

FreeRDP Developer Manual

Marc-André Moreau

Awake Coding Consulting Inc.

Contents

1	Introduction	5
1.1	Glossary	5
1.2	References	5
1.3	Getting Started	5
1.3.1	Linux	5
	Debian / Ubuntu	5
	(Open)SUSE	6
	CentOS / REHL / Fedora	6
2	Build Environment	8
2.1	CMake Build System	8
2.1.1	Installation	8
	Linux	8
	Mac OS X	9
	Windows	9
2.2	Git Version Control	9
2.2.1	Installation	9
2.2.2	Configuration	10
2.3	Native Compilation	10
2.3.1	Linux	10
2.3.2	Mac OS X	10
	SDKs	11
2.3.3	Windows	11

2.4	Cross-Compilation	11
2.4.1	Android	11
	Java Development Kit	12
	Android SDK and NDK	12
	Environment Variables	12
	Preparing Toolchain	13
	Invoking CMake	13
	Android Project	13
	Ant Building	14
	Eclipse Building	14
2.4.2	iOS	15
	iOS SDKs	15
	Invoking CMake	15
2.5	Prerequisites	15
2.5.1	OpenSSL	16
	Linux	16
	Mac OS X	16
	Windows	16
	Android	17
	iOS	19
2.5.2	X11	20
	Linux	20
	Mac OS X	20
	Windows	20
3	Build Configuration	21
3.1	CMake	21
3.1.1	Introduction	21
3.1.2	Generators	22
3.1.3	CMake Cache	22
3.1.4	Configuration Files	23

3.1.5	Out-of-Source Build	24
3.2	Options	24
3.2.1	CMake Options	24
3.2.2	Build Options	25
3.2.3	Mac OS X Options	25
3.2.4	Optimization Options	25
3.2.5	Client Options	25
3.2.6	Server Options	26
3.2.7	Channel Options	26
3.2.8	Feature Options	26
4	Testing	27
4.0.9	CTest	27
	Introduction	27
5	FreeRDS	28
5.1	Getting Started	28
5.2	Building	29
5.3	X11rdp	29
5.3.1	FreeRDS	29
5.4	Running	30
6	API Reference	31
6.1	Input	31
6.1.1	Concepts	31
6.1.2	Keyboard Mapping	33
6.1.3	Synchronize Event	35
6.1.4	Keyboard Event	36
6.1.5	Keyboard Unicode Event	36
6.1.6	Mouse Event	36
6.1.7	Extended Mouse Event	37
6.2	Virtual Channel API	38

6.2.1	VirtualChannelEntry	38
6.2.2	VirtualChannelInit	38
6.2.3	VirtualChannelInitEvent	39
6.2.4	VirtualChannelOpen	39
6.2.5	VirtualChannelOpenEvent	39
6.2.6	VirtualChannelWrite	40
6.2.7	VirtualChannelClose	40

Chapter 1

Introduction

1.1 Glossary

To be expanded.

1.2 References

FreeRDP User Manual¹

FreeRDP Configuration Manual²

FreeRDP Testing Manual³

1.3 Getting Started

1.3.1 Linux

Debian / Ubuntu

FreeRDP:

¹<https://github.com/awakecoding/FreeRDP-Manuals/blob/master/User/FreeRDP-User-Manual.pdf?raw=true>

²<https://github.com/awakecoding/FreeRDP-Manuals/blob/master/Configuration/FreeRDP-Configuration-Manual.pdf?raw=true>

³<https://github.com/awakecoding/FreeRDP-Manuals/blob/master/Testing/FreeRDP-Testing-Manual.pdf?raw=true>

```
sudo apt-get install \  
build-essential git-core cmake \  
libssl-dev \  
libx11-dev libxext-dev libxinerama-dev libxcursor-dev libxkbfile-dev \  
libxv-dev libxi-dev libxdamage-dev libxrender-dev libxrandr-dev \  
libasound2-dev libcups2-dev libpulse-dev \  
libavutil-dev libavcodec-dev \  
libgstreamer0.10-dev libgstreamer-plugins-base0.10-dev
```

FreeRDS:

```
sudo apt-get install \  
libpciaccess-dev libpam0g-dev libpng12-dev libjpeg-dev intltool libexpat1-dev libxml-libxml-  
libtool bison flex xsltproc libfreetype6-dev libfontconfig1-dev libpixman-1-dev xutils-dev
```

(Open)SUSE

FreeRDP:

```
sudo zypper install \  
gcc git cmake \  
libopenssl-devel \  
libX11-devel libXext-devel libXinerama-devel libXcursor-devel libxkbfile-devel \  
libXv-devel libXi-devel libXdamage-devel libXrender-devel libXrandr-devel \  
libpulse-devel alsa-devel cups-devel \  
ffmpeg-devel libavutil-devel \  
gstreamer-0_10-devel gstreamer-0_10-plugins-base-devel
```

ffmpeg and gstreamer development libraries are available in the Restricted Formats⁴ community repository.

FreeRDS:

```
sudo zypper install \  
autoconf automake libtool bison flex libxslt-tools gcc-c++ llvm \  
libpciaccess-devel pam-devel libpng-devel libjpeg-devel intltool libexpat-devel \  
libpixman-1-0-devel perl-libxml-perl freetype2-devel fontconfig-devel \  
protobuf-c protobuf-devel boost-devel
```

CentOS / RHEL / Fedora

FreeRDP:

⁴http://opensuse-community.org/Restricted_formats

```
sudo yum install \  
openssl-devel \  
libX11-devel libXext-devel libXinerama-devel libXcursor-devel libxkbfile-devel \  
libXv-devel libXtst-devel libXi-devel libXdamage-devel libXrandr-devel \  
alsa-lib-devel cups-devel ffmpeg-devel
```

The version of cmake available in the CentOS repositories is too old, use a more recent version:

```
wget http://pkgs.repoforge.org/cmake/cmake-2.8.8-1.el6.rfx.x86_64.rpm  
sudo rpm -i cmake-2.8.8-1.el6.rfx.x86_64.rpm
```

FreeRDS:

```
sudo yum install \  
finger patch gcc gcc-c++ make autoconf libtool automake pkgconfig \  
openssl-devel gettext file pam-devel libX11-devel libXfixes-devel libjpeg-devel \  
flex bison libxslt perl-libxml-perl xorg-x11-font-utils xmlto-tex docbook-utils-pdf
```

Chapter 2

Build Environment

2.1 CMake Build System

FreeRDP uses CMake as its build system. CMake is not an ordinary build system: it generates project files for other build systems to use. This means that we have only one set of scripts to maintain while having the flexibility of using any of the build systems supported by CMake. For instance, FreeRDP developers can use Visual Studio, Xcode, Eclipse, or just plain regular makefiles without having to separately maintain project files for the development tools they are using.

2.1.1 Installation

FreeRDP requires CMake 2.8, but it is highly recommended to use the very latest builds available from the CMake website (www.cmake.org). This is because certain useful features are available only in certain versions that we cannot require since they are too recent. Using an older version of CMake 2.8 will work, but some features taking advantage of functionality present in more recent CMake releases may be disabled. For instance, FreeRDP offers a monolithic build option that combines all libraries into a single one using functionality found in CMake 2.8.9. If you are using CMake 2.8.8, this option will be disabled.

Recent CMake builds are available from the CMake download page:

<http://www.cmake.org/cmake/resources/software.html>

Linux

CMake is packaged by most Linux distributions and should be easy to install. However, since FreeRDP leverages some of the newest CMake features, it is very

likely that the version packaged for your system is not the latest one. In this case, simply download the CMake Linux build from the CMake website and install it manually to the appropriate location.

Mac OS X

The easiest way to install CMake on Mac OS X is by using the latest build available on the CMake website. Alternatively, one can use common Mac OS X package manager such as homebrew to install CMake.

Windows

Installing CMake on Windows is fairly straightforward. Download either the installer or the zip version of the Windows build from the CMake website. Using the installer is simpler, but you might want to consider choosing the zip archive if you need to work from a machine where you do not have sufficient rights to run an installer.

2.2 Git Version Control

FreeRDP is hosted on GitHub, and uses git for version control. Git is one of those tools that is extremely powerful but a bit hard to grasp at first. Luckily, a lot of the complexity of git is made much simpler through the GitHub collaborative coding interface, and provides us with a functional workflow for accepting contributions with minimal overhead.

If you are new to FreeRDP and don't want to bother learning the git basics for now, feel free to simply download a source tarball from GitHub and move on to the next step. However, you will eventually have to learn it, especially if you want to synchronize your work with the latest sources, or contribute your changes back to the community. If you are not so sure about how to contribute changes which have not been done cleanly with git, ping one of the developers in the IRC channel. It's never fun to get contributions in the form of tarball in an email, but we'll take it anyway. Please understand, however, that since this requires much more effort on our side, clean integration of the changes will only happen when time allows for it.

2.2.1 Installation

Git is packaged for most Linux distributions and should be easily obtained. Otherwise, and for all other platforms, pre-built packages are available from the Git download page (<http://git-scm.com/downloads>)

Many other third-party Git packages exist for various operating systems. On Windows, git can be installed in Cygwin <http://www.cygwin.com/>, or through msysgit <http://msysgit.github.com/>

2.2.2 Configuration

Configuring Git can be tricky, especially for the part where one has to generate SSH keys to be used for authentication with GitHub instead of passwords.

One does not necessarily need to use GitHub, but unless you know what you are doing you should simply create a GitHub account and follow the regular GitHub workflow.

GitHub offers extensive help which everyone should refer to: <https://help.github.com/>

Alternatively, the Pro Git book is available for free online: <http://git-scm.com/book>

The Git Reference is concise, small and to the point: <http://gitref.org/>

Along with the GitHub help, there is an excellent step-by-step introduction to GitHub on the GitHub Learn website: <http://learn.github.com/p/setup.html>

If you need more assistance getting started, please come on IRC and we'll help you.

2.3 Native Compilation

2.3.1 Linux

In debian-based distributions, the build-essential package installs GCC and common development tools necessary for C/C++ development.

2.3.2 Mac OS X

Xcode is the easiest way develop for Mac OS X, and is distributed for free on the Mac App Store.

The Xcode command line tools are not installed by default with Xcode and are necessary:

- On the Xcode menu, click Xcode and then Preferences
- Select the Downloads tab from the Preferences panel
- Select the Components sub-tab from the Downloads tab
- Select Command line tools and click install on the right

SDKs

Whenever Xcode is updated through the Mac App Store, it tends to remove SDKs for older versions of Mac OS X to only keep the new ones. In order to provide better backwards compatibility, one can download older releases of Xcode and manually install the missing SDKs in newer versions of Xcode. This process has to be repeated whenever Xcode is updated.

- Download the last release of Xcode containing the SDK you want and mount it.
- Right-click the Xcode package, select “Show Package Contents”.
- Browse to Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs.
- Save a copy of the desired SDK like “MacOSX10.6.sdk” for later.
- Install Mac OS X SDKs to /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/D

Mac OS X 10.6 Snow Leopard SDK: Xcode 4.3.3

Mac OS X 10.7 Lion SDK: Xcode 4.6.3

2.3.3 Windows

FreeRDP supports all versions of Windows starting from Windows XP all the way up to Windows 8. Windows development is normally done using Visual Studio 2010 or Visual Studio 2012. If you do not own a license of Visual Studio, the Express edition is available for free and should work fine.

To build FreeRDP on Windows, you will first need to install OpenSSL (see the “Prerequisites” section), along with the rest of the Windows build environment, including CMake. When installing CMake, it is recommended to select the option that adds it to the system path.

Refer to the “Build Configuration” section to learn the basic usage of CMake. For Windows, common CMake generators are “Visual Studio 2010” and “Visual Studio 2012”. There is also the “NMake Makefiles” generator which may be suitable for more advanced Windows developers who would prefer a build system that does not require the full Visual Studio IDE.

2.4 Cross-Compilation

2.4.1 Android

Android development can be done from Linux, Mac OS X and Windows. There are two development kits for Android: the Software Development Kit (SDK) and the Native Development Kit (NDK). The SDK is for Java, while the NDK is for

C/C++. Since it is not possible to write completely native Android applications (at least not under normal circumstances), we will need both the Android SDK and NDK for FreeRDP development. The NDK alone is enough to build the FreeRDP libraries, but not enough to build a complete application.

Java Development Kit

Before installing the SDK, you will need to install the Java Development Kit. On Windows, the Sun/Oracle JDK is easy to download and install. On Linux, OpenJDK is easier to install from your distribution's package manager.

Android SDK and NDK

Download the Android SDK and Android NDK from the Android developer website:

<http://developer.android.com/tools/>

Extract or install the SDK and NDK to a path which is easy to locate. The Windows installer installs the SDK into “%PROGRAMFILES(x86)%/Android” by default. On Linux, using /opt/android-sdk and /opt/android-ndk is common practice. Regardless of where you choose to install the SDK and NDK, you will need to configure the ANDROID_SDK and ANDROID_NDK environment variables to point to them. It is also very convenient to add *ANDROID_SDK/tools* and *ANDROID_SDK/platform-tools* to your system path in order to have the Android utilities available without specifying the full path.

Once both the SDK and NDK are installed, the first thing to do (it takes some time) is to launch the “android” utility found in \$ANDROID_SDK/tools. On Windows, the same utility can be launched with the “SDK Manager” found in %ANDROID_SDK%. In the SDK manager, select target Android platforms of interest, such as Android 2.3.3 (API 10) or Android 4.0 (API 14), click install, accept the licenses, and be patient because it takes time to download and install everything. The good news is that you can leave the SDK manager open while moving on to setting up the NDK.

Environment Variables

You will need to set a certain number of environment variables to ease Android development. On Linux, this can be done by editing the .bashrc file, and on Mac OS X, this can be done by editing the .profile file.

```
export NDK=/opt/android-ndk
export ANDROID_NDK=$NDK
```

```
export ANDROID_SDK=/opt/android-sdk
export PATH=$ANDROID_SDK/tools:$ANDROID_SDK/platform-tools:$PATH
```

Preparing Toolchain

The prebuilt toolchains that come with the Android NDK should suffice, but if you need to build your own, invoke the `make-standalone-toolchain.sh` script:

```
$NDK/build/tools/make-standalone-toolchain.sh --platform=android-8 --install-dir=$ANDROID_ST
```

Where `$ANDROID_STANDALONE_TOOLCHAIN` is the location where you want to install your standalone toolchain.

Invoking CMake

When invoking CMake, the path to the Android CMake toolchain needs to be passed:

```
cmake -DCMAKE_TOOLCHAIN_FILE=cmake/AndroidToolchain.cmake .
```

The Android CMake toolchain will make use of the environment variables defined in the previous section, so make sure they are correct before invoking CMake.

The Android CMake toolchain included with FreeRDP is based on the `android-cmake` project:

<https://code.google.com/p/android-cmake/>

Android Project

While the native portion of the FreeRDP Android port is built using CMake, the Android Java code is built using `ant` or `Eclipse` with the ADT plugin. The Android project structure is explained here:

<http://developer.android.com/tools/projects/>

This project is generated and maintained using the android command-line tool:

<http://developer.android.com/tools/projects/projects-cmdline.html>

The root of the Android project tree is located in `client/Android`:

src: android Java code (built with `ant` or `Eclipse`)

bin: build output directory (`.apk` files)

jni: android native code (built with `cmake`)

gen: android generated code such as R.java

assets: project assets, added as-is to the .apk

res: contains XML configuration files for menus, layouts, strings, etc

libs: contains third-party libraries for both Java and native code

Since the Android native code is built with cmake from the top-level directory, there is a small issue where the library output path does not match the one from the Android project. CMake outputs libfreerdp-android.so to libs/armeabi-v7a, but the Android project expects them in client/Android/libs/armeabi-v7a. You can either copy libfreerdp-android.so manually every time, or create a symlink to work around the issue:

```
cd client/Android/libs
ln -s ../../../../libs/armeabi-v7a/ armeabi-v7a
```

In order to verify that libfreerdp-android.so gets properly packaged in the .apk, simply unzip the resulting .apk files to see its contents (an .apk file is a regular zip file with a different extension).

Ant Building

Inside client/Android, invoke ant help to list targets:

```
ant help
```

Targets of interest are debug and release, which generate a debug or release .apk file in the bin directory:

```
ant debug
ant release
```

The resulting .apk files can be deployed to actual devices or the Android emulator.

Eclipse Building

There is an Eclipse project located in client/Android, which is different from any potential Eclipse project generated by CMake in the root of the FreeRDP source tree. It can be imported in Eclipse like any other regular Eclipse project. Make sure to configure the Eclipse ADT plugin properly beforehand.

2.4.2 iOS

iOS development requires Xcode and the same development environment as for Mac OS X with the addition of the iOS SDK. Certain features like deploying to real devices are limited to Apple developers only.

iOS SDKs

Just like for Mac OS X SDKs, older versions of the iOS SDKs disappear when you update to a newer version of Xcode.

- Download the last release of Xcode containing the SDK you want and mount it.
- Right-click the Xcode package, select “Show Package Contents”.
- Browse to Contents/Developer/Platforms/iPhoneOS.platform/Developer/SDKs.
- Save a copy of the desired SDK like “iPhoneOS6.1.sdk” for later.
- Install iOS SDKs to /Applications/Xcode.app/Contents/Developer/Platforms/iPhoneOS.platform/Developer/

iOS 6.1 SDK: Xcode 4.6.3

Invoking CMake

FreeRDP includes a modified version of the ios-cmake toolchain in the cmake directory. The original toolchain can be found here:

<http://code.google.com/p/ios-cmake/>

When invoking CMake for iOS development, the iOS toolchain needs to be passed like this:

```
cmake -DCMAKE_TOOLCHAIN_FILE=cmake/iOSToolchain.cmake .
```

2.5 Prerequisites

The general approach towards prerequisites is simple: avoid hard dependencies as much as possible, make soft dependencies modular and optional. This may appear a bit extreme, but it is not when FreeRDP needs to run on a very large variety of platforms with a small footprint. Convenience always comes at a cost, and we have learned from our past mistakes that in our case the extra effort is worth it.

2.5.1 OpenSSL

FreeRDP's only strong dependency is OpenSSL. All other dependencies can be made optional with the side effect of disabling the modules that depend on it. It is within our plans to make FreeRDP less dependent on OpenSSL in the future through the WinPR project: instead of using OpenSSL directly, we will implement the equivalent portions of what we need from the Windows API. The WinPR implementation of the Windows API will make use of OpenSSL and eventually add support for other cryptographic libraries. Not only will this enable OpenSSL-free builds of FreeRDP on Windows, but it will also give us an abstraction layer over OpenSSL, allowing us to add support for alternative libraries such as NSS, PolarSSL and CyaSSL.

OpenSSL is used for TLS, MD4, MD5, HMAC-MD5, SHA1, HMAC-SHA1, RSA, DES3, Base64, BigNum arithmetic and certificate validation. RDP8 support will require DTLS which is also provided by OpenSSL.

Linux

OpenSSL is packaged by most Linux distributions, and is easy to compile from source.

Mac OS X

OpenSSL used to ship with Mac OS X until Mac OS X 10.8 (Mountain Lion). Luckily, OpenSSL ships with Xcode. Otherwise, OpenSSL can be obtained from common Mac OS X package managers such as macports.

Command-line tools are not installed by default with Xcode and are necessary.

Windows

A binary distribution of OpenSSL is maintained by “The Shining Light Productions”:

<http://slproweb.com/products/Win32OpenSSL.html>

In order to use it, download both the Win32 OpenSSL installer (the full one, not the light) and the Visual C++ 2008 Redistributables. Install the Visual C++ redistributables first, since OpenSSL requires it. Install OpenSSL in the default path suggested by the installer (C:\OpenSSL-Win32), otherwise may not find it automatically. If you install OpenSSL in a path not found by CMake, you will need to use the `OPENSSL_ROOT_DIR` option to tell CMake where to find it.

Compilation The Visual C++ 2008 Redistributable requirement is a problem if you want to ship FreeRDP for Windows without an installer. Unfortunately, this requires building OpenSSL manually and modifying the default build scripts to link against the static Visual Studio runtime (/MT switch as opposed to the default /MD switch used by OpenSSL). There are many pitfalls along the way if you want to make a proper build of FreeRDP and OpenSSL using the static Visual Studio runtime.

A modified version of OpenSSL for Windows is available for the sake of convenience here:

<https://github.com/awakecoding/openssl-win32>

The sources were deliberately put on git to allow others to consult the history of modifications from the vanilla sources that allow for an easier Windows build.

In order to build OpenSSL, you will need to install the following prerequisites

ActivePerl (<http://www.activestate.com/activeperl>) NASM for Windows (<http://www.nasm.us/>) Microsoft Visual Studio (2010 or 2012) Raspberry Perl can be used instead of ActivePerl, and other versions of Visual Studio can be used, but the build scripts provided will need to be modified accordingly.

Once prerequisites are installed, download the modified OpenSSL sources from the openssl-win32 GitHub repository. In the root of the openssl-win32 directory, you should find two batch files: win32_env.bat and win32_build.bat. Edit both of them to ensure that the path variables they use are correct for your environment.

Once the batch files are properly edited, open a command prompt and move to the root of the openssl-win32 directory. Run win32_env.bat once in order to configure environment variables. Run win32_build.bat in order to build OpenSSL. When building is complete, you should find the resulting build in the “build” directory. This directory can be renamed to OpenSSL-Win32 and moved to C:\OpenSSL-Win32 to be automatically picked up by CMake. Beware, however, that this build has only been tested at this point with FreeRDP release builds configured with the MSVC_RUNTIME=”static” option.

If your FreeRDP builds start crashing for no obvious reason when freeing memory, this is most likely the side effect of an MSVC runtime mismatch. Don’t waste your time trying to find problems where they are not: even if you think you’ve compiled everything correctly chances are that you’re fighting with the pickiness of OpenSSL builds on Windows.

Android

Even though OpenSSL is used in Android, the headers and libraries are not part the Android NDK. This means that in order to compile FreeRDP for Android, one must first compile OpenSSL manually.

Android uses its own build system which is not supported by OpenSSL out of the box. A modified version of OpenSSL with Android build scripts is available through the Android Open Source Project (AOSP). This is the library which is used for the Android operating system itself, but not made available to application developers through the NDK.

This may look like a good deal, but further modifications to default OpenSSL Android port are required before it can be built for usage with FreeRDP, such as adding an option for a static build.

Usable OpenSSL sources are available from “The Guardian Project”:
<https://github.com/guardianproject/android-external-openssl-ndk-static>

Further modifications to the OpenSSL sources have been made by the FreeRDP developers and can be found here:

```
git://github.com/bmiklautz/android-external-openssl-ndk-static.git
git://github.com/awakecoding/android-external-openssl-ndk-static.git
```

Clone the OpenSSL git repository in external/openssl in the root of the FreeRDP source tree:

```
cd external
git clone git://github.com/awakecoding/android-external-openssl-ndk-static.git openssl
```

Ensure you have the proper environment variables configured for Android development as covered earlier. Move to the root of the OpenSSL source tree use the following command to build:

```
$NDK/ndk-build
```

You can optionally install OpenSSL globally by copying the headers and library files to your target android platform directory:

```
export ANDROID_PLATFORM=android-8
cp -R include/ $NDK/platforms/$ANDROID_PLATFORM/arch-arm/usr/include/
cp obj/local/armeabi/*.a $NDK/platforms/$ANDROID_PLATFORM/arch-arm/usr/lib/
```

If OpenSSL is not installed globally, you will need to manually specify its path when invoking cmake with `FREERDP_ANDROID_EXTERNAL_SSL_PATH`.

```
cmake -DCMAKE_TOOLCHAIN_FILE=cmake/AndroidToolchain.cmake -DANDROID_NDK=/opt/android-ndk -D
```

If `ANDROID_SDK` and `ANDROID_NDK` are set as environment variables, they will be used by cmake, reducing the size of the command. If `openssl` is present in `external/openssl`, it will be used automatically by cmake:

```
export ANDROID_SDK=/opt/android-sdk
export ANDROID_NDK=/opt/android-ndk
cmake -DCMAKE_TOOLCHAIN_FILE=cmake/AndroidToolchain.cmake .
```

iOS

There is no pre-built OpenSSL iOS package, so again we have to build our own. The “iOSPorts” collection of iOS ports simplifies the task of building OpenSSL for iOS.

Download the iOSPorts collection from: <https://github.com/bindle/iOSPorts>

Launch Terminal and move to the root of the iOSPorts source tree. From there, type the following commands:

```
cd ports/security/openssl/
make
cd openssl
../build-aux/configure-wrapper.sh
make install
```

This will download, patch, configure build and install OpenSSL for iOS. The default installation path is in /tmp/iOSPorts. The important files are the include and lib directories from the OpenSSL, which we will need to copy to a location inside the iOS SDK.

The system root for the iOS SDK can be found in the following path:

```
/Applications/Xcode.app/Contents/Developer/Platforms/iPhoneOS.platform/Developer/SDKs/iPhoneOS7.0.sdk
```

Where “iPhoneOS7.0.sdk” is subject to change depending on the SDK version.

To ease installation, configure the following environment variables:

```
export IOS_SDK_ROOT=/Applications/Xcode.app/Contents/Developer/Platforms/iPhoneOS.platform/Devel
export OPENSSL_ROOT=/tmp/iOSPorts
```

Then enter the following commands to install OpenSSL in the iOS SDK:

```
sudo cp -R "$OPENSSL_ROOT/include" "$IOS_SDK_ROOT/usr"
sudo cp "$OPENSSL_ROOT/lib/libssl.a" "$IOS_SDK_ROOT/usr/lib/libssl.a"
sudo cp "$OPENSSL_ROOT/lib/libcrypto.a" "$IOS_SDK_ROOT/usr/lib/libcrypto.a"
```

This procedure only has to be executed once per iOS SDK. After this, OpenSSL should be picked up automatically by CMake.

2.5.2 X11

Linux

For the X11 client, you will need:

```
sudo apt-get install libx11-dev libxcursor-dev libxkbfile-dev libxinerama-dev
```

Additionally, for the X11 server, you will need:

```
sudo apt-get install libxext-dev libxtst-dev libxdamage-dev
```

Mac OS X

If the X11 SDK used to ship with XCode, it was removed from its newer versions. In order to get the X11 dependencies satisfied, one has to obtain it from the XQuartz open source project:

<http://xquartz.macosforge.org/>

Windows

Even if a native Windows port is available, it is possible to compile the X11 FreeRDP client on Windows using Cygwin. Please note, however, that this particular configuration is rarely tested.

Chapter 3

Build Configuration

Projects like FreeRDP need to conform to a large number of requirements at the same time while remaining highly flexible, portable, fast, small, modular, reusable, testable, configurable, and the list goes on. It may not be obvious, but one of the challenges of open source development is having to please everybody at the same time. Different vendors want FreeRDP built in certain ways, with nice extension points for themselves to easily maintain their extensions separately, with just about any option that fits their needs within the upstream sources. It turns out that providing everybody with what they want is possible with enough CMake scripting. Nevertheless, the task of offering build configuration for all these needs makes build configuration management quite complex.

3.1 CMake

Prior to generating build scripts for your environment, you should ensure that you have CMake installed and available within your system path. Installing CMake is covered in the previous section of this manual.

3.1.1 Introduction

To generate build scripts, open a terminal, move to the root of the source tree, and type:

```
cmake .
```

This will tell CMake to generate build scripts using the CMake scripts found in the local directory (“.”). If the command fails, you might need to invoke CMake with additional options which we will cover in the following sections. The path

to the directory where the root CMake script is contained is always the last argument of the `cmake` command.

CMake looks for files with the name `CMakeLists.txt`. These files are the equivalent of makefiles, written in the CMake scripting language. A root `CMakeLists.txt` script can include other such scripts from subdirectories.

3.1.2 Generators

CMake scripts are used to generate platform-specific build scripts with the help of CMake generator. When no generator is specified, CMake attempts finding a suitable default. If it can't find any, generation will fail. Most of the time, you will want to specify the generator of your choice using the `-G "Generator Name"` option. Here is a list of common generators:

- Xcode
- Unix Makefiles
- NMake Makefiles
- Visual Studio 11
- Visual Studio 10
- Visual Studio 9 2008
- Eclipse CDT4 – Unix Makefiles
- Eclipse CDT4 – NMake Makefiles

In order to generate an Eclipse project on Linux, one would type:

```
cmake -G "Eclipse CDT4 - Unix Makefiles" .
```

For simple makefiles without an Eclipse project, “Unix Makefiles” would do it.

On Windows, the generator would likely be “Visual Studio 10” (Visual Studio 2010)

On Mac OS X, the “Xcode” generator would be appropriate.

3.1.3 CMake Cache

When CMake executes, it generates `CMakeCache.txt`, or the CMake cache. The cache stores values from previous executions. Whenever major changes happen, such as installing new dependencies, modifying the CMake scripts, or when something does not work as expected it is advised to simply delete the CMake cache and regenerate from scratch. Clearing the CMake cache is done by deleting the `CMakeCache.txt` file:

```
rm CMakeCache.txt
```

The CMake cache is a simple text file which looks like a simple configuration file with one option per line. You can modify it in a text editor, but it is preferable to use a specialized editor. On Linux, there is `ccmake`, which is ncurses-based (simple text-based GUI inside the terminal). On Windows, there is `cmake-gui`. These tools are very useful in order to visualize and edit the entire set of options available for the current project.

3.1.4 Configuration Files

If you have to frequently regenerate the CMake, the task of entering all the options that you want over and over again will be cumbersome. A flexible build system also means a large number of options, and the process of typing all of them every time is not only slow but also error prone. This is where storing build options in a configuration file becomes useful.

Luckily for us, CMake supports using values from an initial cache when generating a new one. The format of the initial cache is the same as the `CMakeCache.txt` file. The easiest way to use the initial cache is to create a text file and the desired options from a previously generated cache. Since many values in the cache are specific to the current build environment, it is not recommend to copy the entire generated cache to be used as an initial cache.

Here is an example: when developing, I like to keep a working version of FreeRDP installed in a directory separated from other less stable builds. Such a build comes in handy when I need to use FreeRDP rather than develop it, and my current development build is broken. For this build, I use the following options when invoking CMake:

```
cmake -DCMAKE_INSTALL_PREFIX="/opt/freerdp" -DCMAKE_BUILD_TYPE=Release -DSTATIC_CHANNELS=on
```

In the resulting `CMakeCache.txt` file, the options I've passed will look like this:

```
WITH_SERVER:BOOL=on
STATIC_CHANNELS:BOOL=on
MONOLITHIC_BUILD:BOOL=on
CMAKE_BUILD_TYPE:STRING=Release
CMAKE_INSTALL_PREFIX:STRING=/opt/freerdp
```

You will notice that the format in the CMake cache is a bit different that the format used at the command-line. In reality, they turn out to be the same, except that CMake guesses the parameter type when you invoke it. Values like “on/off”

are recognized as `BOOL`, but if you were to explicitly specify the `STRING` data type, CMake would consider “on/off” as a `STRING` rather than a `BOOL`.

The above options can be copied in a text file which I will call “StableBuildCache.txt” for the sake of this example. The next time, instead of retyping everything, all that is needed is passing the initial cache to CMake using the `-C` option:

```
cmake -C StableBuildCache.txt .
```

This will have the same effect as passing all the desired parameters through the command-line, with much less effort.

3.1.5 Out-of-Source Build

One of the powerful features of CMake is the ability to make out-of-source builds. Normally, a build is made in-source, meaning that build scripts and build outputs are produced in the same directories as the sources. An out-of-source build, as the name says, produces build scripts and build outputs outside of the sources.

To generate out-of-source build scripts, invoke CMake from a directory different from the source tree, while providing the path to the directory where the root CMake script is contained. For instance, if I have my sources in `~/src/FreeRDP`, but I want to build everything inside the `~/build/FreeRDP` directory, here is what the command would look like:

```
cmake ../../src/FreeRDP
```

Where the current working directory is `~/build/FreeRDP`.

3.2 Options

CMake options are passed to the `cmake` command, prefixed with `-D` (D for Define). `BOOL` options are either `ON` or `OFF`. `STRING` options are just regular strings, and can be put between quotes.

3.2.1 CMake Options

`CMAKE_BUILD_TYPE` (`STRING` [`Release`]): Defines the build type, usually `Debug` or `Release`.

`BUILD_SHARED_LIBS` (`BOOL` [`OFF`]): Build libraries with an unspecified build type as shared libraries. This option is turned on by default in FreeRDP.

`CMAKE_INSTALL_PREFIX` (`STRING`): Path prefix for installation. On Linux, this is `/usr/local` by default.

3.2.2 Build Options

MONOLITHIC_BUILD (BOOL [OFF]): Combine smaller libraries into single, larger, monolithic libraries. In monolithic build mode, there is a single libfreerdp instead of multiple libraries such as libfreerdp-core, libfreerdp-cache, libfreerdp-utils, etc.

STATIC_CHANNELS (BOOL [ON]): Build all channels statically, including channels normally built as plugins. Client static channels are bundled in libfreerdp-client, and server static channels are bundled in libfreerdp-server.

BUILD_TESTING (BOOL [OFF]): Enable CTest and build unit tests. Unit tests are located in “test” subdirectories, and can be executed with the ctest command.

BUILD_SAMPLE (BOOL [OFF]): Build sample code, such as sample channels, clients or servers.

3.2.3 Mac OS X Options

Mac OS X has specific options for 32-bit, 64-bit and universal binaries. By default, cmake will only build for the native architecture. Target architectures can be specified with the CMAKE_OSX_ARCHITECTURES option:

```
cmake -DCMAKE_OSX_ARCHITECTURES="i386" -G Xcode .
cmake -DCMAKE_OSX_ARCHITECTURES="x86_64" -G Xcode .
cmake -DCMAKE_OSX_ARCHITECTURES="i386;x86_64" -G Xcode .
```

Where i386 is 32-bit, x86_64 is 64-bit. Specifying both 32-bit and 64-bit means building universal binaries.

3.2.4 Optimization Options

WITH_SSE2 (BOOL): Build with SSE2 CPU instructions support (Intel architecture only).

WITH_NEON (BOOL): Build with NEON CPU instructions support (ARM architecture only).

3.2.5 Client Options

WITH_CLIENT (BOOL [ON]): Build clients

WITH_CLIENT_CHANNELS (BOOL [ON]): Build client channels.

3.2.6 Server Options

WITH_SERVER (BOOL [OFF]): Build servers

WITH_SERVER_CHANNELS (BOOL [OFF]): Build server channels.

3.2.7 Channel Options

WITH_CHANNELS (BOOL [ON]): Build channels.

Channel Names:

- AUDIN (Audio Input)
- CLIPRDR (Clipboard Redirection)
- DRIVE (Drive / File System Redirection)
- DRDYNVC (Dynamic Virtual Channel)
- PARALLEL (Parallel Port Redirection)
- PRINTER (Printer Redirection)
- RAIL (Remote Applications)
- RDPDR (Device Redirection)
- RDPSND (Audio Output)
- SERIAL (Serial Port Redirection)
- SMARTCARD (Smart Card Redirection)
- TSMF (Multimedia Redirection)
- URBDRC (USB Redirection)

CHANNEL__ (BOOL): Build channels

CHANNEL_NAME_CLIENT (BOOL): Build client channel. This option is dependent on CHANNEL__.

CHANNEL_NAME_SERVER (BOOL): Build server channel. This option is dependent on CHANNEL__.

3.2.8 Feature Options

WITH_JPEG: Build with JPEG codec support. This feature is not part of the canonical RDP specifications, and is implemented as an alternative codec to RemoteFX or NSCodec.

WITH_WIN8: Build with Windows 8 support. This is used by the Windows FreeRDP server and enables linking against Windows 8 libraries, therefore breaking backwards compatibility.

Chapter 4

Testing

4.0.9 CTest

CTest unit tests are located in “test” subdirectories and are deeply integrated with the rest of the CMake family of software process tools. CTest unit tests do not need a particular test framework. Test code is written as simple programs which return an exit code indicating success or failure.

CTest unit tests can be enabled with the `BUILD_TESTING` option.

Introduction

All unit tests can be executed by invoking `ctest` with no option:

```
ctest .
```

Unit tests can be filtered by name using a regular expression and the `-R` option.

To run all tests with a name beginning with “TestPath”, use:

```
ctest -R "TestPath*" .
```

To execute only the “TestPathCchAppend” test, use:

```
ctest -R "TestPathCchAppend$" .
```

Debug output is turned off by default, but it can be enabled with the `-V` (verbose) option:

```
ctest -V -R "TestPathCchAppend$" .
```

Chapter 5

FreeRDS

These instructions are preliminary for those who want to try FreeRDS as it is being developed.

5.1 Getting Started

Install the dependencies for FreeRDS as instructed in the FreeRDP section of this manual.

Clone the FreeRDP repository with the FreeRDS development branch:

```
mkdir ~/git/FreeRDP
cd ~/git/FreeRDP
git clone git://github.com/FreeRDP/FreeRDP.git
```

Clone the FreeRDS development repository in the server directory of the FreeRDP source tree:

```
cd ~/git/FreeRDP/FreeRDP/server
git clone git://github.com/FreeRDS/FreeRDS.git
```

The FreeRDP git repository automatically ignores the server/FreeRDS directory, such that you can manage the two git repositories independently. Ignoring FreeRDS in the FreeRDP git repository prevents accidental commits where the FreeRDS sources would be included.

Follow the regular instructions for building FreeRDP, with the exception of a few extra options (`WITH_SERVER`, `WITH_X11RDP`). It is currently much easier to deploy FreeRDS to a temporary directory in order to execute it. For the purpose of this example, let's use `/opt/freerds` as an installation prefix:

```
sudo mkdir /opt/freerds
sudo chmod 777 /opt/freerds
```

5.2 Building

5.3 X11rdp

X11rdp is an alternative X11 server like Xvnc, Xephyr, Xnest and Xvb. Normally, these servers are built alongside Xorg with the xorg-server sources as they make use of internal APIs and libraries that are not installed. To build X11rdp, we need to build the xorg-server sources in a known directory such that we can include them in our cmake project.

The generic approach is to install all the distribution-provided packages required to build the xorg-server package from source. You can then obtain the distribution-sources for the xorg-server package or obtain vanilla sources corresponding to the same version. This process is automated with cmake build scripts, but needs to be done prior to generating the FreeRDS project:

```
cd ~/git/FreeRDP/FreeRDP/server/FreeRDS
cd session-manager/module/X11/service/xorg-build
cmake .
make
```

If all goes well, you should have xorg-server built in external/Source/xorg-server. Create a symbolic link to that location from session-manager/module/X11/server/xorg-server:

```
cd ..
ln -s xorg-build/external/Source/xorg-server xorg-server
```

session-manager/module/X11/server/xorg-server will then be picked up by the FreeRDS build system.

5.3.1 FreeRDS

When generating project files with cmake, specify the prefix using `-DCMAKE_INSTALL_PREFIX=/opt/freerds`:

```
cmake -DWITH_SERVER=on -DWITH_X11RDP=on -DCMAKE_INSTALL_PREFIX=/opt/freerds .
```

Then always execute “make install” after building and launch freerds from its installed location. Executing from the source tree may be properly supported in the future but for now it is not recommended.

5.4 Running

You may want to disable the firewall on CentOS:

```
service iptables save
service iptables stop
chkconfig iptables off
```

Open two terminals logged in as root, and change directory to your installation prefix (/opt/freerds).

In the first one, execute freerds:

```
./bin/freerds --nodaemon
```

In the second one, execute freerds-session-manager:

```
./bin/freerds-session-manager
```

There are easier ways of executing freerds but this manual execution method is more flexible for development purposes.

You can then connect locally:

```
./bin/xfreerdp /u:username /p:password /cert-ignore /v:localhost
./bin/xfreerdp /u:username /p:password /cert-ignore /v:localhost /rfx
./bin/xfreerdp /u:username /p:password /cert-ignore /v:localhost /max-fast-path-size:100000
```

Chapter 6

API Reference

6.1 Input

6.1.1 Concepts

Keyboard Type

Physical keyboard type used in Windows keyboard layouts. In the vast majority of cases, type 4 (IBM keyboard) is used. Keyboard type 7 (Japanese 109 keyboard) is used with Japanese keyboard layouts. Most other keyboard types are archaic and are most likely irrelevant today, like type 1 (Olivetti 102 keyboard). Complete keyboard type information also includes a keyboard subtype and number of function keys. Virtual scan code values used in RDP are relative to a specific keyboard type, which is in most cases, but not always, keyboard type 4. This means one should not create static keyboard maps directly involving virtual scan codes, since this approach ignores the keyboard type information and is therefore flawed by design for a minority of cases. Instead, developers are encouraged to create keyboard maps using virtual key codes, and generate on-the-fly a keyboard map to virtual scan codes using the winpr-input module.

Virtual Scan Code

A virtual scan code is a scan code corresponding to a physical key on a given keyboard type. A virtual scan code value ranges from 0 to 127 and is either extended or non-extended. Here are some sample mappings between a keyboard type, virtual scan code, and extended/non-extended flag to help understand the differences:

- (Type4, 0x1C, NonExtended) = VK_RETURN (regular Enter key)
- (Type4, 0x1C, Extended) = VK_RETURN (keypad Enter key)

- (Type4, 0x1E, NonExtended) = VK_KEY_A (physical ‘A’ key when using a qwerty layout)
- (Type4, 0x1E, Extended) = VK_NONE (no corresponding virtual key code)
- (Type7, 0x70, NonExtended) = VK_DBE_KATAKANA (special japanese key)
- (Type4, 0x70, NonExtended) = VK_NONE (no corresponding virtual key code)

Virtual Key Code

Windows virtual key codes are a keyboard-independent codes representing a keyboard key. They are *not* in any way representative of output characters that results from keyboard input. Not all virtual key codes are given names, and information about some of the less frequently used codes is spread out. For the sake of convenience, WinPR provides an exhaustive list of known virtual key codes in winpr/input.h. Since letter keys are not given official names, WinPR defines VK_KEY_LETTER, where each letter key matches the corresponding key on a US qwerty layout. This means that VK_KEY_Q (top left) is still VK_KEY_Q on an azerty keyboard, even if the corresponding physical key has the letter ‘A’ on an azerty layout. Mapping virtual key codes to corresponding output characters is the keyboard layout’s job.

Keyboard Layout

Keyboard layouts on Windows are implemented using DLLs which export some information about them (keyboard type information) and static data structures used for mapping virtual key code input to output characters. A keyboard layout is uniquely identified by an id, such as 0x0000409 for the US keyboard layout.

The process of mapping virtual key code input to output characters is a non-invertible function, meaning that output characters cannot be correctly mapped to the original virtual key code input without loss of information. For instance, the capital letter ‘A’ can be produced by pressing the ‘A’ key while the shift key is down, or when caps locks is toggled. Using only the output character ‘A’, it is impossible to know if the character was produced using the shift key or the caps lock key. The same problem occurs for cases for keys which do not directly produce output characters, such as the control or alt keys. With no output characters to try mapping to original virtual key codes, there’s not much that can be done. Other cases which prevent proper mapping of output characters to input virtual key codes are dead keys, or keys which do not result in a direct character output but affect the next key press. In many non-US keyboard layouts, accented letters are typed by first pressing a key corresponding to the accent and then typing the letter that needs to be accented. In theory, one can get by creating keyboard layout specific maps, but this is a tedious process that needs to be repeated for every keyboard layout in existence, while being broken by design.

6.1.2 Keyboard Mapping

WinPR (winpr-input module) offers reusable keyboard mapping utilities to help implementers with providing exhaustive, accurate keyboard support. Include winpr/input.h to make use of these.

To make debugging and parsing easier, functions are provided to match known key names to their numerical values:

```
char* GetVirtualKeyName(DWORD vkcode);
DWORD GetVirtualKeyCodeFromName(const char* vkname);
DWORD GetVirtualKeyCodeFromXkbKeyName(const char* xkbname);
```

Example: mapping XKB key name to virtual key code and printing out the conversion result:

```
char* xkbname = "RTRN"; /* XKB Enter key */
DWORD vkcode = GetVirtualKeyCodeFromXkbKeyName(xkbname); /* returns VK_RETURN */
printf("XKB key name: %s Virtual Key Code %s (0x%04X)\n", xkbname, GetVirtualKeyName(vkcode));
```

Example: converting textual representation of a virtual key to its corresponding numerical value:

```
char* vkname = "VK_RETURN"; /* this string could be obtained from a configuration file instead */
DWORD vkcode = GetVirtualKeyCodeFromName(vkname);
printf("Virtual Key Code: %s (0x%04X)\n", vkname, vkcode);
```

When mapping keyboard systems, the general approach is to create a map between the native key codes and virtual key codes. However, since RDP uses virtual scan codes, you will still need to map the virtual key codes to virtual scan codes before sending them. The advantage is that virtual scan codes are truly keyboard independent and have proper names, while virtual key codes are highly impractical from the fact that they don't really have names other than their number and are prone to change with regards to the physical keyboard type in use. If the only data you can have from the local input system is output characters, you should attempt using unicode input instead and limit usage of non-unicode keyboard input events to non-character input for keys like shift and num lock.

```
DWORD GetVirtualKeyCodeFromVirtualScanCode(DWORD scancode, DWORD dwKeyboardType);
DWORD GetVirtualScanCodeFromVirtualKeyCode(DWORD vkcode, DWORD dwKeyboardType);
```

Example: mapping virtual key codes to virtual scan codes:

```
DWORD scancode;
```

```
/* Regular Enter key (non-extended) */
```

```
scancode = GetVirtualScanCodeFromVirtualKeyCode(VK_RETURN, 4); /* returns 0x1C */
```

```
/* Keypad Enter key (extended) */
```

```
scancode = GetVirtualScanCodeFromVirtualKeyCode(VK_RETURN | KBDEXT, 4); /* returns (0x1C | KBDEXT) */
```

Example: mapping virtual scan codes to virtual key codes:

```
DWORD vkcode;
```

```
/* Regular Enter key (non-extended) */
```

```
vkcode = GetVirtualKeyCodeFromVirtualScanCode(0x1C, 4); /* returns VK_RETURN */
```

```
/* Keypad Enter key (extended) */
```

```
vkcode = GetVirtualKeyCodeFromVirtualScanCode(0x1C | KBDEXT, 4); /* returns (VK_RETURN | KBDEXT) */
```

Certain keyboard systems use key codes which can be statically mapped to virtual key codes. WinPR currently provides mapping for Linux evdev and Mac OS X keycodes:

```
DWORD GetVirtualKeyCodeFromKeycode(DWORD keycode, DWORD dwFlags);
```

```
DWORD GetKeycodeFromVirtualKeyCode(DWORD vkcode, DWORD dwFlags);
```

```
dwFlags: * KEYCODE_TYPE_APPLE * KEYCODE_TYPE_EVDEV
```

Most modern Linux X11 environments used evdev for input devices. The good news is that evdev uses a static set of X11 keycodes, which is not the case in other X11 environments. When evdev is not in use but XKB is available, a dynamic keyboard map can be generated on-the-fly using the XKB keyboard mapping utilities. One can use the xev utility to test keyboard input and see X11 keycodes. The following is a sample xev output for the 'A' key on a Linux system using evdev:

```
KeyPress event, serial 40, synthetic NO, window 0x5600001,
  root 0xd9, subw 0x0, time 17400104, (4,14), root:(1039,482),
  state 0x2000, keycode 38 (keysym 0x61, a), same_screen YES,
  XLookupString gives 1 bytes: (61) "a"
  XmbLookupString gives 1 bytes: (61) "a"
  XFilterEvent returns: False
```

```
KeyRelease event, serial 40, synthetic NO, window 0x5600001,
  root 0xd9, subw 0x0, time 17400214, (4,14), root:(1039,482),
  state 0x2000, keycode 38 (keysym 0x61, a), same_screen YES,
```

```
XLookupString gives 1 bytes: (61) "a"  
XFilterEvent returns: False
```

Example: mapping Linux evdev keycode to virtual key code back and forth:

```
DWORD vkcode;  
DWORD keycode = 38; /* evdev 'A' key */  
vkcode = GetVirtualKeyCodeFromKeycode(keycode, KEYCODE_TYPE_EVDEV); /* returns VK_KEY_A */  
keycode = GetKeycodeFromVirtualKeyCode(vkcode, KEYCODE_TYPE_EVDEV); /* returns 38 */
```

Mac OS X uses the same set of keycodes in X11 and non-X11 environments. While values from 0 to 7 are unused, the `keyCode` member of the `NSEvent` class is off by 8. This means that when receiving the keycode value from Mac OS X, you should increment it by 8 before passing it to `GetVirtualKeyCodeFromKeycode()`, and decrement the keycode value returned by `GetKeycodeFromVirtualKeyCode()` before passing it to Mac OS X.

Example: mapping Mac OS X keycode to virtual key code back and forth:

```
DWORD vkcode;  
DWORD keycode = APPLE_VK_Return; /* Mac OS X Enter key */  
vkcode = GetVirtualKeyCodeFromKeycode(keycode + 8, KEYCODE_TYPE_APPLE); /* returns VK_RETURN */  
keycode = GetKeycodeFromVirtualKeyCode(vkcode, KEYCODE_TYPE_APPLE) - 8; /* returns APPLE_VK_RETURN */
```

6.1.3 Synchronize Event

An RDP synchronize event is used to synchronize keyboard toggle keys, such as caps lock and num lock. Synchronization of toggle keys is usually required whenever the local session window gains focus, as toggle key states may have changed while the window was not focused.

```
void freerdp_input_send_synchronize_event(rdpInput* input, UINT32 flags);
```

flags:

- `KBD_SYNC_SCROLL_LOCK`
- `KBD_SYNC_NUM_LOCK`
- `KBD_SYNC_CAPS_LOCK`
- `KBD_SYNC_KANA_LOCK`

Each flag should be set only if the corresponding key is toggled on the local keyboard. The Kana Lock key is only meaningful for Japanese keyboard input.

6.1.4 Keyboard Event

An RDP keyboard event is a message containing a Windows virtual scan code (extended or non-extended) along with flags to mark the corresponding key as being up or down.

```
void freerdp_input_send_keyboard_event(rdpInput* input, UINT16 flags, UINT16 code);
```

flags:

- KBD_FLAGS_EXTENDED
- KBD_FLAGS_RELEASE
- KBD_FLAGS_DOWN

6.1.5 Keyboard Unicode Event

An RDP unicode keyboard event is a message containing a unicode character along with flags to mark the key as being released.

```
void freerdp_input_send_unicode_keyboard_event(rdpInput* input, UINT16 flags, UINT16 code);
```

flags:

- KBD_FLAGS_RELEASE

Example: sending 'A' unicode character:

```
freerdp_input_send_unicode_keyboard_event(input, 0, 0x0041); /* Unicode 'A' character (key code) */
freerdp_input_send_unicode_keyboard_event(input, KBD_FLAGS_RELEASE, 0x0041); /* Unicode 'A' character (key code) */
```

6.1.6 Mouse Event

An RDP mouse event is used for standard mouse buttons (left, middle and right), mouse movement and mouse wheel scrolling.

```
void freerdp_input_send_mouse_event(rdpInput* input, UINT16 flags, UINT16 x, UINT16 y);
```

flags:

- PTR_FLAGS_DOWN
- PTR_FLAGS_MOVE

- PTR_FLAGS_BUTTON1 (left)
- PTR_FLAGS_BUTTON2 (right)
- PTR_FLAGS_BUTTON3 (middle)
- PTR_FLAGS_WHEEL
- PTR_FLAGS_WHEEL_NEGATIVE
- WheelRotationMask

Example: mouse left-click:

```
freerdp_input_send_mouse_event(input, PTR_FLAGS_BUTTON1 | PTR_FLAGS_DOWN, 100, 200);
freerdp_input_send_mouse_event(input, PTR_FLAGS_BUTTON1, 100, 200);
```

Example: mouse movement:

```
freerdp_input_send_mouse_event(input, PTR_FLAGS_MOVE, 100, 200);
```

Example: mouse wheel scrolling:

```
int rotationUnits = 0x0078;
freerdp_input_send_mouse_event(input, PTR_FLAGS_WHEEL | (rotationUnits & WheelRotationMask), 100, 200);
freerdp_input_send_mouse_event(input, PTR_FLAGS_WHEEL | (rotationUnits & WheelRotationMask), 100, 200);
```

6.1.7 Extended Mouse Event

An RDP extended mouse event is used for extended mouse buttons, such as button 4 and 5. The extended mouse event accepts the same flags as the regular mouse event, with a few extended flags.

```
void freerdp_input_send_extended_mouse_event(rdpInput* input, UINT16 flags, UINT16 x, UINT16 y, INT16 dx, INT16 dy, INT16 wheel);
```

flags:

- PTR_XFLAGS_DOWN
- PTR_XFLAGS_BUTTON1 (button 4)
- PTR_XFLAGS_BUTTON2 (button 5)

Example: extended mouse click

```
freerdp_input_send_extended_mouse_event(input, PTR_XFLAGS_BUTTON1 | PTR_XFLAGS_DOWN, 100, 200, 0, 0, 0);
freerdp_input_send_extended_mouse_event(input, PTR_XFLAGS_BUTTON1, 100, 200);
```

6.2 Virtual Channel API

Virtual Channel Client DLL¹

6.2.1 VirtualChannelEntry

```
typedef BOOL VCAPITYPE VIRTUALCHANNELENTY(PCHANNEL_ENTRY_POINTS pEntryPoints);
```

VirtualChannelEntry is the entry point of static virtual channel clients. This function is called with a pointer to a CHANNEL_ENTRY_POINTS structure containing the portion of the virtual channel API exposed to the virtual channel client:

```
typedef struct tagCHANNEL_ENTRY_POINTS
{
    DWORD cbSize;
    DWORD protocolVersion;
    PVIRTUALCHANNELINIT pVirtualChannelInit;
    PVIRTUALCHANNELOPEN pVirtualChannelOpen;
    PVIRTUALCHANNELCLOSE pVirtualChannelClose;
    PVIRTUALCHANNELWRITE pVirtualChannelWrite;
} CHANNEL_ENTRY_POINTS, *PCHANNEL_ENTRY_POINTS;
```

The data contained within this structure is only valid for the duration of the call to VirtualChannelEntry and should therefore be copied.

6.2.2 VirtualChannelInit

VirtualChannelInit has to be called by the virtual channel client when VirtualChannelEntry is called.

```
typedef UINT VCAPITYPE VIRTUALCHANNELINIT(LPVOID* ppInitHandle, PCHANNEL_DEF pChannel, INT c
```

```
typedef VIRTUALCHANNELINIT *PVIRTUALCHANNELINIT;
```

ppInitHandle receives the init handle which is used to identify the current connection when calling VirtualChannelOpen².

pChannel is an array of CHANNEL_DEF structures, where the number of elements in the array is given by channelCount.

¹<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383580>

²<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383570>

```

typedef struct tagCHANNEL_DEF
{
    char name[CHANNEL_NAME_LEN + 1];
    ULONG options;
} CHANNEL_DEF;
typedef CHANNEL_DEF *PCHANNEL_DEF;
typedef PCHANNEL_DEF *PPCHANNEL_DEF;

```

6.2.3 VirtualChannelInitEvent

VirtualChannelInitEvent³:

```

typedef VOID VCAPITYPE CHANNEL_INIT_EVENT_FN(LPVOID pInitHandle, UINT event, LPVOID pData, U

typedef CHANNEL_INIT_EVENT_FN *PCHANNEL_INIT_EVENT_FN;

```

6.2.4 VirtualChannelOpen

VirtualChannelOpen⁴ should be called for each channel registered by the current virtual channel client when VirtualChannelInitEvent is called with a CHANNEL_EVENT_CONNECTED event.

```

typedef UINT VCAPITYPE VIRTUALCHANNELOPEN(LPVOID pInitHandle, LPDWORD pOpenHandle, PCHAR pCH

typedef VIRTUALCHANNELOPEN *PVIRTUALCHANNELOPEN;

```

VirtualChannelOpen takes the init handle obtained by a previous call to VirtualChannelInit and receives an open handle which will be used in calls to VirtualChannelWrite and VirtualChannelClose. A pointer to a VirtualChannelOpenEvent⁵ function is also passed to VirtualChannelOpen which will be used to receive notifications for read and write operations.

6.2.5 VirtualChannelOpenEvent

VirtualChannelOpenEvent⁶ is used to notify the virtual channel client of data being received or the completion of data write operations.

```

typedef VOID VCAPITYPE CHANNEL_OPEN_EVENT_FN(DWORD openHandle, UINT event, LPVOID pData, UI

typedef CHANNEL_OPEN_EVENT_FN *PCHANNEL_OPEN_EVENT_FN;

```

³<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383568>

⁴<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383570>

⁵<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383573>

⁶<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383573>

6.2.6 VirtualChannelWrite

Virtual channel clients call `VirtualChannelWrite`⁷ with the open handle obtained by calling `VirtualChannelOpen` in order to write data to the virtual channel.

```
typedef UINT VCAPITYPE VIRTUALCHANNELWRITE(DWORD openHandle, LPVOID pData, ULONG dataLength,
                                             DWORD flags);
typedef VIRTUALCHANNELWRITE *PVIRTUALCHANNELWRITE;
```

6.2.7 VirtualChannelClose

`VirtualChannelClose`⁸ is called by the virtual channel close using the open handle returned by the previous call to `VirtualChannelOpen` to close the virtual channel.

```
typedef UINT VCAPITYPE VIRTUALCHANNELCLOSE(DWORD openHandle);
typedef VIRTUALCHANNELCLOSE *PVIRTUALCHANNELCLOSE;
```

⁷<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383576>

⁸<http://msdn.microsoft.com/en-us/library/windows/desktop/aa383556>