How to Become Cloud Backup Provider

Contents

Abstract

In this white-paper we will explain how one can install and configure cloud storage software and use it for backup purposes. Document can be used by an individual or a company who wishes to become a cloud backup provider and utilize its own hardware for cloud object-based storage. We will use Cloudian object-storage software along with CloudBerry Lab products.

Introduction

Object-based storage is considered a leading emerging technology, and many of its features are deemed ideal for cloud storage. Benefits of object storage include massive scalability, geographic independence, multi-tenant features and the ability to use off-the-shelf hardware. More companies are considering object-based storages as a part of their infrastructure, for instance for backup purposes. They choose it for its benefits and price. There is a number of public object-based cloud storages offered on the market, such as Amazon S3, Glacier, Windows Azure, Rackspace, and many others. However, there are still many companies who already have existing hardware and want to transform a part of it or transform it entirely in a cloud storage, utilizing latest technologies. It's also a huge playground for Managed Service Providers (MSP), who are willing to provide cloud storage and backup services to their final customers at the same time. New technologies, such as object-based storage require new tools, like products of CloudBerry Lab.

Purpose

This white-paper may be used by Managed Service Providers who are willing to expand their portfolio by providing cloud backup services to their customers using their own private cloud storage and CloudBerry Managed Backup. There are two ways to provide cloud storage to the final customers from MSPs' point of view. First one is to create new account on public cloud storage provider, like Amazon, for example, and resell it. However, there is a number of MSPs who run their own hardware and want to utilize its capacity to construct their own private cloud storage built on their own resources. In this white paper, we will briefly explain the second way: how this can be done by using Cloudian object-storage software. We will explain how to use CloudBerry Managed Backup to offer online backup service to the final customer. As an add-on, we will show how this cloud storage can be accessed and used by stand-alone products of CloudBerry Lab for object-based storages.

Architecture

It is recommended to install the Cloudian software in a cluster of three or more nodes. For this documentation, we deployed a single-node installation of Cloudian 2.4.3 in a virtual machine running Red Hat Enterprise Linux 6.3.

Result

We were able to install and configure Cloudian object-storage software on a stand-alone Linux server and use it together with CloudBerry Managed Backup, Backup, Explorer and Drive products. Thus, we proved a concept that any individual or company, MSP especially, may use this basic scenario to become a Managed Backup Provider to the cloud storage, built on their own hardware resources. We encourage everyone to go through the following sections to get the basic idea of the technical requirements, tips, and general steps of the installation.

Requirements

OS

Cloudian is officially supported only on these versions of Linux:

- Red Hat Enterprise Linux (RHEL) 5.x up through 6.4, 64-bit

- CentOS 5.x up through 6.4, 64-bit

For best performance and capacity, Cloudian, Inc. recommends version 6.4.

Sizing

As a sizing example, a Supermicro SuperStorage Server with:

- CPU: 2 quad-core Intel Xeon E5-2600

- 32GB RAM

- 12-2TB SAS2 HD (24TB total)

would have these approximate storage limits (presuming that HyperStore functionality is enabled):

- Assume 80% of raw disk is used for data object storage

- Replicated 2x: (24TB * 80%)/2 → **9.6TB data per node**.

- Replicated 3x: (24TB * 80%)/3 → **6.4TB data per node**.

NB!

An ext4 file system on RHEL/CentOS 6.3 has a maximum size of 16TB. However, with the latest *e2fsprogs* package (which provides file system utilities for ext2/3/4), ext4 file systems of up to 100TB are supported.

Third-party software requirements

You need to install (Cloudian v3.0 will install the following software automatically):

- Java 7 JRE 64-bit JVM
    - Cloudian v2.4.3 allows Java 6 and 7, v3.0 and above requires Java 7
- Ruby (Puppet requirement)
- Puppet
- Facter

Installation of Linux server

For this single-node installation we used Red Hat Enterprise Linux. We installed it on the VMware VM. It's better to set up a fix IP address from the beginning, as each time you shut down a node with DHCP you will have to repeat a certain procedures to start Cloudian node again.

Useful tips

Confirm that:

The machine's hostname **is not** set to *localhost*.

The machine's hostname **is not** mapped to the loopback address (127.0.0.1) in */etc/hosts*.

The host **is** assigned a static IP address. Do not use DHCP.

To Set Host Name

vim /etc/sysconfig/network

To Set DNS Server IP

vi /etc/sysconfig/network-scripts/ifcfg-eth0

PEERDNS=no

To disable firewall

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

To set DNS name

```
#cp /etc/resolv.conf /etc/resolv.dnsmasq.conf
#nano /etc/hosts
```

```
#nano /etc/resolv.conf
```

- Comment out all nameserver entries

- Add the line: nameserver 127.0.0.1

Installation of third-party software

To fulfill the pre-requirements for Cloudian installation, we have to install this third-party software:

1. Java 7 JRE 64-bit JVM (Cloudian v2.4.3 allows Java 6 and 7, v3.0 and above requires Java7)
2. EPEL (Extra Packages for Enterprise Linux)
3. Ruby
4. Puppet
5. Facter

NB!

Please note, that it may be installed automatically during Cloudian v3.0 and above installation. Here is a quote from Cloudian official documentation:

1. From *ftp.cloudian.com*, download the Cloudian release package on to your Cloudian host machine and untar the package.
2. As *root*, run the self-extracting script for installing required third party software:

   root# sh ./selfextract_prereq.bin

Useful tips

These tips may be used if manual installation preferred, or if the host doesn't have internet connectivity.

Link to install Java 7 JRE

http://docs.oracle.com/javase/7/docs/webnotes/install/linux/linux-jre.html#install-64

To install EPEL for RHEL/CentOS 6.x

#rpm -Uvh http://dl.fedoraproject.org/pub/epel/6/x86_64/epel-release-6-8.noarch.rpm

Install Ruby

#sudo yum install ruby ruby-libs ruby-shadow ruby-rdoc

Install Puppet

#sudo yum install puppet puppet-server facter

Cloudian installation

Here we won't show the complete procedure, just give a brief idea and some useful tips, which might help you setting up the test cloudian node. You can download the software and manuals/Installation guide from [here](#).

This paper is based on Cloudian software version 2.4.3, the GA version available during the time of integration and development of this paper. The basic steps are the same and the procedure is even easier for the version 3.0.

At a high level, your single-node Cloudian installation will proceed as follows:

1. Download the Cloudian packages and extract them.
2. Review and edit a few installation settings.
3. Execute the installation.
4. Set up Dnsmasq, a domain resolution utility.

Useful tips

While editing/etc/puppet/manifests/extdata/common.csv:

Javahome - you can leave it "**/usr/java/latest**"
Jrehome- you can leave it "**/usr/java/latest**"

To have Dnsmasq started automatically with node reboot, please add:
"cd /opt/dnsmasq; ./dnsmasq-2.55 -C ./dnsmasq.conf" to the /etc/rc.local file.

Please check mts-ui.properties for admin.host details. It should contain your IP address:
#nano /opt/tomcat/webapps/Cloudian/WEB-INF/classes/mts-ui.properties

To get logs
#tail -f /var/log/cloudian/*

To verify connectivity between the CMC host and the admin host and admin port specified in the mts-ui.properties file

#telnet [s3.cloudian.com](#) 18081
*we used default domain s3.cloudian.com, please use yours if you created it in *singlenode.csv

To set up DNS on the main DNS server, you will need 3 entries (in our case 192.168.3.x, hostname Cloudian, default domain s3.cloudian.com):
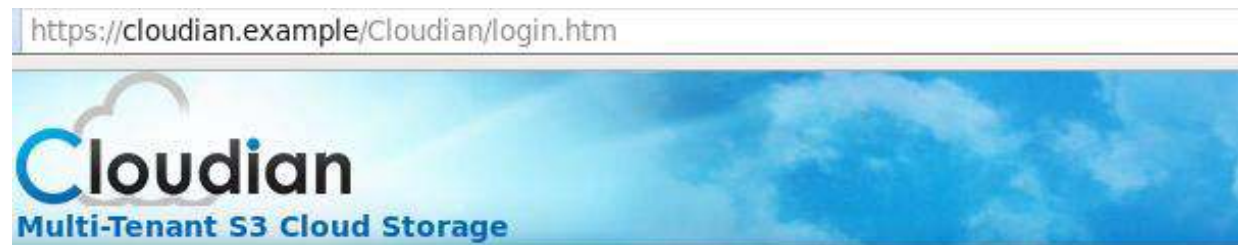
Cloudian IN A 192.168.3.xxx
s3.cloudian.com IN A 192.168.3.xxx
*.s3.cloudian.com IN A 192.168.3.xxx

After the installation check:
/etc/init.d/cloudian-cassandra status
/etc/init.d/cloudian-redis-credentials status
/etc/init.d/cloudian-redis-qos status
/etc/init.d/cloudian-hyperstore status
/etc/init.d/cloudian-s3 status
/etc/init.d/cloudian-redismon status
/etc/init.d/cloudian-cmc status

Configuration of Cloudian software

After installation you will see the login screen (Pic 1), please log in with username *admin@cloudian.com* and password *public*. Cloudian example is the host name of our server.

Pic.1. Initial log in.

Now, you need to create a new group and a new user. Click on the Admin tab and do that by selecting **Manage Groups** → **New Group** and **Manage Users** → **New User** (Pic 2).
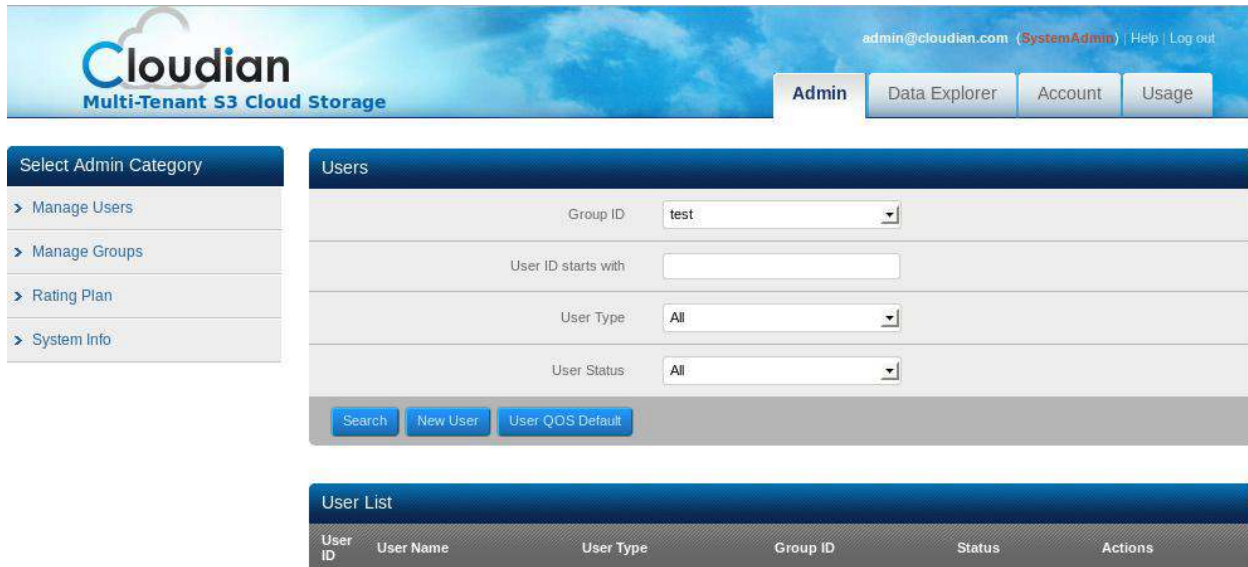
Pic.2. Admin Tab.

For this basic installation we simply created one Group with a name **test**, one user **testuser** and a test bucket **testbucket.** It's pretty easy to set everything up. After doing this please check Account->Security Credentials which we will need later (Pic.3, 4).
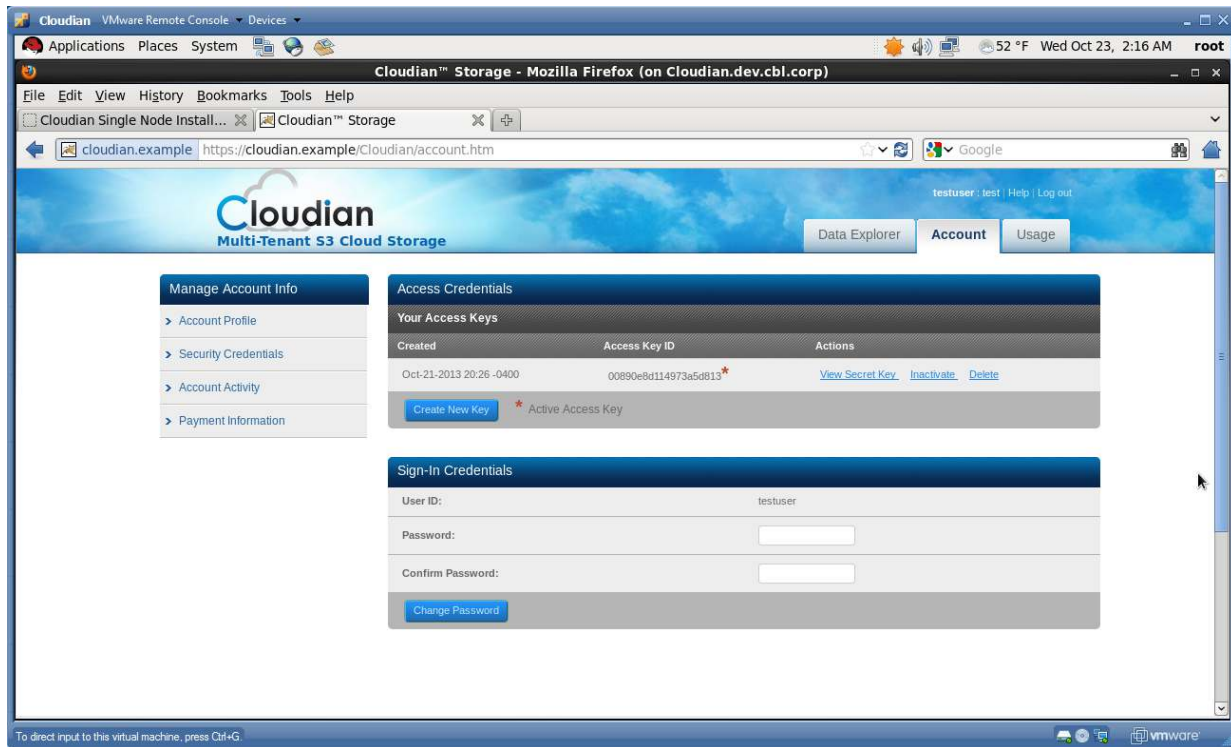

Pic.3. Checking access and secret keys.
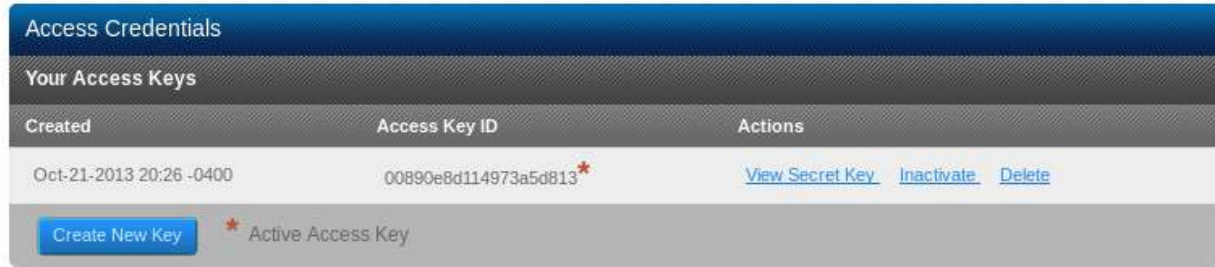
Pic.4. Checking access and secret keys. Closer view.

CloudBerry Managed Backup configuration and usage

You can go to **My Account -> Switch to my own account** and choose Amazon s3 compatible account (Pic 5, 6). Then select storage provider and provide cloud storage credentials (Pic.7) that we've copied from cloudian installation (Pic.3, 4). After finishing the wizard, you will be able to backup your clients to your own cloud storage. To make it happen, please create an arbitrary user (Pic.8), this user will receive a download link for backup client (Pic.9, 10). After installation which is very similar to installation of stand-alone CloudBerry Backup, this user becomes the first user of your managed backup service. Please check Appendix A to get the basic idea how to configure stand-alone products of CloudBerry Lab with this new storage.
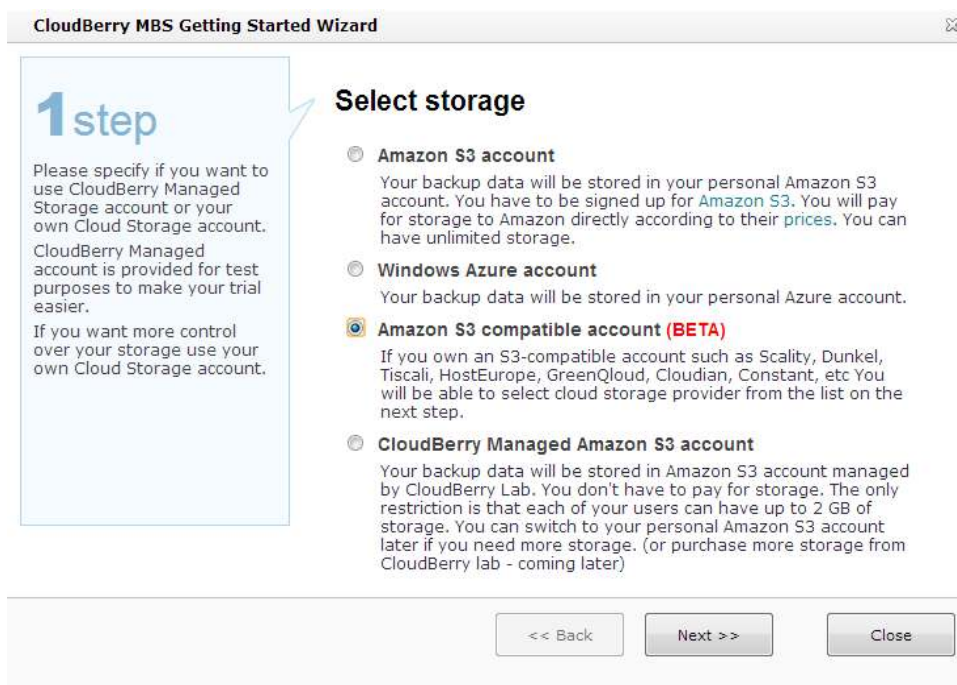
Pic.5. Selecting a private cloud-object storage, based on Cloudian.

Pic.6. Selecting a private cloud-object storage.
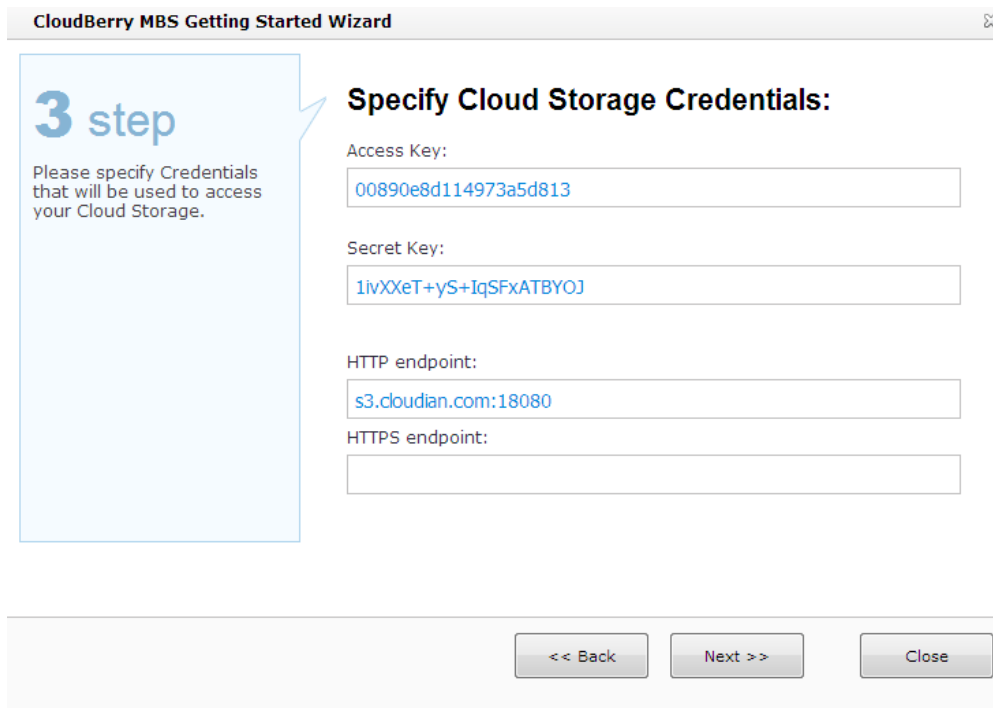
Pic.7. Specifying credentials of a private cloud-object storage for Managed Backup account

Pic. 8. Creating the first customer's profile.

Pic.9. Customer's link to download Backup client.

Pic.10. Customer receives an email with initial password and the link to Backup client.

About Cloudian Inc.

Cloudian is a Foster City, California-based software company specializing in object storage software. With Cloudian®, Enterprise IT and Cloud Service Providers (CSPs) can deploy turnkey multi-tenant private and public object storage systems on commodity IA servers, scaling from two servers to exabytes across multiple datacenters with high reliability and elasticity. Cloudian has a strong record in commercial production at major cloud service providers and enterprises across the U.S., Japan, Europe, and APAC.

About CloudBerry Lab

CloudBerry Lab is established in 2008 by a group of experienced IT professionals with the mission to help organizations in adopting Cloud computing technologies by closing the gap between Cloud vendors propositions and consumer needs through development of innovative low costs solutions.

Appendix A: CloudBerry Lab stand-alone products usage

CloudBerry Lab also develops stand-alone products: Backup - robust backup product, Explorer - file manager, Drive – product, which maps cloud storage as a virtual disk with the Windows environment.

If you set everything up, you can now use your cloud storage with CloudBerry Explorer, Backup and Drive. Please go to **File -> S3 compatible -> Cloudian** (Pic.11) and enter your Service Point name, Access and Secret keys (Pic.12). As the process is very similar for other products, we want to show it.
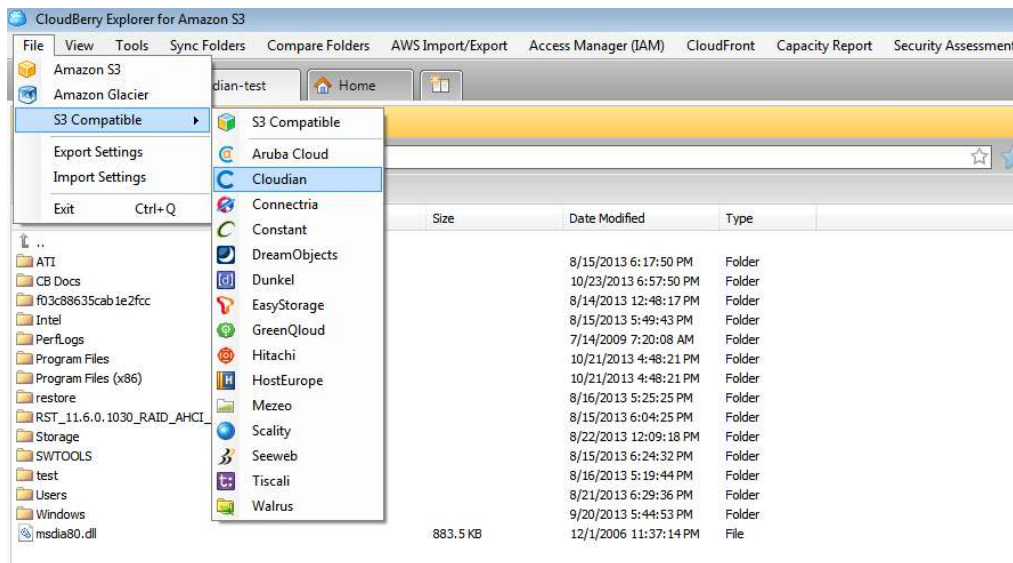
Pic.11. Connecting CloudBerry Explorer to new cloud-object storage.
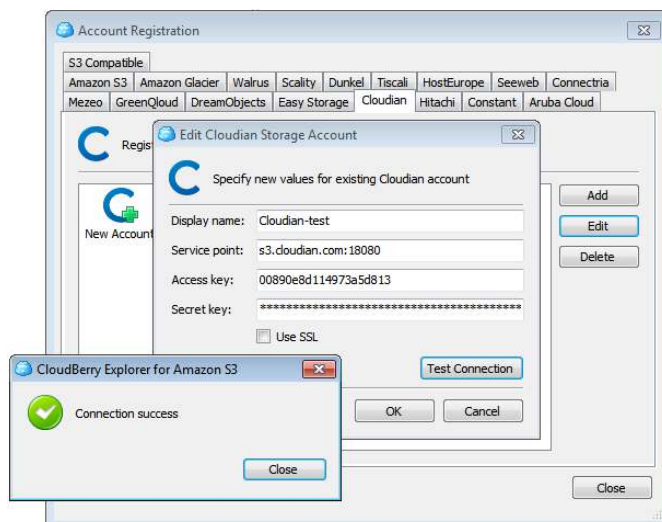
Pic.12. Specifying credentials of a private cloud-object storage for CloudBerry Explorer.