Getting Started on the Cornell Restricted Access Data Center (CRADC)

This version: October 31, 2002

1. Introduction

The CRADC now operates completely behind its own firewall with its own private domain. We made this change to enhance our ability to track movements of files onto and off of the entire system. The firewall permits you to open a window (via the Windows Terminal Services Client or Remote Desktop Client) on the CRADC but it does not permit you to place new files on the system or remove files from the system. These operations are handled through the CRADC custodian, as explained below. We have placed additional software on the system so that you can write papers and prepare presentations without migrating files from the system.

2. Setup

You must sign a CRADC Data Users agreement each year (effective June 1, 2002). We have sent you this agreement as a part of this mailing. Your access to the CRADC nodes will be restored when you have signed and returned this agreement. Cornell University will provide you with a copy that contains both signatures after you return your signed copy.

Your account and the initial password will be sent to you in a separate communication. You will be required to change the password on first access. The rules for password formation are as follows. Passwords must be at least six characters long. Passwords may not contain your user name or any part of your full name. Passwords must contain characters from at least three of the following four classes:

- English upper case letters A, B, C, ... Z
- English lower case letters a, b, c, ... z
- Westernized Arabic numerals 0, 1, 2, ... 9
- Non-alphanumeric ("special characters") Punctuation marks and other symbols

Please sign on (any node) and change your password immediately. This action will reset your domain password for all nodes. Let the CRADC custodian know immediately (cradc_custodian@cornell.edu) if this fails but be sure to follow the rules for a legal password exactly.

All CRADC computing system users are required to change their password periodically. The system is automated to inform users when their account is due for a password change. Caution for users who use multiple CRADC sessions: If you change your password from one node while you have multiple sessions running (disconnected or active) on other nodes, you must logoff those sessions within 10 days. Failure to do so will result in your account lockout. To check if you have sessions running on multiple

CRADC nodes, from a CRADC node go to Start -> Programs -> Administrative Tools -> Terminal Services Manager.

You should configure the IE browser to use http://info.cradc.cornell.edu as its home page. This is the top directory of the CRADC documentation system. There is no access to the internet from CRADC.

In the instructions below [authorized_data] refers to the name of the folder on R:\ that you have been authorized to use and [user_name] refers to your CRADC domain user name.

3. System Access

The CRADC domain now has a total of 10 compute nodes that include an experimental 64bit node. The IP address for the 64bit node is 128.84.158.55. The other nine nodes are referred to as cradc01, cradc02,..., cradc08 and simulator. To access a CRADC node from outside the firewall, use the Windows Terminal Services Client or Remote Desktop Client to connect to [node_name].ciser.cornell.edu. The experimental 64bit node is accessed by using its complete IP address. Inside the firewall, refer to nodes as [node_name].cradc.cornell.edu. Insert the correct name of the node for [node_name].

4. Data Security

Your data use privileges apply to the data in R:\[authorized_data] (\\cradcfs1.cradc.cornell.edu\cradc\[authorized_data]) as indicated in your data user agreement. Your personal files (previously located in R:\Development\[user_name]) have been moved to R:\[authorized_data]\ [user_name], to which only you have access. You create folders under R:\[authorized_data]\programs\projects to contain files that you want to share with other researchers with whom you are collaborating and who also have permission to access R:\[authorized_data].

To request that a file be moved off the system and sent to you, send e-mail (from outside the firewall, there is no e-mail inside) to cradc_custodian@cornell.edu stating the location and name of the file and its contents. Your file movement permissions are governed by the master data use agreement covering the confidential data to which you have been given access. A summary of these rules for your data access privileges was included in your personal data use agreement. The same procedure applies to files that you want to put on CRADC; however, in this case you will have to state that the file contains only programming and non-confidential data before it will be placed on the system. For data files that you wish to place on the system, you will be required to provide a citation that confirms that the data are public use.

5. Using R:\

The R:\ is the primary data storage area for the CRADC system. You have been granted access to folders on R:\ in accordance with your data user agreement. Custom software and other system development activities are conducted in R:\Development. You have read access to this area unless you are doing development programming on CRADC.

6. Using T:\ and other disks

You may use the area T:\[authorized_data]\[user_name] to store temporary files from SAS, Matlab and other programs while you are working. Do not use T:\temp as this area is now reserved for development and system work. Only users authorized to access [authorized_data] may read or write in the folder T:\[authorized_data]\. You should include the line:

```
–work T:\[authorized_data]\[user_name]
```

in your shortcut to SAS. Never use the C:\ drive to store data. Any data that you store outside of R:\[authorized_data] or T:\[authorized_data] will be erased automatically by the system. See info.cradc.cornell.edu for further documentation of the customization of SAS.

7. Using Application Software

In addition to all the programs previously available (SAS, Matlab, Stata, GenStat, GLIM, GAUSS, COMPAQ Visual Fortran 6, Intel Performance Suite, MPIPro, StatTransfer, and DBMSCopy), we have installed the Microsoft Office Suite 2002 and Scientific Workplace 4.0. When using Scientific Workplace, do not store documents in the default area (they will be erased automatically). Use R:\[authorized_data]\docs\[user_name] for documents and bibliographies in Scientific Workplace.

These document preparation programs should allow you to write papers and collaborate with other authorized users without having to move files off of the system as frequently. If you encounter problems, contact Pinky Chandra (the CRADC manager) and she will help you resolve them (pc17@cornell.edu).

8. How to Logon to CRADC Compute Nodes (for New Users)

To connect to a CRADC node from your PC, you can use the Windows Terminal Services Client, Terminal Services Client Software or Remote Desktop Client, as described in step I. Next you respond to a warning message, as shown in step II. In step III you type in your CRADC user name and password to logon to your CRADC desktop.

I. **A: To connect with the Terminal Services via Internet Explorer** go to http://ciserfs1.ciser.cornell.edu/tsweb/ (note: this only works with Internet Explorer).

- In the "Remote Desktop Web Connection" window type in the address for a CRADC node as shown below.
- Leave the "Send logon information... " box blank.
- The "Size" option allows you to select a size for viewing the CRADC desktop. Full-screen mode allows you to view the CRADC.NOTE: After you're logged on you can toggle out and back from Full Screen mode to Window mode by using **Ctrl+Alt+Break**
- Then choose "Connect"

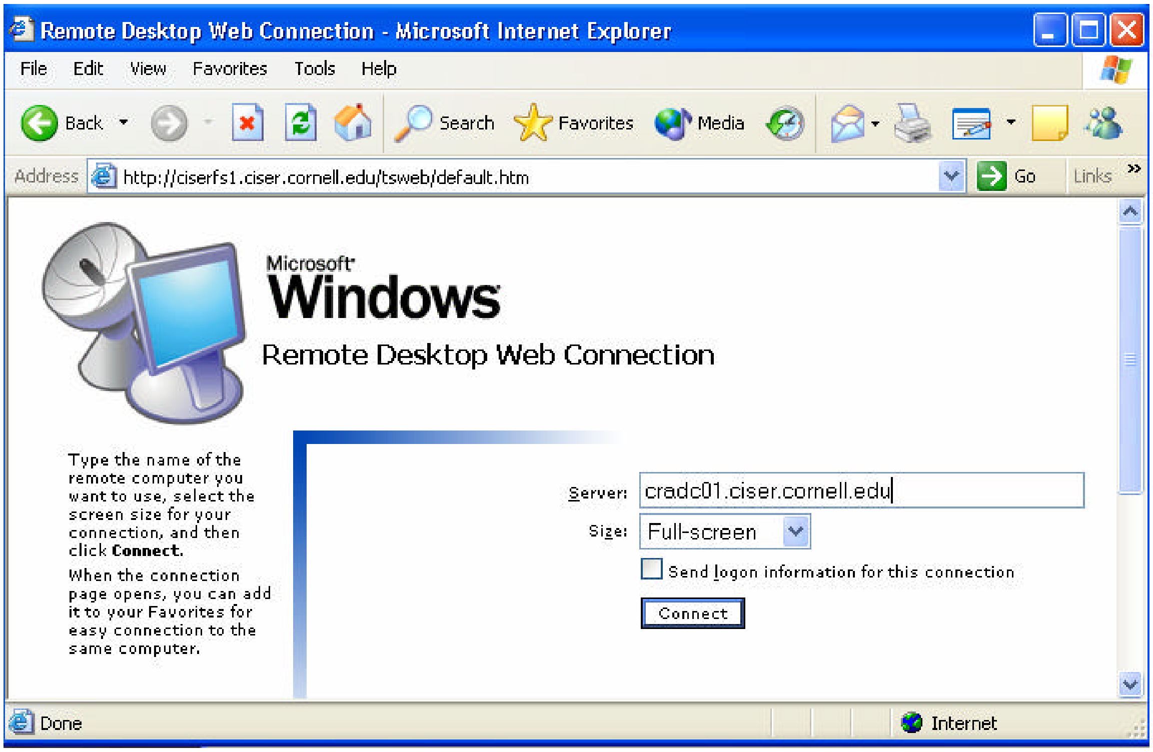

Address: http://ciserfs1.ciser.cornell.edu/tsweb/default.htm

Type the name of the remote computer you want to use, select the screen size for your connection, and then click **Connect.**

When the connection page opens, you can add it to your Favorites for easy connection to the same computer.

Server: cradc01.ciser.cornell.edu

Size: Full-screen

□ Send logon information for this connection

B: To Connect with the TERMINAL SERVICES CLIENT SOFTWARE

- Download and install Terminal Services Client Software on your PC http://ciserfs1.ciser.cornell.edu/public/tsc_32bit_x86.exe. Once the software is installed follow the instructions below for setting up your connections to the CRADC nodes:
- For each CRADC node you can create a CRADC icon on your local desktop as follows:
  - Go to "Start" --> "Programs" --> "Terminal Services Client" --> "Client Connection Manager"
  - The "Client Connection Manager" window will appear and then you should go to "File" --> choose "New Connection"

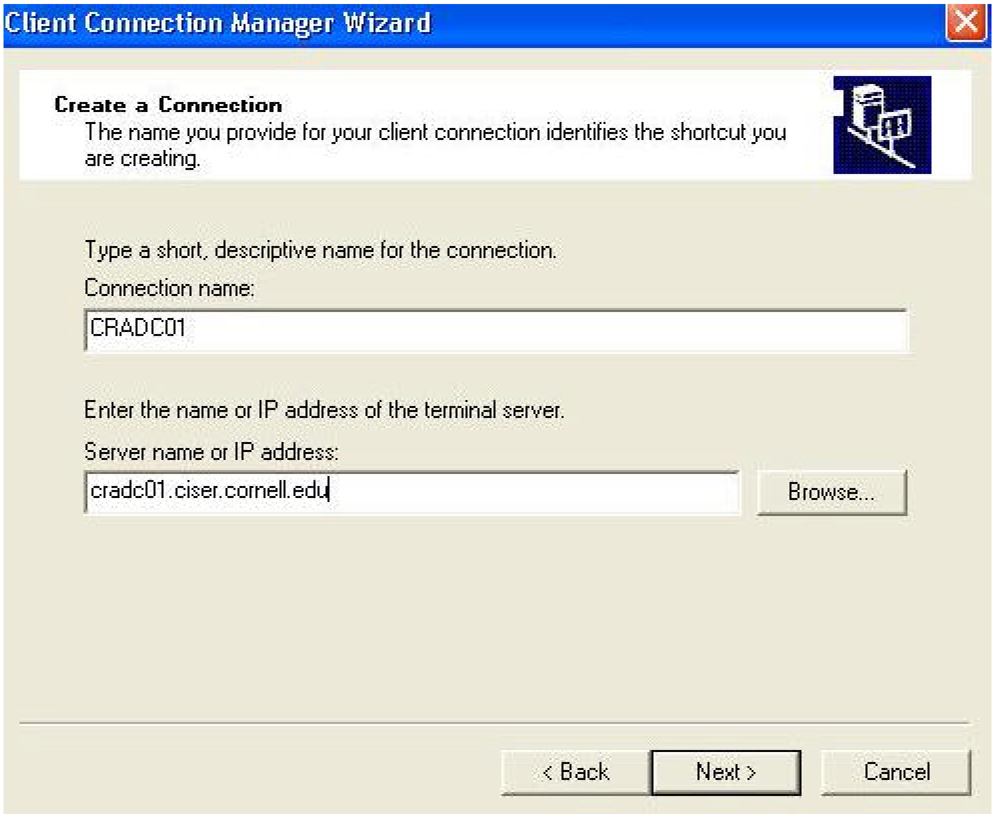

Client Connection Manager Wizard

Create a Connection

The name you provide for your client connection identifies the shortcut you are creating.

Type a short, descriptive name for the connection.

Connection name: CRADC01

Enter the name or IP address of the terminal server.

Server name or IP address: cradc01.ciser.cornell.edu

- Follow the instructions for the following screens.
    1. Automatic logon: Choose "Next"
        - leave this blank since automatic logon will not occur with the CRADC nodes.
    2. Screen Options:
        - You may wish to experiment with screen sizes, or you may choose "Full-screen" and then toggle between Full-screen and Window mode with Ctrl+Alt+Break after you are logged on.
        - Connection Properties: leave "enable data compression" and "cache bitmaps" both unchecked
        - Starting a Program: leave this box unchecked
        - Icon and Program Groups: leave this boxes as they are and check "Next"
    3. Then --> Finish
- A new icon will show up in your "Client Connection Manager Window."
    - You can right click on this icon and choose "create a shortcut on your desktop" if you like.

o **TO LOG ON** to a CRADC node, double click on the icon you have created for that node.

C: To Connect via Remote Desktop Client go to

- Programs -> Accessories -> Communications -> and click on **Remote Desktop Connection**
- Type in the address of a CRADC node (as shown below), and click on Connect.

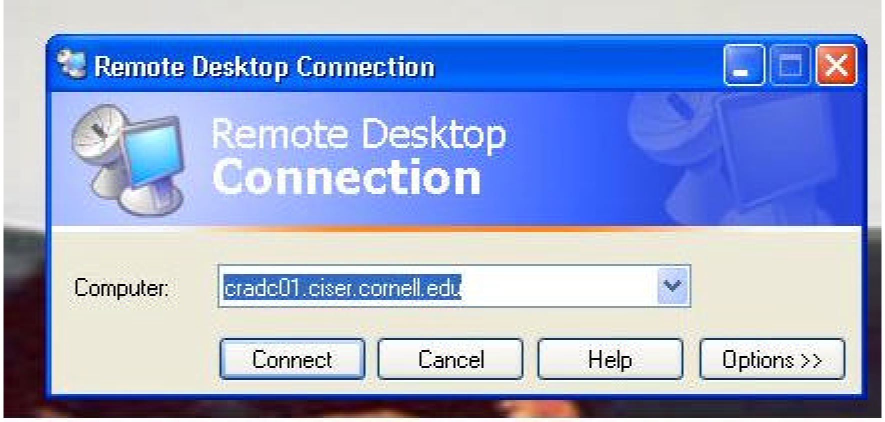

Computer: cradc01.ciser.cornell.edu

II. The system then displays a warning message informing that the system houses restricted-access data, as shown below. Click on “OK” at the bottom of the message.

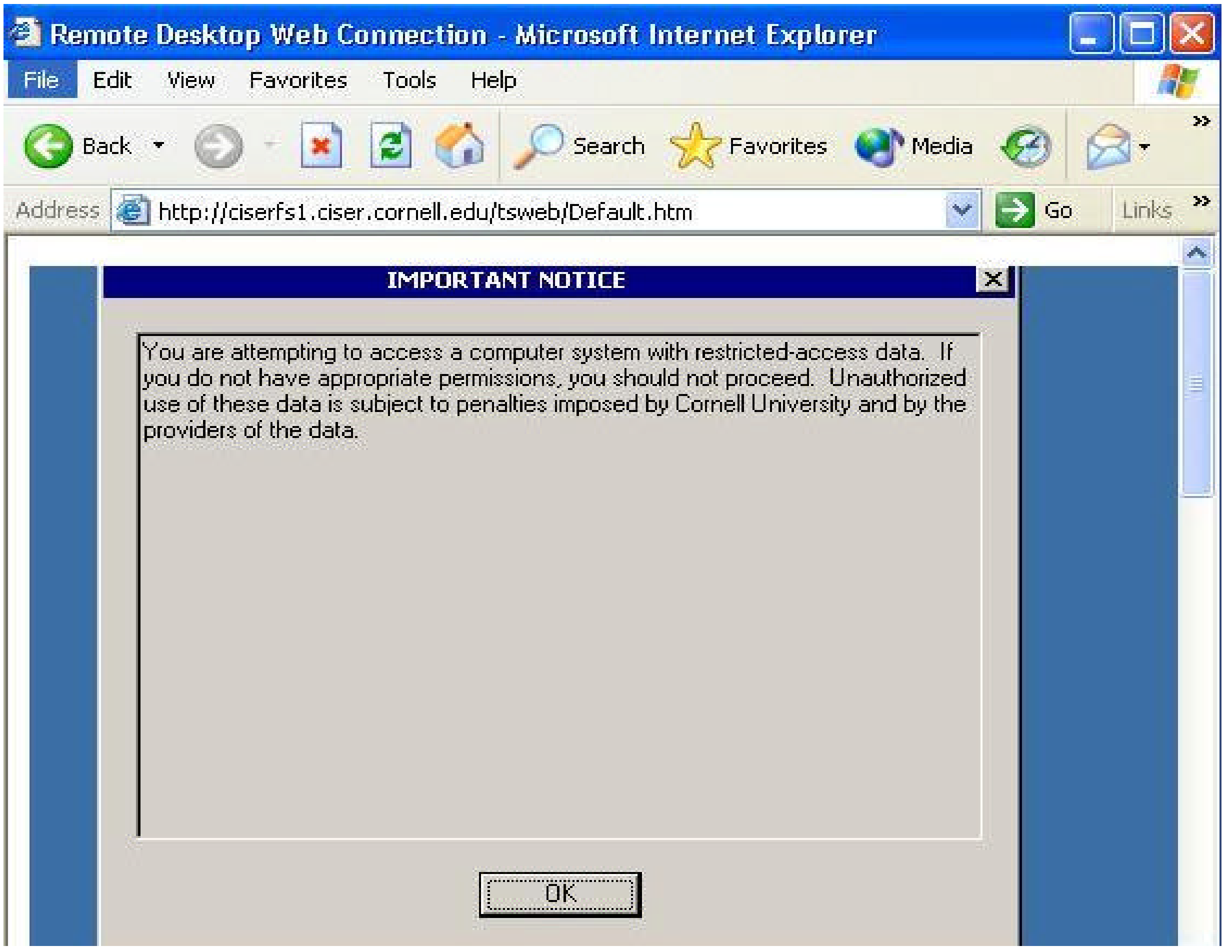

Address: http://ciserfs1.ciser.cornell.edu/tsweb/Default.htm

IMPORTANT NOTICE

You are attempting to access a computer system with restricted-access data. If you do not have appropriate permissions, you should not proceed. Unauthorized use of these data is subject to penalties imposed by Cornell University and by the providers of the data.

III. The system will then provide you with a logon window. Enter your CRADC user name and password and click on OK. A sample logon window is shown below.

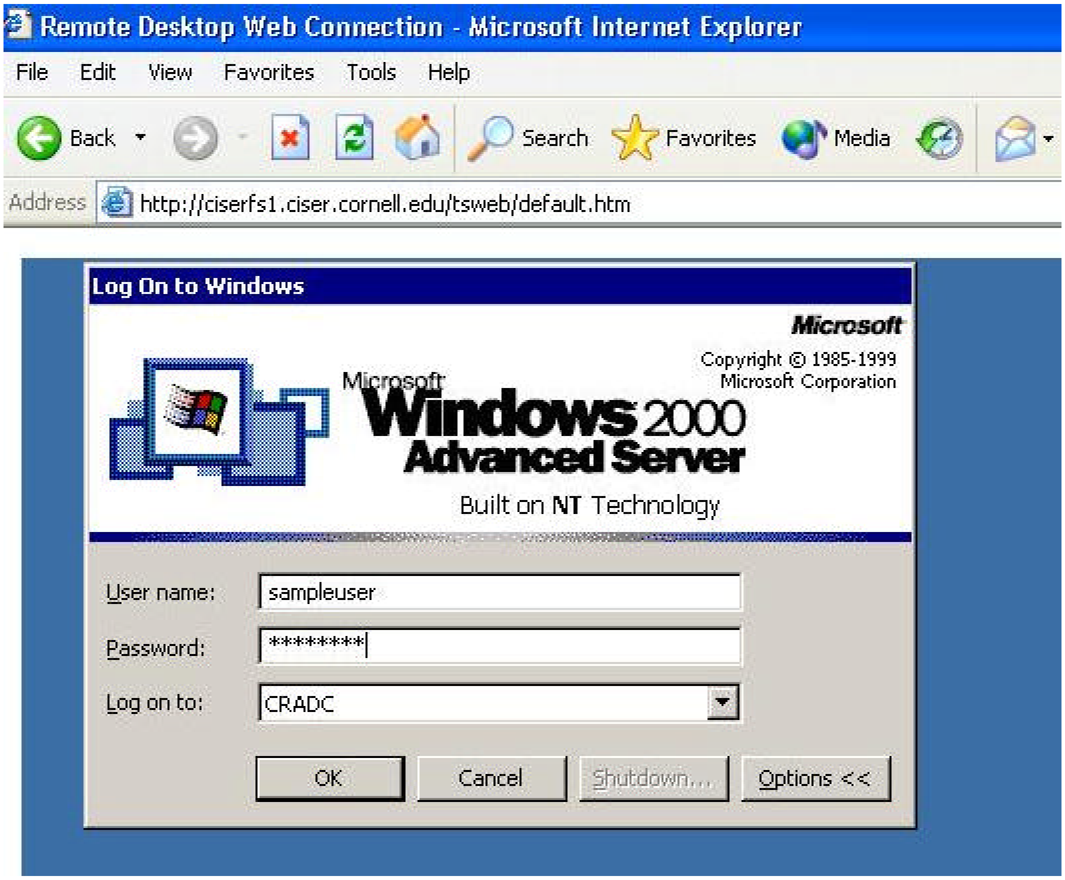

Address: http://ciserfs1.ciser.cornell.edu/tsweb/default.htm

User name: sampleuser

Password: ********

Log on to: CRADC

9. *How to Close a CRADC Session: Logoff vs Disconnect (for New Users)*

- **DISCONNECT:** To exit your CRADC session while your applications are still running.
  - Do this only when you have programs that you wish to continue to run after you close your connection.
  - To disconnect:
    - go to "Start" --> "Shutdown" --> and choose "disconnect ".
    - OR simply close the window by clicking the X box in the upper right corner of the window.
  - You will remain logged on, but disconnected. When you next attempt to connect to that node you will automatically be returned to that session.
  - Note: Disconnected sessions let CRADC users run their programs uninterrupted for days, and also give them the flexibility of preserving their work environments. **However, we strongly recommend loging off your CRADC sessions, whenever possible**. These forgotten disconnected

sessions can result in account lockout problems when the user changes her/his password as required by the system.

- **LOG OFF**: TO END YOUR SESSION
  - Go to " Start" --> "Shutdown" --> and choose "Log off <*your username>"*
  - TIP! To add a "Log off" option to your Start menu on a CRADC node:
    - Go to "Start" --> "Settings" --> "Taskbar & Start Menu". Select the "Advanced " tab. Check the box beside " Display Logoff". Choose " OK".
    - After that you can log off simply by going to "Start" --> "Log off <*your username>"*.

If you encounter problems, contact Pinky Chandra (the CRADC manager) and she will help you resolve them (pc17@cornell.edu).

The following document describes in detail the steps needed to Connect to CRADC from inside a firewall

How to use SSH Windows clients to tunnel to CRADC

Lars Vilhuber

October 11, 2002

1 Quick reference

- Configure an SSH session to *yeehee.ciser.cornell.edu* that allows for tunneling of a TSC session (Section 4 on page 4). This is done only once.
- Set up a TSC session, and save it (Section 4.2 on page 9).
- In every day use, connect to *yeehee.ciser.cornell.edu*, thus establishing a tunnel.
- Then connect to *localhost* using the TSC client.

2 Overview

This document will describe how to use SSH software to exit a firewall-protected computer network with open SSH ports (port 22) in order to allow other applications to connect to another (potentially) protected network. The basic principle described here requires an intermediate SSH server. It is a way to set up a quasi-VPN for TCP-based applications.

We reference specific systems here, but the general setup is generic enough to be used for other systems. In particular, we will be describing the setup needed to allow for a (Windows) Terminal Services Client (TSC) on an arbitrary firewall-protected network to connect to secure CRADC servers called *cradcXX.ciser.cornell.edu* running at Cornell, using the intermediate SSH server *yeehee.ciser.cornell.edu* and the standard TSC port TCP-3389.[1]

[1] The setup as described is not a true VPN, since traffic between *yeehee* and *cradcXX* is not encrypted by the tunnel. It is, however, still encrypted by any encryption used by the application being tunnelled. In the case of CRADC and TSC, this is 128 bit encryption as provided by Windows Terminal Services.

3 Obtaining Software

3.1 SSH

The software needed is the widely available SSH client software. Several implementations are available for free for certain licensees (home and academic users in particular). The following is an (incomplete) list of available software. The first listed is the one to be used in the illustrated examples, your mileage with the other SSH clients may vary. The server on *yeehee* is running SSH Secure Shell Server from SSH Communications, and not OpenSSH.

OS: Windows

1. SSH Secure Shell for Workstations (http://www.ssh.com), Academic and Non-commercial use free, includes SFTP client.
2. SecureCFT (http://www.vandyke.com/), downloadable trial version, no free version, SFTP (SecureFX) client separate. At the present time, there seem to be some problems with this client.
3. Entunnel (also http://www.vandyke.com/), downloadable trial version, no free version, tunneling capability ONLY (no terminal). Not tested.
4. PuTTY (http://www.chiark.greenend.org.uk/ sgtatham/putty/), free for all uses, SFTP client included. Not tested.

OS: Linux/Unix

1. SSH Secure Shell for Workstations (http://www.ssh.com), all uses free, includes (command-line) SFTP client.
2. OpenSSH (http://www.openssh.com/), all uses free, some incompatibilities with SSH Secure Shell for Workstations exits. In particular, the SSH and OpenSSH implementations of SFTP are incompatible.

3.2 Terminal Services Client

OS: Windows

1. Microsoft TSC is often pre-installed in Windows 2000, and always in XP. To run it, click Start, click All Programs, click Accessories, click Communications, and then click Remote Desktop Connection. If not, a full client can be downloaded from
   http://www.microsoft.com, at
   http://www.microsoft.com/windowsxp/pro/downloads/rdclientdl.asp.
2. An ActiveX applet (for those without Administrator privileges) can also be downloaded from
   http://www.microsoft.com/windows2000/downloads/recommended/TSAC/default.asp.
3. Online: http://ciserfs1.cornell.edu/tsweb/Default.htm

OS: Linux

1. WinConnect (http://www.thincomputinginc.com/) works fine, and allows to use arbitrary TSC ports.
2. rdesktop (http://www.rdesktop.org) is a free RDP client. Mileage varies, and it may not support (as of version 1.1.0) the high encryption used on CRADC servers.

4 Setting up the tunnel

The following assumes that

- a TSC connection is to be tunnelled
- from *client.anywhere.com*, short: *client*
- to *cradc01.ciser.cornell.edu*, short: *cradc01*
- via an accessible SSH server on *yeehee.ciser.cornell.edu*, short: *yeehee*
- Port 22 (ssh standard) is open for at least outgoing connections on the firewall surrounding *client*.
- Port 3389 is open for incoming connections on the firewall surrounding *cradc01*

The following setup implies that

- Only one TSC connection can be open at any time on *client* using tunnelling[2]
- The connection is protected by SSH encryption between *client* and *yeehee*, and protected by native TSC encryption between *yeehee* and *cradc01*.

[2] This is a restriction of the Microsoft TSC software, not of the procedure described here, and it does not apply when using the WinConnect Linux client.

4.1 Initiating a session to YEEHEE

- First, set up a standard session to *yeehee*. In your SSH client click on the "Quick Connect" button (Figure 1 for SSH Secure Shell Client)

Figure 1: The Quick Connect button, Secure Shell Client

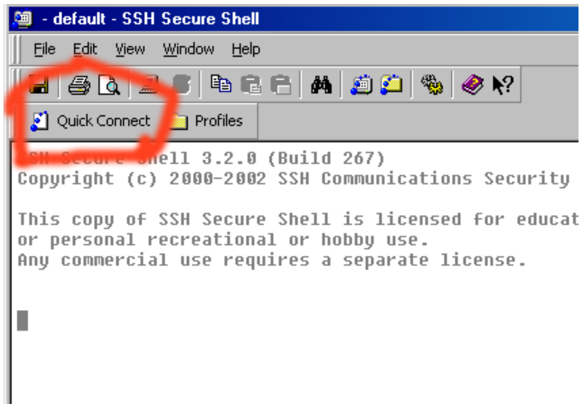

```
SSH Secure Shell 3.2.0 (Build 267)
Copyright (c) 2000-2002 SSH Communications Security

This copy of SSH Secure Shell is licensed for educat
or personal recreational or hobby use.
Any commercial use requires a separate license.
```

- Fill in the information (Host Name is "yeehee.ciser.cornell.edu", your User Name, and the rest as is) (Figure 2).

Figure 2: The Quick Connect dialog

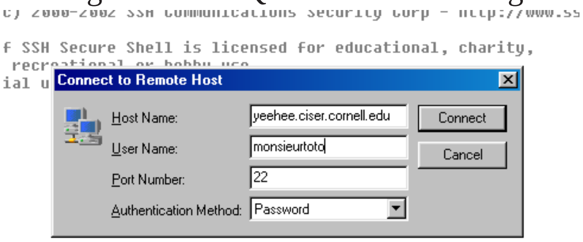

Host Name: yeehee.ciser.cornell.edu

User Name: monsieurtoto

Port Number: 22

Authentication Method: Password

- If this is the first time you connect to *yeehee.ciser.cornell.edu*, then you will be prompted if you want to accept the identifying key provided by the server (Figure 3). This is fine the first time around, but if this occurs in later connection attempts, then it means either that changes have occurred on the server, or that you are not connecting to the server you thought you were connecting to (a so-called "man-in-the-middle" attack). The numbers in Figure 3 are the RIGHT numbers, you can check that when you connect.

Figure 3: Verifying the identity of the server

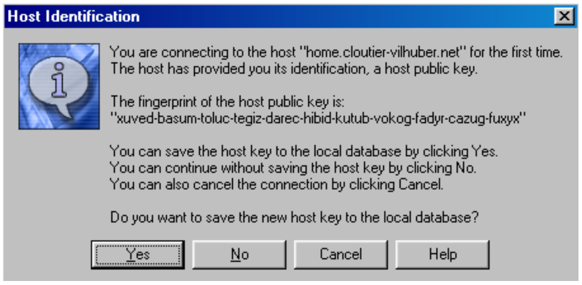

You are connecting to the host "home.cloutier-vilhuber.net" for the first time.
The host has provided you its identification, a host public key.

The fingerprint of the host public key is:
"xuved-basum-toluc-tegiz-darec-hibid-kutub-vokog-fadyr-cazug-fuxyx"

You can save the host key to the local database by clicking Yes.
You can continue without saving the host key by clicking No.
You can also cancel the connection by clicking Cancel.

Do you want to save the new host key to the local database?

- Once connected, you have the option to save the settings to a "profile". Give the profile a name ("yeehee" would be fine), and save. (Figure 4)

Figure 4: Saving the profile

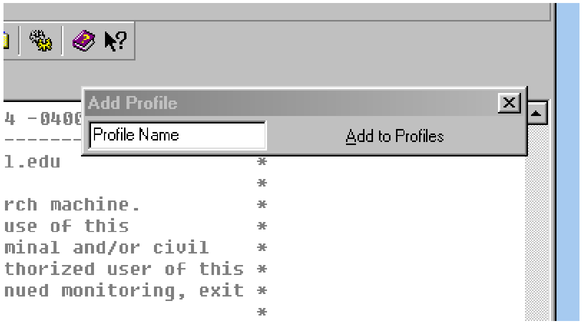

```
4 -0400
-------
l.edu                             *
                                  *
rch machine.                      *
use of this                       *
minal and/or civil                *
thorized user of this             *
nued monitoring, exit             *
                                  *
```

- After disconnecting from this session (by typing “exit”), click on the “Profile” button, and choose “Edit Profile...” (Figure 5).

Figure 5: Editing the profile

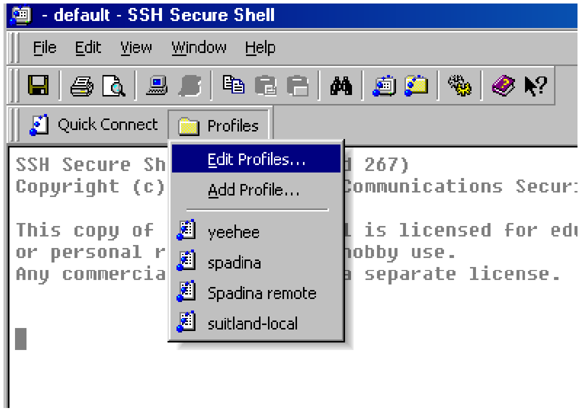

```
SSH Secure Sh                  l 267)
Copyright (c)                 Communications Secur

This copy of                   l is licensed for edu
or personal r                 hobby use.
Any commercia                  a separate license.
```

- Choose the profile just saved (i.e. “yeehee”) in the left panel, and click on the “Tunneling” tab in the right panel (Figure 6).

Figure 6: Creating a tunnel

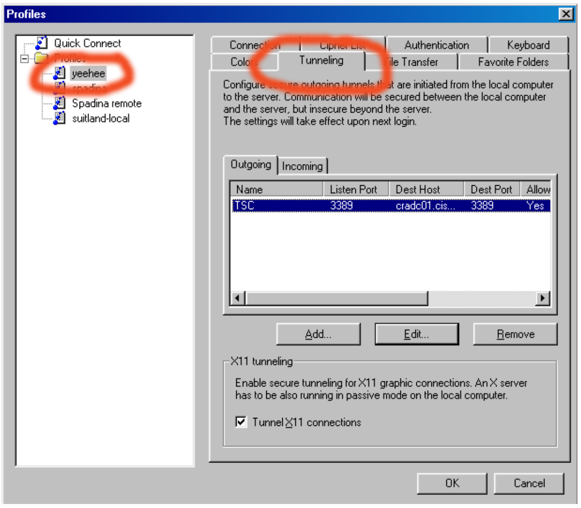

Configure secure outgoing tunnels that are initiated from the local computer to the server. Communication will be secured between the local computer and the server, but insecure beyond the server.
The settings will take effect upon next login.

Name	Listen Port	Dest Host	Dest Port	Allow
TSC	3389	cradc01.cis...	3389	Yes

Enable secure tunneling for X11 graphic connections. An X server has to be also running in passive mode on the local computer.

☑ Tunnel X11 connections

- Click on "Add..." to add another tunnel connection. In the resulting dialog (Figure 7), fill in the required information:
  - Display Name: An arbitrary name to identify this tunnel (i.e. "TSC")
  - Type: Choose "TCP" (the default)
  - Listen Port: "3389" (the standard TSC port)[3]
  - Destination Host: "cradc01.ciser.cornell.edu" (or any other connection where you want to connect to)
  - Destination Port: "3389" (the standard TSC port)[4]

  Click "OK" to accept, and "OK" on the profile display to finish editing this profile.

Figure 7: Edit dialog

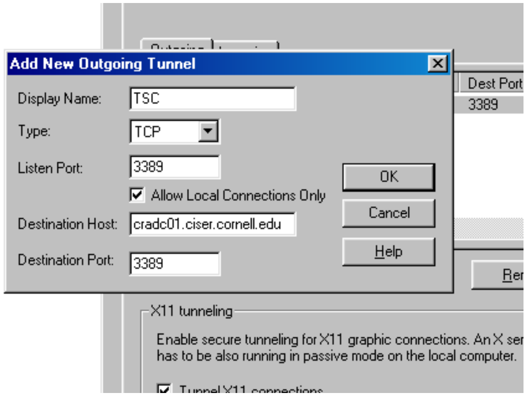
Add New Outgoing Tunnel

Display Name: TSC

Type: TCP

Listen Port: 3389

☑ Allow Local Connections Only

Destination Host: cradc01.ciser.cornell.edu

Destination Port: 3389

Dest Port 3389

X11 tunneling

Enable secure tunneling for X11 graphic connections. An X ser
has to be also running in passive mode on the local computer.

☑ Tunnel X11 connections

- You are now set to initiate the tunnel. Click again on the "Profile" button (see Figure 5 on the preceding page), and choose the profile you just modified "yeehee". Log in using your password, and the tunnel is established.

[3] For WinConnect, this can be any other port, see Section 4.2 on the next page, footnote.

[4] This MUST be 3389 for ANY client, except if you are connecting to a TSC server you know to be listening on other ports.

4.2 Configuring TSC connection

We now need to initiate the TSC connection. The following example will be based on the Windows full client.

- Normally, i.e. in absence of a firewall, the connection is straightforward: enter the name of the server and connect (Figure 8). More options, such as a chosen connection speed, the screen size, and other options, are accessible by clicking on the "Options¿¿" button (Figure 9

Figure 8: Standard TSC connection

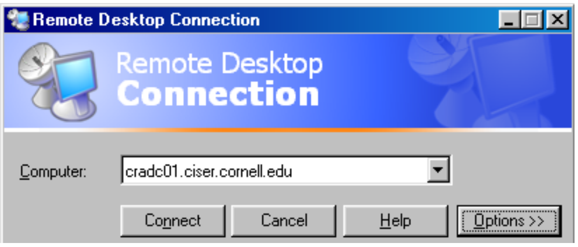

Computer: cradc01.ciser.cornell.edu

Figure 9: Standard TSC connection with options

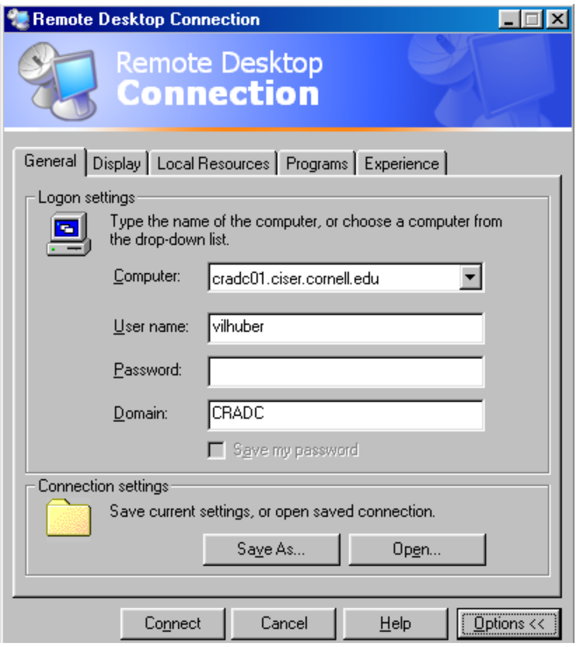

Logon settings

Type the name of the computer, or choose a computer from the drop-down list.

Computer: cradc01.ciser.cornell.edu

User name: vilhuber

Password:

Domain: CRADC

□ Save my password

Connection settings

Save current settings, or open saved connection.

- For the case of tunnelling, all options are accessible, and the only item to be changed is the destination host ("Computer"). Instead of the final destination, you will enter "localhost" in that dialog. Click on "Connect", and the connection will be established via the tunnel previously created.

- For ease of access, this session can be saved (in the expanded “Options” dialog).
- In essence, the TSC software *thinks* it is connecting to a computer with the named address “localhost”, which is actually redirected to the computer in the tunnel created above (in this case, *cradc01.ciser.cornell.edu*.