Markdown document.

The location of R defaults to whatever is in the path. Customize the `Server.RVersion` setting to use a specific R installation. See Chapter 12 for details.

10.1 Sandboxing

The RStudio Connect process runs as the `root` user. It needs escalated privileges to allow binding to protected ports and to create “unshare” environments that contain the R processes.

RStudio Connect runs its R processes as an unprivileged user; both a system default and content-specific overrides are supported. See Section 10.3 for details.

The “unshare” environment created for R execution involves first establishing a number of bind mounts and then switching to the target unprivileged user. RStudio Connect 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.

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.

The following locations are masked during R execution:

- The `Server.DataDir` directory containing all variable data used by RStudio Connect.

- The `Database.Dir` directory, which can optionally be placed outside the data directory.
- Configuration directories, including `/etc/rstudio-connect`.
- The `/tmp` and `/var/tmp` directories.

The following information is exposed during R execution:

- The `packrat` data directory (read-only except when installing packages).
- The R data directory (only when installing packages).
- The directory containing the unpackaged R code (Shiny and R Markdown).
- The document rendering destination directory (only for R Markdown).
- A per-process temporary directory (exposed over the original `/tmp` and `/var/tmp`).

When `Applications.HomeMounting` is enabled, the contents of `/home` are masked by an additional bind mount as follows:

- The contents of `/home` are masked by the home directory of the `RunAs` user.
- If the `RunAs` does not have a home directory, an empty directory masks `/home`.

The path to the home directory is always available through the `HOME` environment variable. With `Applications.HomeMounting`, the mounted path to the `HOME` directory is subject to change. Avoid hard-coding paths to either `/home` and `/home/username`.

Shiny applications have write access to the directory containing the unpackaged R code. This application directory is the working directory when launching Shiny. Data written here will be visible to all processes associated with that Shiny application but are not visible to other R processes. Application directory data remains available until that application is next deployed to RStudio Connect. A deployment creates a new application directory containing only the deployed content.

RStudio Connect may launch multiple processes to service requests for an application. There is no coordination between these processes. Shiny applications that write to local files could experience problems when different processes attempt to write to a single file. We recommend against using the file system for data persistence.

R Markdown documents have write access to the rendering destination directory and read access to a directory containing the unpackaged R code. The source directory is the working directory when calling `rmarkdown::render`. The destination directory is passed as the `output_dir` while a temporary directory is passed as the `intermediates_dir`. The intermediate directory is transient and not available after rendering completes. A new output directory is created whenever the document is rendered. Data created during one rendering is *not* visible to another.

R Markdown multi-document sites have a slightly different rendering pipeline than standalone documents. RStudio Connect uses the `rmarkdown::render_site` function, which does its rendering in-place. The content from the source directory is copied into the rendering destination directory in preparation for rendering. Site rendering has write access to the destination directory. Access to the original source directory is not provided because the source content is duplicated in the destination directory.

The `rmarkdown::render_site` call usually places its output into a subdirectory (typically, `'_site'`). The contents of this output subdirectory will be moved to the root of the rendering destination directory, replacing any other content. No post-rendering file movement occurs if `rmarkdown::render_site` is instructed to render into the current directory instead of a subdirectory. This means that both source and output files will be available for serving.

We recommend against configuring `rmarkdown::render_site` to write its output into the current directory. Rendering the site into a subdirectory (the default) allows RStudio Connect to remove source from the output directory.

RStudio Connect serves rendered content from the document output directory. This content remains available until a subsequent rendering is successful and activated (if requested). Neither incomplete nor unsuccessful document renderings affect the availability of previously rendered content.

10.2 Shiny Applications

Most of the R processes started by RStudio Connect are batch-oriented tasks. R is invoked, does a narrow set of work, and then exits. Shiny applications are different and may see an R process handle many requests for many users over its lifetime.

RStudio Connect launches an R process tied to a Shiny application when the first request arrives for that application. That R process will continue to service requests until it becomes idle and eventually terminated. If there is sufficient traffic against that Shiny application, RStudio Connect may launch additional processes to service those requests.

There are a number of configuration parameters which control the conditions under which processes for Shiny applications are launched and eventually reaped. The default values are appropriate for most applications but occasionally need customization in specialized environments. Section 15.13 explains each of the options.

We recommend that adjustment to these runtime properties be done gradually.

10.3 User Account for R Processes

The RStudio Connect installation creates a local `rstudio-connect` user account. This account runs all the R processes; `root` does not invoke R. If you would like a different user to run R, customize the `Applications.RunAs` property.

Administrators can customize the `RunAs` user on a content-specific level. This means that different Shiny applications and R Markdown reports can be run using different Unix accounts. This setting can be found on the *Access* tab when editing content settings. Publishers and Viewers are prohibited from changing the `RunAs` user on a content-specific level.

If you choose to specify a custom `RunAs` user for content, that user *must* be a member of the Unix group that is the primary group of the `Applications.RunAs` user.

The `rstudio-connect` user, for example, has a primary group also named `rstudio-connect`. Any Unix account configured as a custom `RunAs` user for a Shiny application or R Markdown report *must* be a member of the `rstudio-connect` group.

Installation of R packages always happens as the `Application.RunAs` user. A Shiny application or R Markdown report may override its `RunAs` setting; this alters how the deployed code is executed and does not impact package installation. See Section 10.1 for more information about process sandboxing.

10.4 Current user execution

RStudio Connect can use a local Unix account associated with the currently logged-in user when executing R. This feature requires that user authentication use PAM.

See Section 8.5 for information about using PAM for user authentication.

The `Applications.RunAsCurrentUser` property specifies that content can be configured to execute as the currently logged-in user.

```
[Applications]
RunAsCurrentUser = true
```

Administrators can now customize the `RunAs` settings to permit current-user execution on a content-specific level. The *Access* content setting tab offers the option of executing using “The Unix account of the current user”.

Content accessed anonymously will execute as the specified fallback `RunAs` user.

See Section 10.3 for more information about `RunAs` customization.

Content execution settings are not altered when `RunAsCurrentUser` is enabled. The `RunAsCurrentUser` setting *permits* current-user execution but by itself does not change how R processes are launched. Each Shiny application or R Markdown report must explicitly request current-user execution.

All Unix accounts used to execute R *must* be members of the Unix group that is the primary group of the `Applications.RunAs` user. Applications are not permitted to launch if the Unix account associated with the logged-in user does not have the proper group membership.

The `Applications.RunAs` setting uses the `rstudio-connect` user by default. This user has a primary group also named `rstudio-connect`. Any Unix account that may be used to execute Shiny applications or R Markdown reports *must* be a member of the `rstudio-connect` group.

10.5 PAM sessions

RStudio Connect can use PAM to establish the environment and resources available for R sessions.

See Section 8.5 for information about using PAM for user authentication.

PAM sessions are enabled with the `PAM.UseSession` setting.

```
[PAM]
UseSession = true
```

The default PAM service name used for PAM sessions is `su`. This gives RStudio Connect the ability to launch processes as the specified user without requiring a password.

You can customize the PAM service name used for PAM sessions by customizing the `PAM.SessionService` setting.

```
[PAM]
SessionService = rstudio-connect-session
```

Any custom PAM service must contain the PAM directive that enables authentication with root privilege.

```
# Allows root to su without passwords (required)
auth sufficient pam_rootok.so
```

10.6 Path Rewriting

The sandboxing used by RStudio Connect involves bind mounts which map physical locations on disk onto different directory structures at runtime. Paths used by your R code use these sandboxed locations. If you need to find the physical file on disk, you will need to undo the path transformation.

This section gives some examples of path rewriting and offer some ways of finding the file you need.

Let's start with an `app.R` file that describes a Shiny application. This file will be in the `apps/XX/YY/` directory underneath the `Server.DataDir` location. The `XX` and `YY` path components correspond to the application ID and bundle (or deployment) ID for this version of your application. This directory is available at runtime as `/opt/rstudio-connect/mnt/app/`.

The directory structure of `/opt/rstudio-connect/mnt/` is just a number of empty directories. The “unshare” environment created during sandboxing allows RStudio Connect to associate different application directories with these mount directories.

Here are some common path transformations that may be helpful. All of the physical paths are beneath the `Server.DataDir` hierarchy that defaults to `/var/lib/rstudio-connect`. All of the sandbox paths are beneath the mount directory `/opt/rstudio-connect/mnt/`. This location is not customizable.

Physical path	Sandbox path
DataDir/apps/XX/YY/	MountDir/app/
DataDir/reports/XX.ZZ	MountDir/report/
DataDir/R	MountDir/R
DataDir/packrat	MountDir/packrat

Here are some actual path transformations using the default `Server.DataDir` location:

```
# A source Shiny application
/var/lib/rstudio-connect/apps/4/7/app.R
=> /opt/rstudio-connect/mnt/app/app.R

# A source R Markdown document
/var/lib/rstudio-connect/apps/8/12/index.Rmd
=> /opt/rstudio-connect/mnt/app/index.Rmd

# An HTML document rendered from that R Markdown document
/var/lib/rstudio-connect/reports/8.2/index.html
=> /opt/rstudio-connect/mnt/report/index.html

# A statically deployed document
/var/lib/rstudio-connect/apps/17/21/index.html
=> /opt/rstudio-connect/mnt/app/index.html

# The Shiny package inside the packrat cache
/var/lib/rstudio-connect/packrat/3.2.5/v2/library/shiny/
28d6903a44dc53bd4823fa43ccdc08e5/shiny
=> /opt/rstudio-connect/mnt/packrat/3.2.5/v2/library/shiny/
28d6903a44dc53bd4823fa43ccdc08e5/shiny
```

10.7 Program Supervisors

You may need to modify the environment or resources available to R processes prior to R being launched. This can be accomplished using a program supervisor using the `Applications.Supervisor` configuration setting.

The supervisor command is provided the full R command-line, which *MUST* be invoked by the supervisor. The process exit code from R *MUST* be returned as the exit code of the supervisor. The file descriptors for standard input, output, and error *MUST NOT* be intercepted by the supervisor.

A supervisor is executed as the appropriate `RunAs` user. Package installation always uses the `Applications.RunAs` user. Other R processes will use the content-specific `RunAs` account, falling back to `Applications.RunAs` if no override was configured. See Section 10.3 for details.

Supervisors run within the sandbox established for any R process. See Section 10.1 for more information about process sandboxes.

RStudio Connect configures the `TMPDIR`, `HOME`, and `RSTUDIO_PANDOC` environment variables for launched R processes. RStudio Connect also manages package installation and references. Avoid altering any of this behavior in program supervisors.

10.7.1 Example Supervisors

Here is a configuration that uses the `nice` command to lower the priority of all R processes. See <http://linux.die.net/man/1/nice> for details about `nice`. Because process supervisors are run as a `RunAs` user and not as `root` or another super-user, you may not be permitted to assign a negative (higher priority) privilege.

```
[Applications]
Supervisor = nice -n 2
```

Here is a configuration that uses a custom script to prepare a custom execution environment before finally running R.

```
[Applications]
Supervisor = /some/script/that/prepares/an/environment.sh
```

Here is an example supervisor that echos its arguments, sets an environment variable, then invokes whatever arguments have been passed.

```
#!/bin/bash

echo arguments: "$@"
echo

export COMPANY_DATA_HOME="/data/resides/here"

exec "$@"
```

Your organization may use shell initialization scripts to establish a particular environment. This environment might not be completely compatible with how RStudio Connect attempts to launch R.

We recommend building supervisor scripts gradually and carefully. Changes to the environment can alter how your content executes or even prevent R from running correctly.

11 Content Management

RStudio Connect provides flexibility over how uploaded content is configured and shared.

11.1 Vanity Paths

All content receives a URL that includes its numerical ID at the time of deployment – something like `https://rsc.company.org/connect/#/apps/982`. Connect administrative users can create “vanity paths” for content which make the content available at an additional, customized URL.

This setting can be found at the bottom of the “Access” tab when editing a piece of content. There you can enter the path at which you want this content to be available and preview the complete URL. Once you “Save” your content, you’ll be able to access your content at the new vanity URL.

Vanity URLs can not be nested inside of one another. So if a vanity URL `/finance/` already exists, you would not be able to create a new vanity URL at `/finance/budget/`. You may create sibling paths: `/finance/budget/` and `/finance/quarterly/` may both exist concurrently.

11.2 Bundle Management

Content published to RStudio Connect is encapsulated in a “bundle” that contains the source code and data necessary to execute the content. An application or report is updated by uploading a new bundle. Old bundles are retained on disk until you reach the limit imposed by `Applications.BundleRetentionLimit` at which point older bundles will be deleted.

Users can manage their own bundles in the dashboard by clicking the “Source Versions” button. Collaborators can delete, download, activate, and view activation logs for their applications’ bundles. Activating a different bundle is a way of “rolling back” or “rolling forward” to an older or newer version of your application, respectively.

Activating an alternative bundle for a Shiny application will cause new incoming users to be directed to the new version of the application but will not interrupt existing users of the application who are viewing the previously activated bundle. For reports, activating an alternate bundle will immediately render the newly activated bundle and promote it to be the authoritative version of that document. For parameterized reports, only the default variant will be rerendered; other instances of the report will not automatically be regenerated, but the next manual or scheduled update will be performed on the newly selected bundle.

12 R

RStudio Connect offers a flexible way of deploying your Shiny applications and R Markdown content against a variety of R versions.

A compatible version of R is identified when content is deployed. That R installation is used any time R is needed for that content. Package installation, starting a Shiny application, and rendering R Markdown documents will all use the version of R discovered at deploy-time.

RStudio Connect allows different content to rely on different versions of R. For example, Alice’s R Markdown document may require version 3.2.4 of R while Bob’s Shiny application needs R version 3.3.1. Those two deployments using different R versions can coexist in RStudio Connect without conflict.

This chapter discusses RStudio Connect can be configured to support more than one version of R and how R version compatibility is determined.

Available R installations are analyzed at startup. Connect logs the discovered R versions. Connect will fail to start if no R installation can be found.

Changing any of the configuration items discussed in this chapter requires a restart of RStudio Connect.

12.1 Installing R

Ubuntu and Red Hat/CentOS Linux distributions offer a version of R; installation of the system R is detailed in sections 2.1.1 and 2.1.2.

If you are attempting to make additional R versions available side-by-side with the system version you typically need to install from source. The RStudio Server documentation has a good reference for building and installing R into alternate locations.

This RStudio Support article also includes R installation instructions.

12.2 R Versions

RStudio Connect supports two ways of discovering R versions: scanning well-known locations and through direct configuration. Connect will use the `PATH` environment variable to discover a version of R if one is not

otherwise found.

12.2.1 Scanning

RStudio Connect can automatically scan for versions of R in the following locations:

```
/usr/lib/R
/usr/lib64/R
/usr/local/lib/R
/usr/local/lib64/R
/opt/local/lib/R
/opt/local/lib64/R
```

In addition, Connect scans all subdirectories of the following directories within `/opt`:

```
/opt/R
/opt/local/R
```

For example, any of the following installed versions of R will be automatically detected:

```
/opt/R/3.1.3
/opt/R/3.2.4
/opt/local/R/3.3.1
```

Scanning happens by default. You can disable version scanning by customizing the `Server.RVersionScanning` property.

```
[Server]
RVersionScanning = false
```

12.2.2 R Versions

The `Server.RVersion` property can be used to specify alternate locations for installations of R. Specify this property once for each R installation directory.

```
[Server]
RVersion = /shared/R/3.3.1
RVersion = /shared/R/3.2.4
RVersion = /shared/R/3.1.3
```

12.2.3 `/etc/rstudio/r-versions`

The `/etc/rstudio/r-versions` file is an alternative way of specifying R versions and is shared with RStudio Server. List your R installations in this file. Note that the `r-versions` file is not created by default and will need to be created.

```
/shared/R/3.3.1
/shared/R/3.2.4
/shared/R/3.1.3
```

12.2.4 Excluding Versions

If you have versions of R that are picked up by automatic scanning but which you would like to exclude, disable R version scanning and explicitly specify all versions you would like to use in the `/etc/rstudio/r-versions` file or with the `Server.RVersion` configuration property.

Here is an example configuration that disables scanning, and specifies precisely two R versions that will be available for use.

```
[Server]
RVersionScanning = false
RVersion = /opt/R/3.3.1
RVersion = /opt/R/3.2.4
```

12.3 R Version Matching

RStudio Connect attempts to find an R installation that is appropriate for your content. By default, it applies a “nearest” matching approach. This algorithm attempts to always find a version of R to use with your content. More deployments will succeed but not always with the same version of R that is used by the author.

If you would prefer a strict association between authored and deployed R versions, you can choose to use an “exact” matching approach.

The “nearest” matching algorithm is the most flexible option and favors publication of your content over precise duplication of the authoring environment.

The “major-minor” algorithm is a middle ground between “nearest” and “exact”. It requires exact `MAJOR.MINOR` matching but is flexible about the patch level. This is a useful option when your desktop and server may occasionally have different update cycles when installing bug fix releases.

An inconsistent version of R occasionally causes problems when installing package dependencies. For the best results, make sure that RStudio Connect has access to the same versions of R used to author content.

The R version matching approach is controlled with the `Server.RVersionMatching` configuration setting.

nearest Find an R installation that is close to the version of R used when authoring the Shiny application or R Markdown document. This algorithm uses the ordered tests when looking for `MAJOR.MINOR.PATCH` version matches:

1. Use exact version match.
2. If there are matching `MAJOR.MINOR` releases, use least-greater version.
3. If there are matching `MAJOR.MINOR` releases, use latest of these.
4. Use least-greater version across all releases.
5. Use latest across all releases.

major-minor Find an R installation that is close to the version of R used when authoring the Shiny application or R Markdown document requiring an exact `MAJOR.MINOR` version match. If a compatible version cannot be found, content will fail to deploy. The algorithm is a constrained “nearest” search:

1. Use exact version match.
2. If there are matching `MAJOR.MINOR` releases, use least-greater version.
3. If there are matching `MAJOR.MINOR` releases, use latest of these.

exact Finds an R installation that exactly matches the version of R used when authoring the deployed content. If a matching version cannot be found, content will fail to deploy.

13 Package Management

13.1 Package Installation

RStudio Connect installs the R package dependencies of Shiny applications and R Markdown documents when that content is deployed. The RStudio IDE uses the `rsconnect` and `packrat` packages to bundle the

relevant source code and document its dependencies. RStudio Connect then uses `packrat` to duplicate those package dependencies on the server.

Packrat attempts to re-use R packages whenever possible. The `shiny` package, for example, should be installed only when the first Shiny application is deployed. Subsequent Shiny applications can use that package and see faster deployments as a result. Packrat also allows multiple versions of a package to exist on a system. Two Shiny applications referencing different versions of `shiny` will reference the correct Shiny installation and these two packages will not conflict with each other.

Resolving which packages need installing and which are already available all happens when you deploy content to RStudio Connect.

13.2 Private Repositories

Packrat records details about how a package was obtained in addition to information about its dependencies. Most public packages will come from a public CRAN mirror. Packrat lets RStudio Connect support alternate repositories in addition to CRAN.

Learn how to create your own custom repository; this directory can then be shared over HTTP or through a shared filesystem.

Here are some reasons why your organization might use an alternate/private repository.

1. Internally developed packages are made available through a corporate repository. This is used in combination with a public CRAN mirror.
2. All packages (private and public) are approved before use and must be obtained through the corporate repository. Public CRAN mirrors are not used.
3. Direct access to a public CRAN mirror is not permitted. A corporate repository is used as a proxy and caches public packages to avoid external network access.

RStudio Connect supports private repositories in these situations given that the deploying instance of R is correctly configured. No adjustment to the RStudio Connect server is needed.

Repository information is configured using the `repos` R option. Your users will need to make sure their desktop R is configured to use your corporate repository.

RStudio IDE version 0.99.1285 or greater is needed when using repositories other than the public CRAN mirrors.

We recommend using an `.Rprofile` file to configure multiple repositories or non-public repositories.

The `.Rprofile` file should be created in a user's home directory.

```
# A sample .Rprofile file with two different package repositories.
local({
  r <- getOption("repos")
  r["CRAN"] <- "https://cran.rstudio.com/"
  r["mycompany"] <- "http://rpackages.mycompany.com/"
  options(repos = r)
})
```

This `.Rprofile` creates a custom `repos` option. It instructs R to attempt package installation first from "CRAN" and then from the "mycompany" repository. R installs a package from the first repository in "repos" containing that package.

With this custom `repos` option, you will be able to install packages from the `mycompany` repository. RStudio Connect will be able to install these packages as code is deployed.

For more information about the `.Rprofile` file, see `help(Startup)` in R. For details about package installation, see `help(install.packages)` and `help(available.packages)`.

13.3 Private Packages

Packages available on CRAN, a private package repository, or a public GitHub repository are automatically downloaded and built when an application is deployed. RStudio Connect cannot automatically obtain packages from private GitHub repositories, but a workaround is available.

We recommend using a private repository to host internal packages when possible. See Section 13.2 for details.

The configuration option `Server.SourcePackageDir` can reference a directory containing additional packages that Connect would not otherwise be able to retrieve. This directory and its contents must be readable by the `Applications.RunAs` user. Connect will look in this directory for packages before attempting to obtain them from a remote location.

This feature has some limitations.

- The package must be tracked in a git repository so that each distinct version has a unique commit hash associated with it.
- The package must have been installed from the git repository using the `devtools` package so that the hash is contained in the `DESCRIPTION` file on the client machine.

If these conditions are met, you may place `.tar.gz` source packages into per-package subdirectories of `SourcePackageDir`. The proper layout of these files is `<package-name>/<full-git-hash>.tar.gz`.

For example, if `Server.SourcePackageDir` is defined as `/opt/R-packages`, source bundles for the `MyPrivatePkg` package are located at `/opt/R-packages/MyPrivatePkg`. A commit hash of `28547e90d17f44f3a2b0274a2aa1ca820fd35b80` needs its source bundle stored at the following path:

```
/opt/R-packages/MyPrivatePkg/28547e90d17f44f3a2b0274a2aa1ca820fd35b80.tar.gz
```

When private package source is arranged in this manner, users of RStudio Connect will be able to use those package versions in their deployed content.

Be aware that this mechanism is specific to the commit hash, so you will either need to make many git revisions of your package available in the `SourcePackageDir` directory hierarchy or standardize to a particular git commit of the package.

14 Historical Metrics

This section describes the configuration and management of historical metrics, enabled with the `Metrics.Enabled` setting.

RStudio Connect uses a separate `rserver-monitor` process to record resource (CPU, memory, etc.) usage over time. It is only active when historical metrics are enabled. The customization settings described in the remainder of this section have no effect when `Metrics.Enabled` is off.

14.1 Historical Metrics Settings

Metrics data is written by default to a set of RRD files. This data is stored by default at `/var/lib/rstudio-connect/metrics`. You can specify an alternate data path by using the `DataPath` setting mentioned in Section 15.15.

The `rserver-monitor` process runs (by default) with the same user account Connect uses to run its R processes. By default, this user account is `rstudio-connect` (see the `RunAs` setting in Section 15.12). You can specify an alternate user account for the `rserver-monitor` process by modifying the `User` setting. See Section 15.15 for details.

RStudio Connect also supports logging of metrics to Graphite, and it supports disabling its default behavior of logging to RRD. Please see Section 15.15 for more options for configuring the historical metrics in Connect.

14.2 Historical Metrics Process Management

Connect automatically spawns a process (`rserver-monitor`) to help maintain historical data. If this process exits, Connect will restart it in an attempt to record as much historical information as possible. Connect will delay restarting `rserver-monitor` if it observes rapid, repeated failures.

Since the `rserver-monitor` needs permission to write data to the metrics data directory, Connect attempts to ensure the necessary permissions at startup. When Connect starts, it grants ownership of the metrics data directory to the user account that will be used to start `rserver-monitor`.

14.3 Historical Metrics Process Logging

The `rserver-monitor` process logs its output to syslog. If the process is unable to run, you can check the system log (e.g., `/var/log/messages` or `/var/log/syslog`) for messages.

15 Appendix: Configuration Options

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

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

The RStudio Connect 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).

15.1 Server

The **Server** section contains configuration properties which apply across the whole of RStudio Connect 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.

Name	Type	Default	Description
<code>DataDir</code>	string	<code>/var/lib/rstudio-connect</code>	The directory where RStudio Connect will store its variable data.
<code>SourcePackageDir</code>	string	<code><empty-string></code>	A directory containing source bundles for packages that are unavailable on either CRAN or a public GitHub repository. Must be readable by the <code>Applications.RunAs</code> user.
<code>EnableSitemap</code>	boolean	<code>false</code>	Specifies if RStudio Connect should provide a <code>/sitemap.xml</code> file enumerating the publicly available apps.
<code>RVersionMatching</code>	string	<code>nearest</code>	Specifies how RStudio Connect attempts to match R version associated with uploaded content with the R versions available on the system. Allows values of <code>nearest</code> or <code>exact</code> .
<code>RVersion</code>	multi-string	<code>unspecified</code>	Path to an R installation root. Multiple definitions can be used to provide multiple locations with R.
<code>RVersionScanning</code>	boolean	<code>true</code>	Scan for R installations in well-known locations.

Name	Type	Default	Description
Address	string	<i><empty-string></i>	A public URL for this RStudio Connect server. Must be configured to enable features like including links to your content in emails.
SenderEmail	string	<i><empty-string></i>	An email address used by RStudio Connect to send outbound email. The system will not be able to send administrative email until this setting is configured.
ViewerKiosk	boolean	false	When enabled, RStudio Connect does not prompt view-only users to request elevated privileges when attempting to access restricted resources.
HideEmailAddresses	boolean	false	When enabled, RStudio Connect will not expose email addresses in API requests or its dashboard.
MailAll	boolean	false	When enabled, RStudio Connect will allow scheduled and on-demand documents to send email to all users of the system.
PublicWarning	string	<i><empty-string></i>	An HTML snippet used to inject a message into the RStudio Connect dashboard welcome pages.
LoggedInWarning	string	<i><empty-string></i>	An HTML snippet used to inject a message into the RStudio Connect recent views.
ContentTypeSniffing	boolean	false	If disabled, sets the X-Content-Type-Options HTTP header to nosniff . When enabled, removes that header, allowing browsers to mime-sniff responses.
ServerName	string	<i><empty-string></i>	By default, Connect sets the Server HTTP header to something like RStudio Connect v1.2.3 . This setting allows you to override that value.
AccessLog	string	<code>/var/log/rstudio-connect.access.log</code>	Path to the file that RStudio Connect will use for its access logs. Disabled when empty.

Name	Type	Default	Description
<code>CustomHeader</code>	multi-string	<i>unspecified</i>	Custom HTTP header that should be added to responses from Connect in the format of <code>key: value</code> . 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 Connect would otherwise have set. Multiple definitions can be used to provide multiple custom headers.
<code>FrameOptionsContent</code>	string	<i><empty-string></i>	The value for the <code>X-Frame-Options</code> HTTP header for all user-uploaded content (Shiny apps, RMDs, etc.). If empty, no header will be added.
<code>FrameOptionsDashboard</code>	string	DENY	The value for the <code>X-Frame-Options</code> HTTP header for the Connect dashboard and all other Connect pages. If empty, no header will be added.

15.2 Http

The `Http` section contains configuration properties which control the ability of RStudio Connect to listen for HTTP requests. RStudio Connect 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.

Name	Type	Default	Description
<code>Listen</code>	string	<i><empty-string></i>	RStudio Connect will listen on this network address for HTTP connections. The network address can be of the form <code>:80</code> or <code>192.168.0.1:80</code> . Either <code>Http.Listen</code> or <code>Https.Listen</code> is required.

15.3 Https

The `Https` section contains configuration properties which control the ability of RStudio Connect to listen for HTTPS requests. RStudio Connect 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.

Name	Type	Default	Description
<code>Listen</code>	string	<code><empty-string></code>	RStudio Connect will listen on this network address for HTTPS connections. The network address can be of the form <code>:443</code> or <code>192.168.0.1:443</code> . Either <code>Http.Listen</code> or <code>Https.Listen</code> is required.
<code>Key</code>	string	<code><empty-string></code>	Path to a private key file corresponding to the certificate specified with <code>Https.Certificate</code> . Required when <code>Https.Certificate</code> is specified.
<code>Certificate</code>	string	<code><empty-string></code>	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 <code>Https.Key</code> .
<code>Permanent</code>	boolean	<code>false</code>	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 <code>secure</code> flag on all session cookies and adds a <code>Strict-Transport-Security</code> HTTP header with a value of 30 days.

15.4 HttpRedirect

The `HttpRedirect` section contains configuration properties which control the ability of RStudio Connect 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.

Name	Type	Default	Description
<code>Listen</code>	string	<code><empty-string></code>	RStudio Connect will listen on this network address for HTTP connection and redirect to either the <code>HttpRedirect.Target</code> or <code>Server.Address</code> target location. The network address can be of the form <code>:8080</code> or <code>192.168.0.1:8080</code> . 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 <code>HttpRedirect.Target</code> or <code>Server.Address</code> .

Name	Type	Default	Description
Target	string	<i><empty-string></i>	The target for redirects when users visit the <code>HttpRedirect.Listen</code> HTTP service. <code>Server.Address</code> is used as a redirect target if this property is not specified.

15.5 Database

The **Database** section contains configuration properties which control the location of and how RStudio Connect interacts with its database.

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

Name	Type	Default	Description
Dir	string	<code>{Server.DataDir}/db</code>	The directory containing the RStudio Connect database.
MaxIdleConnections	integer	0	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.
MaxOpenConnections	integer	0	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.

15.6 Authentication

The **Authentication** section contains configuration properties which control how users will log into RStudio Connect.

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

Name	Type	Default	Description
Provider	string	<code>password</code>	Specifies the type of user authentication. Allows values of <code>password</code> , <code>oauth2</code> , <code>ldap</code> , <code>pam</code> , or <code>proxy</code> .
Name	string	<i><empty-string></i>	Specifies a meaningful name for your authentication provider. This presented on the sign-in page and gives users context about the credentials being requested. If unspecified, RStudio Connect will use a generic name for the chosen provider. Just using your company name is often a good choice.

15.7 OAuth2

The `OAuth2` section contains configuration properties which control how RStudio Connect communicates with Google OAuth2 servers in order to authenticate users.

Section 8.4 contains more information about configuring RStudio Connect to use Google for authentication.

The `DiscoveryEndpoint` property should be configured as:

```
[OAuth2]
DiscoveryEndpoint = https://accounts.google.com/.well-known/openid-configuration
```

These properties must appear after `[OAuth2]` in the configuration file.

Name	Type	Default	Description
<code>DiscoveryEndpoint</code>	string	<code><empty-string></code>	Specifies a URL for the OAuth2 discovery endpoint. Required for all OAuth2 configurations.
<code>ClientId</code>	string	<code><empty-string></code>	Identifier for OAuth2 client. Required for all OAuth2 configurations.
<code>ClientSecret</code>	string	<code><empty-string></code>	Client secret for the configured client ID. Either <code>OAuth2.ClientSecret</code> or <code>OAuth2.ClientSecretFile</code> must be specified when using OAuth2. <code>OAuth2.ClientSecret</code> takes priority if both properties are set.
<code>ClientSecretFile</code>	string	<code><empty-string></code>	Path to file containing the client secret. Either <code>OAuth2.ClientSecret</code> or <code>OAuth2.ClientSecretFile</code> must be specified when using OAuth2. <code>OAuth2.ClientSecret</code> takes priority if both properties are set.
<code>AllowedDomains</code>	string	<code><empty-string></code>	Space-separated list of domains permitted to authenticate.
<code>AllowedEmails</code>	string	<code><empty-string></code>	Space-separated list of email addresses permitted to authenticate. When used without <code>OAuth2.AllowedDomains</code> , only the email addresses listed here will be allowed access. When used with <code>OAuth2.AllowedDomains</code> , the email addresses listed here will be added to those with valid domains.

15.8 LDAP

The `LDAP` section contains configuration properties which control how RStudio Connect communicates with an LDAP or Active Directory server.

See Section 8.3 for details about how to configure RStudio Connect with LDAP authentication. Section 16 contains many configuration examples.

The `LDAP` section is different from many other configuration sections, as it allows multiple, distinctly named configuration instances. This name is *case sensitive*. A named section looks like:

[LDAP "European LDAP Server"]

All of the LDAP configuration properties must appear after [LDAP "<name>"] in the configuration file.

Name	Type	Default	Description
ServerAddress	string	<empty-string>	Specifies the location of the LDAP/AD server. This should be of the form <host>:<port>. The host may be either an IP or DNS address. Most LDAP/AD servers operate on port 389 or 636 .
TLS	boolean	false	When enabled, all connections to your LDAP/AD server will use TLS (SSL).
StartTLS	boolean	false	When enabled, the connection will initially be made on an insecure port then the channel will be upgraded to TLS using StartTLS.
ServerTLSInsecure	boolean	false	This option controls if RStudio connect will verify the server's certificate chain and host name. When enabled, RStudio Connect will accept any certificate presented by the server and any host name in that certificate. Setting to true is susceptible to man-in-the-middle attacks but is required in some circumstances, such as when using a self-signed certificate.
TLSCACertificate	string	<empty-string>	Path to a certificate authority used to connect an LDAP server.
UserObjectClass	string	<empty-string>	The name of the LDAP objectClass used to define users.
UserFirstNameAttribute	string	<empty-string>	The LDAP user attribute containing a user's first name. This is often the givenName attribute.
UserLastNameAttribute	string	<empty-string>	The LDAP user attribute containing a user's last name. The sn attribute will usually contain last name.
UserEmailAttribute	string	<empty-string>	The LDAP user attribute containing a user's email address. Many systems use the mail attribute.
UsernameAttribute	string	<empty-string>	The LDAP user attribute containing a user's username. Commonly used attributes include uid , cn , and samaccountname .
BindDN	string	<empty-string>	A DN for a read-only admin account that is used during double-bind authentication and for certain operations that do not occur during the login sequence (such as searching). Must be paired with BindPassword .
BindPassword	string	<empty-string>	The password for the BindDN account.

Name	Type	Default	Description
<code>AnonymousBind</code>	boolean	<code>false</code>	Enable anonymous bind. An anonymous user must have rights to search and view all pertinent groups, group memberships, and users.
<code>UserSearchBaseDN</code>	string	<i><empty-string></i>	The base DN used when performing user searches.

15.9 PAM

The `PAM` section contains configuration properties which control how RStudio Connect interacts with the PAM (Pluggable Authentication Module) API.

See Section 8.5 for details about configuring an appropriate PAM authentication profile for RStudio connect.

See Section 10.5 for information about using PAM sessions when launching R processes.

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

Name	Type	Default	Description
<code>Service</code>	string	<code>rstudio-connect</code>	Specifies the PAM service name that RStudio Connect will use when authenticating users.
<code>UseSession</code>	boolean	<code>false</code>	Use PAM sessions when launching R processes.
<code>SessionService</code>	string	<code>su</code>	Specifies the PAM service name that RStudio Connect will use for PAM sessions.

15.10 Proxied Authentication

The `ProxyAuth` section contains configuration properties which control how RStudio Connect utilizes an external authentication server which proxies all requests.

See Section 8.6 for details about configuring an appropriate proxied authentication for RStudio connect.

Name	Type	Default	Description
<code>UsernameHeader</code>	string	<code>X-Auth-Username</code>	Specifies the name of the header that will contain a username provided by the proxy.

15.11 Authorization

The `Authorization` section contains configuration properties which control permissions and privileges when accessing RStudio Connect.

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

Name	Type	Default	Description
Name	Type	Default	Description
DefaultUserRole	string	publisher	Specifies what abilities given to a newly created user. Allows values viewer , publisher , or administrator .

15.12 Applications

The `Applications` section contains configuration properties which control how RStudio Connect communicates with R processes.

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

Name	Type	Default	Description
RunAs	string	<code>rstudio-connect</code>	User used to invoke R.
RunAsCurrentUser	boolean	<code>false</code>	Allows content to execute as the logged-in user when using PAM authentication.
Supervisor	string	<code><empty-string></code>	Specifies a command to wrap the execution of R.
HomeMounting	boolean	<code>false</code>	Specifies that the contents of <code>/home</code> should be hidden from R processes with additional bind mounts. The existing <code>/home</code> will have the home directory of the <code>RunAs</code> user mounted over it. If <code>RunAs</code> does not have a home directory, an empty temporary directory will mask <code>/home</code> instead. Launched R processes can discover this location through the the <code>HOME</code> environment variable.
ViewerOnDemandReports	boolean	<code>false</code>	Allow report viewers to generate an ad-hoc rendering. The <code>ViewerCustomizedReports</code> property is implicitly disabled when this property is disabled.
ViewerCustomizedReports	boolean	<code>false</code>	Allow report viewers to customize the parameters of an ad-hoc rendering.
BundleReapFrequency	duration	<code>24 hours</code>	Time between the worker that deletes filesystem data for bundles in excess of our retention limit.
BundleRetentionLimit	integer	<code>0</code>	Maximum number of bundles per app for which we want to retain filesystem data. The default is 0, which means retain everything.
ScheduleConcurrency	integer	<code>2</code>	Number of scheduled reports permitted to execute in parallel

Name	Type	Default	Description
ConnectionTimeout	duration	1h	Maximum time allowed without data sent or received across a client connection. A value of 0 means connections will never time-out (not recommended).
ReadTimeout	duration	1h	Maximum time allowed without data received from a client connection. A value of 0 means a lack of client (browser) interaction will never cause the connection to close. This is useful when deploying dashboard applications which send regular updates but have no need for interactivity.
DisabledProtocols	string	<empty-string>	List of comma-delimited protocols to disable on the SockJS client. Allows values of <code>websocket</code> , <code>xhr-streaming</code> , <code>iframe-eventsources</code> , <code>iframe-htmfile</code> , <code>xhr-polling</code> , <code>iframe-xhr-polling</code> , or <code>jsonp-polling</code>

15.13 Performance/Scheduler

The `Scheduler` section contains configuration properties which control how RStudio Connect manages R processes for deployed Shiny applications. These properties are managed on an individual application under the `Performance` tab.

RStudio Connect makes a determination on each new client connection about whether or not it needs to spawn an additional R process. That computation analyzes the number of current R processes and the number of active connections against those processes. If a substantial percentage of connections are consumed, RStudio Connect will create a new process rather than causing the existing processes to become more busy. That percentage of connection use is called the “load factor”.

The algorithm that considers the current load factor looks like the following pseudocode.

```
// Given:
//   numProcesses
//     - The number of R processes for the current application.
//   numConnections
//     - The number of connections across all R processes associated
//       with the current application.
allowedConnections = numProcesses * Scheduler.MaxConnsPerProcess
currentLoadFactor = numConnections / allowedConnections
if currentLoadFactor > Scheduler.LoadFactor {
  // Create a new process if the new process will not exceed
  // Scheduler.MaxProcesses
}
```

The `Scheduler.InitTimeout` and `Scheduler.IdleTimeout` properties may need adjusting when a Shiny application takes a very long time to startup. Increasing `InitTimeout` will allow more time for the Shiny application to start. An increase to `IdleTimeout` lets idle R processes linger longer so they are available the next time a request arrives - avoiding the startup penalty.

The scheduler properties can be changed in the configuration file and apply to all Shiny applications. The RStudio Connect dashboard allows custom scheduler settings for individual applications.

We recommend that **Scheduler** property adjustment be done gradually.

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

Name	Type	Default	Description
<code>MaxProcesses</code>	integer	3	Specifies the total number of concurrent R processes allowed for a single application.
<code>MaxConnsPerProcess</code>	integer	20	Specifies the maximum number of client connections allowed to an individual R process. Incoming connections which will exceed this limit are routed to a new R process or rejected.
<code>LoadFactor</code>	decimal	0.5	Controls how aggressively new R processes will be spawned.
<code>InitTimeout</code>	duration	60s	Maximum time to wait for an app to start.
<code>IdleTimeout</code>	duration	5s	Minimum time to keep a worker process alive after it goes idle.
<code>MinProcessesLimit</code>	integer	20	Maximum value allowed for the <code>MinProcesses</code> setting on an application level. All applications default to <code>MinProcesses=0</code> , but <code>MinProcesses</code> can be increased to this limit per application.

15.14 Jobs

The **Jobs** section contains configuration properties which control the retention of metadata associated with R process execution.

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

Name	Type	Default	Description
<code>MaxCompleted</code>	integer	100	The maximum number of completed jobs preserved on disk for any one application. When this limit is reached, the oldest completed jobs for an application will be deleted as new jobs are launched. On-disk job metadata is removed if either the <code>MaxCompleted</code> or <code>OldestCompleted</code> restrictions are violated.
<code>OldestCompleted</code>	duration	720h	The maximum age of a completed job retained on disk. Jobs older than this setting will be deleted. Set to zero to remove restrictions on the age of a completed job. On-disk job metadata is removed if either the <code>MaxCompleted</code> or <code>OldestCompleted</code> restrictions are violated.

15.15 Historical Metrics

The `Metrics` section contains configuration properties which control how RStudio Connect manages the `rserver-monitor` process for monitoring the use of resources (CPU, memory, etc.) for historical metrics.

See Section 14 for more details about historical metrics in Connect.

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

Name	Type	Default	Description
<code>Enabled</code>	boolean	<code>true</code>	Specifies whether or not the <code>rserver-monitor</code> process that collects historical metrics will be started.
<code>User</code>	string	<code>{Applications.RunAs}</code>	The user for the <code>rserver-monitor</code> process.
<code>DataPath</code>	string	<code>{Server.DataDir}/metrics</code>	The path for writing log entries and RRD database files.
<code>Interval</code>	duration	<code>60s</code>	The frequency of historical metrics collection.
<code>RRDEnabled</code>	boolean	<code>true</code>	Enable logging of historical metrics to RRD.
<code>GraphiteEnabled</code>	boolean	<code>false</code>	Enable logging of historical metrics to Graphite.
<code>GraphiteHost</code>	string	<code>127.0.0.1</code>	Host to which to send Graphite historical metrics.
<code>GraphitePort</code>	integer	<code>2003</code>	Port to which to send Graphite historical metrics.
<code>GraphiteClientId</code>	string	<code><empty-string></code>	Optional Client ID to include along with Graphite historical metrics.

16 Appendix: LDAP/AD Configuration Examples

Sample LDIF file and associated RStudio Connect configuration section.

16.0.1 connect.cfg

Here is an example RStudio Connect configurations for the LDIF file below. We've given you two examples, one that uses double bind, and one that uses single bind. (double bind is recommended)

```
# using single bind
[LDAP "myLDAPserverSingle"]
ServerAddress = 127.0.0.1:389
UserSearchBaseDN = "ou=People,dc=company,dc=com"
UserObjectClass = posixAccount
UserFirstNameAttribute = givenName
UserLastNameAttribute = sn
UserEmailAttribute = mail
UsernameAttribute = uid
```

```
# using double bind
[LDAP "myLDAPserver"]
ServerAddress = 127.0.0.1:389
BindDN = cn=admin,dc=company,dc=com"
```



```
BindPassword = "password"
UserSearchBaseDN = "ou=People,dc=company,dc=com"
UserObjectClass = posixAccount
UserFirstNameAttribute = givenName
UserLastNameAttribute = sn
UserEmailAttribute = mail
UsernameAttribute = uid
```

16.0.2 contents.ldif

```
dn: ou=People,dc=company,dc=com
objectClass: organizationalUnit
```

```
dn: ou=Groups,dc=company,dc=com
objectClass: organizationalUnit
```

```
dn: cn=membera-grp,ou=Groups,dc=suba,dc=company,dc=com
objectClass: posixGroup
cn: membera-grp
gidNumber: 50000
memberUid: membera
```

```
dn: cn=memberb-grp,ou=Groups,dc=subb,dc=company,dc=com
objectClass: posixGroup
cn: memberb-grp
gidNumber: 50001
memberUid: memberb
```

```
dn: cn=memberc-grp,ou=Groups,dc=subc,dc=company,dc=com
objectClass: posixGroup
cn: memberc-grp
gidNumber: 50002
memberUid: memberc
```

```
dn: uid=membera,ou=People,dc=suba,dc=company,dc=com
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
uid: membera
sn: A
givenName: Member
cn: Member A
displayName: Member A
uidNumber: 20000
gidNumber: 50000
userPassword: memberaldap
gecos: MemberA
loginShell: /bin/bash
homeDirectory: /home/membera
mail: membera@company.com
```

```
dn: uid=memberb,ou=People,dc=subb,dc=company,dc=com
objectClass: inetOrgPerson
```

```
objectClass: posixAccount
objectClass: shadowAccount
uid: memberb
sn: B
givenName: Member
cn: Member B
displayName: Member B
uidNumber: 20001
gidNumber: 50001
userPassword: memberbldap
gecos: MemberB
loginShell: /bin/bash
homeDirectory: /home/memberb
mail: memberb@company.com
```

```
dn: uid=memberc,ou=People,dc=subc,dc=company,dc=com
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
uid: memberc
sn: C
givenName: Member
cn: Member C
displayName: Member C
uidNumber: 20002
gidNumber: 50002
userPassword: membercldap
gecos: MemberC
loginShell: /bin/bash
homeDirectory: /home/memberc
mail: memberc@company.com
```

17 Appendix: RStudio Connect Deployment Guide

17.1 Overview

This guide will cover the details of the deployment process in RStudio Connect. For most users, these details can be safely ignored, as the details are handled automatically via push-button publishing. However, some users may want to programmatically publish content using the `rsconnect` package or may have run into an error during deployment.

17.2 Programmatic Deployment

To programmatically publish content to RStudio Connect, use the functions `deployDoc`, `deployApp`, and `deploySite` from the `rsconnect` package. Each of these functions will require a user account and a connected server. To setup an account on a server use `addConnectServer` and `connectUser`. To view currently configured accounts use `accounts`. For more details visit the `rsconnect` reference pages.

Each of the deployment functions listed above can be supplied with optional arguments. If additional arguments are not supplied, defaults are determined based on the content being deployed. All of the deployment functions follow a similar, underlying process. This appendix explains the process in detail.

17.3 Step 1: Building the Bundle

Connect builds an application bundle for the deployed content. The bundle contains the source code, any data files, and a manifest (JSON file) with metadata about the bundle and environment.

17.3.1 Application Metadata

rsconnect infers a number of attributes about the content including:

1. `appMode`: static, shiny, rmd-static, rmd-shiny
2. `hasParameters`: whether or not the R Markdown file includes parameters

In the case of an R Markdown document the YAML is parsed. Otherwise, .R files are flagged as shiny applications, html files and pdf files are flagged as static. (When a plot is published, the plot is wrapped in an html file).

17.3.2 List of Target Files

Next, rsconnect identifies the relevant files for the application. `appFiles` or `appFileManifest` can be passed as arguments to `deployApp` to specify the required files. Otherwise, rsconnect attempts to identify the required files using a number of heuristics.

For R Markdown documents and static HTML files, external dependencies are discovered using the rmarkdown function `find_external_resources`. This function searches for dependencies in the R Markdown file and the rendered HTML file. The function is able to identify files in the YAML header (if a parameter is a file), logos, images, data files used within R code chunks, and HTML dependencies. This process includes a minimal, client-side “render” of the document (the Rmd is not rendered, it is converted to plain markdown and then rendered to HTML without running any R code). Think of this rendering as creating a skeleton of the final HTML document. During push-button deployment, this initial “render” will show up in the IDE R Markdown tab.

The dependencies for R Markdown websites are identified uniquely. Websites should be deployed by calling `deploySite`.

Troubleshooting: To avoid client side rendering, deploy the content directly using `deployApp` with `appFiles` or `appFileManifest`.

For Shiny applications, rsconnect adds all the files in the project directory and subdirectories with a few exceptions: .Rproj file, the packrat directory, and the rsconnect directory. Files are added up to the specified max bundle size: `getOption("rsconnect.max.bundle.size")`.

Troubleshooting: try `rsconnect::listBundleFiles(appDir)` to see the identified dependencies

17.3.3 Lint

After identifying the target file and dependency files, rsconnect applies a series of linters. The rsconnect linters attempt to identify common problems that might prevent an application that works locally from working after deployment. These checks ensure the application code does not contain:

1. absolute paths
2. invalid relative paths
3. inconsistent capitalization among paths (the Connect server has case sensitive file paths)

The linters currently **do not** check for database connections.

Troubleshooting: You can disable the linters by passing `lint=FALSE` to the deployment function.

17.3.4 Create Temporary Folder

If the files pass the linters, RStudio Connect creates the initial bundle by copying all of the files to a temporary directory.

17.3.5 Library Dependencies

Next, `rsconnect` attempts to identify the package dependencies required by the app. (This step is skipped for static content). `rsconnect` does this by using `packrat`. `Packrat` is a dependency management tool for R designed to keep projects isolated, portable, and reproducible. `rsconnect` deployment does not use all of `packrat`'s functionalities. (For example, the package sources are not installed on the client in the project's `packrat` subdirectory). For more information visit: <https://github.com/rstudio/packrat>

`Packrat` looks through the R code and makes note of any `library()` or `require()` calls. `Packrat` creates a list of the required packages and saves the list in the `packrat.lock` file. This lock file includes the package version and package dependencies. This process is recursive. In addition, the lock file also includes information on the version of R being used, the type of repository containing the package, and the specific URI for each type of repository. A few notes about this process:

Packrat searches in the order of `.libPaths`

For example, if the code includes `library(babynames)`, `Packrat` will look for `babynames` inside the first library in `.libPaths`. Imagine there are two libraries: A and B and `.libPaths(A,B)`. In A, `babynames` is version 1.0. In B, `babynames` is version 2.0. `Packrat` will assume the app depends on version 1.0. To understand this behavior, recall that a library is just a folder containing an installed R package. The most common scenario where this occurs is when the target directory is part of an existing `packrat` project.

Repositories

Most packages come from CRAN. In the `packrat` lockfile, `packrat` will record the names of packages originating from CRAN as well as a specific URL for CRAN (i.e. `CRAN='https:cran.rstudio.com'`). The url is determined by the state of `options("repos")` during deployment. The same process is used for other repositories: Github, BioConductor, and local repositories. In the case of a local repository, the repository URI may be a location on disk.

For the edge case of an internal package from a local repository, be sure the package's `Repository` option (found in the package's `Description` file) is mapped to a repo URI in the current `options("repos")`. For example, imagine a package called `myPackage` is stored in a local repo called `myRepo`. The `myPackage` `Description` file should include `repository:myRepo`. `options("repos")` should define a URI for `myRepo` during deployment runtime, i.e. `options(repos = list(myRepo="file://path_to_private_repo"))`.

Troubleshooting: try `rsconnect:::performPackratSnapshot(appDir)`. This command will create the `packrat` lock file helping to identify the dependencies, corresponding repos, and URLs expected for deployment.

Once the lock file is created, `rsconnect` proceeds to copy all of the description files for the packages listed in the `packrat` lock file. The files are copied into `packrat/desc`. Normally, a `packrat` lockfile would be enough to fully reproduce the package environment. This additional step is necessary just in case the version of `packrat` on the client is significantly different from the version on the server.

17.3.6 Manifest

Next, `rsconnect` generates the actual manifest. This manifest includes a list of the relevant source code, package dependencies, and other metadata including the R version, the locale, the app mode, content category, etc. The R version is determined while building the manifest. The R version listed in the manifest will later be used by `Connect` to attempt to re-create a server-side environment consistent with the client. While

creating the manifest, `rsconnect` will also attempt to determine the primary document (if not already listed). Checksums are stored for each file, including the packrat description files. Finally, the manifest is copied to the temporary bundle directory alongside the code and packrat directory.

For example, a target directory with the structure:

```
targetDir
- app.R
+ dataDir
  - data.csv
```

where `app.R` includes:

```
library(babynames)
library(shiny)
```

The final bundle will contain:

```
bundleDir
- app.R
- manifest.json
- index.htm
+ dataDir
  - data.csv
+ packrat
  - packrat.lock
  + desc
  - babynames
  - shiny
  ...
```

The `manifest.json` file will include:

```
{
  "version" : 1,
  "locale" : "en_US",
  "platform" : "3.2.5",
  "metadata" : {
    "appmode" : "shiny",
    "primary_rmd" : null,
    "primary_html" : null,
    "content_category" : "application",
    "has_parameters" : false
  },
  "packages" : {
    ...
  },
  "files" : {
    "app.R" : {
      "checksum" : "bc81fad5645566fe5d228abf57bba444"
    },
    "packrat/desc/babynames" : {
      "checksum" : "ee14db463dc57f078fea1c3d74628104"
    },
    ...
  },
}
```

The `packages` entry will contain a version of each package's DESCRIPTION file. The `files` entry will include a checksum for each package description file.

Troubleshooting: try `rsconnect::bundleApp(appDir, appFiles=rsconnect::listBundleFiles(appDir), ...)`. This command will generate a tarball containing the application bundle.

17.4 Step 2: Push Bundle to Connect

In step 2 `rsconnect` publishes the bundle to the server. This is done with a POST request to an HTTP endpoint. `rsconnect` supports multiple protocols for making HTTP requests. `rsconnect` looks for the server address and account information created when the IDE is linked to Connect. Publisher privileges are required for a user to link the IDE to Connect and publish content. These privileges are checked when the user sets up an account for publishing (this process creates a public-secret key pair unique to the user and Connect server).

Troubleshooting: try `rsconnect::accounts()`

When an application bundle is successfully deployed, `rsconnect` generates a folder in the original target directory called `rsconnect`. This folder contains a DCF file with information on the deployed content (i.e. the name, title, server address, account, URL, and time). If you re-deploy the same directory, `rsconnect` checks for this file allowing the deployed content to be updated. Redeployments will replace the original bundle with the new bundle. At this time there is no way for bundles to be versioned or to download a previous copy of a bundle. Redeployment will only work if the document is the same content type. For instance, you can not redeploy an R Markdown document after adding `runtime:shiny`. Instead, deploy the document to a new endpoint by changing the `appName`.

Currently, each deployed application is tied to an account. For example, imagine `user1` deploys an app and shares the code with `user2`. If `user2` deploys the app, a new copy of the app would be deployed. This is true even if `user1` shares the `rsconnect` folder. (The only way for a different collaborator to deploy to the same app is for both collaborators to use a service account where the username and password are shared by both users. Both users would also need to go through the steps that link the IDE to Connect - generating the public-private keypair).

In some occasions, a single user will have multiple accounts on one server, or an account on multiple servers. To deploy a bundle to a different server or under a different account, specify the account and user parameters in the `deployApp` function. After successful deployment, a new DCF file will be added to the `rsconnect` folder. If you deploy the same content from a new machine to the server, using the same account, `rsconnect` will prompt you asking whether or not the content is a redeploy. This occurs even if the `rsconnect` folder does not exist on the new machine.

17.5 Step 3: Bundle is deployed on Connect

Once the bundle is published to the server, Connect prepares the content to be deployed. This process follows a number of steps:

17.5.1 Parse the Manifest

The bundle is uncompressed at a unique location (assigned based on `appid` and `bundle id`). The manifest from the uncompressed bundle is parsed to determine the type of content. The R version is also identified and matched based on the available R versions on the Connect server. You can find more details here. Files are checked against the checksum listed in the manifest to ensure content was not lost or corrupted during transfer.

17.5.2 Packrat Restore

Packrat is used to ensure the required packages are available. For every package identified in the manifest: Packrat checks to see if the required package is available in the global cache. (The cache is specific to the version of R matched previously).

If the package is available, a symlink is created that points to the package within the global cache. If a symlink is not possible, the package will be copied from the global cache.

If the package is not available, packrat attempts to install the package. The package is requested from the repo URL identified during bundling. The package is installed and built from source and the installed package is added to the global cache.

Many R packages have system-level dependencies (Java, openssl, etc). If the package fails to install, be sure these system dependencies are installed and available.

All packages are installed as the default `[Server].RunAs` user (typically `rstudio-connect`). Connect ensures that the package libraries and uncompressed bundle have the appropriate permissions based on the application specific `runas` user.

17.5.3 R Markdown Render

If the deployed content is an R Markdown document (excluding documents with `runtime:shiny`) the Rmd file is rendered on the server. If the document is parameterized, the default parameters are used.

The application is presented as deployed. User input is currently required to publish the application and specify any server-side attributes (such as performance tuning settings, permissions, etc).

17.6 Other Frequently Asked Questions

1. My app deployed but does not run?

If the application is deployed but does not run, the error message will be caught and displayed in the application log (visible at the app url in Connect on the logs panel).

2. Can I get more details about the deployment failure?

Yes, set the option “Show diagnostic information after publishing” in `Tools -> Global Options -> Publishing`

3. Will database connections work once deployed?

Database connections will only work if the same drivers (and potentially DSNs) are available on the client and on Connect. At this time there is not a linter to check for connection strings.

4. I use a specific distribution of R (i.e. MRO). Will matching work?

The version of R written to the manifest will be the version used during runtime.

On the server side, Connect attempts to match the version of R in the manifest as described here.

Currently Connect only matches based on the version - no other supplemental information (such as distribution) is maintained. For that reason, to ensure a specific distribution is used on the server, install only that distribution for the desired version.

5. Are bundles compressed?

Bundles are not be compressed. Bundles do not need to be read completely into RAM during deployment. Typically the only bottleneck is upload speed. You can specify a maximum bundle size using: `getOption("rsconnect.max.bundle.size")`.

18 Using Continuous Integration to Deploy Content

18.1 Overview

It is possible to use the `rsconnect` R package to programmatically deploy content to a Connect server. This is particularly useful when combined with a continuous integration (CI) server that builds and deploys your content.

18.2 Prerequisites

Currently, it is only feasible to use a CI server to update content that you originally published from the same server. You cannot update content that you published from elsewhere. To clarify, the CI server must perform both the initial deployment and subsequent updates of the application.

Configuring a CI server to deploy content with `rsconnect` requires that you log in to the CI server with the credentials the CI server uses to run `rsconnect`. For example, if your CI server uses the `jenkins` account, you need to log in as `jenkins` to configure `rsconnect` for the CI server.

You must be familiar with deploying content with `rsconnect`. Please see 17 for more information.

18.3 Configuring a CI Server to Deploy Content to Connect

18.3.1 Installing `rsconnect`

The `rsconnect` package is used to deploy content to Connect. Install it with the following command in the R console. In practice, `rsconnect` may already be available.

```
install.packages("rsconnect")
```

18.3.2 Configuring `rsconnect`

Configuring `rsconnect` requires a user home directory. In this use case, a valid home directory is required for the `jenkins` user account.

You must configure `rsconnect` for the user account that will be used by the CI server to deploy content with `rsconnect`. In this document, we assume that this user is `jenkins`.

```
sudo su jenkins
```

Next, while running as `jenkins`, run R and issue the following commands in the R console:

```
library(rsconnect)
addConnectServer("http://myserveraddress:3939", "mylocaldeployserver")
connectUser(server="mylocaldeployserver")
```

The `rsconnect` server name, `mylocaldeployserver`, is an arbitrary name that is used to identify a Connect server when using `rsconnect`. You can choose any name you wish.

After the last command, you will see output similar to this:

```
A browser window should open; if it doesn't, you may authenticate manually by visiting
http://myserveraddress:3939/___login___?url=http%3A%2F%2Fmyserveraddress%3A3939%
2Fconnect%2F%23%2Ftokens%2FTc8f636c59ffff521eef4888b163dcf64%2Factivate&user_id=0.
```

```
Waiting for authentication...
```


Copy the URL in the output above, then paste it into a Web browser and authenticate with the Connect user credentials for your CI server. In this example, we assume that you wish to deploy content with the `ci-server` Connect account.

After successfully connecting the `ci-server` Connect account to `rsconnect`, you will see this message at the R console:

```
Account registered successfully: CI Server (ci-server)
```

The server and account information are persisted to configuration files on the server in the `jenkins` user's home directory:

```
/home/jenkins/.config/R/connect/servers/mylocaldeployserver.dcf
/home/jenkins/.config/R/connect/accounts/mylocaldeployserver/connectuser.dcf
```

18.3.3 Deploying Content with `rsconnect`

Now `rsconnect` is configured to use the `ci-server` Connect account when running with the `jenkins` server account.

18.3.4 Package and R Version Compatibility

`rsconnect` will use the package libraries and the R installation available on the CI server to create the manifest used by Connect. It is crucial that the environment on the CI server is compatible with the content you are deploying. Ideally, you should maintain the same R version, the same available packages, and the same package versions that you use in development.

18.3.5 CRAN Note

If you don't already have it in an `.Rprofile`, be sure to specify a default CRAN repository in your application before issuing the `rsconnect` command to deploy content. For example:

```
options(repos=c(CRAN="https://cran.rstudio.com"))
deployDoc(doc="out.Rmd", appName="ServerDeployedDoc",
          account="ci-server", server="mylocaldeployserver")
```

Please note that Connect content must be published before it is publicly available. This means that you must log in to Connect and publish the content after the initial deployment. Subsequent automated deployments of the same content are automatically published and require no manual intervention.

18.4 Warning and Security Information

A CI server account that is configured to deploy content to Connect can deploy additional content to Connect without further authentication.

For example, Bob logs in to a server console as Unix user `jenkins`, which is the account used by his CI server. Bob then configures `rsconnect` to deploy content. During the authorization step, Bob signs in to Connect as a publisher with user name `ci-server`. Now, any other CI processes running on this server under the `jenkins` user account can deploy additional content using the Connect user `ci-server`.

19 Appendix: Programmatic Deployment with `rsconnect`

19.1 Overview

It is possible to use the `rsconnect` R package to programmatically deploy content to a Connect server. Furthermore, Connect-hosted content can use `rsconnect` to deploy additional content to itself or to another Connect server.

Configuring Connect to deploy content with `rsconnect` requires:

1. administrator privileges for Connect, and
2. `sudo` or root privileges on the server where Connect is installed.

19.2 Use Case: A Shiny Application

Here we present a use case that explains how to configure Connect for programmatic deployment. Please see 19.4 for an example Shiny application for this use case.

19.2.1 Use Case Scenario

Bob White develops a Shiny application (see 19.4) that:

1. Renders an R Markdown document.
2. Deploys the generated document using `rsconnect`

Bob deploys his Shiny application to Connect. The application, as noted above, can automatically deploy documents it generates to Connect. However, the Connect server must first be configured to authorize deployment from `rsconnect`.

19.2.2 Installing `rsconnect`

The `rsconnect` package is not yet available on Bob's Connect server, so Bob installs it by running R as root (`sudo R`) and issuing the following command in the R console. In practice, `rsconnect` may already be available.

```
install.packages("rsconnect")
```

19.2.3 Configuring a Custom “RunAs” User

Since Bob does not want to allow arbitrary Connect users to deploy content using `rsconnect`, he configures a custom `RunAs` user, `robert`, for his Shiny application. See Section 10.3 for configuring the `RunAs` user on a per-application basis in Connect.

19.2.4 Configuring `rsconnect`

Important Note: `rsconnect` configuration requires a user home directory. In this use case, a valid home directory is required for the `robert` user account.

Since Bob's Shiny application will be running as the `robert` user, Bob (at a server console) switches to the `robert` user:

```
sudo su robert
```

Next, while running as `robert`, Bob runs R and issues the following commands in the R console:

```
library(rsconnect)
rsconnect::addConnectServer('http://myserveraddress:3939', 'mylocaldeployserver')
rsconnect::connectUser(server='mylocaldeployserver')
```

NOTE: the `rsconnect` server name, `mylocaldeployserver`, is an arbitrary name that is used to identify a Connect server when using `rsconnect`. You can choose any name you wish.

After the last command, Bob sees the following output:

A browser window should open; if it doesn't, you may authenticate manually by visiting `http://myserveraddress:3939/___login___?url=http%3A%2F%2Fmyserveraddress%3A3939%2Fconnect%2F%23%2Ftokens%2FTc8f636c59fff521eef4888b163dcf64%2Factivate&user_id=0`.

Waiting for authentication...

Bob copies the URL in the output above and pastes it into a Web browser. Then Bob authenticates with his Connect user credentials. Bob's Connect user name (with publishing privileges) is `rwhite`.

After successfully connecting his Connect account to `rsconnect`, Bob sees this message at the R console:

Account registered successfully: Bob White (rwhite)

The server and account information are persisted to configuration files on the server in Bob's home directory:

```
/home/robert/.config/R/connect/servers/mylocaldeployserver.dcf
/home/robert/.config/R/connect/accounts/mylocaldeployserver/connectuser.dcf
```

19.2.5 Deploying Content with `rsconnect`

Now `rsconnect` is configured to use the `rwhite` Connect account when running with the `robert` server account. If Bob's Shiny application uses `robert` as its `RunAs` user, it can deploy content using `rsconnect`.

19.2.6 CRAN Note

If you don't already have it in an `RProfile`, be sure to specify a default CRAN repository in your application before issuing the `rsconnect` command to deploy content. For example:

```
options(repos=c(CRAN="https://cran.rstudio.com"))
rsconnect::deployDoc(doc="out.Rmd", appName="ServerDeployedDoc",
                    account="rwhite", server="mylocaldeployserver")
```

19.3 Warning and Security Information

Please restrict access to any Connect content that can deploy arbitrary content via `rsconnect`. The Connect Dashboard's "Permissions" document provides details on securing content in Connect.

Do not enable deployment via `rsconnect` for the default `Applications.RunAs` user; if you do so, all your Connect users will be able to deploy content using your `rsconnect` credentials.

Once a Connect user authorizes `rsconnect` to deploy content under a particular server account, any content that runs under that server account can use `rsconnect` to deploy content without further authentication.

For example, Bob logs in to a server console as Unix user `robert`. Bob then configures `rsconnect` to deploy content. During the authorization step, Bob signs in to Connect as a publisher with user name `rwhite`. Now, any Connect application that is configured with a `RunAs` user of `robert` can deploy additional content using the Connect user `rwhite`, regardless of who owns the application.

19.4 Example Shiny Application

Below is an example Shiny application that knits R Markdown text and deploys the resulting content using `rsconnect`.

```
library(knitr)
library(rsconnect)
library(shiny)
library(shinyAce)
library(rmarkdown)

# Default text for editor
defaultMarkdown <- '
### Sample R Markdown
This is some markdown text. It may also have embedded R code
which will be executed.
'

# A Shiny UI for editing R Markdown
ui <- shinyUI(
  bootstrapPage(
    headerPanel("Embedded Deployment Example"),
    div(
      class="container-fluid",
      div(class="row-fluid",
        div(class="col-sm-6",
          h2("Source R-Markdown"),
          aceEditor("rmd", mode="markdown", value=defaultMarkdown),
          actionButton("eval", "Update")
        ),
        div(class="col-sm-6",
          h2("Knitted Output"),
          htmlOutput("knitDoc")
        )
      )
    )
  )
)

# A Shiny application that generates and deploys R Markdown content
server <- shinyServer(function(input, output, session) {

  # Only update and deploy when the 'Update' button is clicked
  rmd <- eventReactive(input$eval, {
    input$rmd
  })

  output$knitDoc <- renderUI({
    writeLines(rmd(), "out.Rmd")
    knit2html(input="out.Rmd", fragment.only = TRUE, quiet = TRUE)
    options(repos=c(CRAN="https://cran.rstudio.com"))
    rsconnect::deployDoc(doc="out.Rmd", appName="GeneratedDoc",
                        account="rwhite", server="mylocaldeployserver")
    return(isolate(HTML(
```

```
    readLines("out.html")
  )))
})
})

# Run the application
shinyApp(ui = ui, server = server)
```