

RStudio Connect: Admin Guide

Version 1.7.2-7

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	4
1.1	System Requirements	5
2	Getting Started	5
2.1	Need Help?	5
2.2	Installation	6
2.3	Initial Configuration	9
3	License Management	13
3.1	Capabilities	13
3.2	Notification of Expiration	13
3.3	Product Activation	13
3.4	Connectivity Requirements	14
3.5	Evaluations	15
3.6	Licensing Errors	16
3.7	Floating Licenses	16
4	Files & Directories	19
4.1	Program Files	19
4.2	Configuration	19
4.3	Server Log	19
4.4	Access Logs	19
4.5	Application Logs	20
4.6	Variable Data	20
4.7	Backups	22
4.8	Server Migrations	22
5	Server Management	23
5.1	Stopping and Starting	23
5.2	System Messages	25
5.3	Health-Check	25
5.4	Upgrading	25
5.5	Purging RStudio Connect	26
5.6	Docker	26
6	High Availability and Load Balancing	30
6.1	HA Checklist	30
6.2	HA Limitations	30
6.3	Updating HA Nodes	32
6.4	Downgrading	32
6.5	Backups, Snapshots, and Data Integrity	32
6.6	HA Details	33

7	Running with a Proxy	34
7.1	Nginx Configuration	35
7.2	Apache Configuration	36
8	Security & Auditing	37
8.1	API Security	37
8.2	Browser Security	38
8.3	Audit Logs	41
8.4	Application Environment Variables	42
9	Database	42
9.1	SQLite	42
9.2	PostgreSQL	44
9.3	Changing Database Provider	45
9.4	Authentication Provider	46
10	Authentication	46
10.1	Changing Authentication Provider	46
10.2	Session Management	49
10.3	Username requirements	49
10.4	User Attribute Editability	50
10.5	Password	52
10.6	LDAP and Active Directory	53
10.7	SAML (Beta)	61
10.8	OAuth2 (Google)	65
10.9	PAM	67
10.10	Proxied Authentication	69
11	User Management	81
11.1	User Roles	81
11.2	User Provisioning	82
11.3	User Permissions	84
11.4	Administrator Capabilities	85
11.5	Locked Accounts	85
11.6	Username Requirements	86
11.7	User Renaming	86
11.8	Changing Authentication Provider	86
11.9	Removing Accounts	86
11.10	Command-Line Interface	87
12	Group Management	87
12.1	Remote Group Management	87
12.2	Group Ownership	87
12.3	Proxied Authentication Group Membership Management	88
12.4	LDAP Groups	88
12.5	OAuth2 Group Members	89
12.6	Command-Line Interface	89
13	Process Management	89
13.1	Sandboxing	89
13.2	Temporary Directory	91
13.3	Shiny Applications & Plumber APIs	91
13.4	TensorFlow Model APIs	91
13.5	User Account for Processes	92
13.6	Current user execution	92

13.7 PAM sessions	93
13.8 Path Rewriting	94
13.9 Program Supervisors	95
13.10 Using the <code>config</code> Package	96
14 Content Management	97
14.1 Sharing Settings	97
14.2 Vanity Paths	98
14.3 Tags	98
14.4 Bundle Management	99
14.5 Content Images	99
14.6 API Keys	99
15 R	100
15.1 Installing R	100
15.2 Upgrading R	101
15.3 R Versions	101
15.4 R Version Matching	103
16 Python	104
16.1 Enabling Python Support	104
16.2 Installing Python	105
16.3 Upgrading Python	106
16.4 Python Version Matching	107
16.5 Excluding Python Packages	107
16.6 Configuring pip	108
17 Package Management	108
17.1 Package Installation	108
17.2 Private Repositories	109
17.3 Private Packages	110
18 Historical Information	111
18.1 Historical Metrics	111
18.2 Historical Events	111
Appendix	114
A Configuration Options	114
A.1 Configuration Basics	114
A.2 Server	117
A.3 HTTP	120
A.4 HTTPS	120
A.5 HTTPRedirect	121
A.6 Licensing	122
A.7 Database	123
A.8 SQLite	123
A.9 Postgres	124
A.10 Authentication	124
A.11 Password	126
A.12 OAuth2	126
A.13 LDAP	127
A.14 SAML (Beta)	131
A.15 PAM	133
A.16 Proxied Authentication	134

A.17	Authorization	136
A.18	Applications	137
A.19	Packages	140
A.20	Client	141
A.21	Runtime/Scheduler	141
A.22	Jobs	143
A.23	Historical Information	143
A.24	Load Balancing	145
A.25	Python	145
B	Command-Line Interface	146
B.1	User Management	146
B.2	Migration Utility	154
C	Using a Custom Landing Page	156
C.1	Overview	156
C.2	Configuration	156
C.3	Custom Landing Page Assets	156
C.4	Example	156
D	LDAP/AD Configuration Examples	156
D.1	Single Bind	157
D.2	Double Bind	157
D.3	LDIF	157
E	RStudio Connect Deployment Guide	159
E.1	Overview	159
E.2	Programmatic Deployment	159
E.3	Step 1: Building the Bundle	159
E.4	Step 2: Push Bundle to Connect	162
E.5	Step 3: Bundle is deployed on Connect	163
E.6	Other Frequently Asked Questions	164
F	Using Continuous Integration to Deploy Content	164
F.1	Overview	164
F.2	Prerequisites	164
F.3	Configuring a CI Server to Deploy Content to Connect	165
F.4	Warning and Security Information	166
G	Programmatic Deployment with rsconnect	166
G.1	Overview	166
G.2	Use Case: A Shiny Application	167
G.3	Warning and Security Information	168
G.4	Example Shiny Application	168

1 Introduction

RStudio Connect allows users to share and collaborate on the results they produce with R such as R Markdown documents, Shiny applications, Plumber APIs, 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 processes
- Logs for all 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+
- SUSE Linux Enterprise Server 12 SP3+
- Ubuntu 14.04
- Ubuntu 16.04
- Ubuntu 18.04

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

RStudio Connect can be used with R versions 3.1.0 or higher.

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

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

2 Getting Started

This chapter helps you install RStudio Connect on Ubuntu (Section 2.2.3) or Red Hat Enterprise Linux (Section 2.2.4)/CentOS Linux (Section 2.2.4)/ SUSE Linux Enterprise (Section 2.2.4), learn to manage the server (Section 5.1), and perform some initial configuration (Section 2.3).

We built this checklist to guide you through that process.

1. Install R - Ubuntu 2.2.3, Red Hat/CentOS/SUSE 2.2.4
2. Download RStudio Connect installer
3. Install RStudio Connect - Ubuntu 2.2.3, Red Hat/CentOS/SUSE 2.2.4,
4. Set `Server.SenderEmail` - 2.3.1
5. Set `Server.Address` - 2.3.1
6. Configure Authentication - 2.3.2, 10
7. Restart RStudio Connect - 5.1
8. Sign into RStudio Connect - 2.3.3
9. Configure email sending - 2.3.4
10. Install TensorFlow Dependencies (Optional) - 2.3.5
11. Install Python (Optional) - 16

2.1 Need Help?

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

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

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

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

1. Running the diagnostics script:

```
/opt/rstudio-connect/scripts/run-diagnostics.sh
```

Note: if you have not been able to download the software, skip this step.

2. Email support@rstudio.com with:
 - Your name and company
 - What you are trying to do
 - What you have tried
 - Any errors you have received
 - The diagnostics file produced from step 1

Our Support Agreement contains important details about the support we provide.

2.2 Installation

This section explains how to install R on an Ubuntu or Red Hat/CentOS/SUSE server. You must install R version 3.1.0 or higher.

Administrators should install the versions of R that they wish to support from source so that user content is run in an environment as close as possible to the development environment. This allows maintenance of multiple versions of R simultaneously and mitigates the risk associated with updating the version of R.

2.2.1 Package Managers

The best practice is to install R from source rather than with a package manager like `apt`, `yum`, or `zypper`. Installing R with a package manager means that R will be updated when the package manager update command is run and a new version is released. Any content that was uploaded and built with the old R version will need to be rebuilt for the new R version. This is not guaranteed to succeed.

2.2.2 Build R from Source

To build R from source, first, acquire the build dependencies for R.

In Ubuntu, you can install build dependencies with

```
apt-get build-dep r-base
```

Or determine and install dependencies independently by using the following command:

```
apt-rdepends --build-depends r-base
```

In RedHat Enterprise Linux and CentOS, you can install build dependencies with

```
yum-builddep R
```

Or determine and install dependencies independently by using the following command:

```
yum deplist R
```

In SUSE, build dependencies should be installed directly with `zypper install <dependency>`. A partial list of the libraries that you might need is below.

```
zypper install \  
  gcc \  
  gcc-c++ \  
  gcc-fortran \  
  readline-devel \  
  xorg-x11-devel \  
  liblzma5 \  
  xz-devel \  
  pcre-devel \  
  libcurl-devel \  
  make
```

Second, you should download and unpack the source tarball for the version of R that you want to install from CRAN. To install R-3.4.3, this might look like the following:

```
# Download and extract source code  
wget https://cran.r-project.org/src/base/R-3/R-3.4.3.tar.gz  
tar -xzvf R-3.4.3.tar.gz  
cd R-3.4.3
```

The final step is to configure, make, and install R. The recommended install location in this case is at `/opt/R/3.4.3`. If you would like to choose a different location, you can change `--prefix=`. To learn more about how Connect looks for R installations, see section 15.3.1.

```
# Build R from source  
./configure \  
  --prefix=/opt/R/$(cat VERSION) \  
  --enable-memory-profiling \  
  --enable-R-shlib \  
  --with-blas \  
  --with-lapack  
  
make  
sudo make install
```

It is very important that the R installation folder is *not moved* once it is installed from source. Libraries are statically linked, so moving the folder will break the installation of R.

To test that the installation went smoothly, execute:

```
/opt/R/3.4.3/bin/R --version
```

Once you have R installed on the server, it is important to understand how Connect discovers R and chooses the R version that an asset will use. Chapter 15 explains more about this process.

2.2.3 Ubuntu (14.04+)

RStudio Connect requires an installation of R version 3.1.0 or higher. See section 2.2.2 for details on installing R from source.

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 a `.deb` installer for RStudio Connect. It can be downloaded from the RStudio website. 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_1.7.2-7_amd64.deb
```

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

```
sudo gdebi rstudio-connect_1.7.2-7_amd64.deb
```

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

You can now configure the server following the instructions in Section 2.3.1. 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.

```
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
```


2.2.4 Red Hat Enterprise Linux/SUSE Linux Enterprise (12 SP3+)/CentOS Linux (6.0+)

Prerequisites

RStudio Connect requires an installation of R version 3.1.0 or higher. See section 2.2.2 for details on installing R from source.

For RedHat/CentOS, 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, 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.

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/zypper`. If you have only a link to the RPM file, you can use `wget` to download the file to the current directory.

```
wget https://download-url/rstudio-connect-1.7.2-7.x86_64.rpm
```

For RedHat/CentOS:

```
sudo yum install --nogpgcheck rstudio-connect-1.7.2-7.x86_64.rpm
```

For SUSE:

```
sudo zypper --no-gpg-checks install rstudio-connect-1.7.2-7.x86_64.rpm
```

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

You can now configure the server following the instructions in Section 2.3.1. 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, and SUSE Enterprise 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
```

2.3 Initial Configuration

RStudio Connect is installed, but requires additional configuration before it is ready for use. This section will help you specify the public URL of your server, configure authentication, and validate that RStudio Connect is able to send email.

2.3.1 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 A for details about this file, its syntax, and the available settings.

Start by 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 2.3.4 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. This setting enables Connect 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.

Important Note

Please use a publicly available URL (like the one you set in `Server.Address`) when connecting `rsconnect` or the RStudio IDE to your RStudio Connect server. If a non-public address (e.g., `localhost`) is used for publishing content, `rsconnect` will not be able to automatically open the published content in the user's browser.

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.

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

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

Here is a sample configuration with RStudio Connect using a non-standard port:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
SenderEmail = rstudio-connect@company.com
Address = http://rstudio-connect.company.com:3939/
```

This example has RStudio Connect available underneath a specific URL path. You may have this situation when Connect is accessed through a proxy.

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

By default the version and build number for RStudio Connect is displayed within the dashboard. If you wish to hide this information you can do so by configuring the `Server.HideVersion` setting. When set to `true` the version and build will be hidden from users.

Restart RStudio Connect after altering the `rstudio-connect.gcfg` configuration file.

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

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

2.3.2 Authentication

It is important that you specify the correct style of authentication for your organization. RStudio Connect includes a built-in authentication mechanism and supports a number of external authentication integrations, which are detailed in Section 10.

It's **strongly** recommended that the definitive form of authentication be established before using a permanent RStudio Connect installation. Migrating from one style of authentication to another can be done but it is a challenging process that will cause some temporary disruption.

The means to migrate between authentication providers are described in the section 10.1.

2.3.3 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. These settings are necessary in order for Connect to be able to distribute reports and notify users of errors running their content. Connect also sends confirmation messages when using the default password auth provider.

2.3.4 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!

At this point, RStudio Connect is installed and ready for use. The rest of the administration guide covers additional configuration options.

2.3.5 Install TensorFlow Dependencies (Optional)

Note: `libtensorflow.so` is required for TensorFlow Model API deployment support, and an up-to-date `libtensorflow.so` is required to run TensorFlow Model APIs at the latest version. If users are experiencing difficulty deploying TensorFlow Models to RStudio Connect, try updating `libtensorflow.so` to the latest version following the instructions below.

To install the TensorFlow dependencies, first verify that your license supports TensorFlow Model API deployment. If you do not have a Standard or Enterprise license, please contact RStudio Sales for information on upgrading your license.

Second, verify that your platform is supported by TensorFlow. TensorFlow requires libraries that are unavailable on CentOS 6, Red Hat Enterprise Linux 6, and SUSE Linux Enterprise Server 12. Accordingly, RStudio Connect does not provide support for running TensorFlow Model APIs on these platforms.

Next, determine what version of TensorFlow your users will be deploying their models with. RStudio Connect supports TensorFlow versions 1.9.0 and earlier. Later releases and development TensorFlow builds are not supported.

Finally, there are two options for installing the required libraries:

- You could obtain a `libtensorflow` binary package. This binary package would not be optimized for your platform, but would provide a working runtime for TensorFlow Model APIs on connect. See 2.3.5.1
- You could compile `libtensorflow` from source. This would allow you to make various compile-time optimizations for your hardware. See 2.3.5.2

2.3.5.1 TensorFlow Binary Packages

The guide `Installing TensorFlow For C` should contain a section reading “Download and extract the TensorFlow C library”. Follow the instructions up to that point, download and extract the binary, and then:

- On Ubuntu, the libraries may need to be installed to `/usr/lib` rather than `/usr/local/lib` as specified in the instructions.
- On CentOS7/RHEL7, the libraries need to be installed to `/usr/lib64` rather than `/usr/local/lib` as specified in the instructions.
- On ALL PLATFORMS: Ensure that both `libtensorflow.so` and `libtensorflow_framework.so` are in your shared library search path.

2.3.5.2 Compiling TensorFlow Libraries From Source

NOTE: These instructions have been tested as of TensorFlow 1.7, please refer to the TensorFlow Project for up-to-date information on building TensorFlow Libraries.

NOTE: Compiling TensorFlow from source may be difficult and is highly dependent on your host environment. If you encounter issues compiling `libtensorflow.so` from source, you should raise an issue at the TensorFlow project’s Github Issues Page

1. Install Bazel per the Bazel Installation Instructions
2. Download the TensorFlow Source either using `git clone` or by downloading a source tarball.
3. Python, Swig, and NumPy may be required on the build platform. - On Ubuntu, they can be installed with `sudo apt-get install python swig python-numpy` - On CentOS, they can be installed with `sudo yum install python swig numpy`
4. Enter the source directory and run `./configure`. Answer the questions about your hardware to the best of your knowledge. Application logs will contain a warning if the platform’s tensorflow libraries lack usable optimizations.
5. Execute `bazel build --config opt //tensorflow:libtensorflow.so`. This will compile tensorflow for your platform
6. Copy `libtensorflow.so` and `libtensorflow_framework.so` from `bazel-bin/tensorflow` into your platform’s shared library path

3 License Management

3.1 Capabilities

An RStudio Connect product key limits usage 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. Additionally, users that have not recently been active on the server are not counted against this quota. Users are deemed “inactive” after 365 days without visiting RStudio Connect, though this value may vary for certain licenses.
- 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.
- Whether or not API hosting is supported.

The settings of each metric depends on the license purchased from RStudio.

3.2 Notification of Expiration

RStudio Connect will attempt to send email to administrators when the license key is sixty days from expiration. You can disable this behavior with the `Licensing.ExpirationEmail` setting.

The Connect Dashboard will display a notification to admins and publishers when the license is within 15 days of expiration. You can disable this behavior with the `Licensing.ExpirationUIWarning` setting.

3.3 Product Activation

3.3.1 Activation Basics

When RStudio Connect is first installed on a system it operates in 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’ll receive a license key that can be used to activate the license on a given system.

You can activate your license key with the command:

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

Note that you need to restart RStudio Connect after activation for licensing changes to take effect.

Your platform may use 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 with the command:

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

3.4 Connectivity Requirements

In order to activate or deactivate RStudio Connect, internet connectivity is required for communication with the licensing server. If your server is behind an internet proxy or not connected to the Internet at all this section describes what's required to successfully activate.

Additionally, your server should have a synchronized system clock, using `ntp` or some other clock syncing service. If the server's clock is sufficiently incorrect, licensing verification will fail.

3.4.1 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.

3.4.2 Offline Activation

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

To activate your license offline, 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 enter into our offline activation application or 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>  
sudo systemctl restart rstudio-connect
```

Note that you need to restart RStudio Connect after activation for licensing changes to take effect.

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

If you are renewing your license or 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 enter into the offline activation application 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
```

3.5 Evaluations

3.5.1 Extending Evaluations

If you are unable to evaluate RStudio Connect during the initial evaluation period, you can obtain a key for extending the evaluation period from RStudio customer support (support@rstudio.com). Once you have the key, supply it to the RStudio Connect license manager using the `extend-evaluation` command.

```
sudo /opt/rstudio-connect/bin/license-manager extend-evaluation <key>
```

If you are performing the evaluation on a physical machine (not on virtualized hardware or containers) without a network connection, you may also request an offline evaluation extension key, which does not require an internet connection. This key may be supplied to the license manager as follows:

```
sudo /opt/rstudio-connect/bin/license-manager extend-evaluation-offline <key>
```

Note that offline evaluation extension keys are valid *only* on machines which do not have Internet access and are not virtualized. For most offline evaluation extensions, you will need to generate an offline evaluation request (see below for details).

3.5.2 Connectivity Requirements

3.5.2.1 Beginning Evaluations

Generally speaking, there are no network requirements during the evaluation period. Inside virtual machines or sandboxes (such as Docker), however, Internet access is required to begin the evaluation period.

If you have a proxy, you can supply it using the `--proxy` argument as described above. If however you have no means of connecting to the Internet from inside the virtual environment, you can begin the evaluation as follows:

```
sudo /opt/rstudio-connect/bin/license-manager begin-evaluation-request
```

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 begin the evaluation offline as follows:

```
sudo /opt/rstudio-connect/bin/license-manager \  
  begin-evaluation-offline <evaluation-file>  
sudo systemctl restart rstudio-connect
```

Note that you need to restart RStudio Connect after starting your offline evaluation for licensing changes to take effect.

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

3.5.2.2 Extending Evaluations

You may extend evaluations offline using the same pattern described above (just use `extend-evaluation-request` and `extend-evaluation-offline`):

```
sudo /opt/rstudio-connect/bin/license-manager extend-evaluation-request
```

Then, when you've received the evaluation file:

```
sudo /opt/rstudio-connect/bin/license-manager
  extend-evaluation-offline <evaluation-file>
sudo systemctl restart rstudio-connect
```

Note that you need to restart RStudio Connect after extending your evaluation for licensing changes to take effect.

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

3.6 Licensing Errors

RStudio Connect uses the `license-manager` to determine if a valid license is available. Should an error occur when interacting with the license manager, Connect indicates that problem in the `/var/log/rstudio-connect.log` log. The license manager sends details about the error to the system messages (`syslog`). You should consult both locations if Connect cannot obtain a license.

3.7 Floating Licenses

If you stop and start RStudio Connect instances frequently, for instance because you're running them inside virtual machines or containers, you may wish to use floating licensing instead of traditional licensing.

To use floating licensing, you run a small, lightweight server, which holds a license that grants you the right to run a certain number of concurrent RStudio Connect instances.

When RStudio Connect starts, it will connect to the license server and obtain a temporary lease, releasing it when RStudio Connect is stopped. Using this method, you can have any number of RStudio Connect instances, so long as you do not run more instances at once than specified in your license.

3.7.1 The RStudio Connect License Server

The RStudio License Server site contains license server downloads for all RStudio products. Download and install the license server for RStudio Connect. You then activate your license key with the command:

```
sudo dpkg -i connect-license-server-1.0.3-x86_64.deb
sudo connect-license-server activate <product-key>
sudo connect-license-server start
```

A license key which distributes floating license leases is not the same as a traditional license key, and the two cannot be used interchangeably. If you have purchased traditional license keys and wish to exchange them for a floating license key, or vice versa, please get in touch with RStudio customer support (support@rstudio.com).

The file `/etc/connect-license-server.conf` contains configuration settings for the RStudio Connect License server, including the network port to listen on and any proxy settings required for connecting to the Internet.

3.7.2 License Server Offline Activation

The `connect-license-server activate` command requires an internet connection. If your license server has no connection to the Internet it's also possible to perform an offline activation. The process for doing this on the license server is identical to the process used to activate RStudio Connect offline. Generate an offline activation request as follows:

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

3.7.3 Using Floating Licensing

Once your license server is up and running, you need to tell RStudio Connect to use floating licensing instead of traditional licensing.

```
/etc/rstudio-connect/rstudio-connect.gcfg
; /etc/rstudio-connect/rstudio-connect.gcfg
[Licensing]
LicenseType = remote
```

The value `remote` indicates that RStudio Connect should connect to a remote licensing server to obtain a license; the value `local` can be used to explicitly specify traditional (local) activation.

Then, tell RStudio Connect which licensing server to connect to:

```
sudo /opt/rstudio-connect/bin/license-manager license-server <server-hostname-or-ip>
sudo systemctl restart rstudio-connect
```

Note that you need to restart RStudio Connect after configuring a remote license server for licensing changes to take effect.

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

You only need to run the `license-server` command once; RStudio Connect saves the server name and will use it on each subsequent startup.

By default, the RStudio Connect License Server listens on port 8999. If you wish to use a different port, you will need to specify the port in `/etc/connect-license-server.conf`, and specify `license-server` to the license manager as `<server-hostname-or-ip:port>`.

Depending on your system configuration, it is possible that the RStudio Connect service will be started before the service which allows hostname resolution (this is known to be the case for example on some Amazon EC2 systems). If this is the case, you'll want to specify the license server using a private IP address rather than a hostname, so that RStudio Connect can acquire a license immediately when starting up.

3.7.4 Configuring License Leases

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

```
/etc/connect-license-server.conf
```

```
<lease length="1800"/>
```

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

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

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

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

3.7.5 Lease Expiration and Renewal

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

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

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

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

3.7.6 Troubleshooting Floating Licensing

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

```
sudo connect-license-server activation-status
```

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

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

Then, restart the license server, tail the licensing log, and start your RStudio Connect instances.

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

At the `notification` level, the licensing log will tell you the total number of licenses associated with your key, and how many are currently in use. It will also notify you when RStudio Connect instances acquire leases,

and when those leases are released, renewed, or expired. No rotation is done for this log, so it's recommended to use the `warning` level in production.

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`. This file is initially owned by `root` with permissions `0600`. 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 remain 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. Similarly, the initial installation of RStudio Connect will not overwrite the `rstudio-connect.gcfg` file if it already exists.

Restart RStudio Connect after altering the `rstudio-connect.gcfg` configuration file using the instructions in Section 5.1.

Configuration settings marked as “reloadable” do not require a full restart. See Section A to learn which properties are reloadable. You can find a “reload” command for your operating system in Section 5.1.

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 stored alongside the original with a numeric extension, `.1`, `.2`, etc. The rotated log files are compressed after one day. The `.1` log file is retained uncompressed, but older logs are compressed. Most systems use `gzip` for compression, giving log files with extensions like `.2.gz`, `.3.gz`. Logs will be maintained for 30 days.

The manual for `logrotate` has more information.

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 be maintained for 30 days.

4.5 Application Logs

Each 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 process. Standard output and standard error are captured and available.

4.6 Variable Data

RStudio Connect manages uploaded Shiny applications, Plumber and TensorFlow APIs, 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.
- Application images as uploaded by the user.

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.
- 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 `Server.DataDir` in your configuration file.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
DataDir = /mnt/rstudio-connect
```

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
DataDir = /mnt/rstudio-connect

[SQLite]
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 18.1.1.

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 `reports` tree contains a nested directory structure of the form: `v2/A_ID/V_ID/R_ID`. The `A_ID` directory corresponds to a content deployment (an R Markdown document). The `V_ID` directory represents a configuration of that document (a set of parameter values). The `R_ID` contains a single rendering of that document with the associated parameters. The directories `v2`, `A_ID`, and `V_ID` are all owned by `root` with `0711` permissions. The final directory, `R_ID` contains the actual rendered output and is owned by `rstudio-connect` with `0700` permissions. Files contained in the `R_ID` directory will have `0600` permissions.

The `bookmarks` directory contains a bookmarking state subdirectory for each Shiny application. The top-level directory is owned by `root` with `0711` permissions. Each `bookmarks/A_ID` subdirectory is owned by `rstudio-connect` and the `rstudio-connect` group with `0770` permissions.

Learn more about server-stored Shiny bookmarking state in this article.

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 13.5.

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 and `reports/v2/A_ID/V_ID/R_ID` directories are 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` and `reports/v2/A_ID/V_ID/R_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 systemctl stop rstudio-connect
# Run appropriate backup steps here.
sudo systemctl start rstudio-connect
```

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

4.8 Server Migrations

There are a number of factors that must be considered before migrating your RStudio Connect installation from one server to another. We recommend making as few changes as possible during the initial migration. If, for instance, you will be migrating to a new server, upgrading to a new default version of R, and altering the default `Applications.RunAs` user, complete the migration first. Then upgrade R and alter the `Application.RunAs` user in subsequent steps.

In order to migrate a server, you will follow the same steps as when you perform a backup in order to obtain a consistent copy of the data in the necessary directories. These directories can then be copied to the new server.

1. Install RStudio Connect on the new server, then stop the service. RStudio Connect v1.5.6 introduced features that make server migrations more reliable; migrating servers in older versions is not supported, so the new server should have v1.5.6 or later.
2. Mirror the Unix accounts used by RStudio Connect on the existing server to the new server. Consider the `Applications.RunAs` user and any other users that might have been selected as the user responsible for running any content on the server. These Unix accounts must all exist on the new server and continue to be members of the default `Applications.RunAs` user's primary group as discussed in 13.5.
3. Copy the config and data directories while preserving the *permissions* and file *ownership*. Not all file transfer clients are able to preserve these attributes, so consider using `rsync` with the `-a` flag to copy the data. Bear in mind that certain applications may have overridden settings that alter how their files are stored on disk (for instance, by customizing the user account that runs their processes - e.g. Shiny application), so it is critical that ownership and permissions be preserved exactly during the migration.
4. Update your `/etc/rstudio-connect/rstudio-connect.gcfg` file if you've changed settings like the path to your data directory.
5. Sanity-check the permissions and ownership of the content working directories using the `migrate repair-content-permissions` command as documented in B.2.
6. Install the same version(s) of R on the new server to mimic existing behavior. If you need additional versions or support for multiple versions of R, please see 15.
7. On the new server, install any system dependency that may be used by an R package on the existing server. A list of recommended packages are available in 2.2. Whichever packages you chose to install on your existing server to support the R packages that users have deployed should also be installed on the new server. Otherwise, RStudio Connect will not be able to rebuild users' deployed packages.
8. Run `migrate rebuild-packrat --force` to delete the Packrat cache and rebuild it. This cache likely includes binaries that were compiled against particular versions of libraries on your existing system. Be aware that **this step may take a very long time** (easily multiple hours for large deployments with lots of content). It is recommended you start this before you start RStudio Connect, but you can start

RStudio Connect once it starts. If any application or report is executed, the packrat directory for that application will be rebuilt at runtime.

If you are also migrating to a different database provider, see 9.3.

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.
- Processes serving applications and APIs are stopped.
- Processes rendering documents run through completion.
- In-progress deployments will fail. Processes running as part of the deployment may run to completion.

Start:

- RStudio Connect process is resumed.
- Applications and APIs with a minimum number of processes are started.
- Scheduled document render updates missed during system downtime are run at most once.

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

5.1.1 `systemd`

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

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

The RStudio 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
```

You can trigger a configuration reload with the `systemctl reload` command, which reads the changed configuration from disk:

```
sudo systemctl reload 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.

Use `systemctl reload` to send a HUP signal when your configuration changes are limited to properties marked as “reloadable”. See Appendix A to learn which settings may be reloaded. Perform a full restart of RStudio Connect when changing other properties.

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 14.04, Red Hat 6)

Distributions using Upstart include Red Hat/CentOS 6 and Ubuntu 14.04.

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.*

You can trigger a configuration reload with the `reload` command, which reads the changed configuration from disk:

```
sudo reload 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.

Use `reload` to send a HUP signal when your configuration changes are limited to properties marked as “reloadable”. See Appendix A to learn which settings may be reloaded. Perform a full restart of RStudio Connect when changing other properties.

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.

The `Server.PublicWarning` setting 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:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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 `systemd` (Red Hat/CentOS 7, SUSE 12, Ubuntu 16.04, and Ubuntu 18.04):

```
sudo systemctl reload rstudio-connect
```

or Upstart (Ubuntu 14.04, Red Hat 6):

```
sudo reload rstudio-connect
```

5.3 Health-Check

RStudio Connect provides a simple health-check endpoint that can be used to test if Connect is up/listening. Point your browser to `myserveraddress:myserverport/__/ping__`, which returns an empty JSON response and an HTTP 200 status.

```
curl -I -X GET http://myserveraddress:myserverport/__/ping__
```

5.4 Upgrading

Upgrading RStudio Connect requires limited downtime. Scheduled document renders are not interrupted. Connections to running applications and APIs 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: bash sudo gdebi <rstudio-connect-version.deb>
```

```
Red Hat/CentOS: bash sudo yum install --nogpgcheck <rstudio-connect-version.rpm>
```

```
SUSE: bash sudo zypper --no-gpg-checks install <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.

5.5 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
```

SUSE:

```
sudo zypper remove rstudio-connect
```

3. Remove logs from `/var/log/rstudio-connect*`
4. Purge the database
 - When using SQLite, remove the `SQLite.Dir` directory. This has a default location of `/var/lib/rstudio-connect/db`.
 - When using PostgreSQL, drop the database used by Connect. You may also wish to remove the PostgreSQL user associated with Connect.
5. Remove the `Server.DataDir` directory. By default, this is `/var/lib/rstudio-connect`.
6. Remove configuration files from `/etc/rstudio-connect` if they still exist.

5.6 Docker

Some organizations are starting to deploy RStudio Connect within a Docker container. Connect will manage and run your processes within that container. This section describes the requirements for a container-based deployment.

If you are using multiple Docker containers to run RStudio Connect, Section 6 contains requirements for a load-balanced, multi-host configuration.

Please contact the RStudio Solutions Engineering team (sol-eng@rstudio.com) for guidance deploying RStudio Connect into your particular environment.

5.6.1 Distributions

RStudio Connect is supported on the Linux distributions given in Section 1.1. Docker-based deployments need to use one of these operating systems.

5.6.2 R

R version 3.1.0 or higher must be installed within your Docker container. We recommend installing a curated set of R versions from source rather than using R installed by a package manager. The instructions in Section 2.2.2 contain the commands you will use to install R when building a Docker image.

5.6.3 Privileged containers

RStudio Connect runs processes within an “unshare” environment using bind mounts as described in Section 13.1. These bind mounts help Connect hide and isolate parts of the filesystem from executed user processes. Bind mounts are not permitted when running within an unprivileged container.

Docker containers are allowed mount permissions with the `SYS_ADMIN` capability. You can give these permissions to your container using either the `--cap-add=SYS_ADMIN` or `--privileged=true` option.

Learn more about Linux capabilities and `CAP_SYS_ADMIN` in the `capabilities` man page.

The Docker container running RStudio Connect must be started with either `--cap-add=SYS_ADMIN` or `--privileged=true`.

5.6.4 Licensing

We recommend using a floating license server and having RStudio Connect use that license server. Section 3.7) explains how to obtain and configure the license server as associating RStudio Connect with the floating license server.

We recommend using `docker stop` rather than `docker kill` to terminate your RStudio Connect Docker containers. Using `docker stop` allows floating license leases to be gracefully released. License leases not returned to the license server need to wait for lease expiration before they can be reused.

Please contact the RStudio Solutions Engineering team (sol-eng@rstudio.com) if your organization is unable to run a persistent floating license server.

5.6.5 Storage

Mount a persistent storage location into the Docker container and configure that location as your `Server.DataDir`. Do *not* mount this to the default `/var/lib/rstudio-connect` location. See Section 4.6 for more information on the variable data managed by RStudio Connect.

5.6.6 Database

The default SQLite database provider may be used when RStudio Connect runs inside a *single* Docker container and your `Server.DataDir` is a mounted volume that is NOT part of a networked (NFS) share.

You *must* use PostgreSQL in all multiple-container deployments or if your persistent storage location is a networked location.

5.6.7 Examples

This section contains an `rstudio-connect.gcfg` configuration file and `Dockerfile` that will help get you started building your own environment.

With the `Dockerfile` and `rstudio-connect.gcfg` files in a directory, you can build a Docker image tagged with `rstudio/connect-docker:latest` with the command:

```
docker build -t rstudio/connect-docker .
```

Once the image is built, it can be run with the command:

```
docker run -d --privileged=true --rm \  
  -p :9999:3939 \  
  -v $(pwd)/data:/data \  
  rstudio/connect-docker:latest
```

This launches RStudio Connect within a Docker container. The Connect instance is available on port :9999 (published from :3939 within the container). This sample mounts the `./data` directory as our persistent storage and will need updating for your environment.

5.6.7.1 Configuration

This `rstudio-connect.gcfg` file is copied into the Docker image by a `COPY` command in the `Dockerfile` that follows. This file is only an initial configuration and will need customization for your organization.

```
; /etc/rstudio-connect/rstudio-connect.gcfg  
  
; RStudio Connect sample configuration  
[Server]  
; SenderEmail is 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.  
SenderEmail = account@company.com  
  
; The public URL to this RStudio Connect container. This might be the address for  
; a proxy or the host running the Docker container.  
Address = https://rstudio-connect.company.com  
  
; The persistent data directory mounted into our container.  
DataDir = /data  
  
[Licensing]  
LicenseType = remote  
  
; Use and configure our PostgreSQL database.  
[Database]  
Provider = postgres  
  
[Postgres]  
URL = "postgres://username:password@db.company.com/connect"  
  
[HTTP]  
; RStudio Connect will listen on this network address for HTTP connections.  
Listen = :3939  
  
[Authentication]  
; Specifies the type of user authentication.  
Provider = password
```

5.6.7.2 Dockerfile

Here is a very simple `Dockerfile` that can be used to run RStudio Connect within a Docker container. This simple example uses `apt-get` to install `r-base` from the Ubuntu 18.04 (bionic) repository.

```
# RStudio Connect sample Dockerfile  
FROM ubuntu:bionic
```

```

# R from Ubuntu + necessary tools.
RUN export DEBIAN_FRONTEND=noninteractive && \
  apt-get update && \
  apt-get install -y \
    curl \
    gdebi-core \
    r-base

ARG CONNECT_VERSION=1.7.2-7
ARG CONNECT_URL_BASE=https://s3.amazonaws.com/rstudio-connect
ARG CONNECT_PACKAGE=rstudio-connect_${CONNECT_VERSION}_amd64.deb
ARG CONNECT_URL=${CONNECT_URL_BASE}/${CONNECT_PACKAGE}
RUN curl -sL -o rstudio-connect.deb ${CONNECT_URL} && \
  gdebi -n rstudio-connect.deb && \
  rm rstudio-connect.deb

# Copy our configuration over the default install configuration
COPY rstudio-connect.gcfg /etc/rstudio-connect/rstudio-connect.gcfg

# Use a remote license server issuing floating licenses
RUN /opt/rstudio-connect/bin/license-manager license-server licensing.company.com

# Expose the configured listen port.
EXPOSE 3939

# Launch Connect.
CMD ["--config", "/etc/rstudio-connect/rstudio-connect.gcfg"]
ENTRYPOINT ["/opt/rstudio-connect/bin/connect"]

```

Install R from source (see Section 2.2.2) by updating that Dockerfile with the following:

```

# Install dependencies needed to build R from source + necessary tools.
RUN export DEBIAN_FRONTEND=noninteractive && \
  apt-get update && \
  apt-get install -y curl gdebi-core && \
  apt-get build-dep -y r-base

# Fetch and install R
ARG R_VERSION=3.4.3
ARG R_ARCHIVE=R-${R_VERSION}.tar.gz
ARG R_URL=https://cran.r-project.org/src/base/R-3/${R_ARCHIVE}
RUN curl -sL -O ${R_URL} && \
  tar xzf ${R_ARCHIVE} && \
  cd R-${R_VERSION} && \
  ./configure \
    --prefix=/opt/R/${R_VERSION} \
    --enable-memory-profiling \
    --enable-R-shlib \
    --with-blas \
    --with-lapack && \
  make && make install && \
  cd .. && \
  rm -r R-${R_VERSION}

```

This snippet installs only a single R version but may be extended to install a number of different R versions. Learn more about how RStudio Connect can use multiple R installations in Section 15.

Your organization may have a shared volume containing all of the available R installations. You could mount those installations into the container as it is started. This sample command mounts a shared `/shared/software/R` location into the well-known `/opt/R` location within the container:

```
docker run -d --privileged=true --rm \
  -p :9999:3939 \
  -v $(pwd)/data:/data \
  -v /shared/software/R:/opt/R \
  rstudio/connect-docker:latest
```

6 High Availability and Load Balancing

Multiple instances of RStudio Connect can share the same data in highly available (HA) and load-balanced configurations. We refer to these configurations as “HA” for brevity.

6.1 HA Checklist

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

1. Install and Configure the same version of RStudio Connect on each node - 2
2. Migrate to a PostgreSQL database (if running SQLite) - 9.3. All nodes in the cluster must use the same PostgreSQL database.
3. Configure each server's `Server.DataDir` to point to the same shared location - 4.6 and 6.2.5
4. If the `Database.Dir` setting has been customized, ensure that it points to a consistent, shared location on each server - 4.6 and 6.2.5
5. Configure each server's `Server.LandingDir` to point to the same shared location (if using a custom landing page) - C and 6.2.5
6. Configure each server's `Metrics.DataPath` directory to point to a unique-per-server location - A.23. Alternatively, you may also wish to consider using Graphite to write all metrics to a single location - 6.2.6
7. Configure `LoadBalancing.EnforceMinRsconnectVersion` on each server to ensure that clients use a compatible version of `rsconnect` - 6.2.8.
8. Configure your load balancer to route traffic to your RStudio Connect nodes with sticky sessions - 6.2.8

6.2 HA Limitations

6.2.1 Clock Synchronization

All nodes in an RStudio Connect HA Configuration MUST have their clocks synchronized, preferably using `ntp`. Failure to synchronize system clocks between nodes can lead to undefined behavior, including loss of data.

6.2.2 Consistent Users and Groups

RStudio Connect executes your content using one or more target accounts. The `rstudio-connect` user and its primary group (also named `rstudio-connect`) are created when RStudio Connect is installed. The `rstudio-connect` user serves as the default for `Applications.RunAs` and is used when executing content.

You must ensure consistent UID/GID for all Unix user accounts that may be used to execute your deployed content.

The `id` command is one way to check the user and group identifiers for a single Unix username. Your systems administrator can probably help configure accounts in a uniform way across your cluster of hosts.

```
id rstudio-connect
# => uid=998(rstudio-connect) gid=998(rstudio-connect) groups=998(rstudio-connect)
```

6.2.3 Node Management

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

6.2.4 Database Requirements

RStudio Connect only supports HA when using a PostgreSQL database. If you are using SQLite, please switch to PostgreSQL. See 9.3.

6.2.5 Shared Data Directory Requirements

RStudio Connect manages uploaded content within the server's data directory. This data directory must be a shared location. The `Server.DataDir` configuration on each node must point to the same shared location. See 4.6 for more information on the server's data directory. We recommend and support NFS version 3 for file sharing.

If you configure `Database.Dir` (not required), this also must point to the same shared location.

6.2.6 Metrics Requirements

By default, RStudio Connect writes metrics to a set of RRD files. We do not support metrics aggregation, and each server must maintain a separate set of RRD files to avoid conflicts. The Connect dashboard for a specific node will only show metrics for that node. See A.23 for information on configuring a unique `Metrics.DataPath` for each server.

RStudio Connect includes optional support for writing metrics to Graphite. If you wish to aggregate metrics, consider using Graphite or any monitoring tool compatible with Carbon protocol. See 18.1 for more information.

6.2.7 Shiny Applications

Shiny applications depend on a persistent connection to a single server. Please configure your load-balancer to use cookie-based sticky sessions to ensure that Shiny applications function properly when using HA.

6.2.8 rsconnect Cookie Support

For cookie-based sticky session support, you will need to ensure that your R clients (including the RStudio IDE) use `rsconnect` version 0.8.3 or later. Versions of `rsconnect` prior to 0.8.3 did not include support for cookies. Enable `LoadBalancing.EnforceMinRsconnectVersion` to ensure that clients use a version of `rsconnect` with cookie support.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LoadBalancing]
EnforceMinRsconnectVersion = true
```

If you cannot enforce a minimum `rsconnect` version, you can consider alternatives like:

- Non-cookie-based sticky sessions, or
- Providing a separate host name for deployment from `rsconnect` to a single node in the cluster. Content deployed to a specific node will be available to the cluster assuming the database and shared storage are appropriately configured.

6.3 Updating HA Nodes

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

1. Stop all RStudio Connect nodes in your cluster.
2. Upgrade one RStudio Connect node. The first update will upgrade the database schema (if necessary) and start RStudio Connect on that instance - 5.4.
3. Upgrade the remaining nodes.

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

6.4 Downgrading

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

1. Stop all Connect services on all nodes
2. Reconfigure your network to route traffic directly to one of the nodes, unless you wish to continue using a load balancer.
3. If you wish to move all shared file data to the node, then
 1. Configure the server's `Server.DataDir` to point to a location on the node, and copy all the data from the NFS share to this location - 4.6
 2. If using a custom landing page, configure `Server.LandingDir` to point to a location on the node, and copy the custom landing page data from the NFS share to this location - C
 3. Configure the server's `Metrics.DataPath` directory to point to an appropriate location. If necessary, copy the data from the NFS share to this location. - 6.2.6
4. If you wish to move the database to this node, install PostgreSQL on the node and copy the data. Moving the PostgreSQL database from one server to another is beyond the scope of this guide. Please note that we do not support migrating from PostgreSQL back to SQLite.
5. Start the Connect process - 5.1

6.5 Backups, Snapshots, and Data Integrity

NOTE: Do not attempt to restore snapshots taken while RStudio Connect is running. Incremental snapshots such as these may be useful for security auditing purposes. Additionally, the `bundles` directory has compressed archives that could be redeployed from the RStudio IDE if necessary.

RStudio Connect has two data sources that should be consistent with each other for normal operation: the data directory, and the database.

In a single-node configuration, the data directory is usually some directory on the host node and the database is a SQLite file also located on the host node.

In a HA, multiple-node configuration, the data directory is generally located on an NFS share, and the database is a postgres database, often hosted elsewhere.

When backing up a single node, it is sufficient to stop the node, copy the data directory and database file, and then restart the node.

For a multiple node configuration, you can also stop all nodes, copy the data directory, dump the database using `pg_dump`, and then restart all the nodes again. This is the only way to guarantee that your RStudio Connect environment can be restored without any inconsistency.

In order to address concerns about day-to-day data loss, some users may wish to take incremental snapshots of the database and data directory, using the same method, while Connect is running. **These are not guaranteed to be consistent and you should not attempt to restore an RStudio Connect installation with them.** Notably:

- Depending on the timing between when both snapshots are taken, the database may have records that aren't reflected in your data directory snapshot, or vice versa.
- Because the networked filesystem may lock certain files that are currently in use, the snapshot process may finish much later than it began, leading to a snapshot that isn't consistent with any state. - Example: Imagine app X v0.1 was copied successfully, and app Y is locked on NFS. While the backup script waits for app Y to unlock, app X v0.2 is uploaded. Then app Z is created. The snapshot would reflect apps X v0.1, Y, and Z, despite that the only consistent states were "app X v0.1, Y" and "app X v0.2, Y, Z"
- Copying your data directory and snapshotting your database will create contention on both the networked filesystem and the database, which could cause service slowness or interruption for users. Consider warning users and/or performing incremental snapshots during times of low user activity.

Recognizing that you will not be able to restore the entire backup, it may be useful to only snapshot those parts of the data directory or database that you want to retain. For example, you could use `rsync` to find new bundles added since the last backup was performed, and copy them to a snapshot directory.

6.6 HA Details

6.6.1 Concurrent Scheduled Document Rendering

The `Applications.ScheduleConcurrency` configuration setting specifies the number of scheduled jobs that can run concurrently on a host. By default, two scheduled jobs can run simultaneously.

Hosts with substantial processing power can increase this setting. This can be helpful if your environment has many long-running reports.

A particular host can disable processing of scheduled jobs by setting `Applications.ScheduleConcurrency` to zero.

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
ScheduleConcurrency = 0
```

Note: No schedule job will execute if every host sets `Applications.ScheduleConcurrency` to zero.

The `Applications.ScheduleConcurrency` setting does not affect ad-hoc rendering requests, hosted APIs, or Shiny applications.

6.6.2 Concurrent Shiny Applications and Ad-Hoc Rendering

Each process associated with Shiny applications, hosted APIs, ad-hoc rendering requests, and bundle deployments runs on the server where the request was initiated. We depend on your load balancer to distribute these requests to an appropriate Connect node. The minimum and maximum process limits for Shiny applications are enforced per server. For example, if a Shiny application allows a maximum of 10 processes, a maximum of 10 process per server will be enforced. See A.21 for more information.

6.6.3 Polling

RStudio Connect nodes poll the data directory for new scheduled jobs:

- Every 5 seconds, and
- After every completed scheduled job.

6.6.4 Abandoned Processes

While processing a scheduled job, the RStudio Connect node periodically updates the job’s metadata in the database with a “heartbeat”. If the node goes offline and the “heartbeat” ceases, another node will eventually claim the abandoned job and run it again. Hence, if a server goes offline or the Connect process gets shut down while a scheduled report is running, it is possible that the scheduled job could run twice.

6.6.5 Abandoned Shiny Applications

A Shiny applications depends on a persistent connection to a single server. If the server associated with a particular Shiny application session goes down, the Shiny application will fail. However, simply refreshing the application should result in a new session on an available server, assuming your load balancer detects the failed node and points you to a working one.

Shiny applications that support client-side reconnects using the `session$allowReconnect(TRUE)` feature will automatically reconnect the Shiny application to a working node. See <https://shiny.rstudio.com/articles/reconnecting.html>

7 Running with a Proxy

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 or Apache HTTPD.

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 recommend adding 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`)

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

If your proxy secures traffic with SSL, but uses an insecure (HTTP) connection to Connect, Connect Dashboard users will see a warning about accessing Connect over an insecure connection. You can disable this warning with the `HTTP.NoWarning` configuration setting.

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[HTTP]
NoWarning = true
```

Users are able to upload an image to each application. By default the image is limited to 10MB in size. Note that some proxies may be configured to reject uploads of excessive size. Modify the image size limitation with the `Applications.MaxAppImageSize` configuration setting.

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
;; Use a 5MB limit
MaxAppImageSize = 5000000
```

7.1 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
```

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

```
sudo zypper install nginx
```

To enable an instance of Nginx running on the same server to act as a front-end proxy to RStudio 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;

        client_max_body_size 0; # Disables checking of client request body size

        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;
            proxy_buffering off; # Required for XHR-streaming
        }
    }
}
```

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

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

        client_max_body_size 0; # Disables checking of client request body size

        location /rsconnect/ {
            rewrite ^/rsconnect/(.*)$ /$1 break;
            proxy_set_header X-RSC-Request $scheme://$host:$server_port$request_uri;
            proxy_pass http://localhost:3939;
            proxy_redirect / /rsconnect/;
            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, SUSE 12, Ubuntu 16.04, and Ubuntu 18.04):

```
sudo systemctl restart nginx
```

On `upstart` systems (Ubuntu 14.04, Red Hat 6):

```
sudo restart nginx
```

7.2 Apache Configuration

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

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

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

```
Listen 80
```

```
<VirtualHost *:80>
```

```
    RewriteEngine on
```

```
    # store variable values with dummy rewrite rules
```

```
    RewriteRule . - [E=req_scheme:%{REQUEST_SCHEME}]
```

```
    RewriteRule . - [E=http_host:%{HTTP_HOST}]
```

```
    RewriteRule . - [E=req_uri:%{REQUEST_URI}]
```

```
    # set header with variables
```

```
    RequestHeader set X-RSC-Request "%{req_scheme}e://%{http_host}e%{req_uri}e"
```

```

RewriteCond %{HTTP:Upgrade} =websocket
RewriteRule /(.*) ws://172.17.0.1:3939/$1 [P,L]
RewriteCond %{HTTP:Upgrade} !=websocket
RewriteRule /(.*) http://172.17.0.1:3939/$1 [P,L]
ProxyPass / http://172.17.0.1:3939/
ProxyPassReverse / http://172.17.0.1:3939/
</VirtualHost>

```

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

```
Listen 80
```

```

<VirtualHost *:80>
  RewriteEngine on
  RedirectMatch ^/rsconnect$ /rsconnect/

  # store variable values with dummy rewrite rules
  RewriteRule . - [E=req_scheme:%{REQUEST_SCHEME}]
  RewriteRule . - [E=h_host:%{HTTP_HOST}]
  RewriteRule . - [E=req_uri:%{REQUEST_URI}]
  # set header with variables
  RequestHeader set X-RSC-Request "%{req_scheme}e://%{h_host}e%{req_uri}e"

  RewriteCond %{HTTP:Upgrade} =websocket
  RewriteRule /rsconnect/(.*) ws://172.17.0.1:3939/$1 [P,L]
  RewriteCond %{HTTP:Upgrade} !=websocket
  RewriteRule /rsconnect/(.*) http://172.17.0.1:3939/$1 [P,L]
  ProxyPass /rsconnect/ http://172.17.0.1:3939/
  ProxyPassReverse /rsconnect/ http://172.17.0.1:3939/
  Header edit Location ^/ /rsconnect/
</VirtualHost>

```

8 Security & Auditing

8.1 API Security

8.1.1 Preventing Brute Force & Dictionary Attacks

By default, RStudio Connect allows as many login attempts as it can handle from any source when using the PAM, LDAP, and Password authentication providers. Users will be able to log in directly by entering their user name and password.

Setting the `Authentication.ChallengeResponseEnabled` flag to `true` enables a CAPTCHA form in the login screen, and requires that CAPTCHA be solved in order to authenticate. Both visual and audio CAPTCHA challenges are provided for accessibility needs.

8.1.2 Hiding RStudio Connect Server Version and Build Information

By default RStudio Connect sets the `Server HTTP` header to the Connect version and build number, which looks something like `RStudio Connect v1.2.3-1234`. If you would like to hide this information you can configure `Server.ServerName` to be any string. If you would also like to hide the version and build information within the Connect dashboard you may configure the `Server.HideVersion` setting to be `true`.

8.2 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.

8.2.1 Web Sudo Mode

When a user performs a sensitive operation in a web browser (such as creating a new API key), they will be asked to reenter their login credentials. Once the user enters their password properly, their session will enter a privileged state internally referred to as “web sudo mode” which will allow them to perform these sensitive operations for a certain window of time without entering their password again. Note that this privileged mode is a notion purely internal to RStudio Connect and entirely unrelated to the server’s actual `sudo` or PAM settings.

This feature is not available on servers configured to use proxied authentication or Google OAuth2, as these providers don’t have a mechanism for reprompting the user for their password.

For all other authentication providers, the `WebSudoMode` and `WebSudoModeDuration` configuration options are available within the section pertaining to that provider. If `WebSudoMode` is set to `false`, then this protection is disabled; effectively, all authenticated users are always in privileged mode. `WebSudoModeDuration` controls the length of time for which a user stays in this privileged mode. From within the section pertaining to your authentication provider (such as `[Password]`, `[PAM]`, or `[LDAP]`), you could configure:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
; This example is using password authentication. Use PAM or LDAP as appropriate.
[Password]
WebSudoMode = true
WebSudoModeDuration = 10m
```

In this case, users would be prompted for their password before performing sensitive actions, then would be allowed to continue performing sensitive actions for up to 10 minutes without additional prompts. After that point, any sensitive actions would require entering their password again.

Regardless of the setting, web sudo mode will never affect calls made outside of a browser using a token or an API key.

8.2.2 Limiting Session Lifetime

By default RStudio Connect uses a session cookie that will keep a user logged in for 30 days. You may adjust this session lifetime by configuring the `Authentication.Lifetime` configuration option. It is a good security practice to request user credentials daily to ensure the user still has access to Connect. Use a value of 24 hours for `Authentication.Lifetime` to follow this recommendation.

NOTE: If a user’s session expires with an application left running in their browser that application’s process will be terminated.

You may further restrict a user’s session lifetime by configuring the setting `Authentication.Inactivity`. When this is set to a non-zero value that is less than `Authentication.Lifetime` the user’s session will expire after they have not been actively using the Connect dashboard for the configured amount of time. The security practice is to make the inactivity period as short as possible. Our current recommendation is setting it to 8 hours, which allows most users to go through their office hours without being asked for credentials a second time.

8.2.3 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.

8.2.4 Strong HTTPS

Even with HTTPS, you may wish even more constraints on transport settings. Two settings are available to constrain HTTPS even further:

NOTE: Some Windows builds of the `RCurl` package are not compatible with TLS 1.1 or TLS 1.2. Windows users should execute `options(rsconnect.http = 'curl')` to use the local `curl` binary, if installed, instead of the outdated `RCurl` package.

NOTE: Older versions of the RStudio IDE may use webviews that are not compatible with TLS 1.1 or TLS 1.2. In this case, attempting to pair with RStudio Connect without TLS 1.0 would open a blank screen instead of a login window. Open the blue link from the pairing window in the browser, or install a newer version of RStudio Connect as a workaround.

1. `HTTPS.MinimumTLS`: This setting sets the minimum TLS version. By default, the minimum TLS version supported by RStudio Connect is TLS 1.0. You can choose to set this configuration setting to "1.1" or "1.2"; before doing so, you should check the SSL Labs User Agent List to ensure your clients will be compatible with the version you select.
2. `HTTPS.ExcludedCiphers`: This setting removes a cipher from the list of ciphers available. This may be useful if your organization has a security policy that disallows certain ciphers to be used.

An example exclusion list could be:

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

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

8.2.5 Using a Secure Proxy

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

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

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

8.2.6 Content Sniffing

The `Server.ContentTypeSniffing` setting can be used to configure HTTP responses with the `X-Content-Type-Options` header. This header can protect your users from a certain class of malicious uploads.

When `Server.ContentTypeSniffing` is disabled (the default), the `X-Content-Type-Options` HTTP header is given the value of `nosniff`. This instructs browsers not to inspect the content in an attempt to identify its type.

When `Server.ContentTypeSniffing` is enabled, the `X-Content-Type-Options` HTTP header is not sent; browsers are free to analyze the content to detect its type.

8.2.7 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. The dashboard is any of the internal services or assets that are shipped with RStudio Connect. User content is anything uploaded by a user (reports, Shiny applications, Plumber APIs, 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 Connect 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.

8.2.8 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, by Plumber, 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:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
CustomHeader = "HeaderA: some value"
CustomHeader = "HeaderB: another value"
```

8.3 Audit Logs

RStudio Connect records information about changes to the system. These events can later be retrieved for auditing purposes.

This table contains the event types logged by the auditing system:

Event	Description
<code>add_user</code>	Create a user
<code>edit_user</code>	Change an existing user
<code>remove_user</code>	*Delete a 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. Content may need 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>edit_group</code>	*Change an existing group
<code>remove_group</code>	Delete a group
<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

Event	Description
<code>clear_app_viewer_acl</code>	Change from a specific list of viewers to “just me”
<code>add_api_key</code>	Add API key
<code>edit_api_key</code>	*Change an existing API key
<code>remove_api_key</code>	Remove 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
<code>add_tag</code>	Create a tag/category
<code>remove_tag</code>	Delete a tag/category
<code>update_tag</code>	Update a tag/category
<code>assign_tag_to_parent</code>	Associate a tag with some parent tag/category
<code>add_app_tag</code>	Associate a tag with content
<code>remove_app_tag</code>	Disassociate a tag with content
<code>updated_environment_variables</code>	Change to application environment variables

Note: The items marked with an asterisk can only be done via the `usermanager` CLI.

8.3.1 Audit Logs Command-Line Interface

See Appendix B for more information on using the `usermanager` CLI to dump audit logs.

8.4 Application Environment Variables

User-specified environment variables for applications are encrypted on-disk and in-memory. They are decrypted only when a process is about to be started.

9 Database

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

Customize the `Database.Provider` property with a database scheme appropriate for your organization. See Section A.7 for details

Here is a partial configuration which chooses to use SQLite

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = sqlite
```

9.1 SQLite

SQLite is the default database provider.

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = sqlite
```

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SQLite]
Dir = /mnt/connect/sqlite
```

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

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

9.1.1 SQLite Backups

RStudio Connect can be configured to periodically back up its database while running.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SQLite]
Backup = true
BackupFrequency = 24h
BackupRetentionLimit = 7
```

The above config will execute an online backup every 24 hours. RStudio Connect will retain up to 7 of those backups. If an eighth backup is created, the oldest of the previous backups will be deleted. This provides a grace period for an administrator implementing a data retention policy, such as a practice of copying backups to tape periodically.

Backups are stored in the same directory as the main database file: `/var/lib/rstudio-connect/db` by default. Backups are lexically sortable, because they are timestamped with the UNIX epoch time padded to 11 digits. For example, if the server's database file is `/var/lib/rstudio-connect/db/connect.db`, a backup of that database could be `/var/lib/rstudio-connect/db/connect.db.01508526538`.

The same process is applied to the instrumentation database file. It will be backed up based on the same settings as for the main database file.

Automatic SQLite backups are **not** a complete backup solution for RStudio Connect. You should also make regular backups of the `Server.DataDir` directory. This is especially important because the `Server.DataDir` directory is expected to be kept in sync with the database.

Restoring a SQLite backup is straightforward:

- Ensure that the backup is valid with the command `sqlite3 <backup file name> "PRAGMA integrity_check;"`
- Stop the RStudio Connect service
- Copy the current database file as well as its `.wal` file, if any. (If you wish to analyze them or send a copy to RStudio Support)
- Replace the current database file with the backup
- **If they exist**, delete any `.wal` and `.shm` files associated with the previous database. Failure to do this could lead to further downtime and possible data corruption.
- Start the RStudio Connect service

Note also that RStudio Connect has no way of restoring applications deployed or changes made since the last backup. Restoring the backup file will cause these changes to be lost permanently.

9.2 PostgreSQL

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

You must provide your own Postgres server which will likely be a separate box from your RStudio Connect server (but not required). We currently support any 9.x, 10.x or 11.x version greater than or equal to 9.4. Your Postgres server does not have to be dedicated to RStudio Connect, but it must have its own dedicated database.

Use your Postgres server instead of the default SQLite database by configuring `Database.Provider` as `postgres` and setting `Postgres.URL` to the fully qualified Postgres URL of your server.

The user credentials supplied in the `Database.Provider` URL must be able to create and alter database tables, in addition to read/write permissions in the database referenced as the path of the URL. A blank database with the given name **MUST** already exist.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/connect"
```

9.2.1 SSL

RStudio Connect assumes by default that SSL is enabled on the Postgres server. If SSL is not enabled on your Postgres database, add `?sslmode=disable` to the `Postgres.URL`.

Here is an example configuration using a Postgres database with SSL connections:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/connect"
```

Here is an example configuration using a Postgres database without SSL connections:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/connect?sslmode=disable"
```

9.2.2 Schemas

If you need to, you can tell Connect to restrict itself to keeping its tables within a specific schema. You control this by giving PostgreSQL a search path as part of the URL by adding `options=-csearch_path=<schema-name>` to the URL. If it's the only item you're adding, separate it from the rest of the URL with `?` (just like the `sslmode` item above). Otherwise, separate it from other items with `'&'`.

Here is an example configuration that specifies a schema to use, with SSL:

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/connect?options=-csearch_path=connect_schema"

```

Here is an example configuration that specifies a schema to use, with SSL disabled:

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[Database]
Provider = postgres

[Postgres]
URL = "postgres://username:password@db.seed.co/connect?sslmode=disable&options=-csearch_path=connect_schema"

```

RStudio Connect will refuse to start when given a schema that does not already exist. The schema must be owned by the connecting user or by a group that contains the connecting user.

9.3 Changing Database Provider

Connect includes a `migrate` command for migrating data from one database to another.

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

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

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

Note: Migration from PostgreSQL to SQLite is not supported at this time.

Note: Migration of instrumentation event data is not supported at this time.

If you are also migrating your RStudio Connect installation to a new server, see 4.8.

9.3.1 Database Migration Checklist

Use this checklist to guide your migration process:

1. Shut down Connect - 5.1
2. Back up your data - 4.7
3. Ensure that you have a `Postgres` configuration section - 9.2
4. Run the migration - B.2
5. Update the `Database.Provider` configuration setting to point to the new database - A.7
6. Restart Connect - 5.1

9.3.2 Configuration Requirements

When migrating data, the configuration file must contain valid configuration sections for both `SQLite` and `Postgres`. The migration utility will connect to the SQLite and PostgreSQL databases specified in the configuration.

9.4 Authentication Provider

The migration between authentication providers is detailed in section 10.1.

10 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 10.5.

When signing into RStudio Connect, a browser cookie is used to keep the user logged in. See section 8.2.2 for details about the default session lifetime and how to change it.

External authentication is available through the following integrations:

- LDAP and Active Directory (Section 10.6)
- SAML (Section 10.7)
- OAuth 2.0 using Google Apps accounts (Section 10.8)
- PAM (Section 10.9)
- Proxied Authentication (Section 10.10)

Customize the `Authentication.Provider` property with an authentication scheme appropriate for your organization.

Here is a partial configuration which chooses to use LDAP.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
Provider = ldap
```

10.1 Changing Authentication Provider

The decision to change which authentication provider is configured for RStudio Connect should not be entered into lightly. Doing so in an inappropriate way can lead to loss of access to a user's content, subscriptions and other assets (note that only the *access* to the assets is affected; the assets themselves will still exist).

There are use cases, however, where changing the authentication provider is necessary. One of the most common ones is when RStudio Connect has been using the default password provider for a proof of concept and needs to be switched to a corporate standard like LDAP once it has been decided to make use of RStudio Connect in a production environment.

If maintaining existing content is not important, the simplest approach will be to completely purge and reinstall RStudio Connect. See Section 5.5 for instructions.

10.1.1 Uniquely Identifying Users

Among the data that RStudio Connect maintains about a user is a *Unique ID*. Even though the `username` (the identifier used for logging in and displayed in the Connect dashboard) must also be unique, the Unique ID is the real marker for uniquely identifying a user. All authentication providers require some form of this. Some, like the default password provider, invent their own when users are created. External providers (proxy, LDAP, etc.) generally have their own notion of this. LDAP, for example, uses the *distinguished name* (or DN) by default as a Unique ID, even though the recommendation is using values that don't change over time, usually vendor-specific (such as `objectGUID` for AD or `entryUUID` for OpenLDAP).

The success of changing the authentication provider is completely dependent on how these Unique IDs are mapped from the old provider to the new one. There are two approaches to consider when changing the

provider, the distinguishing point being whether a second user record gets created the first time a user logs in with the new auth provider or, similarly, if there is a reason for using the RStudio Connect Server API to create new user records ahead of time.

Avoiding this user creation is the most straightforward choice, but will require the most knowledge about what the new provider will use for the user's Unique ID. If this is discoverable ahead of time, these steps may be followed:

1. Stop RStudio Connect.
2. Use the `list` command of the `usermanager` tool to discover the identifier for each user in RStudio Connect. The GUID is the recommended ID to use.
3. Using tools available from the source behind the new authentication provider, such as LDAP, determine the Unique ID that the provider will use for each user.
4. For each user, use the user's RStudio Connect identifier to assign the new Unique ID to the user using the `usermanager` tool. For example,

```
sudo /opt/rstudio-connect/bin/usermanager alter --user-guid <guid> \  
--new-unique-id <provider-unique-id>
```

5. Reconfigure RStudio Connect to use the new authentication provider.
6. Restart RStudio Connect.

These steps are generic and may need adjustment to fit with your situation. They can be followed for performing a bulk update of multiple users or for one user at a time. Since the `usermanager` tool can only be run while RStudio Connect is not, doing bulk updates is recommended to minimize downtime.

Note: The Unique ID is stored as it comes from the authentication provider for SAML, OAuth2 and for LDAP using Distinguished Name (DN). LDAP uses a base64-encoded value when not using DN and since these may be sent to RStudio Connect as binary values. Proxied authentication when relying on headers for the Unique ID instead of usernames also stores Unique ID as a base64-encoded value, even when the header value is plain text. Currently, Password, PAM and Proxied authentication relying on usernames use internal formats. If you are migrating from other providers to these three, this first approach is not recommended.

If determining the Unique ID that the new authentication provider will use for a user is problematic, there are two other alternatives.

First, once RStudio Connect has been reconfigured with the new authentication provider, use the RStudio Connect Server API to create new user records yourself. For most authentication providers, the Unique ID will be properly determined so you don't have to work out what it is. Then, you can transfer all the assets from the user's old record to the new one with the `usermanager` tool. Doing this has the advantage of making the user's assets immediately available upon their first login under the new authentication provider.

Or, you can require that the user log in to RStudio Connect under the new authentication provider. This will create a new user record to which you can transfer all the user's assets from the old record.

Note: If you have configured RStudio Connect so that new user records are *not* automatically created, you will need to use the API to create new user records. See the `RegisterOnFirstLogin` configuration option under any provider, such as `OAuth2.RegisterOnFirstLogin`

To use these approaches, these steps may be followed:

1. Stop RStudio Connect.
2. Reconfigure RStudio Connect to use the new authentication provider.
3. Restart RStudio Connect.
4. Either create the new user record via the RStudio Connect Server API or require the user to log in.

5. Stop RStudio Connect.
6. Use the `list` command of the `usermanager` tool to discover the identifier for the old record for the user in RStudio Connect. The GUID is the recommended ID to use.
7. Use the same command to discover the identifier for the new record for the user. If you used the API to create the user, the response to the create endpoint will contain the GUID.
8. Use the `transfer` command of the `usermanager` tool to transfer all user assets from the old user identifier to the new one. You may optionally delete the old RStudio Connect user data as part of the transfer. For example,

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-guid <old-guid> \  
--target-guid <new-guid> --memberships --permissions --api-keys --delete
```

9. Restart RStudio Connect.

Again, these steps are generic and may need adjustment to fit with your situation. Also, you may want to update more than one user at a time here to minimize downtime.

See Appendix B for more information on using the `usermanager` CLI to manage users.

10.1.2 Updating Groups

LDAP and Proxied authentication providers can manage their own groups automatically. When this option is used the groups created in RStudio Connect do not have owners.

When switching to another authentication provider that requires users to manage groups either manually or via the Connect Server API these groups without owners will not be fully effective. For example, group members cannot be manually added or removed if the groups was previously managed by the authentication provider.

RStudio Connect will issue a warning on startup if the ownership of the existing groups does not match the configuration of the current authentication provider. Additionally, the Connect dashboard will label these groups as managed by the provider.

The opposite is also true. Groups always have owners when created under Password, PAM, SAML, OAuth2 and Proxied authentication with automatic group provisioning disabled and these groups will not work with the LDAP provider or Proxied authentication with group provisioning enabled.

In the case above a warning will be issued since the Connect dashboard does not support manual management of groups managed by the authentication providers.

In these situations, either the group needs to be removed or the group ownership needs to be adjusted. This can be done with the `usermanager` CLI tool with the `alter` command using the `--group-owner` switch. Groups themselves can be removed via the Connect dashboard (if not managed by users), via the Connect Server API or via the `usermanager` CLI tool with the `delete` command. See B.1.

The goal is that all groups automatically provisioned should be assigned an owner if you intend to manage them in RStudio Connect. Conversely if you intend to have all groups managed by the authentication provider their owners should be removed. Also, you should remove any groups that do not make sense in the new configuration.

Note: In addition to group ownership switching authentication providers may also require adjustments to the Unique ID as done for users. This is true when switching from any other authentication provider to use LDAP when preserving the groups.

There are situations where a new authentication provider brings its own groups and the existing groups have no function other than their association to content. It is possible to leverage the `usermanager` CLI tool with

the `transfer` command to transfer all access to content for one group to a different group. Group members can also be transferred using this option but only for groups managed by users. See B.1 for more details.

10.1.3 Unsupported Authentication Changes

- Changing user UNIX accounts under PAM is **NOT SUPPORTED**.

Why? The user's Unix account name is considered the source of truth for user identity. This means that a clustered configuration does not need to ensure usernames have the same UIDs on every node.

So, for example, if a user had the Unix account `jim` but now has the account `james`, RStudio will see this as a new user. You would need to use the `usermanager transfer` command to move everything from `jim` to `james` in this case.

- Changing `X-Auth-Username` responses for a user under Proxied authentication is **NOT SUPPORTED** with one exception.

Why? The `X-Auth-Username` response is considered the source of truth for user identity. To be able to update usernames without creating new ones, an alternative and persistent identification must be provided. By using `ProxyAuth.UniqueIdHeader` it is possible to send a user identity to RStudio Connect which is separate from the username.

Note: Any existing users *MUST* be updated to have the new identity before `ProxyAuth.UniqueIdHeader` is changed or configured for the first time. Use the `usermanager` CLI tool to alter user identity.

10.2 Session Management

Sessions are stored on the server in RStudio Connect. Encrypted session cookies stored only on the client are deprecated, as they provide inferior security.

The server will periodically check the data store for expired cookies and remove them. This happens once per hour by default, but is configurable using the `Authentication.CookieSweepDuration` configuration setting. This does not affect the lifetime of web sessions, which is controlled by the `Authentication.Lifetime` configuration setting.

10.3 Username requirements

When using OAuth2 or Password authentication, users are required to choose a valid username when first logging in. Usernames must:

- be 3-64 characters in length,
- start with a letter, and
- contain only alphanumeric characters, underscores, and periods.

By default, the LDAP, PAM, SAML, and Proxied authentication providers require that a valid username is received from the provider. LDAP additionally requires a valid user email. If the provider cannot resolve any of these, an error will be thrown.

Note that there is a known issue: if `LDAP.UsernameAttribute` is blank for the given LDAP credentials, the user will not be able to authenticate.

The LDAP, PAM, SAML, and Proxied authentication providers use relaxed username requirements. These providers permit any username, except the following prohibited usernames:

- connect
- apps

- users
- groups
- setpassword
- user-completion
- confirm
- recent
- reports
- plots
- unpublished
- settings
- metrics
- tokens
- help
- login
- welcome
- register
- resetpassword
- content

10.4 User Attribute Editability

Depending on your authentication provider, certain attributes of a user are editable and certain ones are not. Furthermore, who is able to make edits to which user can change depending on configuration.

10.4.1 Password

RStudio Connect considers all user attributes editable. The setting `Authorization.UserInfoEditableBy` has a default value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. Configure `Authorization.UserInfoEditableBy` with `Admin` if profile editing should be restricted to administrators. It is recommended that if you disable `Password.SelfRegistration` with a value of `false`, that you also set `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means users created by the administrator cannot be changed by non-administrators.

10.4.2 OAuth2 (Google Apps with OpenID Connect)

For OAuth2, first name, last name, and email address are provided by the provider. Therefore, it is not editable by any user. If you need to change/update this information, update your OAuth2 provider and the changes will be synced into RStudio Connect on next user login.

Usernames are derived from the email address (without the domain) at the first login or if the user is added to content as a collaborator or a viewer before the first login in RStudio Connect. Duplicate or prohibited usernames have a sequence number suffix added.

The setting `Authorization.UserInfoEditableBy` controls whether or not your users will be permitted to manage their username and role. The default configuration has `Authorization.UserInfoEditableBy` with a value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. A value of `Admin` permits only administrators to edit username and role.

It is recommended that if you disable `OAuth2.RegisterOnFirstLogin` with a value of `false`, that you also configure `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means that users created by the administrator, or in some cases a publisher, cannot be changed by non-administrators.

10.4.3 LDAP

A username is the primary means of authentication and therefore it is not editable.

The user's first name, last name, and email address are considered editable *only* if not configured to be set by the provider. The setting `Authorization.UserInfoEditableBy` has a default value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. Configure `Authorization.UserInfoEditableBy` with `Admin` if profile editing should be restricted to administrators.

The user attributes configured to be set by the provider are updated in RStudio Connect during every new user login.

It is recommended that if you disable `LDAP.RegisterOnFirstLogin` with a value of `false`, that you also configure `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means that users created by the administrator, or in some cases a publisher, cannot be changed by non-administrators.

Please be aware that any changes to the value that uniquely identifies the user on LDAP will cause RStudio Connect to see it as a brand new user creating a new account and the individual will not be able to log into their old account.

The section 10.6 has more information on how to avoid such issues by using vendor-specific values to identify users.

Additionally, please refer to the appendix B for more information on using the `usermanager` CLI to repair users in the situation above.

10.4.4 SAML

The Unique ID of a user requires a mapping to either the Name ID of the SAML subject or to an Identity Provider (or IdP) assertion attribute and, as such, is not editable in RStudio Connect. Please be aware that any changes to this value that uniquely identifies the user in the SAML IdP will cause RStudio Connect to see it as a brand new user, thus creating a new account, and that the individual will not be able to log into their old account.

Which other user fields are editable is based on which fields are mapped from SAML assertion attributes. If a mapping exists for a field (username, first or last name, or email), it is considered managed by the SAML IdP. In that case, they will not be editable within RStudio Connect; changes to those fields should be made in the IdP; they will be correctly propagated to RStudio Connect the next time the user logs in.

For fields with no attribute mapping, the `Authorization.UserInfoEditableBy` will control who can edit them. That option has a default value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. Configure `Authorization.UserInfoEditableBy` with `Admin` if profile editing should be restricted to administrators.

It is recommended that if you disable `OAuth2.RegisterOnFirstLogin` with a value of `false`, that you also configure `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means that users created by the administrator, or in some cases a publisher, cannot be changed by non-administrators.

10.4.5 Proxied Authentication

The user's first name, last name, and email address are considered editable *only* if not provided by the proxy via configured headers. The setting `Authorization.UserInfoEditableBy` has a default value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. Configure `Authorization.UserInfoEditableBy` with `Admin` if profile editing should be restricted to administrators.

It is recommended that if you disable `ProxyAuth.RegisterOnFirstLogin` with a value of `false`, that you also configure `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means that users created via the API by the administrator cannot be changed by non-administrators.

A username provided by the authenticating proxy as the primary means of identification for a given user and therefore it is not editable. In this situation, changing the username given by the proxy will result in a new user being created in RStudio Connect the next time the user logs in.

Alternatively, the authenticating proxy can provide a user identification separate from the username. Use the setting `ProxyAuth.UniqueIdHeader` for that purpose. The proxy can supply a different username in this situation, as long as the user identification remains the same over time. Even in this scenario the username is still not editable since the proxy is responsible for always providing a username.

10.4.6 PAM

RStudio Connect is always supplied with a username corresponding to the UNIX account username and therefore it is not editable. Changing the username of the UNIX account will result in a new user being created in RStudio Connect the next time the user logs in.

The user's first name, last name, and email address are not provided by the hosting system and are considered editable. The setting `Authorization.UserInfoEditableBy` has a default value of `AdminAndSelf`, permitting users and administrators to manage these user profile attributes. Configure `Authorization.UserInfoEditableBy` with `Admin` if profile editing should be restricted to administrators.

It is recommended that if you disable `PAM.RegisterOnFirstLogin` with a value of `false`, that you also configure `Authorization.UserInfoEditableBy` to `Admin`. A value of `Admin` means that users created via the API by the administrator cannot be changed by non-administrators.

10.5 Password

Password authentication is the default authentication provider used by RStudio Connect. It uses local user accounts backed by the RStudio Connect database and is not integrated with any external, third-party authentication services.

Users will be able to create accounts when they first visit the system and will provide profile details, like username, name and email, at that time. An administrator will also be able to create new accounts. The first user that logs in automatically becomes an administrator.

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.

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
Provider = password
```

10.5.1 Groups

Groups are supported when using Password authentication. They can be managed manually in the Connect dashboard or via the Connect Server API.

10.5.2 Limitations to Remember

Using password authentication makes it easy to get RStudio Connect up and running quickly but is really only appropriate when using RStudio Connect in one of these situations:

- A demonstration or proof-of-concept in which RStudio Connect is being evaluated or explored.
- Training users to use RStudio Connect.
- Using RStudio Connect to do some small-scale testing.

- A small group of users without, or isolated from, a centralized IT system.

Though proper security policies are followed by the password authentication provider so far as it goes, it is important to remember that there are several important policies it does not support. For example, the provider does *not*:

- provide for password expiration,
- provide automatic user lockout on multiple login failures,
- enforce any particular rules for password strength,
- enforce any particular limits on password length or classes of characters.

If your security needs require more fully fledged authentication capabilities, you will need to use an alternate authentication provider such as PAM, OAuth2, LDAP, SAML or Proxied authentication.

10.6 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:

- A user is uniquely identified by the attribute defined in `LDAP.UniqueIdAttribute`. This defaults to "DN" (Distinguished Name) which offers a broad compatibility with different LDAP products but has some known limitations. For information on how to configure this setting for your specific LDAP server, see 10.6.3.8
- The same limitation applies to LDAP groups which use `LDAP.GroupUniqueIdAttribute` to define the attribute which uniquely identifies them, also with the default of "DN".
- The following user information is required: first name, last name, email address, and username. Different LDAP/AD attributes may be used for each of these fields.
- Changes to a user (e.g. their name, email address, or username) will not propagate to RStudio Connect until the next time the user logs in.
- When using single bind, the DN of a user must contain their username (i.e. must utilize the `LDAP.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 enabling `LDAP.Logging` in any LDAP configuration section.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "any"]
Logging = true
```

10.6.1 Defining an LDAP or AD section

Start an LDAP or Active Directory configuration with a LDAP section header like the following:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "European AD Server"]
...
```

An LDAP or Active Directory configuration section header is always of the form [LDAP "EFFECTIVE NAME"]. The “effective name” is a name that is meaningful to your organization ("European AD Server" in the example).

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

Make sure that each LDAP configuration section has a unique effective name. If multiple LDAP sections have the same name, they will be combined as described in Appendix A.

The LDAP section name is treated *case sensitively*.

Here is a sample configuration using two LDAP sections.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "European LDAP Server"]
...

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

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

10.6.2 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 `LDAP.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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample OpenLDAP Configuration"]
ServerAddress = 127.0.0.1:389
BindDN = "cn=admin,dc=example-openldap"
BindPassword = "XXXXXXXX"
UserSearchBaseDN = "ou=People,dc=example-openldap"
UniqueIdAttribute = "entryUUID"
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
entryUUID: c1afa3d6-d688-1037-99f0-d7682e3458da
```

More LDAP configuration scenarios can be found in Appendix D.

10.6.3 LDAP or AD Configuration Settings

10.6.3.1 ServerAddress

LDAP.`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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
ServerAddress = 127.0.0.1:389
ServerAddress = ldap.company.com:389
ServerAddress = ldaps.company.com:636
ServerAddress = private.internal.local:7554
```

10.6.3.2 TLS

LDAP.`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 LDAP.`StartTLS` is `true`.

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
TLS = true
TLS = false
```

10.6.3.3 StartTLS

LDAP.`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 “LDAPS” or “Opportunistic TLS”. The default value for this is `false`. This cannot be enabled if `TLS` is `true`.

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Debug]
Log = ldap
```

10.6.3.4 TLSCACertificate

LDAP.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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
TLSCACertificate= /etc/ssl/cert/ca.pem
```

10.6.3.5 ServerTLSInsecure

LDAP.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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
ServerTLSInsecure = true
ServerTLSInsecure = false
```

10.6.3.6 BindDN and BindPassword

LDAP.BindDN and LDAP.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.

Quote passwords and DN's that contain a literal double-quote (") or one of the ; or # comment characters.

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
; Example DN
BindDN = uid=john,ou=People,dc=company,dc=com
BindPassword = "johnpassword"
```

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
; Example UPN
```



```
BindDN = admin@company.com
BindPassword = "adminpassword"

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
; Example NT-style login
BindDN = COMPANY\\admin # we use double slashes (\\) to character escape the last slash
BindPassword = "adminpassword"
```

10.6.3.7 AnonymousBind

LDAP.`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 LDAP.`BindDN` and LDAP.`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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
AnonymousBind = true
```

10.6.3.8 UniqueIdAttribute

UniqueIdAttribute (optional, default "DN") is the vendor-specific object attribute present for all users which has a unique value that identifies the object persistently. Commonly used attributes include:

- objectGUID (Microsoft AD)
- entryUUID (OpenLDAP)
- orclGuid (Oracle OID)
- ibm-entryUUID (IBM RACF)
- GUID (Novell eDirectory)
- nsUniqueId (389 Directory Server)

Please refer to your LDAP vendor documentation for the correct value.

The default value of DN is stored in RStudio Connect in plain text. However, any other attribute will be stored encoded as base64 since some LDAP vendors (i.e. Microsoft) send binary values for user identities.

WARNING: any change to this setting automatically invalidate all existing users. RStudio Connect will try to detect this and it will not start if problems are found. Before starting up RStudio Connect with the new configuration, all users must be adjusted to use the updated attribute value via the `usermanager` CLI: B.1

Note: The use of the default value of "DN" is not recommended and existing installations should use a different value if possible. RStudio Connect will issue a warning on startup if this condition is detected. The Distinguished Name (DN) is not a persistent value in LDAP/AD and it may change in certain situations. RStudio Connect will no longer be able to refer to a user once their DN changes. You can repair an affected user by updating the DN via the `usermanager` CLI: B.1

10.6.3.9 UserSearchBaseDN

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

Examples

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserSearchBaseDN = dc=company,dc=com
UserSearchBaseDN = ou=People,dc=company,dc=com

```

10.6.3.10 UserObjectClass

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

UserObjectClass is used to define LDAP.UserFilterBase if that property is not explicitly configured. Either UserObjectClass or UserFilterBase must be defined. If both have values, only UserFilterBase will be used.

Examples

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserObjectClass = user
UserObjectClass = posixAccount

```

10.6.3.11 UserFilterBase

The LDAP.UserFilterBase attribute allows more flexible when searching for user objects in complicated LDAP hierarchies.

Either LDAP.UserObjectClass or UserFilterBase must be defined. If UserFilterBase is unset, it is given a default value of objectClass={UserObjectClass}.

If users are identified by two separate objectClass values, you might use the configuration:

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserFilterBase = &(objectClass=user)(objectClass=statistician)

```

You can disqualify an objectClass value with the configuration:

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserFilterBase = &(objectClass=user)(!(objectClass=computer))

```

10.6.3.12 UsernameAttribute

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

Examples

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UsernameAttribute = uid
UsernameAttribute = sAMAccountName

```

10.6.3.13 UserFirstNameAttribute

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

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserFirstNameAttribute = givenName
```

10.6.3.14 UserLastNameAttribute

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

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserLastNameAttribute = sn
```

10.6.3.15 UserEmailAttribute

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

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
UserEmailAttribute = mail
```

10.6.3.16 PermittedLoginGroup

LDAP.PermittedLoginGroup 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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
PermittedLoginGroup = cn=admins,ou=group,dc=company,dc=com
PermittedLoginGroup = cn=scientists,ou=group,dc=company,dc=com
```

10.6.3.17 GroupUniqueIdAttribute

LDAP.GroupUniqueIdAttribute (optional, default "DN") is the vendor-specific object attribute present for all groups which has a unique value that identifies the object persistently. Please refer to your LDAP vendor documentation for the correct value. In general, users and groups use the same value. See 10.6.3.8

WARNING: any change to this setting automatically invalidate all existing groups. RStudio Connect will try to detect this and it will not start if problems are found. Before starting up RStudio Connect with the new configuration, all groups must be adjusted to use the updated attribute value via the `usermanager` CLI: B.1

Note: The use of the default value of "DN" is not recommended and existing installations should use a different value if possible. RStudio Connect will issue a warning on startup if this condition is detected. The Distinguished Name (DN) is not a persistent value in LDAP/AD and it may change in certain situations. RStudio Connect will no longer be able to refer to a group once their DN changes. You can repair an affected groups by updating the DN via the `usermanager` CLI: B.1

10.6.3.18 GroupObjectClass

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

GroupObjectClass is used to define LDAP.GroupFilterBase if that property is not explicitly configured. Either GroupObjectClass or GroupFilterBase may be defined. If both have values, only GroupFilterBase will be used.

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
GroupObjectClass = group
GroupObjectClass = posixGroup
```

10.6.3.19 GroupFilterBase

The LDAP.GroupFilterBase attribute allows more flexible when searching for group objects in complicated LDAP hierarchies.

Either LDAP.GroupObjectClass or GroupFilterBase may be defined. If GroupFilterBase is unset, it is given a default value of objectClass={GroupObjectClass}.

If groups are identified by two separate objectClass values, you might use the configuration:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
GroupFilterBase = &(objectClass=group)(objectClass=club)
```

You can disqualify an objectClass value with the configuration:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
GroupFilterBase = &(objectClass=group)(!(objectClass=flock))
```

10.6.3.20 GroupNameAttribute

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

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
GroupNameAttribute = cn
GroupNameAttribute = sAMAccountName
```

10.6.3.21 GroupSearchBaseDN

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

Examples

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "Sample LDAP Configuration"]
GroupSearchBaseDN = dc=company,dc=com
GroupSearchBaseDN = ou=Groups,dc=company,dc=com
```

10.7 SAML (Beta)

SAML integration is a Beta feature of RStudio Connect. Beta features are supported and unlikely to face breaking changes in a future release. Any issues found in the feature will be addressed during the regular release schedule; they will not result in immediate patches or hotfixes. We encourage customers to try these features and welcome any feedback, but recommend the feature not be used in production until it is in general availability.

RStudio Connect can integrate with the SAML Identity Provider (or IdP) of your company's choice to perform user authentication. In the SAML world, RStudio Connect fulfills the role of service provider (or SP). Note that user searching in RStudio Connect will search local users only as SAML provides no such user search capability.

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
Provider = saml
```

Appendix A.14 contains information about each SAML configuration option.

Note: Though it is not typically required, the `Server.Address` configuration option *is* required when using SAML as the authentication provider.

SAML support in RStudio Connect has the following constraints:

- A user is uniquely identified in Connect by the attribute defined in `SAML.UniqueIDAttribute`. This defaults to "NameID" which is a standard identifier for the authenticated Subject which offers a broad compatibility with different SAML Identity Providers. If your IdP is set up to use the *transient* format for Name IDs, then this value must be specified to something other than NameID. > **Note:** The value that identifies the user must be immutable during the user's > entire lifetime. Otherwise Connect will not be able to identify the user and > will create a new one. For this reason, usernames and email addresses should be > avoided as UniqueIDs unless it is the only option, since those can change over time.
- A user's username, first name, last name, and email address may be mapped from any SAML assertion attribute. If a mapping is not configured, the corresponding field will be editable in RStudio Connect, subject to the `Authorization.UserInfoEditableBy` configuration option.
- Changes to fields mapped to SAML assertion attributes for a user (e.g. their name, email address, or username) will not propagate to RStudio Connect until the next time that user logs in.
- It is highly recommended that RStudio Connect be configured to listen on a port secured with TLS to make sure SAML login assertions are posted to the RStudio Connect server in a secure manner. See Section A.4 for a list of the options available when configuring TLS.

When attempting to troubleshoot a problem relating to SAML, you can enable more verbose logging by enabling `SAML.Logging` in the SAML section.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
Logging = true
```

10.7.1 Defining a SAML Section

Start a SAML configuration with a SAML section header like the following:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
...
```

10.7.2 Identity Provider Configuration

There are two ways to configure the information RStudio Connect needs about your SAML Identity Provider (or IdP), either by the IdP's metadata or by using specific configuration options.

Note: RStudio Connect requires that the SAML IdP be configured to sign all the messages it sends. As such, the metadata for the IdP must contain a signing certificate. Or, if you are using the specific options, the `SAML.IdPSigningCertificate` option is required.

10.7.2.1 Using an IdP's Metadata

To configure an IdP by its metadata file, use the `SAML.IdPMetaData` option. This may be set to either an HTTPS (or HTTP) URL from which the metadata may be retrieved or to the path to a local file that contains the metadata.

For example,

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
IdPMetaData = https://idp.example.com/saml/metadata
```

tells RStudio Connect to pull the metadata for the IdP from the given URL, whereas

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
IdPMetaData = /a/local/dir/idp.xml
```

tells RStudio Connect to read the metadata from the named, locally accessible file.

The retrieved, or read, metadata is cached while RStudio Connect is running. If the metadata expires during that time, the first login attempt will cause the metadata to be retrieved or read again. If the metadata is *still* expired, the login (and all subsequent logins) will fail. Some IdPs do this to account for the issuing of new signing and/or encryption certificates so this is generally not a problem for URL-based metadata. If you are using file-based metadata, the error will persist until you update the metadata file.

10.7.2.2 Providing Specifics Yourself

If there is some reason that an IdP's metadata cannot be made available to the RStudio Connect server, you can provide the information yourself by using the `SAML.IdPEntityID`, `SAML.IdPSingleSignOnServiceURL`, `SAML.IdPSigningCertificate` and `SAML.IdPEncryptionCertificate` configuration options.

The `SAML.IdPEntityID` option is used to identify the IdP. Since all identity providers must have a public name, this entity ID, it is used to verify that SAML messages are being received from the correct system.

The `SAML.IdPSingleSignOnServiceURL` option is used to tell RStudio Connect where to send (via POST or redirect) authentication requests for a user attempting to log in. The IdP system must be listening for this URL for SP-initiated logins. In the case where RStudio Connect is configured for IdP-initiated logins only, this must be the IdP's user login page.

The `SAML.IdPSigningCertificate` specifies the public side of the certificate that the IdP will use to sign the messages it sends to RStudio Connect. It may be the path to a PEM file or the raw, base64 encoded certificate. When not configuring IdP information from metadata, this option **must** be specified.

The `SAML.IdPEncryptionCertificate` specifies the public side of the certificate that the IdP will use to encrypt the messages it sends to RStudio Connect. It may be the path to a PEM file or the raw, base64 encoded certificate.

If you use these options, you *must not* specify the `SAML.IdPMetaData` option, as that will override these. A warning is logged during RStudio Connect startup if this condition exists.

10.7.3 SAML Binding

Use the `SAML.IdPSingleSignOnPostBinding` option to note whether you want RStudio Connect to use POST binding (`true`) or redirect binding (`false`, the default) when RStudio Connect is to initiate a single sign-on operation with your IdP. Redirect is the most common choice.

If the IdP is configured by its metadata, RStudio Connect will verify that the IdP makes an SSO location available using the binding requested here. If the IdP metadata indicates multiple bindings are supported, the `SAML.IdPSingleSignOnPostBinding` option controls which one will be used.

10.7.4 Name ID Format

When a SAML IdP identifies an authenticated user to a service provider, that user is (typically) identified by what SAML calls a “NameID”. There are four possible formats this name ID may be in. The IdP may support multiple values for this. Use the `SAML.NameIDFormat` option to tell RStudio Connect which of them your IdP will use. This may be set to `transient`, `persistent`, `emailAddress` or `unspecified`; the last will be used if this option is not provided.

If the IdP is configured by its metadata, RStudio Connect will verify that the value specified by the `SAML.NameIDFormat` option is supported by the IdP.

10.7.5 Single Sign-On Initiation

The SAML specification allows for a single sign-on flow to be initiated by either the IdP or SP. Use the `SAML.SSOInitiated` option to control this. There are three possible values.

- **IdP** – This indicates that single sign-on may be initiated at the IdP only. In this mode, an end user must visit the sign-on page of their IdP and select RStudio Connect as the service provider to use.

Note: In this mode, you will not be able to use the IdP’s metadata to configure the IdP’s information; you will need to use the separate configuration options. When the IdP is used to initiate authentication, this causes a problem when publishing from the RStudio IDE. Since the IDE can only know about RStudio Connect, RStudio Connect must respond by returning a “blind” redirect to to the IdP’s sign-on page. This must be configured in the `SAML.IdPSingleSignOnServiceURL` configuration option.

- **IdPAndSP** – This indicates that single sign-on may be initiated by either the IdP or by RStudio Connect. In this case, note that the RStudio IDE will always use the service provider initiated flow. This is the default.
- **SP** – This indicates that single sign-on may be initiated only by RStudio Connect. In this mode, an IdP-initiated login will fail, since there would be no matching authentication request that RStudio Connect has sent out.

10.7.6 User Attribute Mapping

When a user is authenticated by an IdP, certain data about the user may be included in the authentication response. These are most commonly called *attributes*, *assertion attributes* or *claims*. So that RStudio Connect may manage its user information correctly, it needs to understand these.

Note: SAML attributes in the authentication response are case sensitive.

Use the `SAML.UsernameAttribute`, `SAML.FirstNameAttribute`, `SAML.LastNameAttribute` and `SAML.EmailAttribute` configuration options to set the names of the assertions to look for in obtaining the related values. For each of these that are *not* specified, RStudio Connect will assume management of the fields, making them editable in the Connect dashboard.

The ways `SAML.UniqueIDAttribute` is configured to identify users in RStudio Connect are detailed above.

Note: The value obtained from `SAML.UniqueIDAttribute` must match exactly the field `unique_id` used for creating new users via the Connect Server API. Once the user is created, the RStudio Connect identity (`guid`) should be used for subsequent API operations with that user.

10.7.7 Metadata Configuration Example

Here is a complete SAML configuration as an example. It tells RStudio Connect to communicate to an IdP by retrieving its metadata from the given URL. Single sign-ons can be initiated by either party. The IdP will use the persistent NameID format. The Unique ID will be the name ID from the IdP.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
IdPMetaData = https://idp.example.com/saml2/metadata
NameIDFormat = persistent
UsernameAttribute = "Username"
FirstnameAttribute = "FirstName"
LastnameAttribute = "LastName"
EmailAttribute = "Email"
```

10.7.8 Individual Settings Configuration Example

Here is a complete SAML configuration as an example but now using individual settings when a metadata is not available or cannot be used. Single sign-ons can be initiated by the IdP only, Connect will redirect to a login page on the IdP if the authentication is attempted first in Connect (which will happen with the RStudio IDE), the signing certificate (shortened for brevity) is a raw base64 value, while the encryption certificate is a path to a PEM file and because `transient` is being used the `UniqueIDAttribute` is something other than the `NameID`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
SSOInitiated = IdP
IdPEntityID = https://idp.example.com/saml2/metadata
IdPSingleSignOnServiceURL = https://idp.example.com/login
IdPSigningCertificate = MIADHIUfdhfkhdshkf454sdfe5e64ff45FF5456456d4d65 ... fd378==
IdPEncryptionCertificate = /path/to/pem.cer
NameIDFormat = transient
UniqueIDAttribute = "GUID"
UsernameAttribute = "Username"
FirstnameAttribute = "FirstName"
LastnameAttribute = "LastName"
EmailAttribute = "Email"
```

10.7.9 Groups

Groups are supported when using SAML authentication. They can be managed manually in the Connect dashboard or via the Connect Server API.

Note: These groups are local to RStudio Connect and have no relation to any organization groups present in the IdP.

10.7.10 Service Provider Configuration for the IdP

Though the number of available IdPs is large and each one has its own configuration process, there are some common things to know.

Almost all IdPs will require you to define and configure the service providers that they should interact with, including RStudio Connect. They may use different terms, but there are three pieces of information you will need to provide to the IdP.

- **Service Provider EntityID** – This is a string that uniquely identifies RStudio Connect as a service provider to the IdP. This might be labeled in other ways such as, *Entity ID*, *Name*, *Audience*, *Audience Restriction*, etc. This is typically a URL to the metadata provided by the SP but doesn't need to be.

You should use a value of `https://<hostname>[:port]/__login__/saml` for RStudio Connect, replacing `hostname` and `port` as appropriate. These should match what is specified in the `Server.Address` configuration option.

It is not required that the IdP be able to access this if it's a URL.

- **Assertion Consumption Service** – This is a URL in the service provider where the IdP will post authentication responses. This is important when initiating an authentication at the IdP. This may also be labeled as *Single sign on URL*, *SSO URL*, *ACS URL*, etc.

You should use a value of `https://<hostname>[:port]/__login__/saml/acs` for RStudio Connect, replacing `hostname` and `port` as appropriate. These should match what is specified in the `Server.Address` configuration option.

- **Attribute Mapping** – This is the mapping of fields known within the IdP to the attribute names that will be used for them in authentication responses. They are usually specified in name/value pairs though they may include a name format specification.

Each name is the string by which its value will be known in authentication responses and must match the values specified in the `SAML.UsernameAttribute`, `SAML.FirstNameAttribute`, `SAML.LastNameAttribute` and `SAML.EmailAttribute` configuration options. The matching of these names is case sensitive.

10.8 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`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
Provider = oauth2
```

Appendix A.12 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.

When attempting to troubleshoot a problem relating to OAuth2, you can enable more verbose logging with `OAuth2.Logging`:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[OAuth2]
Logging = true
```

10.8.1 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. In the left navigation window, click on “Credentials”, then goto the “OAuth consent screen” tab, fill in the information requested and click “Save”.
3. Once again, click “Credentials” in the left navigation window. Then click the dropdown button “New credentials”, then “OAuth client ID”.
4. 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`).
5. 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[OAuth2]
ClientId = <CLIENT ID>
ClientSecret = <CLIENT SECRET>
```

With `OAuth2.ClientId` and either `OAuth2.ClientSecret` or `OAuth2.ClientSecretFile` configured, you can use your Google Apps account to sign into RStudio Connect!

10.8.2 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[OAuth2]
AllowedEmails = jdoe@company.com asmith@subsidiary.com
```

It is important to understand how the `OAuth2.AllowedDomains` and `OAuth2.AllowedEmails` properties interact.

If only `OAuth2.AllowedDomains` is configured, *only* email addresses with a listed domain will be permitted access.

If only `OAuth2.AllowedEmails` is configured, *only* listed email addresses will be permitted access.

When both `OAuth2.AllowedDomains` and `OAuth2.AllowedEmails` are specified, email addresses given in `OAuth2.AllowedEmails` are permitted access *in addition* to email addresses with a domain listed in `OAuth2.AllowedDomains`.

10.8.3 Usernames

Google OAuth2 has no notion of usernames but certain features in RStudio Connect require a username. On the user's first login or when the user is associated with some content, we derive the initial username from the user's email address without the domain. The user may be able to modify it after the first login if allowed. An administrator is always able to modify OAuth2 usernames.

The process of generating an initial username is limited to what is considered a valid username in RStudio Connect. For example, certain characters such as dashes (-) are replaced with underscores (_), shorter names are also padded with underscores. Resulting usernames longer than 64 characters are truncated as necessary and prohibited or duplicate usernames get an incremental suffix.

10.8.4 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.

10.8.5 Groups

Groups are supported when using OAuth2 authentication. They can be managed manually in the Dashboard or via the Connect Server API.

Note: These groups are local to RStudio Connect and have no relation to any organization groups present in Google Apps, for example.

10.9 PAM

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
Provider = pam
```

See Section 13.7 for information about using PAM sessions when launching processes associated with deployed content.

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`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[PAM]
Service = rstudio-connect
```

Note that there are three types of PAM service that can be configured in the PAM configuration section (See Section 13.7 for more information):

- `PAM.Service` - The PAM service used for authenticating users when logging in.
- `PAM.SessionService` - When `PAM.UseSession` is enabled, the PAM service used for running processes either as the default user or as an arbitrary user. Should not require a password.
- `PAM.AuthenticatedSessionService` - The PAM service used for running processes as the currently logged-in user with the user's password. Requires `PAM.UseSession`, `PAM.ForwardPassword`, and `Applications.RunAsCurrentUser` to be enabled. Useful for Kerberos configurations.

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

When attempting to troubleshoot a problem relating to PAM, you can enable more verbose logging with `PAM.Logging`:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[PAM]
Logging = true
```

10.9.1 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.

10.9.2 Red Hat/CentOS/SUSE

Red Hat/CentOS/SUSE systems may 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.

10.9.3 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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.

10.9.4 Groups

Groups are supported when using PAM authentication. They can be managed manually in the Dashboard or via the Connect Server API.

Note: These groups are local to RStudio Connect and have no relation with Unix/Linux groups present in the host machine where PAM is configured.

10.10 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.

10.10.1 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 custom HTTP headers added before the request is proxied

through to RStudio Connect. The HTTP headers will contain the username and other identifying information of that visitor. RStudio Connect will take the values from the HTTP headers and treat the traffic as the user specified in the headers.

We have no means of validating that this HTTP header was added by your authentication server and not by the user directly. 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 HTTP header identifying the user 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 will return a generic authentication failure if duplicate authentication headers are provided.

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.

10.10.2 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. The minimum required version of `rsconnect` for proxied authentication is 0.8.7.

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

The `X-Auth-Token` header identifies deployment requests. When the `X-Auth-Token` header is present, the request should be allowed to pass through your proxy without attempting to authenticate the user. Further, the following three headers are used in the deployment process and should not be edited or removed by your proxy.

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

10.10.3 Using the RStudio Connect Server API

The `Authorization` header should allow traffic to pass through the proxy without attempting to authenticate the user when:

- The RStudio Connect Server API is being used.
- RStudio Connect is serving hosted APIs (i.e. Plumber, TensorFlow, etc.)

Note: Look for a value containing `Key` followed by the API key for a more specific match of the `Authorization` header used by RStudio Connect.

10.10.4 Configuring Proxied Authentication

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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 2.3.1 for more information about `Server.Address`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
UsernameHeader = X-Auth-Username
```

RStudio Connect can also be configured to receive more complete user profile information from the proxy via additional headers.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
UsernameHeader = X-Auth-Username
FirstNameHeader = X-Auth-First-Name
LastNameHeader = X-Auth-Last-Name
EmailHeader = X-Auth-Email
```

Note: The fields configured to be sent by the proxy are not editable via the Connect dashboard or Connect Server API. Any updated values sent by the proxy will be reflected in the user's profile the next time that user logs in. A non-editable field (one configured to be sent in the proxy headers) can be forcefully updated before a login via the `usermanager` CLI tool. See B.1.

Additionally, a user role in RStudio Connect can be informed via a proxy header. Use the setting `ProxyAuth.RoleHeader` to inform which header will contain the role.

Note: The role header accepts the values `viewer`, `publisher` or `administrator`. Invalid values are defaulted to value of `Authorization.DefaultUserRole` and an error is reported in the logs.

10.10.4.1 Supporting Changes to the Username

By default, the username is the only means of user identity required by RStudio Connect, and every new username gets a new account in RStudio Connect. If you wish to be able to update usernames over time while maintaining a single account for each user, a different form of unique user identity must be provided via the proxy headers.

The setting `ProxyAuth.UniqueIdHeader` determines which header holds the user identity. It's important that the user identity sent over this header by the proxy be: - Unique: It is extremely important that different users do not have the same identity. - Persistent: Ideally this identity must not be attached to the username, email or names, since those may change over time. Values based on UUID/GUID are a great options if your authentication system provides them.

IMPORTANT: Existing customers willing to enable the `ProxyAuth.UniqueIdHeader` *MUST* first update all users to the new identity before configuring the setting. The `usermanager` CLI tool can help in this migration. See B.1

Note: Configuring `ProxyAuth.UniqueIdHeader` also enables the RStudio Connect Server API to accept the field `unique_id` for creating new users. This value *must* match the value sent by the proxy via headers. Internally in Connect the UniqueID is stored encoded as a base64 value. Once the user is created, the RStudio Connect identity (`guid`) should be used for subsequent API operations with that user.

10.10.4.2 Supporting Groups

The setting `ProxyAuth.GroupsHeader` defines the name of the header containing the names of the groups the authenticating user is member of.

The list of group names sent by the proxy is compared with the memberships the user currently has in RStudio Connect and then any missing memberships are added to existing groups and any memberships that no longer apply are removed.

By default it is expected that multiple groups names will come as separate occurrences of the same header, for example:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
GroupsHeader = X-Auth-Groups
```

HTTP Request:

```
X-Auth-Groups: Engineering
X-Auth-Groups: Managers
...
```

However, it is also possible to send group names on a single header. In this case, the group names in the header value must be sent as a separated list. Use the setting `ProxyAuth.GroupsHeaderSeparator` to enable this behavior and set which character should be used to split the group names.

For example:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
GroupsHeader = X-Auth-Groups
GroupsHeaderSeparator = "|"
```

HTTP Request:

```
X-Auth-Groups: Engineering|Managers|...
```

Note: Certain characters cannot be used as separators since they are prohibited in HTTP headers. The RFC 7230 define those.

10.10.4.3 Automatic Group Provisioning

In addition to delegating group membership management to the authentication engine behind the proxy, RStudio Connect can also delegate the management of groups themselves to the proxy. By using `ProxyAuth.GroupsAutoProvision` RStudio Connect will automatically create and optionally remove groups based on the list of group names received from the proxy.

With this option enabled groups are provisioned in RStudio Connect when the first member is added and remain there indefinitely, even after the last member has been removed, so that any access to content is preserved for a future member of those groups.

Note: RStudio Connect can be configured to automatically decommission (remove) groups without members with the additional setting `ProxyAuth.GroupsAutoRemoval`. It also removes the group from all content it has been previously associated with.

In this mode, it is assumed that all group management will be done externally to RStudio Connect and therefore the user interface for group management will be hidden in the Connect dashboard.

10.10.4.4 Switching Between Manual and Automatic Group Provisioning

Groups created by automatic provisioning do not have owners while any groups created manually or via the Connect Server API are always associated with a user.

RStudio Connect will issue a warning on startup if the ownership of the existing groups does not match the provider configuration.

In these situations, either the group needs to be removed or group ownership needs to be adjusted. This can be done with the `usermanager` CLI tool with the `alter` command using the `--new-owner` or `--drop-owner` switches.

The goal is that all groups automatically provisioned should be assigned an owner if you intend to manage them in RStudio Connect. Conversely if you intend to have all groups managed by the authentication provider their owners should be removed. Also, you should remove any groups that do not make sense in the new configuration and this can be done via the Connect dashboard, the Connect Server API or the `usermanager`. See B.1

Note: Additional to group ownership, groups created manually or via the Connect Server API may not match the names of the groups sent by the proxy and therefore those cannot be removed if both auto provisioning and auto removal are enabled. The Connect Server API can also be used to remove these extra groups.

10.10.5 Troubleshooting Proxied Authentication

When attempting to troubleshoot a problem relating to proxied authentication, you can enable more verbose logging with `ProxyAuth.Logging`:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
Logging = true
```

All the contents of headers and the names of the cookies received with the authentication request will be logged. Cookies and headers used by RStudio Connect will be logged as `<internal>`. External cookie values will be logged as `<redacted>`. If an external cookie comes without a value it will be logged as `<empty>`.

- “Rejected insecure proxied authentication attempt” appears in the server logs, users cannot log in
 1. Ensure that the proxy is configured to delete the configured username, names, email and unique identity headers from incoming requests (`X-Auth-Username` by default).
 2. Ensure that users are connecting to RStudio Connect by its proxy, and not directly to the server. As noted above, your network should be configured to make non-proxied connections to RStudio Connect impossible.
- Attempts to deploy to RStudio Connect from the IDE fail because users are redirected to a Single Sign-On page.
 1. Ensure that the proxy is configured to pass through all requests that set the `X-Auth-Token` header.
 2. Ensure that the user has the `rsconnect` package with at least version 0.8.7 installed. If not, and if the package isn’t available from CRAN, it may be installed from the R console using `devtools::install_github('rstudio/rsconnect')`

See 10.10.2

- Attempts to use the RStudio Connect server API fail because users are redirected to a Single Sign-On page.
 1. Ensure that the proxy is configured to pass through all requests that set the `Authentication` header with a `Key` value.
 2. Ensure that the proxy is not blocking other headers that may be required by hosted APIs (i.e. Plumber, TensorFlow, etc.)

See 10.10.3

10.10.6 Example Proxy Configuration

The following examples demonstrate how to intercept and re-route unauthenticated requests to RStudio Connect, which is served via reverse-proxy at `/rsconnect`. The examples assume an authentication server is running at `/auth` which, upon successful authentication, will set a HTTP-only cookie named `verified-user` containing the username of the authenticated user and redirect the user back to RStudio Connect via the URL specified in the `url` query parameter. A similar logic is applied to remaining examples that send additional user information.

IMPORTANT: These examples are intended to aid with an initial configuration of RStudio Connect behind an authenticating proxy. They are not recommended for production environments and have severe functional limitations (i.e. `'` and `;` are not accepted in values because the examples use plain text cookies) and offer no security at all. Customers are urged to adhere to the best security practices and recommended settings for their proxy.

All examples below are based on how to run a proxy in front of RStudio Connect. See 7.

10.10.6.1 Basic configuration example

The bare minimum to allow a user coming from an authentication proxy to enter RStudio Connect. It assumes default values for the existing settings for proxied authentication. See 10.10.4.

These examples send traffic through the proxy if the `verified-user` cookie is present. The examples also allow traffic for deployments from the RStudio IDE or API use if the respective headers are present. See 10.10.3 and 10.10.2.

Nginx:

```
map $http_upgrade $connection_upgrade {
    default upgrade;
    ""      close;
}

map $http_cookie $auth_username {
    default "";
    "~*verified-user=(?<username>[^;]+)" "$username";
}

map "$http_x_auth_token" $has_token {
    default "";
    "~.+" "token";
}

map "$http_authorization" $has_key {
    default "";
    "~Key .+" "key";
}

map "$auth_username:$has_token:$has_key" $requires_auth {
    default 0;
    "::*" 1;
}
```

```

}

server {
    listen 80 default_server;

    location /rsconnect/ {

        if ($requires_auth = 1) {
            return 307 $scheme://$host:$server_port/auth/login?url=$request_uri;
        }

        rewrite          ^/rsconnect/(.*)$ /$1 break;
        proxy_pass       http://rstudio-connect-host:3939;
        proxy_redirect   / /rsconnect/;
        proxy_pass_request_headers on;
        proxy_connect_timeout 5;
        proxy_http_version 1.1;
        proxy_buffering  off; # Required for XHR-streaming

        proxy_set_header    Upgrade $http_upgrade;
        proxy_set_header    Connection $connection_upgrade;
        proxy_set_header    X-RSC-Request $scheme://$host:$server_port$request_uri;
        proxy_set_header    X-Auth-Username $auth_username;
    }
}

```

Apache:

```

RewriteEngine on
ProxyPreserveHost on

```

```

<VirtualHost *:80>
    SetEnvIf Cookie "verified-user=(\[^\;]+\)" AuthUsername=$1
    <If "%{REQUEST_URI} !~ m#^/auth# && -z env('AuthUsername') && -z %{HTTP:X-Auth-Token} && %{HTTP:Autho
        Redirect 307 "%{REQUEST_SCHEME}://%{HTTP_HOST}/auth/login?url=%{REQUEST_URI}"
    </If>

    <Location /rsconnect/>
        ProxyPass http://rstudio-connect-host:3939/ connectiontimeout=5
        ProxyPassReverse /

        # store variable values with dummy rewrite rules
        RewriteRule . - [E=req_scheme:%{REQUEST_SCHEME}]
        RewriteRule . - [E=http_host:%{HTTP_HOST}]
        RewriteRule . - [E=req_uri:%{REQUEST_URI}]
        # set header with variables
        RequestHeader set X-RSC-Request "%{req_scheme}e://%{http_host}e%{req_uri}e"

        <If "-n env('AuthUsername')">
            RequestHeader set X-Auth-Username "%{AuthUsername}e"
        </If>
        <Else>
            RequestHeader unset X-Auth-Username
        </Else>

```

```

RewriteCond %{HTTP:Upgrade} =websocket
RewriteRule /rsconnect/(.*) ws://rstudio-connect-host:3939/$1 [P,L]
RewriteCond %{HTTP:Upgrade} !=websocket
RewriteRule /rsconnect/(.*) http://rstudio-connect-host:3939/$1 [P,L]
</Location>
</VirtualHost>

```

10.10.6.2 Configuration example with user profile support

These configurations extend the basic examples and add support for headers carrying user profile information, such as names and email. The names for the different headers must be also present in your RStudio Connect configuration. See 10.10.4.

Nginx:

```

map $http_upgrade $connection_upgrade {
    default upgrade;
    ""      close;
}

map $http_cookie $auth_username {
    default "";
    "~*verified-user=(?<username>[^;]+)" "$username";
}

map $http_cookie $auth_firstname {
    default "";
    "~*proxyauth-as-first=(?<first>[^;]+)" "$first";
}

map $http_cookie $auth_lastname {
    default "";
    "~*proxyauth-as-last=(?<last>[^;]+)" "$last";
}

map $http_cookie $auth_email {
    default "";
    "~*proxyauth-as-email=(?<email>[^;]+)" "$email";
}

map "$http_x_auth_token" $has_token {
    default "";
    "~.+" "token";
}

map "$http_authorization" $has_key {
    default "";
    "~Key .+" "key";
}

map "$auth_username:$has_token:$has_key" $requires_auth {
    default 0;
    "::$" 1;
}

```

```

server {
    listen 80 default_server;

    location /rsconnect/ {

        if ($requires_auth = 1) {
            return 307 $scheme://$host:$server_port/auth/login?url=$request_uri;
        }

        rewrite                ^/rsconnect/(.*)$ /$1 break;
        proxy_pass              http://rstudio-connect-host:3939;
        proxy_redirect          / /rsconnect/;
        proxy_pass_request_headers on;
        proxy_connect_timeout   5;
        proxy_http_version      1.1;
        proxy_buffering         off; # Required for XHR-streaming

        proxy_set_header        Upgrade $http_upgrade;
        proxy_set_header        Connection $connection_upgrade;
        proxy_set_header        X-RSC-Request $scheme://$host:$server_port$request_uri;
        proxy_set_header        X-Auth-Username $auth_username;
        proxy_set_header        X-Auth-Firstname $auth_firstname;
        proxy_set_header        X-Auth-Lastname $auth_lastname;
        proxy_set_header        X-Auth-Email $auth_email;
    }
}

```

Apache:

```

RewriteEngine on
ProxyPreserveHost on

```

```
<VirtualHost *:80>
```

```

    SetEnvIf Cookie "verified-user=(\[^\;]+\)" AuthUsername=$1
    <If "%{REQUEST_URI} !~ m#^/auth# && -z env('AuthUsername') && -z %{HTTP:X-Auth-Token} && %{HTTP:Autho
        Redirect 307 "%{REQUEST_SCHEME}://%{HTTP_HOST}/auth/login?url=%{REQUEST_URI}"
    </If>

```

```

    SetEnvIf Cookie "proxyauth-as-first=(\[^\;]+\)" AuthFirstName=$1
    SetEnvIf Cookie "proxyauth-as-last=(\[^\;]+\)" AuthLastName=$1
    SetEnvIf Cookie "proxyauth-as-email=(\[^\;]+\)" AuthEmail=$1

```

```
<Location /rsconnect/>
```

```

    ProxyPass http://rstudio-connect-host:3939/ connectiontimeout=5
    ProxyPassReverse /

```

```
    # store variable values with dummy rewrite rules
```

```
    RewriteRule . - [E=req_scheme:%{REQUEST_SCHEME}]
```

```
    RewriteRule . - [E=http_host:%{HTTP_HOST}]
```

```
    RewriteRule . - [E=req_uri:%{REQUEST_URI}]
```

```
    # set header with variables
```

```
    RequestHeader set X-RSC-Request "%{req_scheme}e://%{http_host}e%{req_uri}e"
```

```
    <If "-n env('AuthUsername')">
```

```
        RequestHeader set X-Auth-Username "%{AuthUsername}e"
```

```

</If>
<Else>
    RequestHeader unset X-Auth-Username
</Else>
<If "-n env('AuthFirstName')">
    RequestHeader set X-Auth-Firstname "%{AuthFirstName}e"
</If>
<Else>
    RequestHeader unset X-Auth-Firstname
</Else>
<If "-n env('AuthLastName')">
    RequestHeader set X-Auth-Lastname "%{AuthLastName}e"
</If>
<Else>
    RequestHeader unset X-Auth-Lastname
</Else>
<If "-n env('AuthEmail')">
    RequestHeader set X-Auth-Email "%{AuthEmail}e"
</If>
<Else>
    RequestHeader unset X-Auth-Email
</Else>

    RewriteCond %{HTTP:Upgrade} =websocket
    RewriteRule /rsconnect/(.*) ws://rstudio-connect-host:3939/$1 [P,L]
    RewriteCond %{HTTP:Upgrade} !=websocket
    RewriteRule /rsconnect/(.*) http://rstudio-connect-host:3939/$1 [P,L]
</Location>
</VirtualHost>

```

10.10.6.3 Configuration example with user profile and unique identity support

In these examples, username is considered part of the user profile. An additional header is used to communicate user identity. See 10.10.4.

Ngix:

```

map $http_upgrade $connection_upgrade {
    default upgrade;
    ""      close;
}

map $http_cookie $auth_uniqueid {
    default "";
    "~*proxyauth-as-id=(?<uniqueid>[^;]+)" "$uniqueid";
}

map $http_cookie $auth_username {
    default "";
    "~*verified-user=(?<username>[^;]+)" "$username";
}

map $http_cookie $auth_firstname {
    default "";
    "~*proxyauth-as-first=(?<first>[^;]+)" "$first";
}

```

```

}

map $http_cookie $auth_lastname {
    default "";
    "~*proxyauth-as-last=(?<last>[^;]+)" "$last";
}

map $http_cookie $auth_email {
    default "";
    "~*proxyauth-as-email=(?<email>[^;]+)" "$email";
}

map "$http_x_auth_token" $has_token {
    default "";
    "~.+ "token";
}

map "$http_authorization" $has_key {
    default "";
    "~Key .+" "key";
}

map "$auth_uniqueid:$has_token:$has_key" $requires_auth {
    default 0;
    ":@" 1;
}

server {
    listen 80 default_server;

    location /rsconnect/ {

        if ($requires_auth = 1) {
            return 307 $scheme://$host:$server_port/auth/login?url=$request_uri;
        }

        rewrite          ^/rsconnect/(.*)$ /$1 break;
        proxy_pass        http://rstudio-connect-host:3939;
        proxy_redirect    / /rsconnect/;
        proxy_pass_request_headers on;
        proxy_connect_timeout 5;
        proxy_http_version 1.1;
        proxy_buffering   off; # Required for XHR-streaming

        proxy_set_header   Upgrade $http_upgrade;
        proxy_set_header   Connection $connection_upgrade;
        proxy_set_header   X-RSC-Request $scheme://$host:$server_port$request_uri;
        proxy_set_header   X-Auth-Uniqueid $auth_uniqueid;
        proxy_set_header   X-Auth-Username $auth_username;
        proxy_set_header   X-Auth-Firstname $auth_firstname;
        proxy_set_header   X-Auth-Lastname $auth_lastname;
        proxy_set_header   X-Auth-Email $auth_email;
    }
}

```

Apache:

RewriteEngine on

ProxyPreserveHost on

```
<VirtualHost *:80>
  SetEnvIf Cookie "verified-user=(\[^\;]+)" AuthUniqueId=$1
  <If "%{REQUEST_URI} !~ m#^/auth# && -z env('AuthUniqueId') && -z %{HTTP:X-Auth-Token} && %{HTTP:Autho
    Redirect 307 "%{REQUEST_SCHEME}://%{HTTP_HOST}/auth/login?url=%{REQUEST_URI}"
  </If>

  SetEnvIf Cookie "proxyauth-as-user=(\[^\;]+)" AuthUsername=$1
  SetEnvIf Cookie "proxyauth-as-first=(\[^\;]+)" AuthFirstName=$1
  SetEnvIf Cookie "proxyauth-as-last=(\[^\;]+)" AuthLastName=$1
  SetEnvIf Cookie "proxyauth-as-email=(\[^\;]+)" AuthEmail=$1

  <Location /rsconnect/>
    ProxyPass http://rstudio-connect-host:3939/ connectiontimeout=5
    ProxyPassReverse /

    # store variable values with dummy rewrite rules
    RewriteRule . - [E=req_scheme:%{REQUEST_SCHEME}]
    RewriteRule . - [E=http_host:%{HTTP_HOST}]
    RewriteRule . - [E=req_uri:%{REQUEST_URI}]
    # set header with variables
    RequestHeader set X-RSC-Request "%{req_scheme}e://%{http_host}e%{req_uri}e"

    <If "-n env('AuthUniqueId')">
      RequestHeader set X-Auth-Uniqueid "%{AuthUniqueId}e"
    </If>
    <Else>
      RequestHeader unset X-Auth-Uniqueid
    </Else>
    <If "-n env('AuthUsername')">
      RequestHeader set X-Auth-Username "%{AuthUsername}e"
    </If>
    <Else>
      RequestHeader unset X-Auth-Username
    </Else>
    <If "-n env('AuthFirstName')">
      RequestHeader set X-Auth-Firstname "%{AuthFirstName}e"
    </If>
    <Else>
      RequestHeader unset X-Auth-Firstname
    </Else>
    <If "-n env('AuthLastName')">
      RequestHeader set X-Auth-Lastname "%{AuthLastName}e"
    </If>
    <Else>
      RequestHeader unset X-Auth-Lastname
    </Else>
    <If "-n env('AuthEmail')">
      RequestHeader set X-Auth-Email "%{AuthEmail}e"
    </If>
```



```

<Else>
  RequestHeader unset X-Auth-Email
</Else>

RewriteCond %{HTTP:Upgrade} =websocket
RewriteRule /rsconnect/(.*) ws://rstudio-connect-host:3939/$1 [P,L]
RewriteCond %{HTTP:Upgrade} !=websocket
RewriteRule /rsconnect/(.*) http://rstudio-connect-host:3939/$1 [P,L]
</Location>
</VirtualHost>

```

11 User Management

11.1 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 `Authorization.DefaultUserRole` property specifies the role for new accounts and defaults to `viewer`. `Authorization.DefaultUserRole` may be either `viewer` or `publisher`; new accounts are not permitted to automatically have the `administrator` role. For all authentication providers, the first user is always created as `administrator`.

Note: There are no restrictions regarding roles for the users created via the Connect Server API.

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. The specific capabilities of an administrator are documented here.

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 “Anyone”.

Note: Any logged-in user can list all the other existing users. To limit who can list users use `Authorization.UsersListingMinRole` to specify the minimum role allowed to list users.

For example: `Authorization.UsersListingMinRole` with a value of `administrator` will allow only administrators to list all users while publishers and viewers would only see themselves.

```

; /etc/rstudio-connect/rstudio-connect.gcfg
[Authorization]
UsersListingMinRole = administrator

```

This setting does not affect the ability publishers have to add other users to content as collaborators or viewers.

11.2 User Provisioning

How a user gets created in RStudio Connect depends on the capabilities offered by the configured authentication provider. See 10 for more details.

11.2.1 Password

Note: Sending emails is a prerequisite for password authentication. See 2.3.4

The users can be created by an administrator or they can register themselves through the RStudio Connect dashboard. The Connect Server API can also be used to create users ahead of the first login.

The users created by an administrator or via Connect Server API without a password will receive an email confirmation which should be used to configure a password.

Existing user accounts can reset their passwords through the Connect login page.

11.2.1.1 Self Registration

When using password authentication, users can self-register by clicking the “Sign Up” button on the login page. Self-registered accounts will be created with the role specified in the `Authorization.DefaultUserRole` property (see 11.1).

If you wish to disable self-registration, please give the configuration setting `Password.SelfRegistration` a value of `false`.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Password]
SelfRegistration = false
```

When self-registration is disabled, the first account (the administrator) is still created using self-registration. All other accounts must be created by an administrator.

11.2.2 OAuth2 (Google Apps with OpenID Connect)

OAuth2 users are created in RStudio Connect upon the first successful login attempt. `OAuth2.RegisterOnFirstLogin` can be used to disable this behavior.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[OAuth2]
RegisterOnFirstLogin = false
```

If registration through login is disabled, users must either be a member of an allowed group or added as viewers or collaborators to content in order to be able to login.

The authentication will happen entirely in the Google Auth engine which once completed will return the remote user information to RStudio Connect.

Note: User search support is conditional to your Google Auth infrastructure. See 10.8

New users can also be created ahead of their first login by searching based on their email or name and then associating the remote user found in the results with some content as either a viewer or collaborator.

Note: New users can only be selected as collaborators when `Authorization.DefaultUserRole` is configured as `publisher`.

The Connect Server API can be used to create OAuth2 users. This option enables associating users with content ahead of their first login attempt.

Note: Prior to adding the user to RStudio Connect via the Connect Server API, a separate API call for retrieving the user profile from the provider is needed.

11.2.3 LDAP

LDAP users are created in RStudio Connect upon the first successful login attempt. `LDAP.RegisterOnFirstLogin` can be used to disable this behavior.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "corporate LDAP"]
RegisterOnFirstLogin = false
```

If registration through login is disabled, users must be added as viewers or collaborators to content in order to be able to login.

Note: To restrict which users can login, see the setting `LDAP.PermittedLoginGroup`.

RStudio Connect will forward the entered LDAP credentials to the LDAP or Active Directory server which once authenticated will return the remote users information.

Note: User search support for LDAP requires bind credentials or anonymous bind. See 10.6

New users can also be created ahead of their first login by searching based on their email or name and then associating the remote user found in the results with some content as either a viewer or collaborator.

Note: New users can only be selected as collaborators when `Authorization.DefaultUserRole` is configured as `publisher`.

The Connect Server API can be used to create OAuth2 users. This option enables associating users with content ahead of their first login attempt.

Note: Prior to adding the user to RStudio Connect via the Connect Server API, a separate API call for retrieving the user profile from the provider is needed.

11.2.4 SAML (Beta)

SAML users are created in RStudio Connect upon the first successful login attempt. `SAML.RegisterOnFirstLogin` can be used to disable this behavior.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[SAML]
RegisterOnFirstLogin = false
```

User fields which are mapped from SAML assertion attributes are updated (if they change) on all subsequent successful logins.

The actual authentication will happen entirely in the SAML Identity Provider that RStudio Connect is configured to talk with. See section 10.7 for details on how to do that.

Users are uniquely identified by their *Unique ID* in the IdP system. If RStudio Connect is presented with a Unique ID it has never seen, it will create a new user.

Note: New users can only be selected as collaborators when `Authorization.DefaultUserRole` is configured as `publisher`.

The Connect Server API can be used to create SAML users. For this to work correctly, the user's information must be obtained from the IdP *with the correct Unique ID* and then provided to the RStudio Connect Server API. This enables associating users with content ahead of their first login attempt.

Note: RStudio Connect has no facility for retrieving user data from a SAML IdP; some other means, provided by the IdP, must be used.

11.2.5 Proxied Authentication

Proxied authentication will create users in RStudio Connect upon the first successful login attempt. `ProxyAuth.RegisterOnFirstLogin` can be used to disable this behavior.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[ProxyAuth]
RegisterOnFirstLogin = false
```

If registration through login is disabled, users must first be created via the Connect API in order to be able to login.

The authentication happens entirely in the Proxy placed in front of RStudio Connect which expects to receive from the former a HTTP header containing the username of the authenticated remote user account.

The Connect Server API can be used to create Proxy users. This option enables associating users with content ahead of their first login attempt.

Note: Users created via the API must match exactly the identifier expected to be received through the authentication proxy. By default, this identifier is the username, but it's possible to employ a separate unique identifier. See 10.10.

11.2.6 PAM

PAM authentication will create users in RStudio Connect upon the first successful login attempt. `PAM.RegisterOnFirstLogin` can be used to disable this behavior.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[PAM]
RegisterOnFirstLogin = false
```

If registration through login is disabled, users must first be created via the Connect API in order to be able to login.

PAM authentication relies on local Unix accounts. Users must login to RStudio Connect with their Unix credentials which is authenticated by the Linux system hosting Connect.

The Connect Server API can be used to create PAM users. This option enables associating users with content ahead of their first login attempt.

Note: The username is the user's unique identification in Connect. Users created via the API must match exactly the username for the respective Unix user.

11.3 User Permissions

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

11.3.1 All Content

Anonymous Visitors Anonymous users can access content listed for **Anyone**. 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 **Anyone**, **All users - login required**, 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.

11.3.2 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 **Anyone**, all public versions will be accessible to **Anyone**.

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

11.3.3 Shiny Applications & Plumber APIs

Collaborators “Collaborators” can change the runtime settings for Shiny applications and Plumber APIs.

11.4 Administrator Capabilities

Administrative users on RStudio Connect are empowered to inspect and manage various settings on the server. Regardless of their level of privilege on some piece of content (viewer, collaborator, or neither), administrators can manage collaborators and viewers on content, manage the runtime settings for Shiny applications and Plumber APIs, and adjust the schedules for R Markdown documents. Additionally, only administrators can modify the Vanity Path and RunAs settings for content through the Connect dashboard; they can do so even on content that they don’t have the ability to view the content.

Administrators do not have implicit rights to view content or download the source bundles. If an administrator visits a report without viewership privileges to the report, they will see an error message rather than the report’s content. Despite being unable to see the contents of the report, administrators can still manage the settings for all content. Because an administrator has the ability to manage the collaborators and viewers of others’ content on the system, they can choose to add themselves as a viewer or collaborator on the report to gain access. Administrative overrides of permissions on content require that the administrator take an explicit action which is captured in the audit log.

11.5 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.

11.6 Username Requirements

Connect's username requirements vary depending upon the authentication provider. Please see 10.3 for more information on username requirements.

11.7 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.

11.8 Changing Authentication Provider

There are circumstances where the type of authentication provider must be changed. See section 10.1 for details on how to do this.

11.9 Removing Accounts

Removing accounts from RStudio Connect is considered a last resort option.

Users are kept around so the content and groups they might own, as well as the historical and audit information associated with user is not left without a reference.

For these reasons locking accounts is still the preferred option.

If the decision is to remove the account, use the `usermanager` CLI tool. The removal process requires that the user does not own anything in RStudio Connect. It's possible to transfer the user's assets to another user ahead of the account removal. This transfer is also possible with the `usermanager`.

The historical information will still refer to the removed user account. See section 18 for details about what information is being tracked.

Once the user account is removed, the only place where the user information can still be found is in the audit logs. See section 8.3 for details.

11.10 Command-Line Interface

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

Additionally, you can transfer ownership of content and groups between users as well as memberships to groups, permissions to content, content subscriptions and, under certain circumstances, even API keys. It's also possible to remove users that are not owners of any content or groups in RStudio Connect.

The `usermanager` can adjust the Unique IDs of users in case this identification has been modified in the configured authentication provider or after switching between providers.

See Appendix B for more information on using the `usermanager` CLI to manage users.

12 Group Management

Groups can be used to associate multiple users to content as viewers or collaborators.

An administrator in RStudio Connect can use the dashboard, specifically the “People” tab to create groups and manage their members. Group support is available for all authentication providers and enabled by default.

Note: Groups can also be managed via the Connect Server API.

Group support can be disabled with `Authorization.UserGroups`.

Note: Disabling setting is not effective if groups are still present. RStudio Connect will issue a warning on startup and ignore this setting. In order to use this setting all groups must be removed first.

12.1 Remote Group Management

Some authentication providers can manage their own groups without manual intervention. For example, groups are always managed by the provider for LDAP and they can be managed by the provider optionally using proxied authentication. Under these scenarios, there's no group section under the “People” tab in the Connect dashboard.

12.2 Group Ownership

When a group is created manually in the Connect dashboard, the user that created the group is set as the group's owner. There are two exceptions to this:

1. If you are using the LDAP authentication provider, groups in RStudio Connect will have no owners since they are controlled by the related LDAP system.
2. If you are using the Proxied authentication provider *and* you have enabled automatic group provisioning, the groups created during the login process will have no owners.

Note: The rules above also apply to groups created via Connect Server API.

12.2.1 Publisher Ownership

In older releases of RStudio Connect publishers were allowed to create groups. To do so without proper consent of an administrator has made it more difficult to appropriately manage access to content. As such, by default publishers are no longer allowed to create groups.

The following should be noted about this change.

- Any groups created, and therefore owned, by publishers in an older release of RStudio Connect will still be owned by the same user.
- Publishers will still be able to add members to or remove members from the groups they own.
- Publishers will still be able to delete groups that they own.
- Publishers will not be able to create any new groups going forward.
- Publishers will still be able to remove themselves from groups they don't own. This also applies to viewers.

If there is a reason that publishers *should* be allowed to create groups (i.e., to restore the legacy behavior), set the `Authorization.PublishersCanOwnGroups` configuration option to `true`.

Administrators can take control over existing groups with the help of the `usermanager` CLI tool which allows assigning a new owner to a group with the `alter` command.

12.3 Proxied Authentication Group Membership Management

RStudio Connect can be configured to automatically assign users to existing groups according to the list of group names sent by the proxy. For every login attempt the list of group names received will be compared with the current memberships the user has, adding the user as a member of newly listed groups and removing the user from groups no longer listed by the proxy.

To enable group assignment use the setting `ProxyAuth.GroupsHeader`.

Note: The list of groups sent by proxy will override any memberships defined manually or via the Connect Server API. However, these operations should still be used to keep the group memberships in sync with the authentication engine behind the proxy between login attempts.

12.3.1 Proxied Authentication Provisioned Groups

In addition to memberships, RStudio Connect can also be configured to automatically provision (create) groups according to the list of group names sent by the proxy. This can be enabled by using `ProxyAuth.GroupsAutoProvision`.

This means that groups not yet present will be created in RStudio Connect for a user when the user logs in.

Note: In this mode groups are left in RStudio Connect if they have no members in order to preserve their association to content. If you wish to remove the empty groups use the setting `ProxyAuth.GroupsAutoRemoval`.

In this mode groups must be managed directly through the authentication engine behind the proxy. There's no Connect dashboard support to manage these groups.

Note: An administrator can still use the Connect Server API to manage groups between user login attempts. Care should be taken when removing groups via API. That action will also remove all association a group may have with some content.

12.4 LDAP Groups

RStudio Connect needs to be configured to automatically recognize LDAP groups. See 10.6.

LDAP groups must be managed directly through LDAP or Active Directory. There's no Connect dashboard support to manager these groups.

For LDAP, the group information is stored in RStudio Connect when an LDAP group is associated with some content.

Note: An administrator can still use the Connect Server API to manage groups between user login attempts. Care should be taken when removing groups via API. That action will also remove all association a group may have with some content.

12.5 OAuth2 Group Members

New remote OAuth2 user information is stored in RStudio Connect when a OAuth2 user is associated with a group. This allows a user to be present in RStudio Connect ahead of a first login. This is similar to the association of users with some content.

12.6 Command-Line Interface

Connect includes a `usermanager` command for some basic group management tasks. This utility helps you list groups and modify their attributes. This can be helpful in the event that no one can access a Connect administrative user account.

The tool can also be used to adjust group name or ownership and it can transfer content permissions and members between groups and even remove groups entirely from RStudio Connect.

The `usermanager` can adjust the Unique IDs of groups in case this identification has been modified in the configured authentication provider or after switching between providers.

See Appendix B for more information on using the `usermanager` CLI to manage users.

13 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.
- Running APIs using Plumber
- Running TensorFlow Model APIs

The location of R defaults to whatever is in the path. Customize the `Server.RVersion` setting to use a specific R installation. See Chapter 15 for details.

13.1 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 where content processes are run.

RStudio Connect runs its processes as an unprivileged user; both a system default and content-specific overrides are supported. See Section 13.5 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 `SQLite.Dir` directory, which can optionally be placed outside the data directory.
- Configuration directories, including `/etc/rstudio-connect`.
- The `Server.TempDir` directory contains a `connect-workspaces` sub-directory with per-process temporary 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, Plumber, and R Markdown).
- The document rendering destination directory (only for R Markdown).
- A per-process temporary directory specified in the `TMPDIR` environment variable of the process. This temporary directory is created under `Server.TempDir/connect-workspaces`.

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`.

Running R applications, like Shiny apps and Plumber APIs, have write access to the directory containing the unpackaged R code. This application directory is the working directory when launching an application. Data written here will be visible to all processes associated with that 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. Applications that write to local files could experience problems when different processes attempt to write to a single file.

For example, two different processes writing to the same file may see output incorrectly interleaved or even overwritten.

We **do not** recommend using the file system for data persistence.

R Markdown documents have write access to the rendering destination directory and to a directory containing the unpackaged R code. When RStudio Connect is rendering a document, it first makes a copy of the unpackaged R code into a new, temporary directory so that simultaneous rendering processes are isolated from each other and cannot corrupt each other's output files. This temporary source directory is the working directory when calling `rmarkdown::render`. The destination directory is passed as the `output_dir`. 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.

13.2 Temporary Directory

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

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

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

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

You can learn more about the `noexec` option [here](#).

13.3 Shiny Applications & Plumber APIs

Most of the processes started by RStudio Connect are batch-oriented tasks. A process is invoked, does a narrow set of work, and then exits. Shiny applications and Plumber APIs are different and may see a process handle many requests for many users over their lifetimes. Both Shiny Applications and Plumber APIs are live applications that react to user requests on-demand.

RStudio Connect launches an process tied to a live application when the first request arrives for that application. That process will continue to service requests until it becomes idle and eventually terminated. If there is sufficient traffic against that 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 applications are launched and eventually reaped. The default values are appropriate for most applications but occasionally need customization in specialized environments. Section A.21 explains each of the options.

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

13.4 TensorFlow Model APIs

TensorFlow Model APIs have a similar lifecycle to Plumber APIs. They are live processes that handle user requests on demand. TensorFlow Model API processes do not run R. However, the same per-process and global scheduler settings may still apply to TensorFlow Model API processes.

TensorFlow Model APIs may also be run within a supervisor script if one is provided. The API server requires access to the app's content directory and to shared object files and their dependencies.

13.5 User Account for Processes

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

Administrators can customize the `RunAs` user on a content-specific level. This means that different applications and 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, Plumber API, or R Markdown report *must* be a member of the `rstudio-connect` group.

Installation of R packages always happens as the `Application.RunAs` user. An 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 13.1 for more information about process sandboxing.

13.6 Current user execution

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

See Section 10.9 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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 13.5 for more information about `RunAs` customization.

Content execution settings are not altered when `Applications.RunAsCurrentUser` is enabled. The `Applications.RunAsCurrentUser` setting *permits* current-user execution but by itself does not change how processes are launched. Each Shiny application or Shiny document 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 applications or R Markdown reports *must* be a member of the `rstudio-connect` group.

13.7 PAM sessions

Note: Please see the special instructions at the bottom of this section for running RStudio Connect on Ubuntu 14.04 (Trusty Tahr).

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

See Section 10.9 for information about using PAM for user authentication.

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[PAM]
SessionService = rstudio-connect-session
```

The `PAM.SessionService` must contain the PAM directive that enables authentication with root privileges. Otherwise, processes will not run and will return error code 70.

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

Ubuntu 14.04 (Trusty Tahr) uses `upstart` as init by default, but also uses `systemd-logind` to clean up processes from closed user sessions. There is a known issue where `PAM.UseSession` causes this specific host configuration to rapidly terminate processes, returning `error code 129`.

If you enable `PAM.UseSession`, you also need to edit the `upstart` configuration file at `/etc/init/rstudio-connect.conf`, replacing the line beginning `exec /opt/rstudio-connect/bin/connect` with the following:

```
exec su -s /bin/sh -c 'exec "$0" "$@"' root -- /opt/rstudio-connect/bin/connect \
    --config=/etc/rstudio-connect/rstudio-connect.gcfg >> /var/log/rstudio-connect.log 2>&1
```

After altering `rstudio-connect.conf`, trigger an `upstart` configuration reload and then restart RStudio Connect.

```
sudo initctl reload-configuration
sudo stop rstudio-connect
sudo start rstudio-connect
```

Changing the `rstudio-connect.conf` in this way has considerable side effects because it is the equivalent of opening a `su` session for root and leaving it open for the life cycle of the RStudio Connect daemon.

If this solution is unacceptable, alternative solutions may include:

- Upgrading the host to Ubuntu 16 or later
- Updating `systemd-logind` to be newer than v204
- Altering the init provider to use `systemd` instead of `upstart`
- Disabling `systemd-logind` on the host

13.7.1 PAM Credential Caching (Kerberos)

Note: RStudio Connect’s PAM cache is **encrypted** and **is not stored on disk**. The credentials **must expire** after a certain period of time.

RStudio Connect can be configured to securely cache a user’s PAM credentials when they log in to RStudio Connect. This enables RStudio Connect to let users run R processes as their current UNIX account when the PAM profile requires a user’s credentials, such as when using Kerberos.

The following config settings are required for credential caching to be enabled:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
RunAsCurrentUser = true

[PAM]
UseSession = true ; Enable PAM sessions
ForwardPassword = true ; Forward the current user's password into the PAM session
PasswordLifetime = 12h ; Cache passwords for 12 hours after login
AuthenticatedSessionService = YOUR_PAM_SERVICE_HERE ; PAM service that accepts credentials
```

Replace 12h with the amount of time you would like credentials to be cached. Credential lifetime is counted from the moment the user logs into RStudio Connect. It is not tied to the user’s web session, except that logging in again will restart the timer for that user’s credentials.

The `PAM.AuthenticatedSessionService` setting is similar to `PAM.SessionService`, except that it should accept user credentials and validate them. For example, a PAM service that uses the host’s Kerberos configuration to expose functionality could be:

```
auth        required      pam_krb5.so
account     [default=bad success=ok user_unknown=ignore] pam_krb5.so
password    sufficient    pam_krb5.so use_authtok
session     requisite      pam_krb5.so
```

13.8 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
<code>DataDir/apps/XX/YY/</code>	<code>MountDir/app/ (non-renders)</code>

Physical path	Sandbox path
DataDir/reports/v2/XX/YY/temp.render.TT	MountDir/app/ (renders)
DataDir/reports/v2/XX/VV/RR	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 Plumber API
/var/lib/rstudio-connect/apps/38/10/plumber.R
=> /opt/rstudio-connect/mnt/app/plumber.R

# A source R Markdown document
/var/lib/rstudio-connect/reports/v2/8/12/temp.render.639085504/index.Rmd
=> /opt/rstudio-connect/mnt/app/index.Rmd

# An HTML document rendered from that R Markdown document
/var/lib/rstudio-connect/reports/v2/8/2/17/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
```

13.9 Program Supervisors

You may need to modify the environment or resources available to processes before the processes are 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 processes will use the content-specific `RunAs` account, falling back to `Applications.RunAs` if no override was configured. See Section 13.5 for details.

Supervisors run within the sandbox established for any process. See Section 13.1 for more information about process sandboxes.

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

13.9.1 Example Supervisors

Here is a configuration that uses the `nice` command to lower the priority of executing content. 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
Supervisor = nice -n 2
```

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

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[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"

# Execute the target process after the environment is established.
# All customization must happen before this "exec".
exec "$@"
```

The argument list of the supervisor is the full command-line of the target command. The supervisor *MUST* invoke this target command using `exec` or an equivalent technique.

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.

13.10 Using the `config` Package

The `config` package makes it easy to manage environment specific configuration values in R code. For example, you might want to use one value for a variable locally, and another value when deployed on RStudio Connect. The package vignette contains more information.

The desired configuration is identified to the `config` package by the `R_CONFIG_ACTIVE` environment variable. By default, R processes launched by RStudio Connect set `R_CONFIG_ACTIVE` to `rsconnect`. The value can be changed by modifying the `Applications.RConfigActive` configuration setting. Note that the value of `R_CONFIG_ACTIVE` is not available during package installation.

14 Content Management

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

14.1 Sharing Settings

Each deployment in RStudio Connect can have specific access controls which specify which users are allowed to view and/or edit that content.

14.1.1 Collaborators

The list of collaborators enumerates the users allowed to edit and help manage the settings for a given deployment. The content owner is always included as a collaborator. Collaborators must be either “publisher” or “administrator” accounts.

14.1.2 Viewers

A viewer is able to view content. Any type of account can be made a viewer for a given piece of content. Choose from the following options.

Anyone - no login required Any visitor to RStudio Connect will be able to view this content. This includes anonymous users who are not authenticated with the system.

All users - login required All RStudio Connect accounts are permitted to view this content.

Specific users or groups Specific users (or groups of users) are allowed to view this content. Other users will not have access.

You Only the owner of this content is able to view this content.

14.1.2.1 Limiting allowed viewership

Some organizations want to restrict the types of access that publishers and administrators can assign to content.

The settings `Applications.MostPermissiveAccessType` and `Applications.AdminMostPermissiveAccessType` limit the viewership options allowed to publishers and administrators, respectively.

These settings take the following values:

all The most permissive access type. Allows all viewership permissions.

logged_in Permits only logged-in viewership options. This includes allowing the content owner and specific users and groups.

acl Content must explicitly enumerate specific users and groups.

There is not an access type that only allows the owner to view the content. Owner-only access is implemented as an empty access control list (ACL).

The following example requires publishers to constrain viewership by enumerating specific users and groups in access control lists. Publishers cannot grant access to all logged-in users or to anonymous, unauthenticated users. Administrators are allowed to configure any access type.

These values can be helpful to avoid accidentally exposing content too widely, while still allowing select individuals to publicize approved content.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
MostPermissiveAccessType = acl
AdminMostPermissiveAccessType = all
```

This next example prohibits content from being given unconstrained viewership. Content cannot be configured to allow anonymous, unauthenticated access. The restriction applies to both publishers and administrators. Use this configuration when you want to require user login before accessing any content on RStudio Connect.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Applications]
MostPermissiveAccessType = logged_in
```

The access type of existing content is not automatically altered. The two access type settings limit which permissions may be given to content. This restriction is only enforced when content is created or modified.

14.2 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.

14.3 Tags

You can use tags to organize content and make it easy for users to find content that they’re interested in. To begin, create a tag schema in the “Tags” section of the Connect dashboard by creating one or more tag categories. Define some tags, which can be nested any number of levels deep.

For example, if your data scientists are creating reports covering different geographical areas, you could create a category called “Geographical Area”. Then, you could create tags such as “Americas” or “Asia” and nest the tags “North America” and “South America” under “Americas”.

Only administrators can create and edit the tag schema. Categories and tags can be added, deleted, and renamed. Once a tag or category is deleted, all tags nested under it are also deleted.

Collaborators can associate content with one or more tags in the “Tags” tab of the content settings sidebar. Users can filter by tags to discover content, as long as they have permission to view that content.

For example, if multiple reports analyze the same set of data, those reports could be tagged with some identifier, such as “FY2016 Q3” for the third quarter of the 2016 fiscal year. A report that analyzes the third and fourth quarter could be tagged with “FY2016 Q3” and “FY2016 Q4”, and would appear when a user filters for either “FY2016 Q3” or “FY2016 Q4”.

14.4 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 Connect 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.

When Activating an alternative bundle for a Plumber API, existing requests will be serviced by processes already launched running the old code. New requests will be serviced by new processes running the new code.

14.5 Content Images

Once content has been published to RStudio Connect, the publisher may enter the “info” tab in the application settings pane and upload an image to be associated with the content.

- Application images are stored in the server’s data directory.
- Application images can be viewed by anyone with permission to view the associated content.
- Administrators cannot view images associated with content that they do not have access to. However, all administrators may change the visibility mode, or add themselves to the viewer list, for any piece of content. This ensures that the Administrator’s act of giving themselves viewer privileges can be tracked in the server’s audit logs.
- RStudio Connect will reject uploads of images larger than the `Applications.MaxAppImageSize` configuration setting value. By default, the server will tolerate images up to 10MB in size.
- RStudio Connect does no post-processing; images will be sent back to the user exactly as they were uploaded, and the RStudio Connect dashboard will make its best effort to display them in a consistent manner.
- The RStudio Connect dashboard will attempt to cache images in an efficient and consistent manner. Every time the image is displayed, the dashboard will check the server for a newer version of the image, and accept its cached version if a newer image does not exist.

14.6 API Keys

RStudio Connect allows users to access hosted content outside the web browser by utilizing API Keys - e.g. via shell scripts. API Keys are enabled by default. To change this behavior please see Section 14.6.2.

14.6.1 How this Works

API Keys are associated with user accounts. They provide roughly the same level of access to RStudio Connect as a user logged in via the browser would have.

If a user has a compromised API Key, the Key should be deleted as soon as possible. The administrator may wish to lock the account if the user is having difficulty deleting the API Key.

To retrieve static content or to invoke Plumber endpoints via API Keys an HTTP request must be made to the target URL of the published content. The request must contain an HTTP header whose key is `Authorization` and value is set to `Key API_KEY`.

Authorization: Key ABCDEFGHIJKLMNO

Note: Use the information above to allow API Keys to pass through a Proxy.

API Keys have the same authorization access levels as the user that owns them. Someone who uses an API Key will be able to view all content that the owner of the API Key has access to. API Keys are shared secrets and as such they should be stored securely and only be given to trusted applications. It is advisable that content requests be made securely over HTTPS. If a user believes that an API Key has been compromised, they can revoke just that key by deleting it.

For more details regarding API Keys please see the API Keys section in the User Guide.

To learn how to configure RStudio Connect to listen for HTTPS requests please see Section A.4.

14.6.2 Configuring API Keys

To disallow API Keys, set `Authentication.APIKeyAuth` to `false`.

```
;/etc/rstudio-connect/rstudio-connect.gcfg
[Authentication]
APIKeyAuth = false
```

15 R

RStudio Connect offers a flexible way of deploying your Shiny applications, Plumber APIs, 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 or Plumber API, 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 how 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. RStudio 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.

15.1 Installing R

Ubuntu and Red Hat/CentOS/SUSE Linux distributions offer a version of R; installation of the system R is detailed in sections 2.2.3 and 2.2.4.

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.

15.2 Upgrading R

RStudio Connect supports running multiple versions of R. In most cases, upgrading R should consist of building the new version of R and retaining the previous version. We strongly recommend supporting multiple versions of R instead of upgrading and maintaining a single version of R. Supporting multiple versions of R is the best way to ensure applications or reports published with specific package dependencies will continue to run.

In cases where a single version of R is being managed, R can be upgraded using the system package manager with the following steps:

1. Stop RStudio Connect; see 5.1
2. Follow the instructions to upgrade R.

For Ubuntu, be sure an up-to-date CRAN repo is in the source list, and then run:

```
sudo apt-get update
sudo apt-get install r-base --upgrade
```

For RedHat/CentOS:

```
sudo yum update R
```

For SUSE, be sure to follow the install instructions in the [CRAN SUSE page] (<https://cran.rstudio.com/bin/linux/suse/>) and then run:

```
sudo zypper update R-base
```

3. Start RStudio Connect; see 5.1

Following an upgrade, content dependent on R will be rebuilt on-demand. For example, during the next execution of a scheduled R Markdown document, Connect will automatically reinstall and rebuild all of the necessary packages before rendering the report. For Shiny applications and Plumber APIs, the reinstall and rebuild will occur the first time the application is requested. During the package updates, Connect will present a message and spinner indicating to the end user that the Shiny application will be available after the packages are successfully installed and built for the new version of R.

Not all packages can be reinstalled and rebuilt on newer versions of R. Rebuilding and restoring packages can take a significant amount of time and could delay or prevent the rendering of a report or the availability of a Shiny application.

15.3 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.

15.3.1 Scanning

RStudio Connect can automatically scan for R in some well-known locations.

RStudio Connect scans each of the following directories to determine if they contain a full, single-version installation:

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

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

Each of these locations that contains an R installation will be available for use within RStudio Connect.

Connect also scans directories that often contain multiple R installations:

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

Any child directory that contains an R installation will be available for use within RStudio Connect.

Use the `Server.RVersion` setting or the `r-versions` file to explicitly enumerate R installations if your organization uses a complicated directory hierarchy.

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/R/3.4.2-optimized-build
/opt/local/R/3.3.1
```

Symbolic links from any of these scan locations to alternate locations are permitted.

R version scanning happens by default. Disable version scanning with the `Server.RVersionScanning` property.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
RVersionScanning = false
```

15.3.2 Direct Configuration

RStudio Connect can be instructed to use specific R installations through direct R version configuration.

Direct configuration can be used in combination with scanning to add R installations that are outside the scanned locations. It is most commonly used when R version scanning is disabled - giving complete control over the permitted R versions

15.3.2.1 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
RVersion = /shared/R/3.3.1
RVersion = /shared/R/3.2.4
RVersion = /shared/R/3.1.3
```

The `Server.RVersion` property is permitted to be a symbolic link to an R installation.

15.3.2.2 `/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
```

Paths enumerated in the `r-versions` file are permitted to be symbolic links to R installations.

15.3.2.3 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.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
RVersionScanning = false
RVersion = /opt/R/3.3.1
RVersion = /opt/R/3.2.4
```

15.4 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, Plumber API, 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, Plumber API, 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.

16 Python

RStudio Connect offers a flexible way of deploying Jupyter Notebooks and reticulated R content against a variety of Python versions.

A compatible version of Python is identified when content is deployed. That Python installation is used any time Python is needed for that content. Package installation and rendering content that utilized Python will use the version of Python discovered at deploy-time.

RStudio Connect allows different content to rely on different versions of Python. For example, Alice's Jupyter Notebook may require version 2.7.15 of Python while Bob's Shiny application with reticulate needs Python version 3.7.1. Those two deployments using different Python versions can coexist in RStudio Connect without conflict.

This chapter discusses how to install Python, how RStudio Connect can be configured to support one or more versions of Python, how Python version compatibility is determined, and other configuration settings relevant to Python support in RStudio Connect.

Available Python installations are analyzed at startup. RStudio Connect logs the provided Python versions and whether it can use each installation.

Note that each Python installation is required to have the `pip` and `virtualenv` Python packages installed. `virtualenv` is used to create content-specific environments and `pip` is used to install Python packages.

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

16.1 Enabling Python Support

Python support can be enabled by setting the `Python.Enabled` property to `true`. Note that it is `false` by default.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Python]
Enabled = true
```

RStudio Connect also needs to be configured to use a specific Python installation in order to fully support Python content.

The `Python.Executable` property can be used to specify locations for Python installations. Specify this property once for each Python installation directory.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Python]
Enabled = true
Executable = /shared/Python/2.7.9/bin/python
Executable = /shared/Python/3.5.6/bin/python3
Executable = /shared/Python/3.7.1/bin/python3
```

The `Python.Executable` property is permitted to be a symbolic link to a Python installation.

As with R, the best practice is to install Python from source rather than with a package manager like `apt`, `yum`, or `zypper`. See section 16.2 to learn how to install Python from source.

16.2 Installing Python

The minimum Python 2 version supported is 2.7.9 and the minimum Python 3 supported version is 3.4.0. Each Python installation must additionally have the pip, setuptools, and virtualenv packages installed.

16.2.1 Build Python from Source

To build Python from source, first, acquire the build dependencies for Python.

In Ubuntu, you can install build dependencies with

```
apt-get build-dep python
apt-get install libffi-dev libgdbm-dev libsqlite3-dev
```

In RedHat Enterprise Linux and CentOS, you can install build dependencies with

```
yum-builddep python python-libs
yum install libffi-devel zlib zlib-devel
```

In SUSE, build dependencies should be installed directly with `zypper install <dependency>`. A partial list of the libraries that you might need is below.

```
zypper install \
  automake \
  fdupes \
  gcc \
  gcc-c++ \
  gcc-fortran \
  gdbm-devel \
  gettext-tools \
  gmp-devel \
  intltool \
  libbz2-devel \
  libexpat-devel \
  libffi-devel \
  libnsl-devel \
  lzma-devel \
  make \
  ncurses-devel \
  netcfg \
  openssl-devel \
  pkgconfig \
  readline-devel \
  sqlite-devel \
  xz
```

Second, you should download and unpack the source tarball for the version of Python that you want to install. To install Python 3.7.1, this might look like the following:

```
# Download and extract source code
export VERSION=3.7.1 PYTHON_MAJOR=3
wget https://www.python.org/ftp/python/${VERSION}/Python-${VERSION}.tgz
tar -xzf Python-${VERSION}.tgz
cd Python-${VERSION}
```

The third step is to configure, make, and install Python. The recommended install location in this case is at `/opt/Python/${VERSION}`. If you would like to choose a different location, you can change `--prefix=`.

```

# Build Python from source
./configure \
  --prefix=/opt/Python/${VERSION} \
  --enable-shared \
  --enable-ipv6 \
  LDFLAGS=-Wl,-rpath=/opt/Python/${VERSION}/lib,--disable-new-dtags
make

# Install Python
sudo make install

```

It is very important that the Python installation folder is *not moved* once it is installed from source.

To test that the installation went smoothly, execute:

```
/opt/Python/${VERSION}/bin/python${PYTHON_MAJOR} --version
```

The final step is to install `pip`, `setuptools`, and `virtualenv` into the Python installation.

```

# Install pip
wget https://bootstrap.pypa.io/get-pip.py
sudo /opt/Python/${VERSION}/bin/python${PYTHON_MAJOR} get-pip.py

# Install virtualenv
sudo /opt/Python/${VERSION}/bin/pip install virtualenv

```

Ensure `setuptools` is installed in the Python installation by running `sudo /opt/Python/${VERSION}/bin/pip list | grep setuptools`. If `setuptools` is missing you can install it via `sudo /opt/Python/${VERSION}/bin/pip install setuptools`.

Once you have Python installed on the server, it is important to understand how Connect discovers Python and chooses the Python version that an asset will use. 16.4 explains more about this process.

16.3 Upgrading Python

RStudio Connect supports running multiple versions of Python. In most cases, upgrading Python should consist of building the new version of Python and retaining the previous version. We strongly recommend supporting multiple versions of Python instead of upgrading and maintaining a single version of Python. Supporting multiple versions of Python is the best way to ensure applications or reports published with specific package dependencies will continue to run.

After installing another Python version:

1. Stop RStudio Connect; see 5.1
2. Add an additional `Python.Executable` property to the configuration file that points to the newly built Python installation.
3. Start RStudio Connect; see 5.1

Following an upgrade, content dependent on Python will be rebuilt on-demand. For example, during the next execution of a scheduled Jupyter Notebook, RStudio Connect will automatically reinstall and rebuild all of the necessary packages before rendering the report. During the package updates, RStudio Connect will present a message and spinner indicating to the end user that the rendered Jupyter Notebook will be available after the packages are successfully installed and built for the new version of Python.

Not all packages can be reinstalled and rebuilt on newer versions of Python. Rebuilding and restoring packages can take a significant amount of time and could delay or prevent the rendering of a Jupyter Notebook or reticulated R content.

16.4 Python Version Matching

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

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

The “major-minor” algorithm 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.

e.g. If content was deployed with Python 3.4.0 and RStudio Connect is configured with Python 3.4.3, then Python 3.4.3 will be used as the MAJOR.MINOR numbers match (3.4 and 3.4) but if RStudio Connect is configured with only Python 3.5.6 then the content will fail to deploy as the MAJOR.MINOR versions are different (3.4 v. 3.5).

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

The Python version matching approach is controlled with the `Python.VersionMatching` configuration setting.

major-minor Find a Python installation that is close to the version of Python used when authoring content, requiring an exact MAJOR.MINOR version match. If a compatible version cannot be found, content will fail to deploy.

1. Use exact version match.

e.g. If content was deployed with Python 2.7.10 and RStudio Connect is configured with Python 2.7.9, Python 2.7.10, and Python 2.7.15, then Python 2.7.10 will be used.

2. If there are matching MAJOR.MINOR releases, use least-greater version.

e.g. If content was deployed with Python 2.7.9 and RStudio Connect is configured with Python 2.7.10 and Python 2.7.15, then Python 2.7.10 will be used as it is the lowest matching version greater than 2.7.9.

3. If there are matching MAJOR.MINOR releases, use latest of these.

e.g. If content was deployed with Python 2.7.15 and RStudio Connect is configured with Python 2.7.9 and Python 2.7.10, then Python 2.7.10 will be used as it is the highest matching version less than 2.7.15.

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

16.5 Excluding Python Packages

To exclude a certain Python package from being utilized, set `Python.ProhibitedPackage` to the name of the package. Specify this property once for each Python package that is to be excluded.

This may be used, for example, when deploying Python-enabled content that utilizes OS-specific packages which are unavailable on the OS that RStudio Connect runs on.

RStudio Connect excludes certain Python packages by default. For a list of these packages please see section A.25.

16.6 Configuring pip

Configuration options for `pip` may be set in `/etc/pip.conf`. These options will be read by `pip` when installing Python packages. This may be used, for example, to add additional package indexes or to control connection timeouts.

Refer to the `pip` documentation for further information.

17 Package Management

17.1 Package Installation

RStudio Connect installs the R package dependencies of Shiny applications, Plumber APIs, 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. This installation of `shiny` is placed into the packrat package cache as well as associated with that Shiny application deployment. Subsequent Shiny applications can use that cached package installation 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.

17.1.1 Package compilation

Some packages contain C and C++ code components. That code needs to be compiled during package installation. The `Server.CompilationConcurrency` setting controls the number of concurrent compilation processes used by package installation.

The default value for the `Server.CompilationConcurrency` setting is derived from the number of available CPUs with the formula $\max(1, \min(8, (\text{cpus}-1)/2))$. This property controls the number of concurrent C/C++ compilations during R package installation. This value makes it less likely for package installs to encounter memory capacity issues on lightweight hosts while allowing more concurrency on high-capacity servers.

CPUs	1	2	4	6	8	16	24	32
CompilationConcurrency	1	1	1	2	3	7	8	8

You can customize `Server.CompilationConcurrency` to force a specific level of concurrency.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
CompilationConcurrency = 1
```

17.1.2 External Package Installation

Warning: Adding external packages decreases the reproducibility and isolation of content on RStudio Connect, and should only be done as a last resort.

You can indicate that a system-wide installation of a package should be used instead of one fetched by packrat. The `Packages.External` can be used to enumerate each system-provided package

For example, `RJava` or `ROracle` are large installations, potentially with odd dependencies, such as your choice of JDK and/or Oracle InstantClient. First, you would install these packages in every R installation that RStudio Connect will be using. Then, you would configure RStudio Connect with the following parameters:

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Packages]
External = ROracle
External = RJava
```

This is the same as settings the packrat option `external.packages` to `c("ROracle", "RJava")` using `packrat::set_opts`. The `external.packages` option instructs `packrat::restore` to load certain packages from the user library. See the packrat documentation for more information.

17.1.3 Proxy Configuration

If the `http_proxy` and/or `https_proxy` environment variables are provided to RStudio Connect when the server starts, those variables will be passed to all processes run by RStudio Connect, including the package installation process.

Configuring `Packages.HTTPProxy` and `Packages.HTTPSProxy` will provide their values as the `http_proxy` and `https_proxy` environment variables only when packages are installed during deployment. This could be useful if you have a special proxy just for downloading package dependencies. You could regulate access to unapproved packages in non-CRAN repositories by rejecting certain URL patterns.

17.2 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)`.

17.3 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 17.2 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 `Server.SourcePackageDir` directory hierarchy or standardize to a particular git commit of the package.

18 Historical Information

This section describes the configuration and management of historical information, of which there are two types. Metrics about resource usage over time (trends) are enabled with the `Metrics.Enabled` setting. More discrete usage events are enabled with the `Metrics.Instrumentation` setting. Both settings are `true`, by default; to disable either one, set it to `false` in your configuration file.

18.1 Historical Metrics

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.

18.1.1 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 `Metrics.DataPath` setting.

The `rserver-monitor` process runs (by default) with the same user account Connect uses to run its processes associated with deployed content. By default, this user account is `rstudio-connect` (see the `Applications.RunAs` setting). You can specify an alternate user account for the `rserver-monitor` process by modifying the `Metrics.User` setting.

RStudio Connect also supports logging of metrics to Graphite, and it supports disabling its default behavior of logging to RRD. Please see Section A.23 for more options for configuring the historical metrics in Connect.

18.1.2 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`.

18.1.3 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.

18.2 Historical Events

RStudio Connect can record event-style usage information which is intended to answer questions like, “Who used my Shiny app and for how long?” This information is stored in dedicated tables in the database. When using SQLite, this is handled automatically by creating a second database file named `connect-instrumentation`. For PostgreSQL, a second, full database URL can be provided in the `Postgres.InstrumentationURL` setting. If it is not specified, it will default to the value of `Postgres.URL`. This allows you to store the event data in the same place as the rest of the Connect information, in a different schema, or even a different database, whichever meets your needs best. Please see Section 9.2 for more details about using Postgres.

Note: There is currently no data retention policy so all data will always be kept. Data retention controls will be added in a future release.

Note: This data is not migrated by the `migrate` tool (see 9.3).

18.2.1 Shiny Application Events

When a user opens a Shiny application, an event containing their user information and the length of their session will be logged to the instrumentation database. It's important to note that some configuration settings may affect how the ending time for a session is set.

- The ended time will always include the duration configured in the `Client.ReconnectTimeout` property. This is 15 seconds, by default.
- The `Scheduler.ConnectionTimeout` and `Scheduler.ReadTimeout` values control when an idle session is terminated by RStudio Connect. This will, as a side effect, set the ended time for the session.

These events may be retrieved by making use of the “Get Shiny Application Usage” API. The API returns information in pages and provides URLs in each response that may be used as-is to request the next or previous page of information. All data may be retrieved by first invoking the endpoint without `next` or `previous` parameters to return the first page of results and then repeatedly following the `next` link in each response until that link becomes null.

The API may only be used by administrators and publishers. Additionally, publishers may only retrieve information about the shiny apps that they own.

Optional filters within the request may be used to limit what usage records are returned with each response. Filters are “ANDed” together (i.e. data returned will satisfy all filters).

Application GUIDs may be provided to limit responses to particular applications. A publisher will be implicitly limited to only applications he or she owns. If a publisher asks for information about other content (content owned by someone else), the result will not contain data for that application and will not be reported as an error.

Timestamps may be provided to limit usage information to a more narrow time window of interest. By using the `from` or `to` filters, either independently or together, the information returned will be limited to applications that were being accessed within that window of time. It's worth noting that information for an app will be included in such a request if any portion of its usage by a user falls within the specified time window.

The data returned for a shiny application session includes a version number. As the Connect software has evolved, issues have been identified with how data is recorded. The version number provides an indication of any known issues and are described below.

Version

Issue

0

Extraneous records were recorded under some conditions, notably when protocols other than websockets were used. These can be identified by the same value for `started` and `ended`. This may adversely affect analyses involving counts or session lengths.

1

No known issues.

The `min_data_version` filter may be used to control what data to return. The default minimum data version to return is 1.

More information about using the API is included in the Connect Server API reference and in the Cookbook chapter of the RStudio Connect User Guide.

18.2.2 Rendered and Static Content Visit Events

When a user visits an R Markdown document (or other such rendered content) or static page (such as a plot), an event will be logged to the instrumentation database noting the time of the visit, the visitor and the content visited.

These events may be retrieved by making use of the “Get Content Visits” API. The API returns information in pages and provides URLs in each response that may be used as-is to request the next or previous page of information. All data may be retrieved by first invoking the endpoint without `next` or `previous` parameters to return the first page of results and then repeatedly following the `next` link in each response until that link becomes null.

The API may only be used by administrators and publishers. Additionally, publishers may only retrieve information about the content that they own.

Optional filters within the request may be used to limit what visit records are returned with each response. Filters are “ANDed” together (i.e. data returned will satisfy all filters).

Application GUIDs may be provided to limit responses to particular content. A publisher will be implicitly limited to only content he or she owns. If a publisher asks for information about other content (content owned by someone else), the result will not contain data for that content and will not be reported as an error.

Timestamps may be provided to limit visit information to a more narrow time window of interest. By using the `from` or `to` filters, either independently or together, the information returned will be limited to content that was visited within that window of time.

The data returned for a visit to content includes a version number. As the Connect software has evolved, issues have been identified with how data is recorded. The version number provides an indication of any known issues and are described below.

Version

Issue

0

Extraneous records were recorded under some conditions, notably when content is not rendered to a self-contained page *and* refers to images, CSS, JavaScript, and the like as files external to the page but within the content.

1

No known issues.

The `min_data_version` filter may be used to control what data to return. The default minimum data version to return is 1.

More information about using the API is included in the Connect Server API reference and in the Cookbook chapter of the RStudio Connect User Guide.

18.2.3 User Login Events

When a user logs in to the Connect dashboard their user information will be logged to the instrumentation database.

This event information is not currently presented in the dashboard or via an API.

18.2.4 Server Node Session Events

Note: Node session events may be erroneous if you have multiple nodes with the same hostname and do not reconfigure your node name, as explained below.

When a node is started, an event will be logged to the database containing a node name, the server start time, and a periodically-updating heartbeat timestamp indicating the length of the node's running session. A node that exits cleanly will log `true` to the `exited_cleanly` column for its session.

If `exited_cleanly` is `false`, it means either:

- The node is still running. In this case, the heartbeat will continue to update.
- The node's `rstudio-connect` process was terminated with `SIGKILL`, or the system lost power while the process was running. Confirm this by cross-referencing the node's log file with the row for the session at issue. The log for that session will end abruptly in this case.
- The node was terminated with `SIGTERM` or `SIGINT`, but did not successfully write to the database before being terminated with `SIGKILL` or before the system lost power. Confirm this by cross-referencing the node's log file, looking specifically for the line beginning with `Caught SIGINT/SIGTERM`. The log for that session will end abruptly after that line in this case.
- The node was terminated with `SIGTERM` or `SIGINT`, but couldn't write to the database for some other reason. Confirm this by cross-referencing the node's log file, looking specifically for `Error storing server exit time`. The log for that session will contain that line in this case.

The node name defaults to the node's hostname, but can be changed using the `Server.NodeName` configuration setting or the `RSTUDIO_CONNECT_NODE_NAME` environment variable. The node name **MUST** be unique for every node in your cluster. RStudio Connect cannot detect duplicate node names at this time, including the situation where multiple nodes have the same hostname.

The heartbeat timestamp can be changed from its default of 30m by setting `Metrics.InstrumentationServerHeartbeat` to another duration.

This event information is not currently presented in the dashboard or via an API.

A Configuration Options

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

A.1 Configuration Basics

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.

A.1.1 File Format

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 (;) or hash (#) 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.

A.1.2 Multi-value Properties

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.

A.1.3 Property Types

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"`. Strings containing the comment characters ; and # also need to be quoted.

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), `h` (hours), and `d` (days).

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

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

A.1.4 Reloadable Properties

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

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

A.2 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.

A.2.1 Server Settings

A.2.1.1 DataDir

The directory where RStudio Connect will store its variable data.

Type: string Default: `/var/lib/rstudio-connect`

A.2.1.2 TempDir

The directory that will contain all temporary directories needed by R processes. If `TMPDIR` environment variable is not defined then `/tmp`.

Type: string Default: `$TMPDIR`

A.2.1.3 LandingDir

Specifies an optional path from which a customized landing page is served to logged-out users. See `examples/landing-page` for a directory containing a sample landing page.

Type: string Default: `<empty-string>`

A.2.1.4 EnableSitemap

Specifies if RStudio Connect should provide a `/sitemap.xml` file enumerating the publicly available apps.

Type: boolean Default: `false`

A.2.1.5 RVersionMatching

Specifies how RStudio Connect attempts to match R version associated with uploaded content with the R versions available on the system. Allows values of `nearest` or `exact`.

Type: string Default: `nearest`

A.2.1.6 RVersion

Path to an R installation root. Multiple definitions can be used to provide multiple locations with R.

Type: multi-string Default: `unspecified`

A.2.1.7 RVersionScanning

Scan for R installations in well-known locations.

Type: boolean Default: `true`

A.2.1.8 CompilationConcurrency

The amount of parallelism allowed to `make` during R package installs. This value is passed to `make` as the value given to the `-jNUM` flag. The value `DEFAULT_CPUS` is derived from the number of available CPUs with the following calculation: $\max(1, \min(8, (\text{numcpu}-1)/2))$.

Type: integer Default: `#DEFAULT_CPUS`

A.2.1.9 SourcePackageDir

A directory containing source bundles for packages that are unavailable on either CRAN or a public GitHub repository. Must be readable by the `Applications.RunAs` user.

Type: string Default: `<empty-string>`

A.2.1.10 Address

A public URL for this RStudio Connect server such as `https://connect.company.com/`. Must be configured to enable features like including links to your content in emails.

Type: string Default: `<empty-string>`

A.2.1.11 SenderEmail

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.

Type: string Default: `<empty-string>`

A.2.1.12 EmailSubjectPrefix

A leading subject prefix for all mail sent by RStudio Connect.

Type: string Default: `[RStudio Connect]`

A.2.1.13 ViewerKiosk

When enabled, RStudio Connect does not prompt view-only users to request elevated privileges when attempting to access restricted resources.

Type: boolean Default: `false`

A.2.1.14 HideEmailAddresses

When enabled, RStudio Connect will not expose email addresses to non-administrators in API requests or its dashboard.

Type: boolean Default: `false`

A.2.1.15 MailAll

When enabled, RStudio Connect will allow scheduled and on-demand documents to send email to all users of the system.

Type: boolean Default: `false`

A.2.1.16 PublicWarning

An HTML snippet used to inject a message into the RStudio Connect dashboard welcome pages.

Type: string Default: *<empty-string>* Reloadable: true

A.2.1.17 LoggedInWarning

An HTML snippet used to inject a message into the RStudio Connect recent views.

Type: string Default: *<empty-string>* Reloadable: true

A.2.1.18 ContentTypeSniffing

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

Type: boolean Default: `false`

A.2.1.19 ServerName

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

Type: string Default: `RStudio Connect vX.Y.Z-NNNN`

A.2.1.20 AccessLog

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

Type: string Default: `/var/log/rstudio-connect.access.log`

A.2.1.21 CustomHeader

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

Type: multi-string Default: *unspecified*

A.2.1.22 FrameOptionsContent

The value for the `X-Frame-Options` HTTP header for all user-uploaded content (Shiny apps, RMDs, etc.). If empty, no header will be added.

Type: string Default: *<empty-string>*

A.2.1.23 FrameOptionsDashboard

The value for the `X-Frame-Options` HTTP header for the Connect dashboard and all other Connect pages. If empty, no header will be added.

Type: string Default: `DENY`

A.2.1.24 HideVersion

If `true`, will suppress display of the server version and build numbers.

Type: boolean Default: `false`

A.2.1.25 NodeName

The name of this node. If not specified, the contents of the `RSTUDIO_CONNECT_NODE_NAME` environment variable or the special file `/proc/sys/kernel/hostname` are used, in that order.

Type: string Default: `$RSTUDIO_CONNECT_NODE_NAME`

A.3 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.

A.3.1 HTTP Settings

A.3.1.1 Listen

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

Type: string Default: `<empty-string>`

A.3.1.2 NoWarning

Disables warnings about insecure (HTTP) connections.

Type: boolean Default: `false`

A.3.1.3 ForceSecure

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

Type: boolean Default: `false`

A.4 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.

A.4.1 HTTPS Settings

A.4.1.1 Listen

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

Type: string Default: *<empty-string>*

A.4.1.2 Key

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

Type: string Default: *<empty-string>*

A.4.1.3 Certificate

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

Type: string Default: *<empty-string>*

A.4.1.4 Permanent

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

Type: boolean Default: `false`

A.4.1.5 MinimumTLS

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

Type: string Default: `1.1`

A.4.1.6 ExcludedCiphers

List of IANA strings for ciphers to exclude. See <https://www.iana.org/assignments/tls-parameters/tls-parameters.xml> for a list of all known cipher strings. e.g.: `TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA` `TLS_RSA_WITH_3DES_EDE_CBC_SHA` excludes DES ciphers.

Type: multi-string Default: *unspecified*

A.5 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.

A.5.1 HTTPRedirect Settings

A.5.1.1 Listen

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

Type: string Default: *<empty-string>*

A.5.1.2 Target

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

Type: string Default: *<empty-string>*

A.6 Licensing

The `Licensing` section contains configuration properties which control how RStudio Connect interacts with its licensing system.

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

A.6.1 Licensing Settings

A.6.1.1 LicenseType

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

Type: string Default: `local`

A.6.1.2 RemoteRetryFrequency

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

Type: duration Default: `10s`

A.6.1.3 ExpirationEmail

Enables sending of email when the license approaches expiration.

Type: boolean Default: `true`

A.6.1.4 ExpirationUIWarning

Enables display of a warning to administrators and publishers when the license approaches expiration.

Type: boolean Default: `true`

A.7 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.

A.7.1 Database Settings

A.7.1.1 Provider

The type of database. Either `sqlite` or `postgres`.

Type: string (case-insensitive) Default: `sqlite`

A.7.1.2 Dir

The directory to store database support files.

Type: string Default: `{Server.DataDir}/db`

A.7.1.3 MaxIdleConnections

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

Type: integer Default: 5

A.7.1.4 MaxOpenConnections

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

Type: integer Default: 10

A.7.1.5 ConnectionMaxLifetime

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

Type: duration Default: 60s

A.8 SQLite

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

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

A.8.1 SQLite Settings

A.8.1.1 Dir

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

Type: string Default: `{Database.Dir}`

A.8.1.2 Backup

When enabled and Provider is `sqlite`, periodically backs up the database

Type: boolean Default: `true`

A.8.1.3 BackupFrequency

How often to back up the SQLite database

Type: duration Default: `24h`

A.8.1.4 BackupRetentionLimit

Connect will periodically delete backups older than this number. Set to 0 to disable sweeping, for example, if you plan to manage your backups with an external task.

Type: integer Default: `3`

A.9 Postgres

The `Postgres` section contains configuration properties which control the location of and how RStudio Connect interacts with its postgres.

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

A.9.1 Postgres Settings

A.9.1.1 URL

The fully qualified URL to connect to a Postgres database

Type: string Default: `<empty-string>`

A.9.1.2 InstrumentationURL

The fully qualified URL to connect to a Postgres database where instrumentation data will be written.

Type: string Default: `<empty-string>`

A.10 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.

A.10.1 Authentication Settings

A.10.1.1 Provider

Specifies the type of user authentication. Allows values of `password`, `ldap`, `saml`, `oauth2`, `pam`, or `proxy`.

Type: string (case-insensitive) Default: `password`

A.10.1.2 Name

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.

Type: string Default: `RStudio Connect`

A.10.1.3 Lifetime

The lifetime of an authenticated session.

Type: duration Default: `30d`

A.10.1.4 Inactivity

The period of time after which inactive sessions are considered idle and therefore invalid. Effective only when non-zero and less than `Authentication.Lifetime`.

Type: duration Default: `0`

A.10.1.5 Notice

Text of notice displayed to the user on the login, sign up, and user completion pages. Useful for displaying notices and reminders to users about the system and the data they enter.

Type: string Default: `<empty-string>`

A.10.1.6 APIKeyAuth

Whether API key authentication is enabled.

Type: boolean Default: `true`

A.10.1.7 ChallengeResponseEnabled

Whether a second factor challenge-response is enabled

Type: boolean Default: `false`

A.10.1.8 CookieSweepDuration

Duration between sweeps of expired cookies

Type: duration Default: `1h`

A.11 Password

The `Password` section contains configuration properties which control how RStudio Connect's default password authentication provider behaves.

See Section 10.5 for details about configuring password authentication for RStudio connect.

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

A.11.1 Password Settings

A.11.1.1 SelfRegistration

Allow users to self-register. Self-registered users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: `true`

A.11.1.2 WebSudoMode

Whether web-sudo mode is enabled. Web-sudo mode will ensure that users are prompted to re-enter their password before performing sensitive actions.

Type: boolean Default: `true`

A.11.1.3 WebSudoModeDuration

The lifetime of web-sudo mode. This is the amount of time before a user will be prompted to enter their password again when performing sensitive actions.

Type: duration Default: `5m`

A.12 OAuth2

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

Section 10.8 contains more information about configuring RStudio Connect to use Google for authentication.

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

A.12.1 OAuth2 Settings

A.12.1.1 ClientId

Identifier for OAuth2 client. Required for all OAuth2 configurations.

Type: string Default: `<empty-string>`

A.12.1.2 ClientSecret

Client secret for the configured client ID. One of `OAuth2.ClientSecret` or `OAuth2.ClientSecretFile` must be specified when using OAuth2.

Type: string Default: `<empty-string>`

A.12.1.3 ClientSecretFile

Path to file containing the client secret for the configured client ID. One of `OAuth2.ClientSecret` or `OAuth2.ClientSecretFile` must be specified when using `OAuth2`.

Type: string Default: `<empty-string>`

A.12.1.4 AllowedDomains

Space-separated list of domains permitted to authenticate.

Type: string Default: `<empty-string>` Reloadable: true

A.12.1.5 AllowedEmails

Space-separated list of email addresses permitted to authenticate. When used without `OAuth2.AllowedDomains`, only the email addresses listed here will be allowed access. When used with `OAuth2.AllowedDomains`, the email addresses listed here will be **added** to those with valid domains.

Type: string Default: `<empty-string>` Reloadable: true

A.12.1.6 DemoteSearchErrors

Disables logging of user search errors. Recommended only for public servers where gmail.com accounts or accounts from multiple, unrelated Google Apps domains appear.

Type: boolean Default: `false`

A.12.1.7 Logging

Enables logging for all `OAuth2` operations.

Type: boolean Default: `false`

A.12.1.8 RegisterOnFirstLogin

Allow users to be created on their first login attempt. Otherwise users must be created ahead of the first login which varies across authentication providers. Users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: `true`

A.13 LDAP

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

See Section 10.6 for details about how to configure RStudio Connect with LDAP authentication. Section D 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.

A.13.1 LDAP Settings

A.13.1.1 ServerAddress

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.

Type: string Default: *<empty-string>*

A.13.1.2 TLS

When enabled, all connections to your LDAP/AD server will use TLS (SSL).

Type: boolean Default: `false`

A.13.1.3 StartTLS

When enabled, the connection will initially be made on an insecure port then the channel will be upgraded to TLS using StartTLS.

Type: boolean Default: `false`

A.13.1.4 ServerTLSInsecure

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.

Type: boolean Default: `false`

A.13.1.5 TLSCACertificate

Path to a certificate authority used to connect an LDAP server.

Type: string Default: *<empty-string>*

A.13.1.6 UserObjectClass

The name of the LDAP `objectClass` used to define users.

Type: string Default: *<empty-string>*

A.13.1.7 UserFilterBase

An LDAP filter clause used to select user objects. Defaults to `objectClass={UserObjectClass}`.

Type: string Default: *<empty-string>*

A.13.1.8 UserFirstNameAttribute

The LDAP user attribute containing a user's first name. This is often the `givenName` attribute.

Type: string Default: *<empty-string>*

A.13.1.9 UserLastNameAttribute

The LDAP user attribute containing a user's last name. The `sn` attribute will usually contain last name.

Type: string Default: *<empty-string>*

A.13.1.10 UserEmailAttribute

The LDAP user attribute containing a user's email address. Many systems use the `mail` attribute.

Type: string Default: *<empty-string>*

A.13.1.11 UsernameAttribute

The LDAP user attribute containing a user's username. Commonly used attributes include `uid`, `cn`, and `sAMAccountName`.

Type: string Default: *<empty-string>*

A.13.1.12 UniqueIdAttribute

The LDAP attribute with an immutable value which uniquely identifies objects. Commonly used attributes include `objectGUID` (Microsoft AD), `entryUUID` (OpenLDAP), `orclGuid` (Oracle OID), `ibm-entryUUID` (IBM RACF), `GUID` (Novell eDirectory), `nsUniqueID` (389 Directory Server). Please refer to your LDAP vendor documentation for the correct value.

Type: string Default: DN

A.13.1.13 BindDN

A DN for a read-only administrator 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`.

Type: string Default: *<empty-string>*

A.13.1.14 BindPassword

The password for the `BindDN` account.

Type: string Default: *<empty-string>*

A.13.1.15 BindPasswordFile

Path to file containing the bind password. Either `BindPassword` or `BindPasswordFile` may be specified when using LDAP, but if both are set, it is an error.

Type: string Default: *<empty-string>*

A.13.1.16 AnonymousBind

Enable anonymous bind. An anonymous user must have rights to search and view all pertinent groups, group memberships, and users.

Type: boolean Default: `false`

A.13.1.17 UserSearchBaseDN

The base DN used when performing user searches.

Type: string Default: *<empty-string>*

A.13.1.18 PermittedLoginGroup

Limit who can log into Connect by specifying a group DN. Multiple definitions can be used to provide multiple groups.

Type: multi-string Default: *unspecified*

A.13.1.19 GroupObjectClass

The name of the LDAP `objectClass` used to define groups. Commonly this is `group` or `posixGroup`.

Type: string Default: *<empty-string>*

A.13.1.20 GroupSearchBaseDN

The base DN used when performing group searches.

Type: string Default: *<empty-string>*

A.13.1.21 GroupFilterBase

An LDAP filter clause used to select group objects. Defaults to `objectClass={GroupObjectClass}`.

Type: string Default: *<empty-string>*

A.13.1.22 GroupNameAttribute

The LDAP group attribute containing a group's name. Commonly this is `cn` or `sAMAccountName`.

Type: string Default: *<empty-string>*

A.13.1.23 GroupUniqueIdAttribute

The LDAP attribute with an immutable value which uniquely identifies groups. In general, users and groups use the same value. Please refer to your LDAP vendor documentation for the correct value for groups.

Type: string Default: DN

A.13.1.24 Logging

Enables logging for all LDAP operations.

Type: boolean Default: `false`

A.13.1.25 WebSudoMode

Whether web-sudo mode is enabled. Web-sudo mode will ensure that users are prompted to re-enter their password before performing sensitive actions.

Type: boolean Default: `true`

A.13.1.26 WebSudoModeDuration

The lifetime of web-sudo mode. This is the amount of time before a user will be prompted to enter their password again when performing sensitive actions.

Type: duration Default: 5m

A.13.1.27 RegisterOnFirstLogin

Allow users to be created on their first login attempt. Otherwise users must be created ahead of the first login which varies across authentication providers. Users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: `true`

A.14 SAML (Beta)

The **SAML** section contains configuration properties which control how RStudio Connect communicates with a SAML Identity Provider server.

See Section 10.7 for details about how to configure RStudio Connect with SAML authentication.

All of the SAML configuration properties must appear after `[SAML]` in the configuration file.

A.14.1 SAML Settings

A.14.1.1 IdPMetaData

The URL or file name of the metadata for the Identity Provider (IdP). This overrides the other IdP-related configuration properties.

Type: string Default: `<empty-string>`

A.14.1.2 IdPEntityID

The entity identifier (name or URI) of the Identity Provider (IdP).

Type: string Default: `<empty-string>`

A.14.1.3 IdPSingleSignOnServiceURL

The HTTP endpoint able to receive authentication requests on the Identity Provider (IdP).

Type: string Default: `<empty-string>`

A.14.1.4 IdPSigningCertificate

The path to a PEM file or the base64-encoded body of the public certificate to be used for validating the digital signature of a SAML response.

Type: string Default: `<empty-string>`

A.14.1.5 IdPEncryptionCertificate

The path to a PEM file or the base64-encoded body of the public certificate to be used for decrypting a SAML response.

Type: string Default: *<empty-string>*

A.14.1.6 IdPSingleSignOnPostBinding

When true uses HTTP POST for authentication requests to the Identity Provider (IdP). Uses HTTP redirect otherwise.

Type: boolean Default: `false`

A.14.1.7 NameIDFormat

The expected format for the user's NameID. One of `unspecified`, `persistent`, `transient` or `emailAddress`. This will be what RStudio Connect will request when initiating a single sign-on from the IdP.

Type: string Default: *<empty-string>*

A.14.1.8 UniqueIDAttribute

The name or URI of the SAML attribute containing the user's unique identifier. If `NameIDFormat = transient` the value cannot be the default of `NameID`.

Type: string Default: `NameID`

A.14.1.9 UsernameAttribute

The name or URI of the SAML attribute containing the username. Use `NameID` to refer to its value instead of an attribute.

Type: string Default: *<empty-string>*

A.14.1.10 FirstNameAttribute

The name or URI of the SAML attribute containing the user's first name.

Type: string Default: *<empty-string>*

A.14.1.11 LastNameAttribute

The name or URI of the SAML attribute containing the user's last name.

Type: string Default: *<empty-string>*

A.14.1.12 EmailAttribute

The name or URI of the SAML attribute containing the user's email address. Use `NameID` to refer to its value instead of an attribute.

Type: string Default: *<empty-string>*

A.14.1.13 SSOInitiated

Indicates who (IdP and/or SP) can initiate a login sequence. Valid values are IdP, SP, or IdPAndSP. It will default to IdPAndSP.

Type: string Default: IdPAndSP

A.14.1.14 Logging

Enables logging for all SAML operations.

Type: boolean Default: false

A.14.1.15 RegisterOnFirstLogin

Allow users to be created on their first login attempt. Otherwise users must be created ahead of the first login which varies across authentication providers. Users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: true

A.15 PAM

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

See Section 10.9 for details about configuring an appropriate PAM authentication profile for RStudio connect.

See Section 13.7 for information about using PAM sessions when launching processes.

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

A.15.1 PAM Settings

A.15.1.1 Service

Specifies the PAM service name that RStudio Connect will use when authenticating users.

Type: string Default: `rstudio-connect`

A.15.1.2 UseSession

Use PAM sessions when launching R processes.

Type: boolean Default: false

A.15.1.3 SessionService

Specifies the PAM service name that RStudio Connect will use for running R processes. This PAM service cannot require user credentials when executed by root, or all R processes run by Connect will fail yielding error code 70.

Type: string Default: `su`

A.15.1.4 AuthenticatedSessionService

Specifies the PAM service name that RStudio Connect will use for running R processes with cached user credentials. This can be used with a Kerberized PAM service if Kerberos exposes certain resources to the R process.

Type: string Default: `su`

A.15.1.5 ForwardPassword

If true, will retain the user's password when they login so that it can be provided to PAM when spawning proceses on this user's behalf. This can be used where PAM features such as Kerberos or `pam_mount` which require the user's password are in use.

Type: boolean Default: `false`

A.15.1.6 PasswordLifetime

The duration of time for which a password should be retained in memory. Only used if `ForwardPassword` is `true`.

Type: duration Default: `24h`

A.15.1.7 Logging

Enables logging for all PAM operations.

Type: boolean Default: `false`

A.15.1.8 WebSudoMode

Whether web-sudo mode is enabled. Web-sudo mode will ensure that users are prompted to re-enter their password before performing sensitive actions.

Type: boolean Default: `true`

A.15.1.9 WebSudoModeDuration

The lifetime of web-sudo mode. This is the amount of time before a user will be prompted to enter their password again when performing sensitive actions.

Type: duration Default: `5m`

A.15.1.10 RegisterOnFirstLogin

Allow users to be created on their first login attempt. Otherwise users must be created ahead of the first login which varies across authentication providers. Users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: `true`

A.16 Proxied Authentication

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

See Section 10.10 for details about configuring an appropriate proxied authentication for RStudio connect.

A.16.1 ProxyAuth Settings

A.16.1.1 UniqueIdHeader

Specifies the name of the header that will contain a user unique identifier provided by the proxy. By using this setting the RStudio Connect Server API will require an `unique_id` value to create users.

Type: string Default: *<empty-string>*

A.16.1.2 UsernameHeader

Specifies the name of the header that will contain a username provided by the proxy.

Type: string Default: `X-Auth-Username`

A.16.1.3 FirstNameHeader

Specifies the name of the header that will contain the user's first name provided by the proxy.

Type: string Default: *<empty-string>*

A.16.1.4 LastNameHeader

Specifies the name of the header that will contain the user's last name provided by the proxy.

Type: string Default: *<empty-string>*

A.16.1.5 EmailHeader

Specifies the name of the header that will contain the user's email provided by the proxy.

Type: string Default: *<empty-string>*

A.16.1.6 RoleHeader

Specifies the name of the header that will contain a user role provided by the proxy.

Type: string Default: *<empty-string>*

A.16.1.7 Logging

Enables logging for all Proxy Auth operations.

Type: boolean Default: `false`

A.16.1.8 GroupsHeader

Requires `Authorization.UserGroups`. Specifies the name of the header that will contain the group or groups of which the authenticated user is a member.

Type: string Default: *<empty-string>*

A.16.1.9 GroupsHeaderSeparator

Requires `GroupsHeader`. If not empty expects a single header value representing a list of group names separated by the character(s) defined here. By default, the header in `GroupsHeader` is expected to be received multiple times, one for each group name.

Type: string Default: *<empty-string>*

A.16.1.10 GroupsAutoProvision

Requires `GroupsHeader`. Enables automatic provisioning and removal of local groups as needed to remain in sync with proxy-specified memberships.

Type: boolean Default: `false`

A.16.1.11 GroupsAutoRemoval

Requires `GroupsAutoProvision`. When enabled provisioned groups are removed when there are no more members.

Type: boolean Default: `false`

A.16.1.12 RegisterOnFirstLogin

Allow users to be created on their first login attempt. Otherwise users must be created ahead of the first login which varies across authentication providers. Users will be created using the role specified in the `Authorization.DefaultUserRole` setting.

Type: boolean Default: `true`

A.17 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.

A.17.1 Authorization Settings

A.17.1.1 UserGroups

Enable the group support in Connect using locally-stored groups. This option is also required for associating remote groups to Connect users. Connect will fail to start if this setting is disabled and groups exist.

Type: boolean Default: `true`

A.17.1.2 UserInfoEditableBy

Specifies who is able to modify user attributes. Accepted configurations are `Admin` or `AdminAndSelf`. If `Admin` is specified, only administrators can edit user attributes. `AdminAndSelf` expands the permissions by allowing a user to edit their own user attributes.

Type: string (case-insensitive) Default: `AdminAndSelf`

A.17.1.3 DefaultUserRole

Specifies what abilities given to a newly created user. Allows values `viewer` or `publisher`.

Type: string (case-insensitive) Default: `viewer`

A.17.1.4 UsersListingMinRole

Specifies the minimum role level required to list users. Allows values `viewer`, `publisher`, or `administrator`.

Type: string (case-insensitive) Default: `viewer`

A.17.1.5 PublishersCanOwnGroups

When disabled only administrator can own and manage groups.

Type: boolean Default: `false`

A.18 Applications

The `Applications` section contains configuration properties which control how RStudio Connect communicates with the processes associated with deployed content.

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

A.18.1 Applications Settings

A.18.1.1 RunAs

User used to invoke R.

Type: string Default: `rstudio-connect`

A.18.1.2 RunAsCurrentUser

Allows content to execute as the logged-in user when using PAM authentication.

Type: boolean Default: `false`

A.18.1.3 RConfigActive

Specifies a value for the `R_CONFIG_ACTIVE` environment variable for R processes; supported by the `config` package.

Type: string Default: `rsconnect`

A.18.1.4 Supervisor

Specifies a command to wrap the execution of R.

Type: string Default: `<empty-string>`

A.18.1.5 HomeMounting

Specifies that the contents of `/home` should be hidden from R processes with additional bind mounts. The existing `/home` will have the home directory of the `RunAs` user mounted over it. If `RunAs` does not have a home directory, an empty temporary directory will mask `/home` instead. Launched R processes can discover this location through the the `HOME` environment variable.

Type: boolean Default: `false`

A.18.1.6 TempMounting

Specifies that per-process temporary directories are hidden from other R processes with additional bind mounts. If disabled, R processes may inspect or modify the temporary data of other processes.

Type: boolean Default: `true`

A.18.1.7 ProhibitedEnvironment

Environment variables that users are not allowed to override. Environment variables specified here will be added to the default set.

Type: multi-string Default: `CONNECT_SHARED_SECRET, HOME, LOGNAME, PWD, R_CONFIG_ACTIVE, RETICULATE_PYTHON, RSC_EMAIL_SUBJECT, RSC_REPORT_NAME, RSC_REPORT_RENDERING_URL, RSC_REPORT_SUBSCRIPTION_URL, RSC_REPORT_URL, RSTUDIO_PANDOC, TMPDIR, USER, USERNAME, VIRTUALENV, VIRTUALENV_HOME, WORKON_HOME`

A.18.1.8 ShinyBookmarking

Toggles support for on-disk Shiny bookmarking state. Configuring Shiny applications to use server bookmarking is described in this article.

Type: boolean Default: `true`

A.18.1.9 ShinyErrorSanitization

Toggles support for Shiny error sanitization as described in this article.

Type: boolean Default: `true`

A.18.1.10 ShinyIsolation

Specifies if a multi-Shiny R Markdown document deployment attempts to serve all documents from a single process. Disable if interactive documents expect in-process shared state.

Type: boolean Default: `true`

A.18.1.11 ExplicitPublishing

Content requires an explicit publication step after creation.

Type: boolean Default: `false`

A.18.1.12 MostPermissiveAccessType

The broadest access type publishers may assign to content. More permissive access types are disallowed. Applies to administrators when `Applications.AdminMostPermissiveAccessType` is not set. Allows values `all`, `logged_in`, or `acl`.

Type: string Default: `all`

A.18.1.13 AdminMostPermissiveAccessType

The broadest access type administrators may assign to content. More permissive access types are disallowed. The `Applications.MostPermissiveAccessType` value applies to administrators when this property is not set.

Type: string Default: all

A.18.1.14 MaxAppImageSize

The maximum file size(in bytes) allowed for an application image.

Type: integer Default: 10000000

A.18.1.15 RenderingSweepLimit

The maximum number of renderings retained for any one application. When this limit is reached, the oldest renderings for an application will be removed. If this setting is not set or is less than 1, Connect will not remove renderings with respect to the number of renderings per application.

Type: integer Default: 30

A.18.1.16 RenderingSweepAge

The maximum age of a rendering retained for any one application. Renderings older than this setting will be removed. If this setting is not set or is of a zero value, Connect will not remove renderings with regards to its age.

Type: duration Default: 30d

A.18.1.17 RenderingSweepFrequency

How often should Connect look for and purge renderings.

Type: duration Default: 1h

A.18.1.18 ViewerOnDemandReports

Allow logged in report viewers to generate an ad-hoc rendering. The `ViewerCustomizedReports` property is implicitly disabled when this property is disabled.

Type: boolean Default: false

A.18.1.19 ViewerCustomizedReports

Allow logged in report viewers to customize the parameters of an ad-hoc rendering.

Type: boolean Default: false

A.18.1.20 BundleReapFrequency

Time between the worker that deletes filesystem data for bundles in excess of our retention limit.

Type: duration Default: 24h

A.18.1.21 BundleRetentionLimit

Maximum number of bundles per app for which we want to retain filesystem data. The default is 0, which means retain everything.

Type: integer Default: 0

A.18.1.22 ScheduleConcurrency

Number of scheduled reports permitted to execute in parallel. A setting of zero disables scheduled execution on this host.

Type: integer Default: 2

A.18.1.23 DisabledProtocols

List of comma-delimited protocols to disable on the SockJS client. Allows values of `websocket`, `xhr-streaming`, `iframe-eventsources`, `iframe-htmfile`, `xhr-polling`, `iframe-xhr-polling`, or `jsonp-polling`. Protocols `xdr-streaming` and `xdr-polling` are always disabled.

Type: string Default: *<empty-string>*

A.19 Packages

The **Packages** section contains configuration properties which alter how R packages are installed and managed. See Section 17 for details.

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

A.19.1 Packages Settings

A.19.1.1 HTTPProxy

Value to be set for the `http_proxy` environment variable during package installation when content is deployed. When set, this will override the `http_proxy` environment variable only when content is built by connect.

Type: string Default: *<empty-string>*

A.19.1.2 HTTPSProxy

Value to be set for the `https_proxy` environment variable during package installation when content is deployed. When set, this will override the `https_proxy` environment variable only when content is built by connect.

Type: string Default: *<empty-string>*

A.19.1.3 External

Package to be excluded from packrat build. This can be provided multiple times, once for each package. You will need this package available in your library path.

Type: multi-string Default: *unspecified*

A.20 Client

The `Client` section contains configuration properties which control the behavior of browsers when interacting with applications. Interactive Shiny applications are the primary example.

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

A.20.1 Client Settings

A.20.1.1 ReconnectTimeout

The amount of time to allow a user connection to be restored. If a zero value, reconnects will be disabled. Disabling reconnects can cause instability with the `session$allowReconnects(TRUE)` feature in Shiny. Note that this amount is always added to a Shiny Application Usage end time.

Type: duration Default: 15s

A.21 Runtime/Scheduler

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

RStudio Connect makes a determination on each new client connection about whether or not it needs to spawn an additional process. That computation analyzes the number of current 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 processes for the current application.
// numConnections
// - The number of connections across all 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 associated idle processes linger longer so they are available the next time a request arrives - avoiding the startup penalty.

Worker processes are terminated when none of the following hold:

- Active client connections to the process
- Client connection during prior `Schedule.IdleTimeout` period
- Process is needed to satisfy the `LoadFactor` application setting
- Process is needed to satisfy the `MinProcesses` application setting

Any one of these conditions will keep a worker process alive. Consider a Shiny application configured with `IdleTimeout=10m` and `MinProcesses=1`. The process continues running if it is the only instance of that application, well beyond its configured `IdleTimeout`. Terminating that process would violate the `MinProcesses=1` constraint.

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.

A.21.1 Scheduler Settings

A.21.1.1 MaxProcesses

Specifies the total number of concurrent processes allowed for a single application.

Type: integer Default: 3

A.21.1.2 MaxConnsPerProcess

Specifies the maximum number of client connections allowed to an individual process. Incoming connections which will exceed this limit are routed to a new process or rejected.

Type: integer Default: 20

A.21.1.3 LoadFactor

Controls how aggressively new processes are spawned.

Type: decimal Default: 0.5

A.21.1.4 InitTimeout

Maximum time to wait for an app to start.

Type: duration Default: 60s

A.21.1.5 IdleTimeout

Minimum time a worker process remains alive after it goes idle (no active connections).

Type: duration Default: 5s

A.21.1.6 MinProcessesLimit

Maximum value allowed for the `MinProcesses` setting on an application level. All applications default to `MinProcesses=0`, but `MinProcesses` can be increased to this limit per application.

Type: integer Default: 10

A.21.1.7 MaxProcessesLimit

Maximum value allowed for the `MaxProcesses` setting on an application level. All applications take the value of `MaxProcesses` as default but can be individually adjusted. This setting prohibits a single application from consuming too many processes.

Type: integer Default: 20

A.21.1.8 ConnectionTimeout

Maximum time allowed without data sent or received across a client connection. A value of 0 means connections will never time-out (not recommended). Once a session is determined to be idle and terminated, the related usage information is marked as ended.

Type: duration Default: 1h

A.21.1.9 ReadTimeout

Maximum time allowed without data received from a client connection. A value of 0 means a lack of client (browser) interaction never causes the connection to close. This is useful when deploying applications which send regular updates but have no need for interactivity. Once a session is determined to be idle and terminated, the related usage information is marked as ended.

Type: duration Default: 1h

A.22 Jobs

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

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

A.22.1 Jobs Settings

A.22.1.1 MaxCompleted

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 `MaxCompleted` or `OldestCompleted` restrictions are violated.

Type: integer Default: 1000

A.22.1.2 OldestCompleted

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 `MaxCompleted` or `OldestCompleted` restrictions are violated.

Type: duration Default: 30d

A.23 Historical Information

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. It also controls how discrete event instrumentation is managed.

See Section 18 for more details about historical information in Connect.
These properties must appear after `[Metrics]` in the configuration file.

A.23.1 Metrics Settings

A.23.1.1 Enabled

Specifies whether or not the `rserver-monitor` process that collects historical metrics will be started.

Type: boolean Default: `true`

A.23.1.2 User

The user for the `rserver-monitor` process.

Type: string Default: `{Applications.RunAs}`

A.23.1.3 DataPath

The path for writing log entries and RRD database files.

Type: string Default: `{Server.DataDir}/metrics`

A.23.1.4 Interval

The frequency of historical metrics collection.

Type: duration Default: `60s`

A.23.1.5 RRDEnabled

Enable logging of historical metrics to RRD.

Type: boolean Default: `true`

A.23.1.6 GraphiteEnabled

Enable logging of historical metrics to Graphite.

Type: boolean Default: `false`

A.23.1.7 GraphiteHost

Host to which to send Graphite historical metrics.

Type: string Default: `127.0.0.1`

A.23.1.8 GraphitePort

Port to which to send Graphite historical metrics.

Type: integer Default: `2003`

A.23.1.9 GraphiteClientId

Optional Client ID to include along with Graphite historical metrics.

Type: string Default: *<empty-string>*

A.23.1.10 Instrumentation

Enable the recording of instrumented data, such as shiny app usage by user.

Type: boolean Default: `true`

A.23.1.11 InstrumentationServerHeartbeat

Controls how often the server will update its lifetime.

Type: duration Default: `30m`

A.24 Load Balancing

The `LoadBalancing` section contains a configuration property that determines whether or not clients must use a version of `rsconnect` with cookie-based sticky session support.

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

A.24.1 LoadBalancing Settings

A.24.1.1 EnforceMinRsconnectVersion

Require a minimum version of `rsconnect` for cookie-based sticky session support.

Type: boolean Default: `false`

A.25 Python

The `Python` section contains configuration properties pertaining to Python support for deployable content that can be optionally enabled in RStudio Connect. See section 16 for further details.

A.25.1 Python Settings

A.25.1.1 Enabled

Python feature toggle.

Type: boolean Default: `false`

A.25.1.2 VersionMatching

Specifies how RStudio Connect attempts to match Python versions associated with uploaded content with the Python versions available on the system. Allows values of `exact` or `major-minor`.

Type: string Default: `major-minor`

A.25.1.3 Executable

Path to a Python executable. Multiple definitions can be used to provide multiple versions of Python.

Type: multi-string Default: *unspecified*

A.25.1.4 ProhibitedPackage

List of prohibited python packages.

Type: multi-string Default: `anaconda-client, anaconda-project, appnope, conda-env, conda, libcxx, libcxxabi, menuinst, mkl-fft, mkl-random, msinttypes, pywin32, pywinpty, rsconnect-jupyter, rsconnect_jupyter, rsconnect, vc-14, vc-9, vs2008_runtime, vs2008_win-32, vs2008_win-64, vs2013_runtime, vs2015_runtime, vs2015_win-32, vs2015_win-64, vs2017_win-64, win-inet-pton, win_inet_pton, win_unicode_console, wincertstore, winkerberos, winpty, xlwings`

B Command-Line Interface

Connect includes tools with a command-line interface (CLI). These tools are typically targeted towards actions that might be performed when the web server is offline or is otherwise inaccessible. Other CLI commands are useful for performing actions against the server in a batch or scripted fashion.

These utilities are installed in `/opt/rstudio-connect/bin/`. They use the configuration defined in `/etc/rstudio-connect/rstudio-connect.gcfg` unless you specify an alternate configuration file with the `--config` flag.

Note: All command flags accept both double (`--`) dashes as indicated here or single (`-`) dashes as indicated in the command built-in help and error messages.

B.1 User Management

This utility helps you list users and modify user attributes such as role, username, name, email, and Unique ID. This can be used to recover if you are unable to access an RStudio Connect administrative account. Similar operations can be done for groups, such as listing and changing the group name, Unique ID or even owner.

Modifications done to users and groups made via `usermanager` are recorded in the audit logs.

Connect's `usermanager` CLI also includes the ability to dump audit logs. By default, the logs are displayed in a formatted table, but you can also choose to output comma-separated values for easy analysis in other tools.

The `usermanager` utility can only be run when Connect is stopped if you use the SQLite database provider. See Section 5.1 for information on stopping and restarting Connect. See Section 9 for information on database providers.

B.1.1 Commands

The `usermanager` utility supports the following commands:

- `list`: Lists users or groups
- `alter`: Alter a user or group
- `transfer`: Transfer user owned items to another user or group relationships to another group
- `delete`: Delete a user or a group
- `tokens`: Lists tokens
- `deactivate`: Deactivates tokens

- **audit:** Dumps audit logs

B.1.2 Flags

Configuration for **usermanager**:

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

Note: The configuration path must be the first flag if not present in the default location.

Flags for the **list** command:

- **--users:** List users (default).
- **--groups:** List groups.
- **--include-locked:** Includes locked user accounts in the list.
- **--detect-unreachable:** Detect LDAP objects (users or groups) no longer visible to Connect.
- **--base64-unique-id:** Detect, decode and attempt to interpret UniqueId values encoded as base64.
- **--force-hex:** When used with **--base64-unique-id** do not attempt to interpret values showing them as hex.
- **--force-ms-guid:** When used with **--base64-unique-id** will attempt to convert GUID/UUID values to the Microsoft format.

Flags for the **alter** command:

- **--users:** Specify operations over users (optional).
- **--groups:** Specify operations over groups (optional).
- **--username:** Specifies the user name of the user to alter.
- **--user-id:** Specifies the user id of the user to alter.
- **--user-guid:** Specifies the user guid of the user to alter.
- **--groupname:** Specifies the group name of the group to alter.
- **--group-id:** Specifies the group id of the group to alter.
- **--group-guid:** Specifies the group guid of the group to alter.
- **--new-owner:** Specifies a new owner to be set on the group (except groups managed by an auth provider). Requires specifying both a user to alter and group to alter.
- **--drop-owner:** Specifies the removal of the group owner (except groups managed by users).
- **--new-groupname:** Specifies the new name for the group (except groups managed by an auth provider).
- **--new-username:** Specifies the new username for the user (only Password, OAuth2, Proxied auth with UniqueID).
- **--new-last-name:** Specifies the new last name for the user (except OAuth2 or managed by Proxied, LDAP, SAML).
- **--new-first-name:** Specifies the new first name for the user (except OAuth2 or managed by Proxied, LDAP, SAML).
- **--new-email:** Specifies the new email address for the user (except OAuth2 or managed by Proxied, LDAP, SAML).
- **--update-ids-using:** Update all Unique IDs for users or groups matching the specified attribute (only LDAP).
- **--new-unique-id:** Specifies a new Unique ID (base64 value or Distinguished Name) for a user or group.
- **--new-role:** Specifies the role to set for the user. Allowed roles are **viewer**, **publisher**, and **administrator**.
- **--lock:** Locks the user. Fails if **--unlock** is also present.
- **--unlock:** Unlocks the user. Fails if **--lock** is also present.
- **--force-demoting:** Force demotion of the last remaining administrator.
- **--yes:** Alter without asking for confirmation. Use with caution.
- **--base64-unique-id:** When used with **--new-unique-id** will encode the input as base64.

- `--force-hex`: When used with `--base64-unique-id` will handle the input as a hex-encoded binary value.
- `--force-ms-guid`: When used with `--force-hex` will first convert GUID/UUID values to the Microsoft format.

Flags for the `transfer` command:

- `--users`: Specify operations over users (optional).
- `--groups`: Specify operations over groups (optional).
- `--source-username`: Specifies the user name of the source user to transfer from.
- `--source-user-id`: Specifies the user id of the source user to transfer from.
- `--source-user-guid`: Specifies the user guid of the source user to transfer from.
- `--target-username`: Specifies the user name of the target user to transfer to.
- `--target-user-id`: Specifies the user id of the target user to transfer to.
- `--target-user-guid`: Specifies the user guid of the target user to transfer to.
- `--source-groupname`: Specifies the group name of the source group to transfer from.
- `--source-group-id`: Specifies the group id of the source group to transfer from.
- `--source-group-guid`: Specifies the group guid of the source group to transfer from.
- `--target-groupname`: Specifies the group name of the target group to transfer to.
- `--target-group-id`: Specifies the group id of the target group to transfer to.
- `--target-group-guid`: Specifies the group guid of the target group to transfer to.
- `--memberships`: Additionally transfer group memberships (except LDAP).
- `--permissions`: Additionally transfer permissions and email subscriptions to content. This is implied for groups.
- `--api-keys`: Additionally transfer API keys.
- `--delete`: Delete the user or group and references to it after the transfer.
- `--yes`: Skip confirmation for both transfer and deletion.
- `--verbose`: Show a detailed list of items affected by the operation before the confirmation.

Flags for the `delete` command:

- `--users`: Specify operations over users (optional).
- `--groups`: Specify operations over groups (optional).
- `--username`: Specifies the user name of the user to delete.
- `--user-id`: Specifies the user id of the user to delete.
- `--user-guid`: Specifies the user guid of the user to delete.
- `--groupname`: Specifies the group name of the group to delete.
- `--group-id`: Specifies the group id of the group to delete.
- `--group-guid`: Specifies the group guid of the group to delete.
- `--yes`: Skip confirmation for deletion.
- `--verbose`: Show a detailed list of items affected by the operation before the confirmation.

Flags for the `token` command:

- `--active`: Show only active tokens.

Flags for the `deactivate` command:

- `--all`: Deactivate all tokens.

Flags for the `audit` command:

- `--csv`: Output comma-separated values.

B.1.3 Examples:

Display help:

```
/opt/rstudio-connect/bin/usermanager help
```

Specify a custom configuration file

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

List unlocked users:

```
sudo /opt/rstudio-connect/bin/usermanager list
```

Or:

```
sudo /opt/rstudio-connect/bin/usermanager list --users
```

Note: This distinguished name field (DN) is included in the results only when LDAP is used.

List all users (locked and unlocked):

```
sudo /opt/rstudio-connect/bin/usermanager list --include-locked
```

Or:

```
sudo /opt/rstudio-connect/bin/usermanager list --users --include-locked
```

Note: The distinguished name field (DN) is included in the results only when LDAP is used.

List all groups:

```
sudo /opt/rstudio-connect/bin/usermanager list --groups
```

Note: The distinguished name field (DN) is included in the results only when LDAP is used.

List all users no longer seen by Connect (LDAP):

```
sudo /opt/rstudio-connect/bin/usermanager list --detect-unreachable
```

Or:

```
sudo /opt/rstudio-connect/bin/usermanager list --users --detect-unreachable
```

List all groups no longer seen by Connect (LDAP):

```
sudo /opt/rstudio-connect/bin/usermanager list --groups --detect-unreachable
```

List users or groups attempting to decode and interpret base64 UniqueIDs:

```
sudo /opt/rstudio-connect/bin/usermanager list --users --base64-unique-id
```

```
sudo /opt/rstudio-connect/bin/usermanager list --groups --base64-unique-id
```

Note: The `--base64-unique-id` option works with `--detect-unreachable` and `--include-locked`.

List users or groups just decoding base64 UniqueIDs and showing values as hex:

```
sudo /opt/rstudio-connect/bin/usermanager list --users --base64-unique-id --force-hex
```

```
sudo /opt/rstudio-connect/bin/usermanager list --groups --base64-unique-id --force-hex
```

List users or groups giving a hint that GUIDs may be in the Microsoft format:

```
sudo /opt/rstudio-connect/bin/usermanager list --users --base64-unique-id --force-ms-guid
```

```
sudo /opt/rstudio-connect/bin/usermanager list --groups --base64-unique-id --force-ms-guid
```

Note: Use this option if the UniqueID is represented by a binary Microsoft GUID (such as objectGUID).

Promote the user `john` to an administrator role

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john --new-role administrator
```

To do the same as above to a user without a username, use `--user-id` to specify the ID of the desired user. IDs can be discovered via `usermanager list`.

```
sudo /opt/rstudio-connect/bin/usermanager alter --user-id 1 --new-role administrator
```

Or with `--user-guid` to specify the GUID of the desired user. GUIDs can be discovered via `usermanager list`.

```
sudo /opt/rstudio-connect/bin/usermanager alter \  
--user-guid 00000000-0000-4000-a000-000000000000 --new-role administrator
```

Note: `--username`, `--user-id` and `--user-guid` can be used interchangeably but not at the same time. Some users may not have names defined in which case only `--user-id` or `--user-guid` will work.

Change the owner of the group `Administrators` to be the user `administrator` (not supported with LDAP groups or automatic group provisioning in Proxied authentication):

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Administrators \  
--new-owner --username admin
```

To do the same as above to a user without a group name, use `--group-id` to specify the ID of the desired group. IDs can be discovered via `usermanager list --groups`.

```
sudo /opt/rstudio-connect/bin/usermanager alter --group-id 1 --new-owner --username admin
```

Or with `--group-guid` to specify the GUID of the desired groups. GUIDs can be discovered via `usermanager list --groups`.

```
sudo /opt/rstudio-connect/bin/usermanager alter \  
--group-guid 00000000-0000-4000-a000-000000000000 --new-owner --username admin
```

Note: `--groupname`, `--group-id` and `--group-guid` can be used interchangeably but not at the same time the same applies to `--username`, `--user-id` and `--user-guid` while specifying the new owner.

Remove the owner of the existing group `Administrator` so it can be used with LDAP or with automatic group provisioning in Proxied authentication:

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Administrators \  
--drop-owner
```

Demote the last remaining administrator to a non-administrative role

```
sudo /opt/rstudio-connect/bin/usermanager alter --username admin --new-role publisher --force-demoting
```

Update all users' Unique IDs currently using DN:

```
sudo /opt/rstudio-connect/bin/usermanager alter --update-ids-using DN
```

Or:

```
sudo /opt/rstudio-connect/bin/usermanager alter --users --update-ids-using DN
```

Use this command after you have configured `LDAP.UniqueIdAttribute` for the first time.

Note: If the current value for the `LDAP.UniqueIdAttribute` setting is "DN" or not present this command does nothing.

Update all groups' Unique IDs currently using DN:

```
sudo /opt/rstudio-connect/bin/usermanager alter --groups --update-ids-using DN
```

Use this command after you have configured `LDAP.GroupUniqueIdAttribute` for the first time.

Note: If the current value for the `LDAP.GroupUniqueIdAttribute` setting is "DN" or not present this command does nothing.

Updating all users' Unique IDs to use another attribute:

```
sudo /opt/rstudio-connect/bin/usermanager alter --update-ids-using < attribute >
```

Or:

```
sudo /opt/rstudio-connect/bin/usermanager alter --users --update-ids-using < attribute >
```

Where `<attribute>` is the case-sensitive name of the LDAP attribute used *previously* in the RStudio Connect configuration for identifying LDAP users such as `objectGUID` (Microsoft AD) or `entryUUID` (OpenLDAP).

Use this command after you have re-configured `LDAP.UniqueIdAttribute` with a different attribute. This command can also be used to revert back to "DN" after removing `LDAP.UniqueIdAttribute`.

Note: If the current value for the `UniqueIdAttribute` setting is the same as `<attribute>` this command does nothing.

Updating all groups' Unique IDs to use another attribute:

```
sudo /opt/rstudio-connect/bin/usermanager alter --groups --update-ids-using < attribute >
```

Where `<attribute>` is the case-sensitive name of the LDAP attribute used *previously* in the RStudio Connect configuration for identifying LDAP groups such as `objectGUID` (Microsoft AD) or `entryUUID` (OpenLDAP).

Use this command after you have re-configured `LDAP.GroupUniqueIdAttribute` with a different attribute. This command can also be used to revert back to "DN" after removing `LDAP.GroupUniqueIdAttribute`.

Note: If the current value for the `GroupUniqueIdAttribute` setting is the same as `<attribute>` this command does nothing.

Change a user's or a group's Unique ID individually:

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john \  
--new-unique-id="CN=John Johnson,OU=Users,DC=example,DC=com"
```

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Admins \  
--new-unique-id="CN=Admins,OU=Groups,DC=example,DC=com"
```

Use this command when the user's or group's DN has changed on the LDAP server and `usermanager --update-ids-using` was not able to reach and recover them.

Change a user's or a group's Unique ID individually, encoding the result as base64:

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john \  
--new-unique-id="f04db6b3-5bab-437b-b46d-18b46d1fd16f" --base64-unique-id
```

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Admins \  
--new-unique-id="f04db6b3-5bab-437b-b46d-18b46d1fd16f" --base64-unique-id
```

Note: Use this option if the UniqueID is represented by a textual value in LDAP (such as `entryUUID`).

Change a user's or a group's Unique ID individually, interpreting hex values as binary then encoding the result as base64:

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john \  
--new-unique-id="0105000000000051500000bed626e0d5fbb2934fb4ee9557040000" \  
--base64-unique-id --force-hex
```

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Admins \  
--new-unique-id="0105000000000051500000bed626e0d5fbb2934fb4ee9557040000" \  
--base64-unique-id --force-hex
```

Note: Use this option if the UniqueID is represented by a binary value in LDAP (such as objectSID).

Change a user's or a group's Unique ID individually, interpreting hex values as binary, forcing the Microsoft representation for GUIDs and then encoding the result as base64:

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john \  
--new-unique-id="f04db6b3-5bab-437b-b46d-18b46d1fd16f" \  
--base64-unique-id --force-hex --force-ms-guid
```

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Admins \  
--new-unique-id="f04db6b3-5bab-437b-b46d-18b46d1fd16f" \  
--base64-unique-id --force-hex --force-ms-guid
```

Note: Use this option if the UniqueID is represented by a binary Microsoft GUID (such as objectGUID).

Change a user's or a group's Unique ID individually with base64 values encoded externally:

```
sudo /opt/rstudio-connect/bin/usermanager alter --username john \  
--new-unique-id="+4FR9WW0+0ecjM20sf30ZA=="
```

```
sudo /opt/rstudio-connect/bin/usermanager alter --groupname Admins \  
--new-unique-id="hYGFGBI987ud+hg7H8dndd=="
```

Use this command when the value for the attribute used as the user's or group's Unique ID has changed on the LDAP server and `usermanager --update-ids-using` was not able to reach and recover them.

Note: An LDIF (LDAP Data Interchange Format) representation of the user object may contain values in the base64 format. This is true for binary values (such as `objectGUID` from Active Directory). On the other hand, string values (such as `entryUUID` from OpenLDAP) must be encoded as base64 before use. Please refer to your LDAP documentation how to obtain an LDIF for a user.

Transfer user-owned groups and content to another user

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane
```

Transfer using ID or GUID to identity source and/or target users

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-id 1 --target-guid 00000000-0000-4000-a000-000000000001
```

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-guid 00000000-0000-4000-a000-000000000001 --target-id 1
```

Transfer user-owned items and group memberships to another user

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane \  
--memberships
```


Transfer user-owned items, API keys, group memberships, content permissions and subscriptions to another user

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane \  
--memberships --permissions --api-keys
```

Note: Transferring API keys is not a recommended practice and it should be used only when it's extremely difficult to replace the API key on the application consuming the Connect Server API. This is a convenience tool with security implications that should be considered carefully. After this operation the audit logs will record that the target user, not the source user, is now the one responsible for the operations done using the transferred key.

Transfer user-owned items and content permissions and subscriptions to another user and then delete the source user

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane \  
--permissions --delete
```

Transfer without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane \  
--yes
```

Transfer and delete source user without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-username john --target-username jane \  
--delete --yes
```

Delete user

```
sudo /opt/rstudio-connect/bin/usermanager delete --username john
```

Delete user by ID

```
sudo /opt/rstudio-connect/bin/usermanager delete --user-id 1
```

Delete user by GUID

```
sudo /opt/rstudio-connect/bin/usermanager delete --user-guid 00000000-0000-4000-a000-000000000000
```

Delete user without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager delete --username john --yes
```

Transfer group content permissions to another group

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-groupname john --target-groupname jane
```

Transfer using ID or GUID to identify source and/or target group

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-id 1 --target-guid 00000000-0000-4000-a000-
```

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-guid 00000000-0000-4000-a000-000000000001 --
```

Transfer group content permissions and members to another group

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-groupname viewers --target-groupname admins \  
--memberships
```

Transfer group content permissions and members to another group and then delete the source group

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-groupname viewers --target-groupname admins \  
--memberships --delete
```

Transfer without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-groupname viewers --target-groupname admins --yes
```

Transfer and delete source group without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager transfer --source-groupname viewers --target-groupname admins --delete --yes
```

Delete group

```
sudo /opt/rstudio-connect/bin/usermanager delete --groupname viewers
```

Delete group by ID

```
sudo /opt/rstudio-connect/bin/usermanager delete --group-id 1
```

Delete group by GUID

```
sudo /opt/rstudio-connect/bin/usermanager delete --group-guid 00000000-0000-4000-a000-000000000000
```

Delete group without confirmation

```
sudo /opt/rstudio-connect/bin/usermanager delete --groupname viewers --yes
```

Dump audit logs to screen

```
sudo /opt/rstudio-connect/bin/usermanager audit
```

Dump audit logs (comma-separated) to a file:

```
sudo /opt/rstudio-connect/bin/usermanager audit --csv > ~/audits.txt
```

B.2 Migration Utility

The `migrate` utility assists system administrators in migrating from one database to another or in transitioning RStudio Connect to a new server. For a high-level overview of the steps necessary to migrate from SQLite to Postgres, see the section on changing database providers. For the high-level steps involved in completing a server migration, see 4.8.

B.2.1 Commands

The `migrate` utility supports four commands

- `db`: Migrate data between databases
- `rebuild-packrat`: Rebuilds the Packrat cache for all content on the server. This command can be used WHILE RStudio Connect is running.
- `rebuild`: alias for `rebuild-packrat`.
- `repair-content-permissions`: Checks and corrects permissions and ownership for the working directories of each deployed piece of content.
- `help`: Displays help

B.2.2 Flags

Configuration for migrate:

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

Flags for the `migrate db` command:

- `--verify`: Verify migration only.
- `--drop-all`: Drop all existing data in the target before migrating.

By default, the `migrate db` command will copy the data from the SQLite database into PostgreSQL, and verify the migration. We assume that the PostgreSQL destination does not contain any data unless the `--drop-all` flag is included.

Data migration copies data from the SQLite database to the PostgreSQL database. Data in the SQLite database remains after the migration; it is not removed. A verification step runs after the data copy completes and confirms the integrity of the migration:

- Row counts for all tables are verified.
- Each record is checked for the correct values.

Data verification will fail if Connect is started prior to the completion of data verification. Please ensure that Connect remains down until the data migration and verification are complete.

Flags for the `rebuild-packrat` command:

- `-force`: Delete the Packrat cache before rebuilding
- `-fast-fail`: Stop when the packages for a single application cannot be installed.

Proactively rebuilds the Packrat cache for all applications on the server. When the `-force` flag is used, the entire existing Packrat cache directory will be deleted first. This command can be used for instances in which the Packrat cache may be incomplete for the current environment. For example, if the system only has one version of R installed and it has been upgraded, the cache will not include packages built on the appropriate version of R. Similarly, if you migrate your RStudio Connect installation to a different server which might have different versions of system libraries, you should delete the cache and rebuild it as discussed in 4.8.

When the `-fast-fail` flag is used, rebuilding the Packrat package cache is halted when the packages for any application cannot be installed or verified as installed.

Flags for the `repair-content-permissions` command:

No flags supported.

Scans for issues with the permissions and ownership of application directories on the server. This command can be used if you have moved some data on disk and need to confirm that the attributes were transferred properly.

B.2.3 Examples

Display help:

```
/opt/rstudio-connect/bin/migrate help
```

Migrate SQLite data to an empty PostgreSQL database:

```
sudo /opt/rstudio-connect/bin/migrate db
```

Migrate SQLite data to a PostgreSQL database, first dropping all data in the PostgreSQL database:

```
sudo /opt/rstudio-connect/bin/migrate db --drop-all
```

Perform data verification only:

```
sudo /opt/rstudio-connect/bin/migrate db --verify
```

Specify a custom configuration file:

```
sudo /opt/rstudio-connect/bin/migrate --config /etc/connect/mycustomconfig.gcfg db
```

Delete the existing Packrat cache and rebuild it by pro-actively rebuilding each application's library.

```
sudo /opt/rstudio-connect/bin/migrate rebuild-packrat --force
```

Check and fix any disk permission errors for applications' working directories.

```
sudo /opt/rstudio-connect/bin/migrate repair-content-permissions
```

C Using a Custom Landing Page

C.1 Overview

It is possible to specify a custom landing page that your anonymous/logged-out users will see when they visit Connect.

C.2 Configuration

Use the `Server.LandingDir` configuration setting to specify the path to a custom landing page. If you do not specify an absolute path, the server will resolve the path starting at your Connect server installation directory (probably `/opt/rstudio-connect`).

Please see A for more information on the `Server.LandingDir` setting.

C.3 Custom Landing Page Assets

Include all assets (JavaScript, CSS, images, etc.) for your custom landing page in the directory you specified in the `Server.LandingDir` configuration setting. Be sure to include an `index.html`, which will be served by default.

C.4 Example

See the `/opt/rstudio-connect/examples/landing-page` directory for an example custom landing page. You can enable this example landing page by adding the following configuration setting and restarting the Connect server.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[Server]
LandingDir = examples/landing-page
```

D LDAP/AD Configuration Examples

This section contains sample RStudio Connect configurations to help you get started with LDAP authentication. We have provided a single bind and a double bind example (double bind is recommended).

The LDIF (LDAP Data Interchange Format) file contained in D.3 describes a LDAP organization used in our examples.

D.1 Single Bind

Here is a partial RStudio Connect configuration file showing how to connect using single-bind LDAP authentication. We are assuming the LDIF contained in D.3 describes the LDAP structure.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "myLDAPserverSingle"]
ServerAddress = 127.0.0.1:389
UniqueIdAttribute = entryUUID
UserSearchBaseDN = "ou=People,dc=company,dc=com"
UserObjectClass = posixAccount
UserFirstNameAttribute = givenName
UserLastNameAttribute = sn
UserEmailAttribute = mail
UsernameAttribute = uid
```

D.2 Double Bind

Here is a partial RStudio Connect configuration file showing how to connect using double bind LDAP authentication. We are assuming the LDIF contained in D.3 describes the LDAP structure.

```
; /etc/rstudio-connect/rstudio-connect.gcfg
[LDAP "myLDAPserver"]
ServerAddress = 127.0.0.1:389
BindDN = cn=admin,dc=company,dc=com"
BindPassword = "password"
UniqueIdAttribute = entryUUID
UserSearchBaseDN = "ou=People,dc=company,dc=com"
UserObjectClass = posixAccount
UserFirstNameAttribute = givenName
UserLastNameAttribute = sn
UserEmailAttribute = mail
UsernameAttribute = uid
```

D.3 LDIF

Here is an LDIF (LDAP Data Interchange Format) file describing a hypothetical organization.

```
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
```

E RStudio Connect Deployment Guide

E.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.

E.2 Programmatic Deployment

To programmatically publish content to RStudio Connect, use the functions `deployDoc`, `deployApp`, `deployAPI`, 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.

E.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.

E.3.1 Application Metadata

`rsconnect` infers a number of attributes about the content including:

1. `appMode`: static, shiny, rmd-static, rmd-shiny, api
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).

E.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 and Plumber APIs, `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

E.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.

E.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.

E.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.

E.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.

E.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 deploy and activate the new

bundle for this application. You may use the “source versions” menu option in the dashboard to revert the application to a previous 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`.

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.

E.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:

E.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. More information is available in Section 15.4. Files are checked against the checksum listed in the manifest to ensure content was not lost or corrupted during transfer.

E.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 `Applications.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.

E.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 tuning runtime settings, permissions, etc).

E.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. More information can be found in Section 15.4.

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")`.

F Using Continuous Integration to Deploy Content

F.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.

RStudio Connect includes experimental HTTP APIs that let you build your own deployment workflows. Those APIs will replace the `rsconnect` package as the preferred way to perform programmatic deployment.

The RStudio Connect API Reference describes all of the content APIs you must use to perform deployments. The RStudio Connect User Guide includes deployment recipes that will help get you started.

F.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 E for more information.

F.3 Configuring a CI Server to Deploy Content to Connect

F.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")
```

F.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=.....
```

```
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
```

F.3.3 Deploying Content with `rsconnect`

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

F.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.

F.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.

F.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`.

G Programmatic Deployment with `rsconnect`

G.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.

RStudio Connect includes experimental HTTP APIs that let you build your own deployment workflows. Those APIs will replace the `rsconnect` package as the preferred way to perform programmatic deployment.

The RStudio Connect API Reference describes all of the content APIs you must use to perform deployments. The RStudio Connect User Guide includes deployment recipes that will help get you started.

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.

G.2 Use Case: A Shiny Application

Here we present a use case that explains how to configure Connect for programmatic deployment. Please see G.4 for an example Shiny application for this use case.

G.2.1 Use Case Scenario

Bob White develops a Shiny application (see G.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`.

G.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")
```

G.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 13.5 for configuring the `RunAs` user on a per-application basis in Connect.

G.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=.....`

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
```

G.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`.

G.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")
```

G.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.

G.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)

```