

## 1. Overview

---

There are 3 components to the system, the storage service, the optimizer and the network accelerator. The storage service should be installed on a server to act as a gateway to the cloud. Clients access the server via the iSCSI protocol. The optimizer should be installed on all iSCSI clients to reduce your data storage requirements. The network accelerator can be installed on all computers within the home or office whether they are clients or not. The accelerator compliments cloud storage with high performance local storage and can greatly increase performance.

## 2. System Requirements

---

|                     |   |
|---------------------|---|
| Operating System    | Windows (XP, 2003, Vista), Linux, VMware Appliance, Mac OS X  |
| CPU                 | As fast as possible and as many as possible.  |
| Memory              | As much as possible. 4 GB RAM + is recommended  |
| Disk                | The service requires as little as 1% of your total cloud data requirements available on disc. A 20 TB installation requires as little as 200 GB disk space. Performance is dependent on the ration of local storage to cloud storage. |
| LAN                 | >= 1 GB networking is recommended   |
| Internet connection | Video can be streamed directly from the cloud with a connection of around 5Mb/s, 100 Mb/s is recommended  |

## 3. Storage Requirements

---

The Cloud Drive Storage Service relies heavily on local discs. There are no strict minimum disc space requirements; however the following table should be consulted for the optimizing performance.

| Cloud Storage | Minimum (1%) | Recommended (5%) | High Performance (10%) |
|---------------|--------------|------------------|------------------------|
| 1 TB          | 10 GB        | 50 GB            | 100 GB                 |
| 10 TB         | 100 GB       | 500 TB           | 1 TB                   |
| 100 TB        | 1 TB         | 5 TB             | 10 TB                  |

For example

- a low performance application such as a 10 TB network file server requires a 100 GB local cache
- A high performance application such as a 100 TB content server requires a 10 TB local cache

## 4. Software Requirements

---

| Operating System      | Dependency   |
|-----------------------|--|
| Windows/Linux         | .NET 3.5   |
| Linux (Ubuntu/Debian) | mono-2.0-service<br>libmono-microsoft-visualbasic8.0-cil<br>open-iscsi<br>samba<br>upstart<br>alien          |
| Linux (RHEL/CentOS 6) | mono-addon-extras<br>mono-addon-basic<br>mono-addon-libgdiplus0<br>iscsi-initiator-utils<br>samba<br>upstart |

## 5. Installation - Linux

---

The following commands, directories and files needed to install, configure and run the software

| Item                       | Command/Directory/File   |
|----------------------------|--|
| Install<br>(Ubuntu/Debian) | <pre>apt-get update apt-get upgrade  apt-get install mono-2.0-service apt-get install libmono-microsoft-visualbasic8.0-cil apt-get install open-iscsi apt-get install samba apt-get install alien  alien -i -c CloudDriveStorageService.noarch.rpm</pre>   |
| Install<br>(RHEL/CentOS 6) | <pre>yum upgrade  Add /etc/yum.repos.d/mono.repo file with the following contents: [mono] name = mono baseurl=http://ftp.novell.com/pub/mono/download-stable/RHEL_5/ enabled=1 gpgcheck=0  yum install mono-addon-extras yum install mono-addon-basic yum install mono-addon-libgdiplus0 yum install samba  rpm -i CloudDriveStorageService.noarch.rpm</pre> |

|                    |   |
|--------------------|---|
| Configure          | CloudDriveStorageService /configure                               |
| Start              | start CloudDriveStorageService                                    |
| Stop               | stop CloudDriveStorageService                                     |
| Status             | status CloudDriveStorageService                                   |
|                    |   |
| Installation dir   | /usr/CloudDriveStorageService                                     |
| Configuration file | /usr/CloudDriveStorageService/CloudDriveStorageService.exe.config |
| Default cache dir  | /CloudDriveStorageService   |
|                    |   |
| Service (init.d)   | /etc/init.d/CloudDriveStorageService                              |
| Service (upstart)  | /etc/init/CloudDriveStorageService.conf                           |
|                    |   |
| Logging            | /etc/var/syslog   |

The Network accelerator follows the same naming conventions

## 6. Installation - VMware

---

The VMware appliance is a 32 bit Linux Ubuntu distribution with 1 GB RAM and 128 GB HD. Knowledge of Linux isn't required; the software is configured graphically in the same way as the Windows version. Simply login with the following username/password clouddrive/clouddrive

| Step                         | Command  |
|------------------------------|--|
| Login                        | clouddrive/clouddrive  |
| Open Terminal                | Applications -> Accessories -> Terminal                              |
| Configure service (headless) | vi /usr/CloudDriveStorageService/CloudDriveStorageService.exe.config |
| Configure service (GUI)      | sudo CloudDriveStorageService /configure                             |
| Configure service            | sudo CloudDriveStorageService /configure                             |
| Restart service              | sudo service CloudDriveStorageService restart                        |
| Exit                         | exit   |

## 7. Installation - Database

---

Cloud Drive can use a database as well as local hard drives to store cache files. The SQL Server schema can be generated by the SQLSERVER.SQL create script

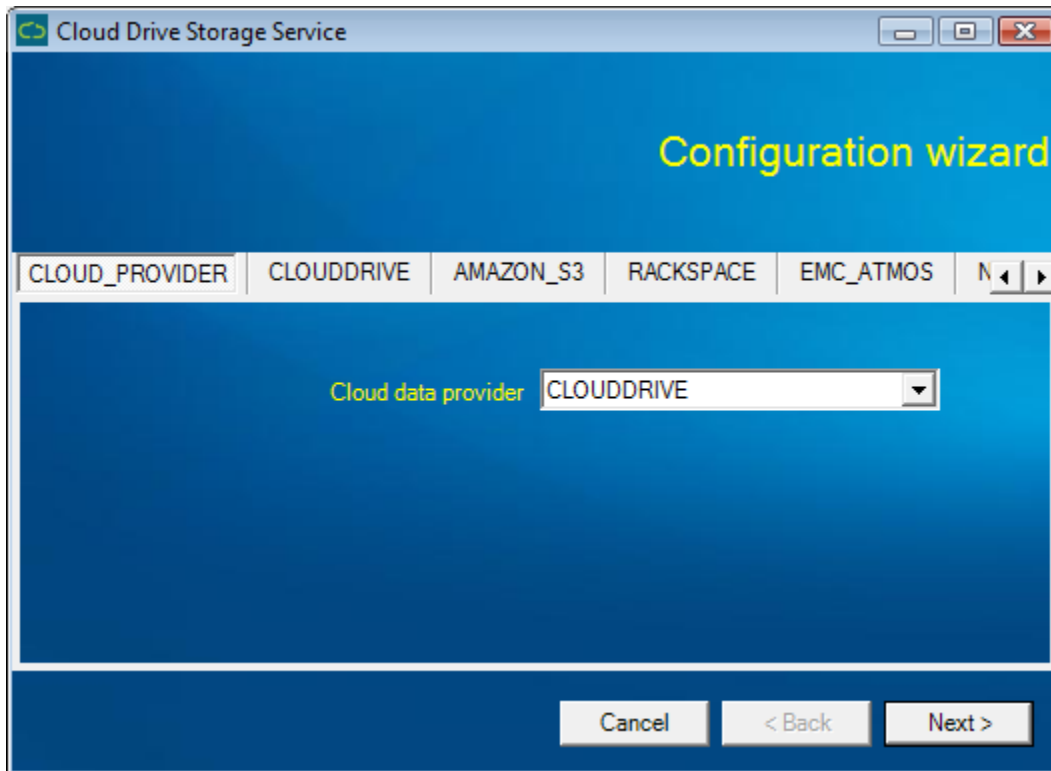
## 8. Configuration

---

After installation you will be prompted to enter your cloud provider details. Do this by simply selecting the tab that relates to your account, “CloudDrive”, “Microsoft Azure“, “Amazon S3”, “Rackspace”, “EMC Atmos”, “Nirvanix” and enter your account details.

The “Windows” option directs the service to store data on any windows file share (via UNC names only). This option is useful if your provider is a “vanilla” windows server such as GoGrid or Amazon EC2. This option is also very useful within the enterprise, storage can be centralized and remote offices with slow network connections will have high speed access to their data even over a slow WAN link. Using this option over the internet requires careful firewall or VPN considerations.

The “iSCSI Server Only” option runs the cloud drive storage service without a cloud provider. In this mode the service is high performance software only SAN device. This is very useful for centralizing data within the office or enterprise.



| Setting      | Use  |
|--------------|--|
| ISCSI_PORT   | The port the service listens to<br>Default 3260  |
| STORAGE_FREE | What percentage of the servers total available disk space should be kept available (not including attached cloud drives)<br>Default 50<br>Type percentage<br>Notes This is the optimal value in terms of |

|                        |  |
|------------------------|--|
|                        | <p>performance and the ability of the service to handle large volumes of new data</p>  |
| STORAGE_SIZE           | <p>The default drive size<br/> Default 2199023255552<br/> Type bytes<br/> Notes Must be a power of 2 Linux and 64 bit windows clients, a larger disc size can be used.<br/> <b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b></p>  |
| STORAGE_CLUSTER_SIZE   | <p>The default cluster size<br/> Default 1048576<br/> Type bytes<br/> Notes Must be a power of 2. A larger cluster size means higher performance when the cluster is cached, at the cost of more latency when the cluster needs to be downloaded<br/> <b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b></p>  |
| STORAGE_CLUSTER_BUFFER | <p>The default cluster buffer size<br/> Default 8388608<br/> Type bytes<br/> Notes Must be a power of 2. Must be a multiple of the cluster size. A larger buffer means better media streaming at the cost of a longer delay in opening or seeking within files. A smaller buffer better opening and seeking within files as the cost of poorer media streaming</p> |
| STORAGE_IMAGE          | <p>The default disc image to format the drive with<br/> Default NTFS<br/> Type string<br/> Notes The service comes with 2 disk images, NTFS and EXFAT. NTFS is more robust and has permissions and transactional capabilities. EXFAT is a lightweight file system potentially faster but less robust.</p>  |
| STORAGE_CLOSE_TIMER    | <p>How long to keep files open<br/> Default 10<br/> Type seconds<br/> Notes How long to keep files open in memory<br/> A larger value means increase performance at the cost of memory usage</p>   |
| STORAGE_ARCHIVE_TIMER  | <p>How often to analyze which clusters should be archived<br/> Default 3600<br/> Type seconds<br/> Notes A smaller value means better decisions on which clusters to archive, at the cost of more disk</p>   |

|                         |  |
|-------------------------|--|
|                         | <p>IO A bigger the value means course grained decisions on which clusters to archive but requires less disk IO</p>   |
| STORAGE_COMPRESSION     | <p>Compress blocks when archiving<br/> Default true<br/> Type Boolean<br/> Notes Compressing blocks when archiving means less storage and bandwidth to the cloud. This has to be set before using the cloud drives and can never be changed.<br/> <b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b></p>  |
| STORAGE_ENCRYPTION      | <p>Encrypt blocks when archiving<br/> Default None<br/> Type string<br/> Notes Encrypt blocks with AES-256 when archiving means nobody can access your data, including us.<br/> <b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b></p>  |
| STORAGE_REMOVABLE       | <p>Present the drive to the SCSI subsystem as a removable drive.<br/> Default false<br/> Type Boolean<br/> Notes This option is used in conjunction with STORAGE_IMAGE to serve light weight EXFAT file images to newly connected clients</p>  |
| STORAGE_RECOVER         | <p>Check cloud provider for missing data. Provides disaster recovery.<br/> Default True<br/> Type Boolean</p>  |
| STORAGE_UPLOAD_SCHEDULE | <p>A string representing hours when the service should upload data to the cloud storage provider. This option is useful if you have cheaper upload data rates during off peak hours. The string represents the 24 hours in a day, starting at midnight, a "0" means don't upload during that hour, anything else means uploading is ok.<br/> Default<br/> "11111111111111111111111111111111"<br/> Type<br/> string<br/> Example<br/> Upload 24 hours a day;<br/> "11111111111111111111111111111111"<br/> Upload between midnight and 12pm;<br/> "1111111111110000000000000000"<br/> Upload between midnight and 6am;</p> |

|                            |   |
|----------------------------|---|
|                            | <p>“111111000000000000000000”<br/> Upload between 6 pm and 6am;<br/> “111111000000000000011111”<br/> Upload between 7 pm and 5am;<br/> “111110000000000000011111”<br/> Upload between 9 pm and 3am;<br/> “111000000000000000000111”</p>   |
| STORAGE_UPLOAD_THREADS     | The number of upload threads  |
| STORAGE_BANDWIDTH          | Bandwidth used by the optimizer to recover unused storage   |
| STORAGE_DRIVER             | Where to store cache files<br>Default FILESYSTEM<br>Example<br>FILESYSTEM<br>DATABASE   |
| STORAGE_FILESYSTEM_INCLUDE | The drives to include<br>Default All hard drives<br>Type string<br>Example<br>Windows “c:\;\server1\disk1;\server1\disk2”<br>Linux “/mnt/sdb:/mnt/sdc”<br><br>Notes Only use these locations to store data. Network locations must be specified as UNC notation<br>Linux The storage service will use all available devices mounted within the file systems. If you mount /home on a separate device cloud drive will use this. |
| STORAGE_FILESYSTEM_EXCLUDE | The drives to exclude<br>Default None<br>Type string<br>Example<br>“D:\;E:\” – Don’t use D or E drives<br>“/;/home” – Don’t use “/” or “/home”  |
| STORAGE_DATABASE_TYPE      | The type of database to store cache files in<br>Default SQLSERVER<br>Type String<br>Example<br>SQLSERVER<br>ODBC  |
| STORAGE_DATABASE_SERVER    | The database server or DSN identifier<br>Default None<br>Type String<br>Example<br>SERVER1<br>192.168.0.10  |

|                           |  |
|---------------------------|--|
|                           | DSNCloudDrive  |
| STORAGE_DATABASE_NAME     | The database name<br>Default None<br>Type String   |
| STORAGE_DATABASE_USERNAME | The database username<br>Default None<br>Type String   |
| STORAGE_DATABASE_PASSWORD | The database password<br>Default None<br>Type String   |
| CACHE_PORT                | The multicast port to communicate with network accelerators on<br>Default 2805<br>Unit integer   |
| CACHE_ADDRESS             | The multicast address to communicate with network accelerators on<br>Default "224.28.05.74"<br>Type string   |
| CACHE_REDUNDANCY          | The number of redundant copies of data the network accelerators will store<br>Default 3<br>Type Integer<br>Notes A larger number means more likely hood that storage requests can be satisfied without needing the storage provider as the cost of more network traffic                  |
| CACHE_TIMEOUT             | The number of milliseconds to wait for a response from the network cache(network accelerators)<br>Default 125<br>Type Integer  |
| CLOUD_PROVIDER            | Which cloud storage provider to use;<br>Default None<br>Type string<br>Example<br>CLOUDDRIVE<br>AMAZON_S3<br>RACKSPACE<br>EMC_ATMOS<br>NIRVANIX<br>WINDOWS<br>MICROSOFT_AZURE<br><b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b> |
| CLOUD_BUCKET              | To resolve clashes with other "bucket" names also hosted on the cloud provider. This prefix disambiguates your data.<br>Default None   |

|                                  |  |
|----------------------------------|--|
|                                  | Type string<br>Example "Lucky"<br><b>Warning Don't change this setting after data has been stored otherwise you may corrupt your data</b>  |
| CLOUD_TIMEOUT                    | The number of milliseconds to wait for a response from the cloud data provider<br>Default 60000<br>Type Integer  |
| CLOUD_PROXY                      | Proxy server to connect to the internet with<br>Default None<br>Type string<br>Example "http://proxyserver:8080"   |
| CLOUD_PROXY_USERNAME             | Proxy username to connect to the internet with   |
| CLOUD_PROXY_PASSWORD             | Proxy password to connect to the internet with   |
| CLOUD_PROXY_DOMAIN               | Proxy user domain to connect to the internet with  |
| CLOUD_AMAZON_S3_ENDPOINT         | Your Amazon S3 Endpoint<br>Default None<br>Example<br>s3.amazonaws.com<br>s3-us-west-1.amazonaws.com<br>s3-ap-southeast-1.amazonaws.com<br>s3-eu-west-1.amazonaws.com<br>dcs.dunkel.de |
| CLOUD_AMAZON_S3_ID               | Your Amazon S3 ID  |
| CLOUD_AMAZON_S3_KEY              | Your Amazon S3 Key   |
| CLOUD_AMAZON_S3_REGION           | Your Amazon S3 Region  |
| CLOUD_RACKSPACE_USERNAME         | Your Rackspace username  |
| CLOUD_RACKSPACE_API_KEY          | Your Rackspace API key   |
| CLOUD EMC_ATMOS_FULL_TOKEN_ID    | Your EMC Atmos full token ID   |
| CLOUD EMC_ATMOS_SHARED_SECRET    | Your EMC Atmos shared secret   |
| CLOUD_NIRVANIX_APP_KEY           | Your Nirvanix appkey   |
| CLOUD_NIRVANIX_USERNAME          | Your Nirvanix username   |
| CLOUD_NIRVANIX_PASSWORD          | Your Nirvanix password   |
| CLOUD_NIRVANIX_APP_NAME          | Your Nirvanix app name   |
| CLOUD_WINDOWS_SHARE              | Your Windows share name  |
| CLOUD_WEBDAV_URL                 | Your Webdav URL  |
| CLOUD_WEBDAV_USERNAME            | Your Webdav username   |
| CLOUD_WEBDAV_PASSWORD            | Your Webdav password   |
| CLOUD_WEBDAV_DOMAIN              | Your Webdav domain   |
| CLOUD_FTP_URL                    | Your FTP URL   |
| CLOUD_FTP_USERNAME               | Your FTP username  |
| CLOUD_FTP_PASSWORD               | Your FTP password  |
| CLOUD_MICROSOFT_AZURE_ACCOUNT    | Your Microsoft Azure account   |
| CLOUD_MICROSOFT_AZURE_ACCESS_KEY | Your Microsoft Azure access key  |
| CLOUD_CLOUDDRIVE_USERNAME        | Your Cloud Drive username  |
| CLOUD_CLOUDDRIVE_PASSWORD        | Your Cloud Drive password  |

## 9. Backup and maintenance

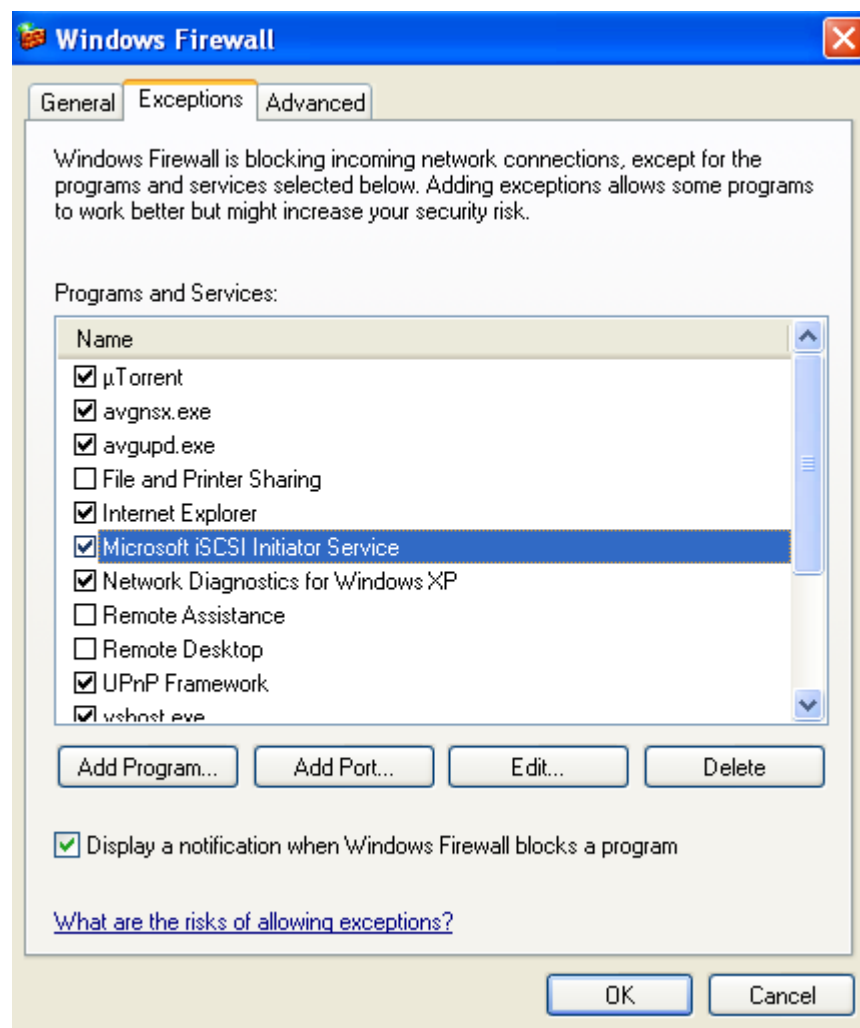
---

Cloud drive doesn't require backups or maintenance. All data is replicated in the cloud redundantly and the optimizer periodically runs "checkdisk". There is no performance gain from defragmenting a cloud drive. Defragmenting a cloud drive will be extremely slow and require downloading all your data from your cloud storage provider.

## 10. Firewall Settings

---

iSCSI clients communicate with the cloud drive storage service on TCP port 3260. Network accelerators communicate with the cloud drive storage service on via both TCP and UDP on port 2805. When this service runs for the first time you will be prompted automatically prompted to allow the port. Alternatively the firewall can be manually configured to allow incoming connections.



## 11.iSCSI Installation

---

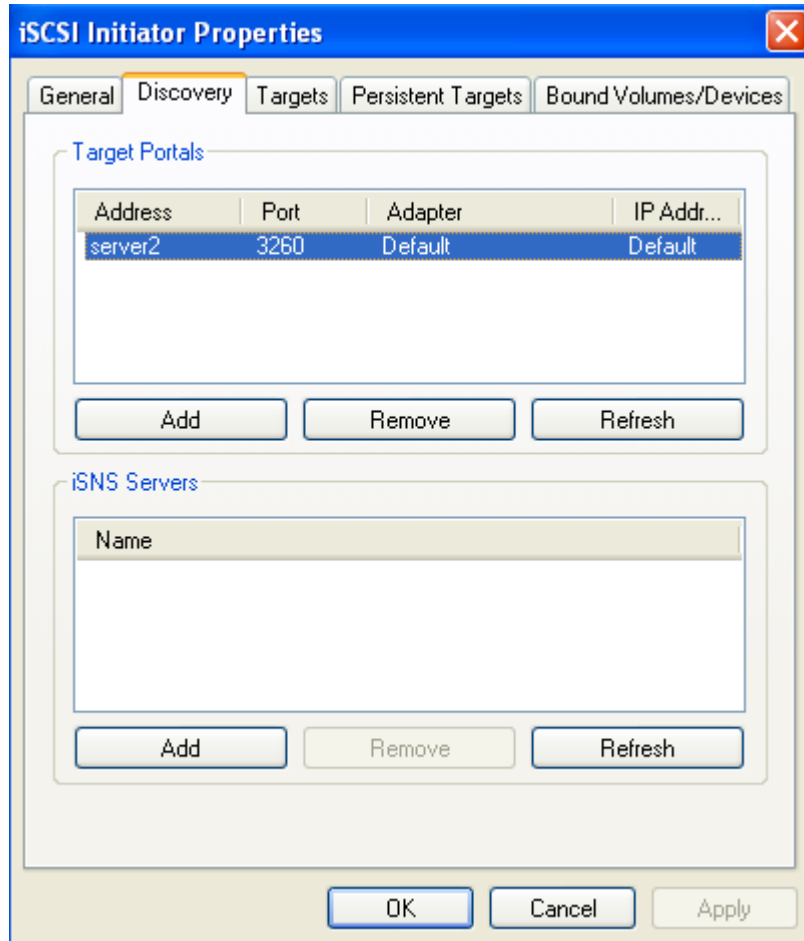
The following table show which iSCSI initiator required to connect to the cloud drive storage service

| Operating System       | Initiator                          |
|------------------------|------------------------------------|
| Vista/2008 Server      | Inbuilt iSCSI initiator            |
| Windows XP/2003 Server | Microsoft iSCSI Software Initiator |
| Linux                  | Open-iSCSI Project                 |
| Mac                    | globalSAN iSCSI Initiator          |

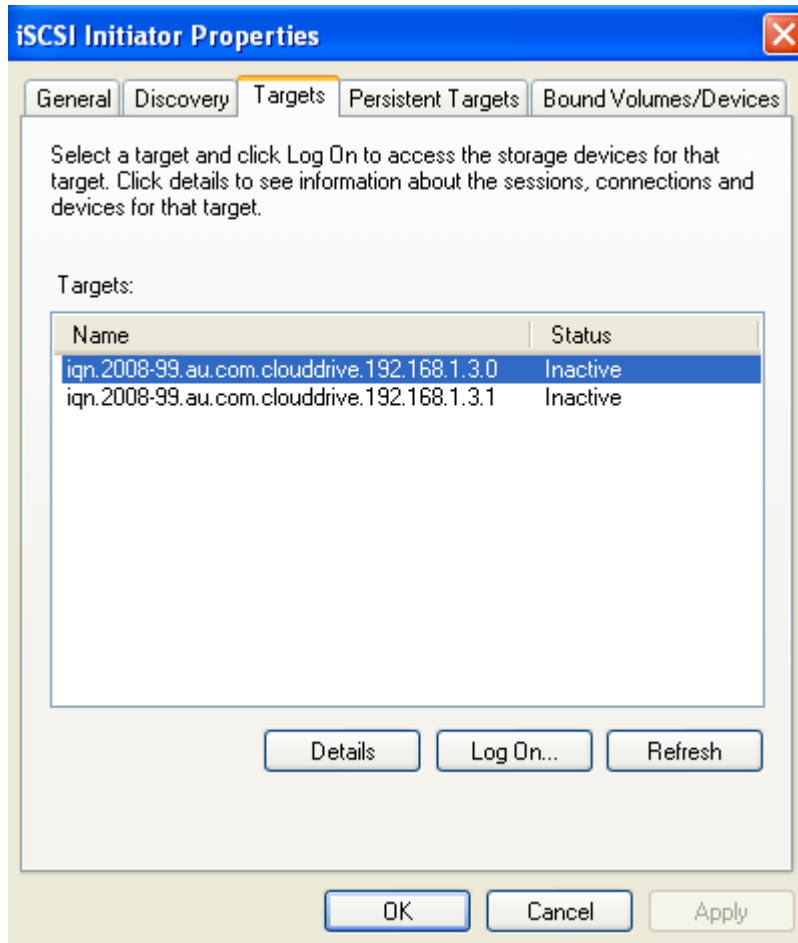
## 12. iSCSI Configuration

---

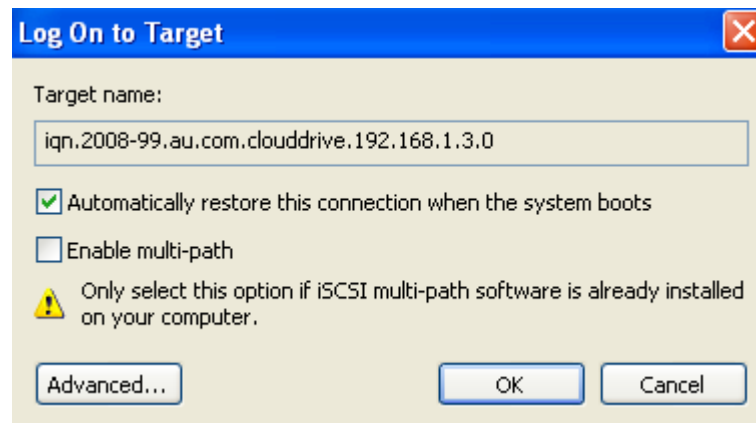
Specify the server where the Cloud Drive Storage Server is located



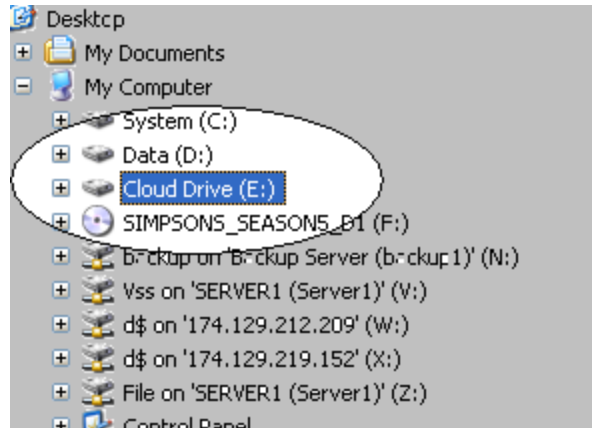
Select the target you wish to log on to



and log on to the iSCSI target



The cloud drive should appear already formatted



You can use all your normal disc tools

