RemotelyAnywhere

Security Considerations

Table of Contents

Introduction

RemotelyAnywhere has security features that enhance the built-in security features of Microsoft Windows. When configured properly, you can minimize security risks. This document has been prepared with a single Windows web server in mind; therefore, there may be concepts that do not apply to your configuration.

Microsoft Windows

When working with Windows computers, it is recommended to take security precautions such as the following:

- Disable NetBIOS on the network adapter that provides Internet access
- Rename any built-in accounts that are in use
- Disable any built-in accounts that are not in use
- Apply the latest patches to the operating system and its components

Since this guide was written with a security conscious system administrator in mind, we are not going into details about using Windows security features that enhance Internet security.

Default Configuration

Having configured the computer to be as safe as possible, you can now turn to securing RemotelyAnywhere. The default RemotelyAnywhere configuration is as follows:

- Accepts Web and Console connections on port 2000 on all network adapters.
- Accepts SSH connections on port 22 on all network adapters.
- All above connections must be authenticated with a username/password pair that identifies a user with Administrator rights.
- Connections are accepted from any Internet address.

Unused Services

Decide which RemotelyAnywhere services you want to use on the computer, and disable the rest. We recommend disabling Telnet and SSH if you do not need access to a command prompt. If you do, disable Telnet and leave SSH running.

Incoming Connections

You can change the IP address on which RemotelyAnywhere listens to incoming connections. Assume you have a web server on the Internet named www.MyWebServer.com and RemotelyAnywhere is installed on your web server on port 2000. The web server can now be remotely managed by typing "http://www.mywebserver.com:2000" in any web browser. However, by adding another IP address to the server for remote administration purposes and restricting RemotelyAnywhere to listen only to that particular address the server can only be accessed from this address. For example, if you add the address 177.246.27.91, the server can only be accessed by typing http://177.246.27.91:2000. Any attempt to access RemotelyAnywhere on the original address (http://www.mywebserver.com:2000) will fail.

If possible, try to use an IP address that is on a different subnet than the IP addresses of the website or websites hosted on the computer. Most attacks begin with a port scan on the target computer. If the attacker sees that port 2000 is open on www.MyWebServer.com they will know that RemotelyAnywhere is installed on the computer. Using a different administrative IP address will hide this fact.

Default Port Numbers

Unlike the default Windows services, such as file sharing, the port providing access to RemotelyAnywhere can be changed to be something the potential intruder cannot detect or is not known. You should, however, choose a port that is easy to remember.

IP Address Filtering

Another way to strengthen security is IP Address Filtering. You can define IP addresses, or ranges of IP addresses, from which RemotelyAnywhere accepts connections. In this way, even if someone took the time to find out what IP address and port to connect to, they cannot do so unless they are on the "trusted" list. Any number of IP addresses and networks can be listed as "trusted".

Similarly, you can define a list of IP addresses and networks from which RemotelyAnywhere will refuse a connection. If you only connect to the computer from a number of known networks, you can tell RemotelyAnywhere to ignore incoming connections from anywhere else.

Suppose that the computer is in a datacenter and you want to access it from the office, from home, and from your wireless PDA. In this case, create an IP address filtering rule that allows connections from these three networks but denies connections from any other computer. You will only need to know the Internet IP address blocks and subnet addresses for the three networks you want to enable.

If you have dialup accounts in any of these locations, contact the ISP's customer service to find out the range of IP addresses that they can assign to you. You will not "trust" everyone who's on your ISP, of course - but even if they have a whole class A subnet it's better to

grant access to it as opposed to not setting up IP address filtering at all. Granting access to a huge class A subnet and denying connections from anywhere else locks out about 99.6% of possible intruders.

Secure Sockets Layer

Set up Secure Sockets Layer on the computer using the **Security > SSL Setup** menu in RemotelyAnywhere. You can also disable unsecured connections in **Preferences > Network**. This way, all traffic between the host and the remote computers will be encrypted with 256-bit ciphers – the same encryption method that major banks use – protecting your passwords and data.

User Accounts and OS Integration

The above methods are effective against unwanted visitors. However, they only prevent possible intruders from getting to the login screen where they must authenticate themselves with a username and a password. The user must enter a valid Windows username and password to RemotelyAnywhere to access the computer. Initially, this must be any user who has administrative credentials.

After configuring RemotelyAnywhere, any user or group can be granted access to the remote administration interface. Not all administrators must be given access, and access is not necessarily limited to administrators. It is up to the administrator to configure RemotelyAnywhere to best suit their needs. There is simply no way around being authenticated by RemotelyAnywhere. If it fails, the user will be denied access. Once the user has been authenticated RemotelyAnywhere impersonates them towards the operating system when servicing requests. This ensures that the user is only able to perform actions that their Windows credentials allow.

There is always the chance that someone will gain access to the password by a brute-force method or by simply guessing until access is granted. There are two very effective measures against this. The first is common sense! The System Administrator must always remember to use a password that is difficult to guess and set a password policy that enforces this on ordinary users as well. Passwords should be long, and preferably contain numbers and special characters, such as punctuation. The second is the IP address lockout feature of RemotelyAnywhere that is discussed below.

IP Address Lockout

RemotelyAnywhere can be set to lock out IP addresses after a pre-defined number of failed login attempts. System administrators can determine how many failed login attempts are allowed within a certain period of time, and for how long the offending IP address should be locked out. Then the IP address will be put on the “distrusted” list and all future connection attempts will fail from that address. For example, a typical configuration is where the IP

address is locked out for 30 minutes (lockout period) after 5 unsuccessful login attempts that takes place within 30 minutes (login period). The login and lockout periods can be as long as 4 billion seconds.

This is fundamentally different from Windows' own lockout mechanism: when Windows detects a certain number of failed login attempts, it disables logins to that account from all network locations as opposed to disabling the offending network address only. The two lockout methods complement each other.

Denial of Service Filter

RemotelyAnywhere also features a Denial of Service Filter. A DoS filter rejects connections if the IP address the request is coming from has made an excessive number of requests without authentication within the observation time window. This is done to protect against someone overloading the host computer by, for example, automatically and very quickly requesting the login page over and over again.

RSA SecurID

To add an extra layer of security over the simple username/password authentication, you can configure RemotelyAnywhere to require RSA SecurID authentication. RemotelyAnywhere was certified by RSA Security as SecureID Ready in 2003. Since that time, LogMeIn has continued to maintain the high level of security consistent with RSA technology.

User Access Controls

System administrators can configure RemotelyAnywhere so that users with certain roles have access only to a subset of tools offered by RemotelyAnywhere; for example, the Helpdesk department can be configured to only view a computer's screen and performance data, but not actually take over the mouse and the keyboard or make any changes to the system configuration. Alternatively, the Sales department might be given full remote control access to their respective computers, but features such as performance monitoring and remote administration would be made unavailable to them.

Using the operating system access token obtained when the user was authenticated, RemotelyAnywhere impersonates the user towards the operating system while performing actions on their behalf. This ensures that RemotelyAnywhere adheres to the operating system's security model, and users have access to the same files and network resources as if they were sitting in front of their computer. Resources unavailable to users in Windows or OS X also remain unavailable via RemotelyAnywhere.

The Human Factor

Any chain is as strong as its weakest link but there are always few steps you can take to make it stronger:

- Always remember to use long, hard to guess (but easy to remember) passwords.
- Always use encrypted connections to RemotelyAnywhere, such as SSL or SSH.
- Always log off when you've finished working.
- Set up the additional security measures detailed above if they apply.

Summary

Reducing or eliminating security threats is a simple but vital process. Following the guidelines and practices detailed in this document will help you ensure a secure implementation of RemotelyAnywhere.