

READYNAS INSTANT STORAGE

# Using Rsync for NAS-to-NAS Backups

---



© Infrant Technologies  
3065 Skyway Court, Fremont CA 94539  
[www.infrant.com](http://www.infrant.com)

# Using Rsync For NAS-To-NAS Backups

You've heard it before, but it's worth repeating—an organization's most valuable asset is the data that comprises its information store. With ever increasing amounts of data, and the need to find an efficient method for ensuring its confidentiality, integrity, and availability, organizations often find themselves struggling with managing data stores.

The ReadyNAS product line (1000S, NV, 600, and X6) of network attached storage devices provides a cost-effective way for organizations to share information without investing in expensive server-class hardware, bloated network operating systems, and seemingly endless client access licenses. Yet, they provide many of the same benefits as traditional file serving methods, including support for multiple sharing protocols as well as share-based, user-based, and domain-based security models. At the same time, because of their large data store capabilities, they need to be backed up as efficiently as possible.

In this guide, we'll focus our attention on implementing NAS-to-NAS backups using Rsync, a native file synchronization service included with the ReadyNAS devices. In our presentation, you'll discover methods for backing up NAS devices on a LAN, between subnets isolated from each other by a firewall, and across a WAN connection using Virtual Private Networking (VPN) technologies. Moreover, each method we'll present builds on the previous one, so you can implement the appropriate backup strategy as your needs change.

## Rsync vs. traditional backups

---

Traditional backup methods typically copy source data into a proprietary file format for archival purposes. If you need to access the data that's been backed up, you need to restore it from the archive. Rsync, which has its roots in the UNIX world, operates differently. Instead of copying data into an archive, it keeps the data in its native format, meaning it's readily available without going through a timely restore process.

The first time you run an Rsync backup job, a full backup is performed, i.e. all the data is copied from the source NAS device to the destination NAS device. When you run the backup job subsequent to the first backup, only incremental changes in the source data are copied to the destination. This makes Rsync an efficient and practical method for ensuring that two locations have the same data store.

## The process

---

In this guide, we'll discuss several different network topologies and walk you through the process of performing NAS-to-NAS backups unique to each one. More specifically, we'll:

- ✓ Set up two ReadyNAS devices on a local LAN segment.
- ✓ Create an Rsync-based backup job and verify its operation.
- ✓ Relocate one ReadyNAS device to a remote location, residing behind a NAT firewall.
- ✓ Configure the NAT device to accept incoming Rsync requests.
- ✓ Create a virtual private network with two VPN endpoints.
- ✓ Define a gateway and network policy to support Rsync backups over a VPN.
- ✓ Discuss snapshots, a feature unique to ReadyNAS devices that allows you to freeze your data a given point in time.
- ✓ Tie everything together by providing an example of a typical backup scenario.

## System requirements

---

To complete the tasks presented in this article, you'll need:

- ✓ A LAN segment with the appropriate switches and cabling.
- ✓ An administrative workstation with a web browser installed.
- ✓ Two ReadyNAS devices.
- ✓ A NAT router.
- ✓ Two VPN endpoint routers.

## Rsync in practice

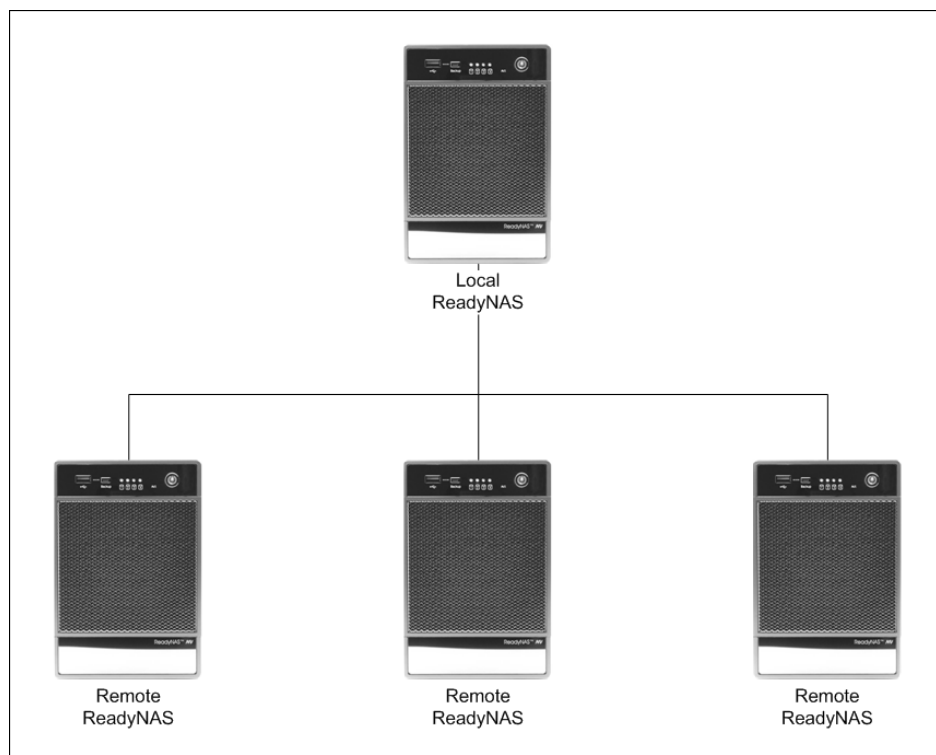
---

Consider the topology illustrated in **Figure A**. In this scenario, a local ReadyNAS can maintain a copy of all the data generated at one or more remote locations. For example, the local device can reside at an off-site location, such as a home office or other facility, and back up the data stored at a remote business site. Rsync can accomplish this goal in two different ways:

- ✓ Each remote device can “push” its data to the local device.
- ✓ The local device can “pull” data from each remote device.

If you choose the “push” solution, you need to create a backup job at each remote site. However, this can add administrative overhead to your organization because you need to maintain separate backup jobs residing at potentially different locations. Choosing the “pull” solution is more efficient because, although you still need to create a backup job for each remote site, all the backup jobs are maintained on the local device.

In general, you can think of each location as a separate network, either internal or external to your organization. For the purpose of this guide, the ReadyNAS device hosting the backup jobs is the local device and the ReadyNAS devices you want to back up via Rsync are the remote devices.




**Figure A:** A ReadyNAS device can be either a local or a remote device.

## Deploy Rsync on your LAN

---

Because each backup job runs initially as a full backup, you'll want to create and test it locally before implementing it across subnets or over a WAN connection. Doing so not only makes the initial backup run more efficiently but also ensures that the job operates as expected. Once you have everything working properly, you can relocate the remote device to its intended location.

Because we want the local device to “pull” data from the remote device, we'll set up the local device as the Rsync client and the remote device as the Rsync server, as shown in **Figure B**. We'll provide the details on accomplishing this in the upcoming sections. For now, let's just set up the devices on a local LAN segment, so that we can create and test a typical backup job.



**Figure B:** *The remote device makes its shares available to the local device.*

### To set up the ReadyNAS devices on your LAN:

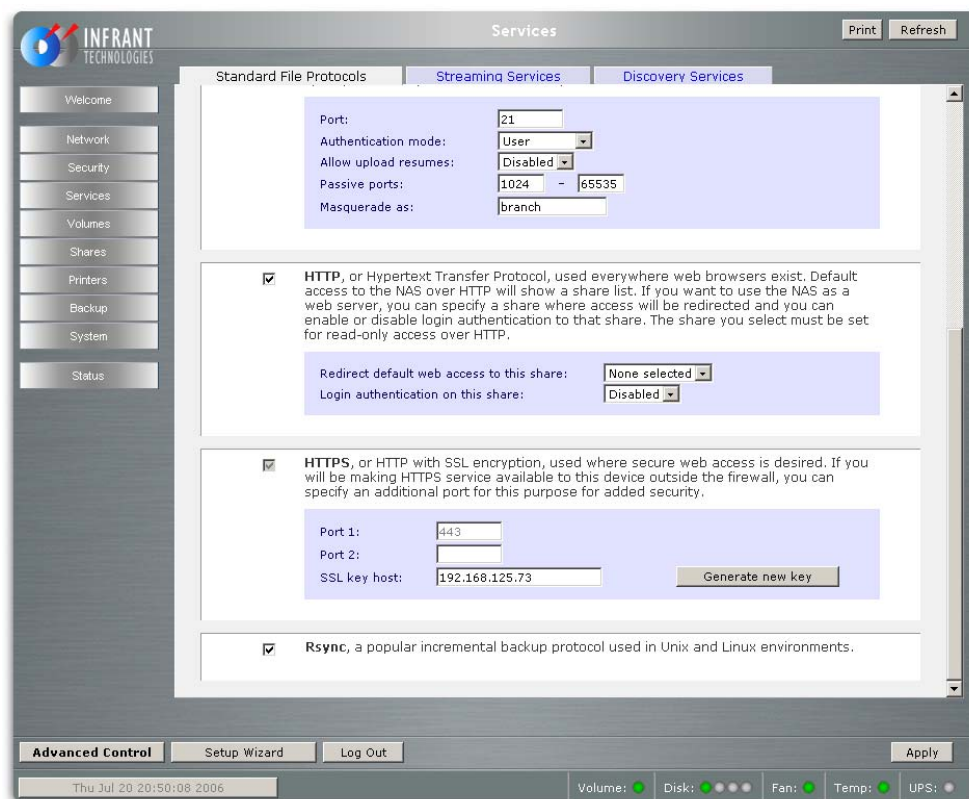
1. Physically connect each device to the same subnet on your LAN.
2. Use the RAIDar discovery utility to complete the initial setup of each device.
3. Name one device *local-nas* and the other device *remote-nas* to make following our examples easier.

## Enable the Rsync service

After you install the ReadyNAS devices, you need to enable the Rsync service on both the local and the remote devices. This allows you to easily reverse the roles of the devices if your needs change. When you enable the Rsync service, you're essentially allowing shares on the ReadyNAS device to act as Rsync servers; Rsync clients can then connect to these servers. As noted previously, the remote device acts as an Rsync server (it makes its data available via Rsync) while the local device acts as an Rsync client (it retrieves data from the Rsync server).

### To enable the Rsync service:

1. Log in to each ReadyNAS device using your administrative credentials.
2. Navigate to the Standard File Protocols tab, as shown in **Figure C**.



**Figure C:** Enable the Rsync service on both ReadyNAS devices.

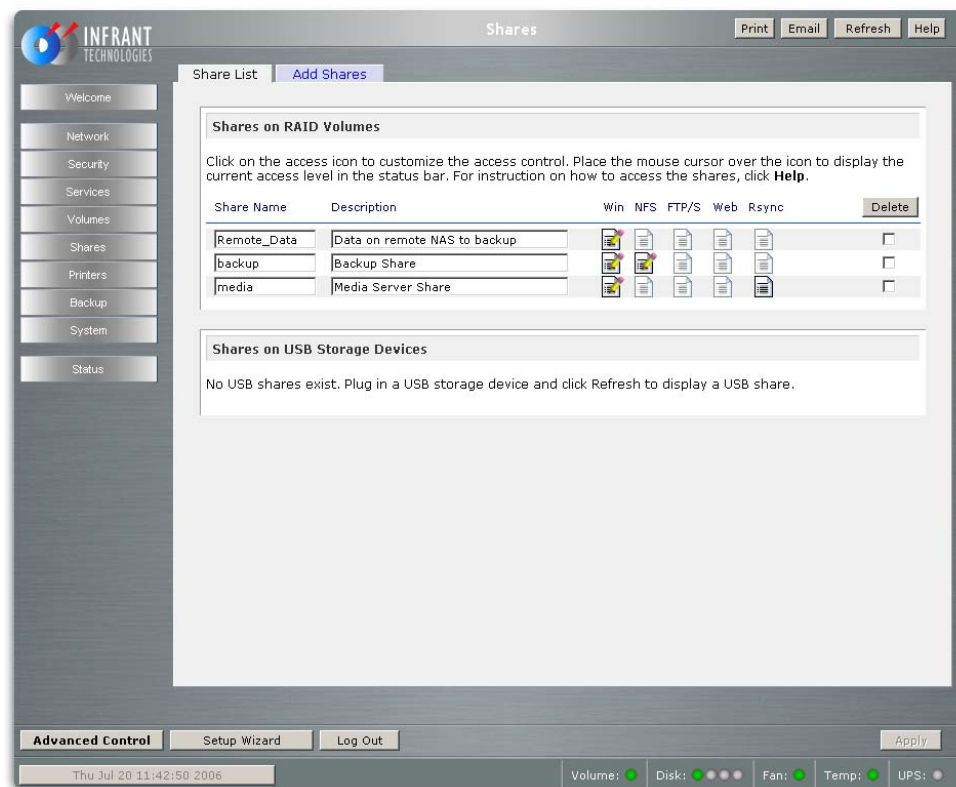
3. Scroll to the bottom of the list of available services, and then enable the Rsync check box.
4. Click the Apply button.

## Configure the Rsync protocol

After enabling the Rsync service, you need to configure the Rsync protocol for each share that you want to back up. Doing so enables the respective share to assume the role of an Rsync server. This involves specifying the share's default access (Read/Write or Read Only) as well as optionally restricting access to it. You can restrict access to the share in two different ways: by specifying allowed hosts and by specifying login credentials.

### To configure the Rsync protocol:

1. Navigate to the remote device's Share list tab, as shown in **Figure D**.



**Figure D:** The Rsync column appears only if you've enabled the Rsync service.

2. Identify the share that you want to back up via Rsync, e.g. *Remote\_Data*.
3. Click the document icon located at the intersection of the share name and the Rsync column.

4. On the Rsync tab that appears and is shown in **Figure E**, complete the required information in the Share Access Restrictions and the Rsync Password Option sections.

The screenshot shows the 'Shares' configuration window for Infrant Technologies. The 'Rsync' tab is selected. The 'Share Name' is 'Remote\_Data' and the 'Default Access' is 'Read-only'. The 'Share Access Restrictions' section has a checkbox for 'Hosts allowed access' checked, with the value '192.168.202' entered. The 'Rsync Password Option' section has a checkbox for 'Enable password protection' checked. Below this, there are three rows for user login and password: 'User login 1: admin' with a password field, 'User login 2:' with an empty password field, and 'User login 3:' with an empty password field. The window includes a sidebar with navigation links, a top bar with 'Print', 'Email', and 'Refresh' buttons, and a bottom bar with 'Advanced Control', 'Setup Wizard', 'Log Out', 'Apply', and 'Cancel' buttons. A status bar at the very bottom shows the date 'Thu Jul 20 12:05:24 2006' and hardware status indicators for Volume, Disk, Fan, Temp, and UPS.

**Figure E:** To protect the source data, specify the Default Access as Read-Only.

5. Click the Apply button.



## Create a backup job

Now that you've enabled the Rsync service and you've configured the Rsync protocol for the share you want to back up, you need to create a backup job. When you create the backup job, you specify the Rsync server as the backup source and the Rsync client as the backup destination.

The path to the Rsync Server takes the form *host::module/path*, where *host* is the name or IP address of the remote device, *module* is the name of the share you want to backup, and */path* is an optional subdirectory of the source share. The path to the Rsync Client takes the form */path*, where */path* is an optional subdirectory of the destination share on the local device.

### To create a backup job:

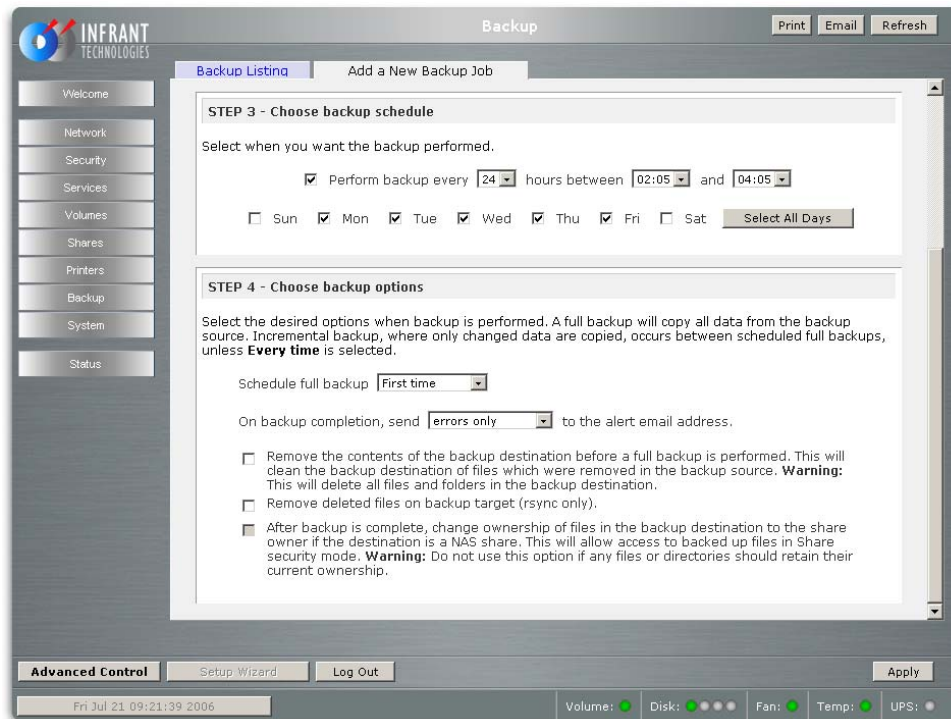
1. On the local device, navigate to the Add a New Backup Job tab.
2. Specify the Backup Source and Backup Destination, as shown in **Figure F**.

The screenshot shows the 'Backup' configuration window for Infrant Technologies. The window has a sidebar on the left with a 'Welcome' message and a list of system components: Network, Security, Services, Volumes, Shares, Printers, Backup, System, and Status. The main area is titled 'Backup' and has tabs for 'Backup Listing' and 'Add a New Backup Job'. The 'Add a New Backup Job' tab is active, showing two steps: 'STEP 1 - Select backup source' and 'STEP 2 - Select backup destination'. Step 1 includes a dropdown for 'Remote: Rsync Server', a 'Path' field with the value '192.168.100.201::Remote\_Data', a 'Login' field with 'admin', and a 'Password' field with masked characters. A 'Test connection' button is below these fields. Step 2 includes a dropdown for 'Share: Local\_Data', a 'Path' field, a 'Login' field, and a 'Password' field. Another 'Test connection' button is below these fields. At the bottom of the window, there are buttons for 'Advanced Control', 'Setup Wizard', 'Log Out', and 'Apply'. A status bar at the very bottom shows the date and time 'Thu Jul 20 12:00:17 2006' and hardware status indicators for Volume, Disk, Fan, Temp, and UPS.

**Figure F:** The remote device is the backup source and the local device is the backup destination.

3. Click the Test Connection button to ensure that the local device can communicate with the remote device.

4. In the Choose Backup Schedule section shown in **Figure G**, schedule your backup job to run during off-peak hours; this avoids server overload and traffic congestion.



**Figure G:** You can schedule backup jobs as well as specify the frequency of full backup.

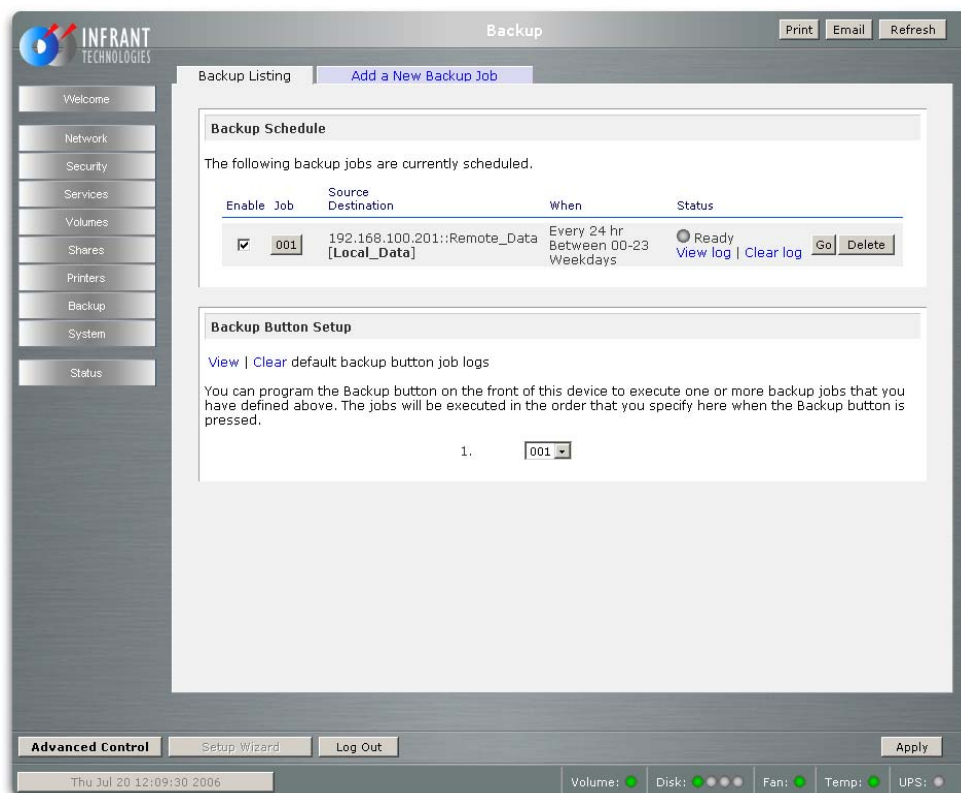
5. In the Choose Backup Options section, choose a time interval to perform a full backup, the desired logging level, and any other options you want.
6. Click the Apply button.

## Test the backup job

After you create the backup job, it appears on the Backup Listing tab. When you run the backup job, a full backup is performed, i.e. all the files from the source are copied to the destination. Each subsequent run copies only those files that have changed on the source device, i.e. incremental changes. You should run the job at least twice—once to perform a full backup and once to perform the first incremental backup, so you can verify that it operates as expected.

### To test the backup job:

1. On the Backup Listing tab shown in **Figure H**, make sure the job you want to run is enabled.



**Figure H:** On the ReadyNAS NV, you can associate a backup job with the backup button located on the front of the device.

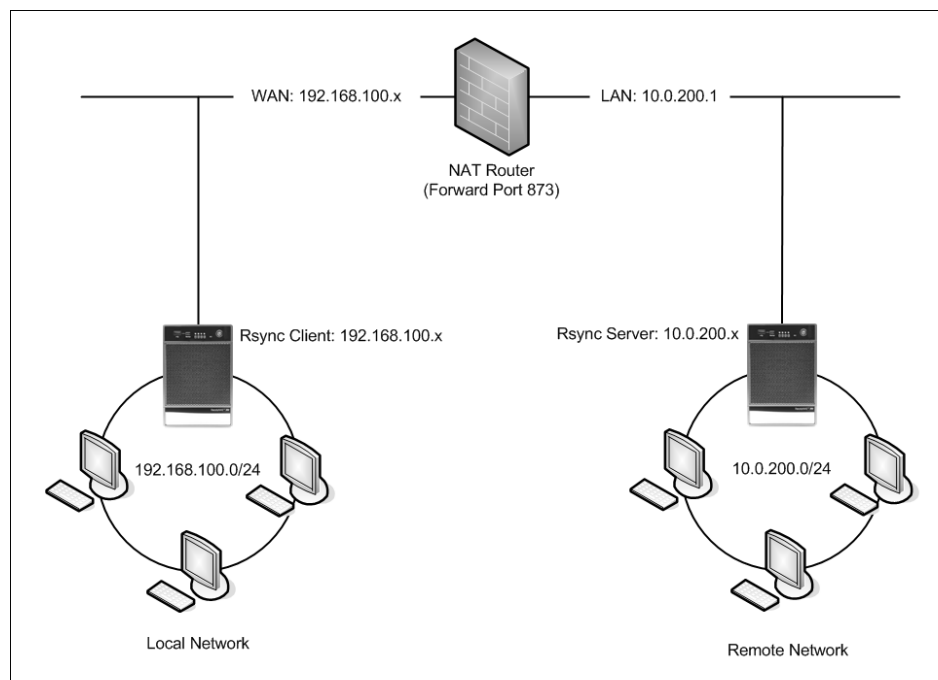
2. Click the Go button and wait for the job to complete.
3. Click on the View Log link to examine the backup job log.
4. If you need to modify the backup job to correct errors, click the job number button, e.g. 001.

## Deploy Rsync across network boundaries

The next scenario we'll examine is deploying Rsync across network boundaries. In particular, we'll examine its operation when a firewall exists between two networks. The two networks could be distinct subnets on your LAN (e.g. different departmental workgroups), or they could be a WAN-LAN pair (e.g. a business site and a home office).

In the topology shown in **Figure I**, the remote network is sitting behind a NAT router with an integrated firewall. As you can see, the Rsync server's IP address and the NAT router's internal IP address are on the same subnet, namely 10.0.200.0/24. You can also see that the IP address of the Rsync client and the NAT router's external IP address are on an entirely different subnet, namely 192.168.100.0/24. It's worth noting that the latter subnet could also be a WAN segment with a public rather than a private IP address.

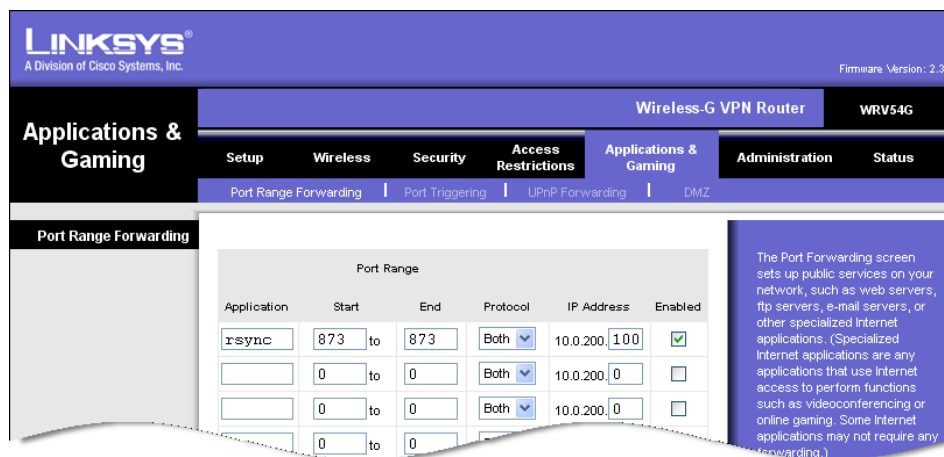
In order for the Rsync client to connect to the Rsync server, the NAT router must forward incoming Rsync requests to the appropriate ReadyNAS device on the remote network. So, you must open TCP/UDP port 873 on the NAT router, and forward Rsync requests to 10.0.200.x in this example.



**Figure I:** You must open port 873 on the NAT router to support Rsync across network boundaries.

## To deploy Rsync across network boundaries:

1. Install the local ReadyNAS device on the WAN side of the firewall.
2. Install the remote ReadyNAS device on the LAN side of the firewall.
3. Log in to the NAT router's management utility.
4. Navigate to the router's Port Forwarding section; consult your router's documentation if necessary.
5. Specify that incoming Rsync requests on port 873 be forwarded to the remote device's IP address, as shown in **Figure J**.
6. Edit the backup job you created earlier such that the IP address of the backup source is the same as the IP address of the router's WAN port.
7. Run the backup job to test its operation.



**Figure J:** Rsync requests received on a WAN port can be forwarded to a ReadyNAS device on your LAN.


## Deploy Rsync over a VPN

While transferring data across networks is a business necessity, so also is ensuring the information doesn't fall into the wrong hands. For example, a malicious user can sniff data in transit and steal confidential information, or he can perform a man-in-the-middle attack and modify critical information during transit. Although Rsync doesn't provide any native facility to encrypt data, it can be used effectively over a VPN connection.

In brief, a VPN allows you to create a secure tunnel (between two endpoints) over a public network, such as the Internet. There are three common types of VPNs, listed here from least to most secure:

- ✓ Point-to-Point Tunneling Protocol (PPTP)
- ✓ Layer 2 Tunneling Protocol (L2TP)
- ✓ IP Security (IPSec).

Each type of VPN offers its own authentication and encryption algorithms, and a complete discussion of each is beyond the scope of this guide. Our discussion focuses primarily on creating the necessary gateway and network policies to support a VPN. And finally, because the configuration details can vary significantly between router manufacturers, we'll provide only a general outline of the procedures. Let's start by examining the scenario shown in **Figure K**.



**Figure K:** Using a VPN, you can ensure that your Rsync backup jobs are secure.

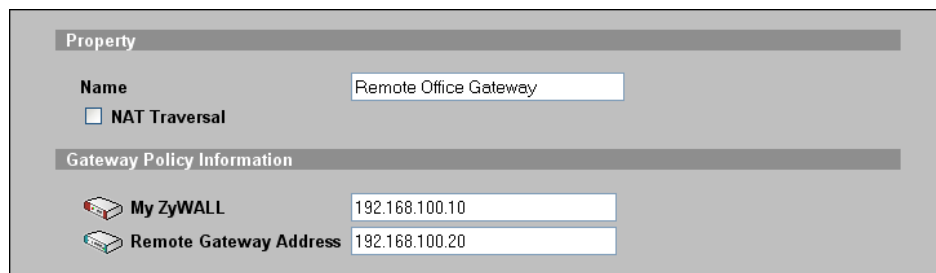
In this topology, note the following:

- ✓ Each ReadyNAS device (an Rsync client and an Rsync server) resides on its own private network.
- ✓ Each private network is connected to an untrusted network (192.168.100.0/24) via a VPN router. Typically, the untrusted network is the Internet, but it can also be parent LAN segment as in our example.
- ✓ A VPN tunnel on the untrusted network ensures a secure connection between the WAN side of each VPN router, i.e. the tunnel endpoints.

When the Rsync client on the local network runs its backup job, the Rsync request passes through the VPN tunnel to the WAN port on the VPN router at the remote network. The remote VPN router forwards the request (via TCP/UDP port 873) to the Rsync Server on the remote network. In order to set this up, you typically need to create a gateway policy, which defines the endpoints of the VPN tunnel, and a network policy, which defines the visibility of each private network.

#### To create a remote gateway policy:

1. Log in to the local VPN router using your administrative credentials.
2. Navigate to the Gateway Policy section, as shown in **Figure L**.



The screenshot shows a web-based configuration interface for a ZyWALL router. It features a 'Property' section with a 'Name' field set to 'Remote Office Gateway' and an unchecked 'NAT Traversal' checkbox. Below this is a 'Gateway Policy Information' section containing two fields: 'My ZyWALL' with the IP address '192.168.100.10' and 'Remote Gateway Address' with the IP address '192.168.100.20'.

**Figure L:** You need to define a remote gateway policy on the local VPN router.

3. Provide a descriptive name for the policy you're creating, e.g. *Remote Office Gateway*.
4. Enter the IP address of the WAN side of the remote VPN router.
5. Enter the IP address of the WAN side of the remote VPN router.
6. Define the authentication method to use, typically a pre-shared key or an SSL certificate.

## To create a remote network policy:

1. Log in to the local VPN router using your administrative credentials.
2. Navigate to the Network Policy section, as shown in **Figure M**.

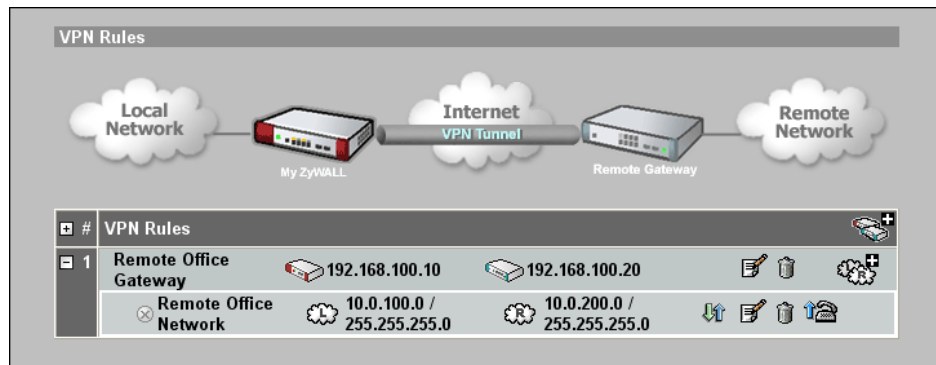
The screenshot shows a web-based configuration interface for a VPN router. It is divided into several sections:   
1. **Property**: Contains checkboxes for 'Active', 'Nailed-Up', 'Allow NetBIOS Traffic Through IPSec Tunnel', and 'Check IPSec Tunnel Connectivity'. The 'Name' field is set to 'Remote Office Network', 'Protocol' is '0', and 'Ping this Address' is '0.0.0.0'.   
2. **Gateway Policy Information**: Shows 'Gateway Policy' set to 'Remote Office Gateway'.   
3. **Local Network**: Includes 'Address Type' (Subnet Address), 'Starting IP Address' (10.0.100.0), 'Ending IP Address / Subnet Mask' (255.255.255.0), and 'Local Port' (Start 0, End 0).   
4. **Remote Network**: Includes 'Address Type' (Subnet Address), 'Starting IP Address' (10.0.200.0), 'Ending IP Address / Subnet Mask' (255.255.255.0), and 'Remote Port' (Start 0, End 0).

**Figure M:** You define a remote network policy on the local VPN router and associate it with the remote gateway policy.

3. Provide a descriptive name for the policy you're creating, e.g. *Remote Office Network*.
4. Associate the network policy with the gateway policy you created earlier.
5. Specify the IP address space on the local network that should be visible to the remote network.
6. Specify the IP address space on the remote network that should be visible to the local network.
7. Define the data encryption method to use.



After you create the gateway and network policies, review the tunnel summary, as shown in **Figure N**. As you can see, the gateway policy defines the endpoints of the VPN tunnel while the network policy defines the two private networks sitting behind the VPN routers.



**Figure N:** Make sure the remote gateway and the network policies are correct before proceeding.

#### To test the backup job:

1. Edit the backup job you created earlier such that the IP address of the backup source is the same as the IP address of the remote VPN router's WAN port.
2. On the local device, open the tunnel by clicking the Dial button.
3. Run the backup job to test its operation.

## Combine backups with snapshots

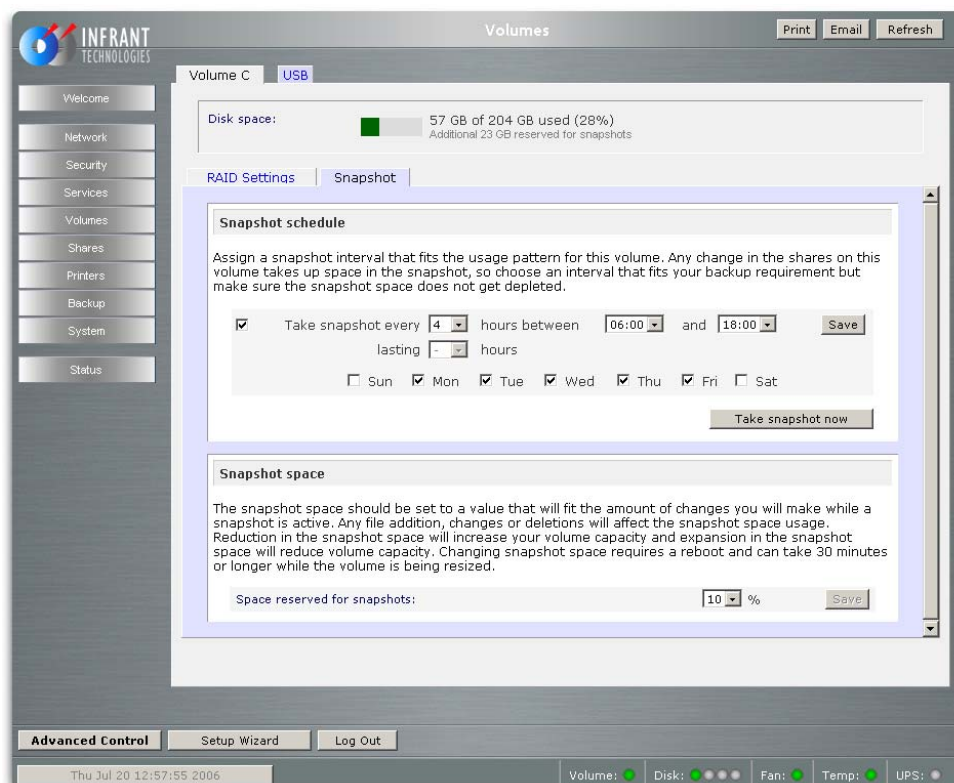
You can combine your backups with snapshots. In brief, a *snapshot* is an image of your data frozen in time. However, because the image isn't the actual data but rather pointers to the it, snapshots are virtually instantaneous while remaining transparent to the user.

You can schedule a snapshot of any volume at a specified frequency and time interval. Why would you want to do that? Well, imagine that a LAN workstation becomes infected with a virus one morning, and that the virus spreads to the files on a ReadyNAS share. If you had scheduled a daily snapshot at midnight, you could go back to the snapshot of the share and restore the version of the files as of midnight.

After you take a snapshot, you can schedule a backup of it. Rather than backing up live shares from the ReadyNAS while file contents are potentially changing, you can backup the snapshot instead, i.e. your data frozen in time. If you coordinate your backup job to start shortly after your snapshot is taken, even if the backup process is still running into the day, you won't have to worry about backing up potentially inconsistent and outdated data.

### To schedule a snapshot:

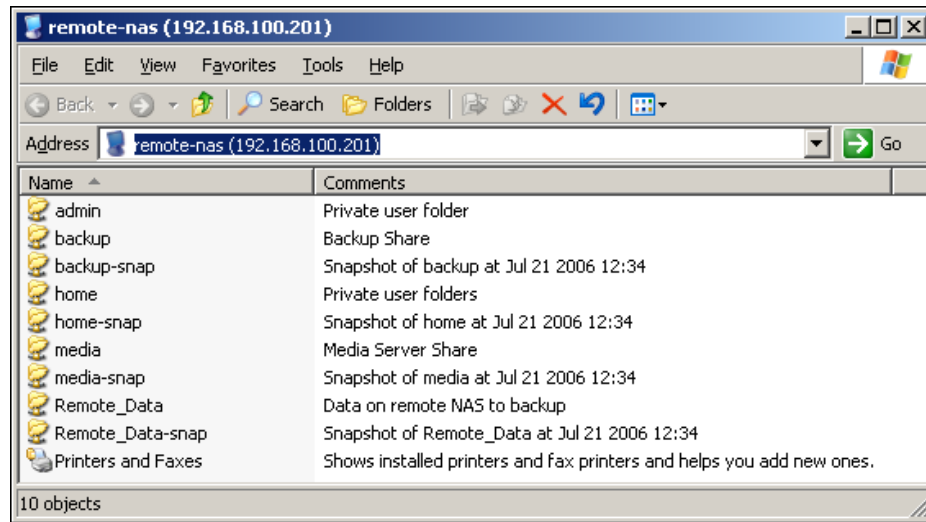
1. In the Snapshot Space section of the Volume tab shown in **Figure O**, specify the percentage of the volume you want to reserve for snapshots.



**Figure O:** The snapshot Schedule section is available only after you reserve snapshot space.

2. Click the Save button, and then follow the prompts to reboot the ReadyNAS device.
3. In the Snapshot Schedule section, specify the frequency and time interval for snapshots.
4. Click the Save button.

Snapshots are identified by *-snap* appended to the original share name, as shown in **Figure P**. As you can see, when you take a snapshot, all the shares on a given volume are included in the snapshot. To back up a snapshot, you simply create a backup job as you normally would, specifying the share with the appended *-snap* as the backup source.



**Figure P:** You can use RAIDar to browse for existing snapshots.

## Putting it all together

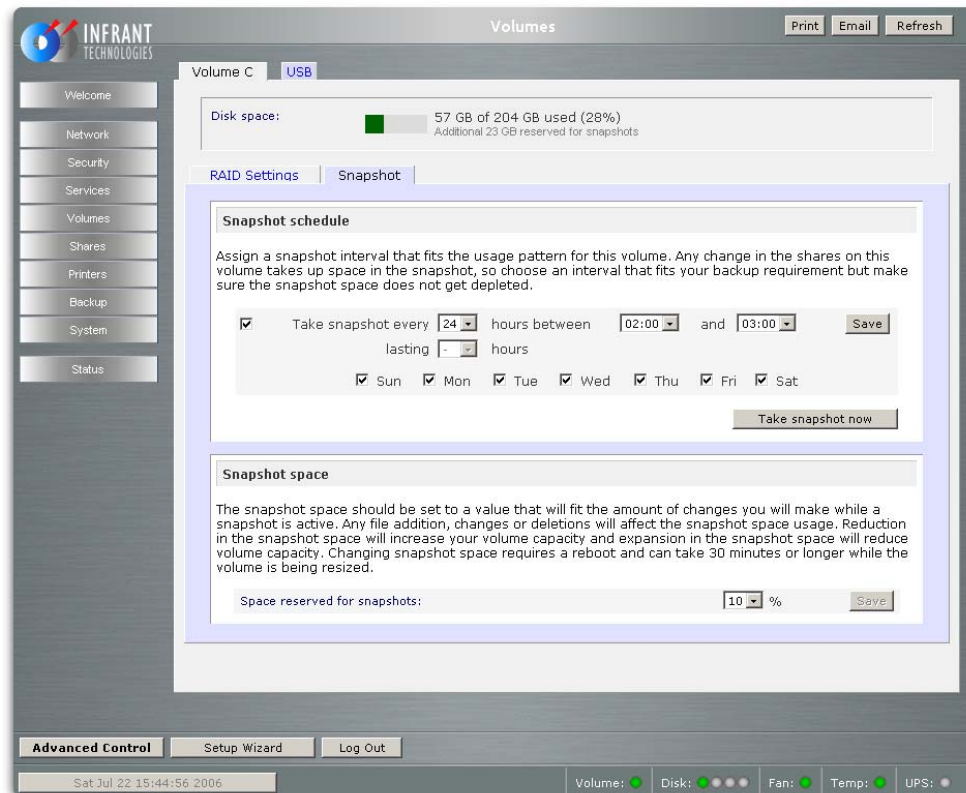
Now that we've discussed several different network topologies as well as scheduling backup jobs and snapshots, let's tie everything together in a final example. For this example, we'll use a local ReadyNAS to host backup jobs and a remote ReadyNAS whose data we want to back up. We'll backup snapshots from the remote device to the local device on a daily basis using Rsync. Furthermore, we want the ability to restore up to a week's worth of changes.

### 1. Set up your network

Use the discussion presented earlier in this guide to choose a network topology. Make sure the devices can communicate properly by testing connections thoroughly.

### 2. Schedule a snapshot

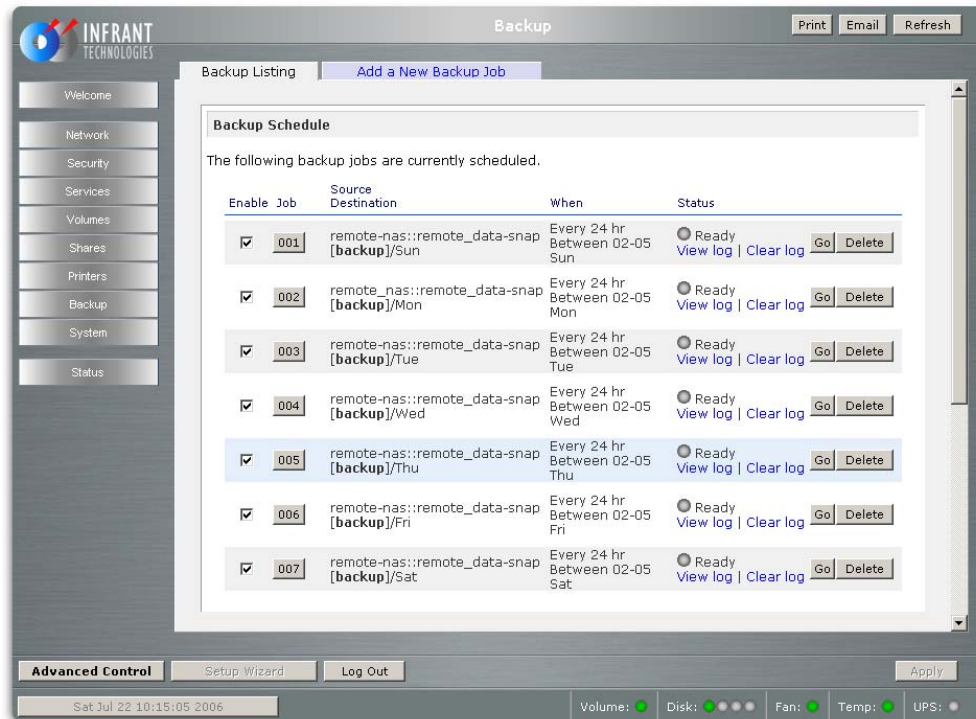
On the remote device, schedule a snapshot to occur once a day on each day of the week, as shown in **Figure Q**.



**Figure Q:** The snapshots are schedule on the remote device.

### 3. Create backup jobs

On the local device, create seven backup jobs—one for each day of the week, as shown in **Figure R**. Notice that the destination for each backup job specifies the path to a folder representing a particular day of the week.



**Figure R:** The backup jobs are scheduled on the local device.