



Catalogic DPX™ 4.4

Using AWS Cloud with DPX

©Catalogic Software, Inc.™, 2016. All rights reserved.

This publication contains proprietary and confidential material, and is only for use by licensees of Catalogic DPX™, Catalogic BEX™, or Catalogic ECX™ proprietary software systems. This publication may not be reproduced in whole or in part, in any form, except with written permission from Catalogic Software.

Catalogic, Catalogic Software, DPX, BEX, ECX, and NSB are trademarks of Catalogic Software, Inc. Backup Express is a registered trademark of Catalogic Software, Inc. All other company and product names used herein may be the trademarks of their respective owners

Using AWS Cloud with DPX

Introduction

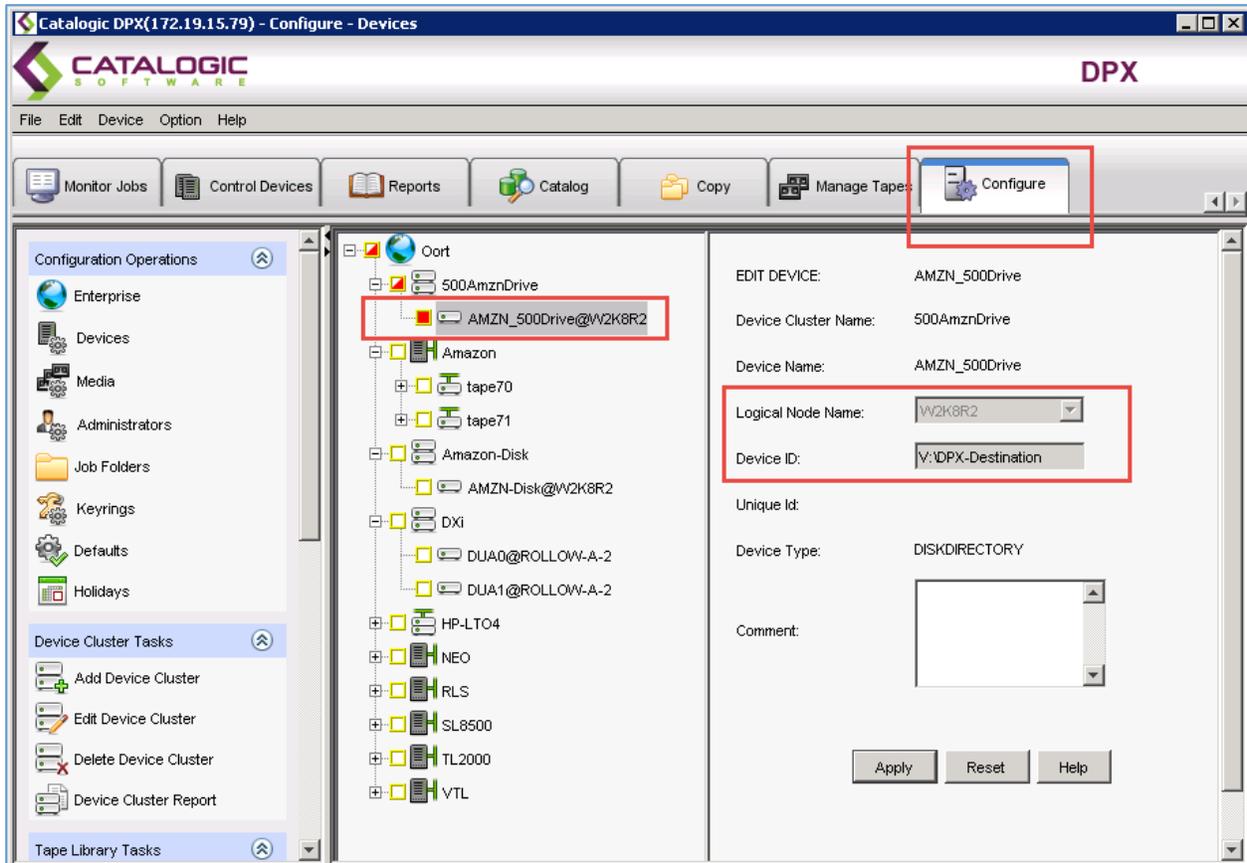
DPX is a Catalogic Data Protection suite which facilitates protection of heterogeneous sources (virtual and physical servers running different Operating System and variety of applications). DPX supports destination servers and devices comprised of tape libraries, standalone tape drives, disk arrays, NetApp filers and more. Diversity of destinations provides to the Catalogic customer community a convenient way to use available resources and to plan protection solutions based on projected data size, data growth, suitable fabric and other internal and external factors. Recently Catalogic has certified and added a cloud to the mixture of possible destinations. Cloud availability facilitates and widens DPX customer use cases and creates a virtually infinite data array for storing and retrieving corporate data.

DPX Data Mover Architecture

As stated above, DPX supports disk directory as a destination (target) for backups. Disk directory device is a predefined location (directory) defined on a volume visible from a “device server” and defined within DPX as a part of a “Device Cluster” which is a hierarchical subdivision within the DPX Enterprise.

As illustrated on the screenshot below, the Definition includes:

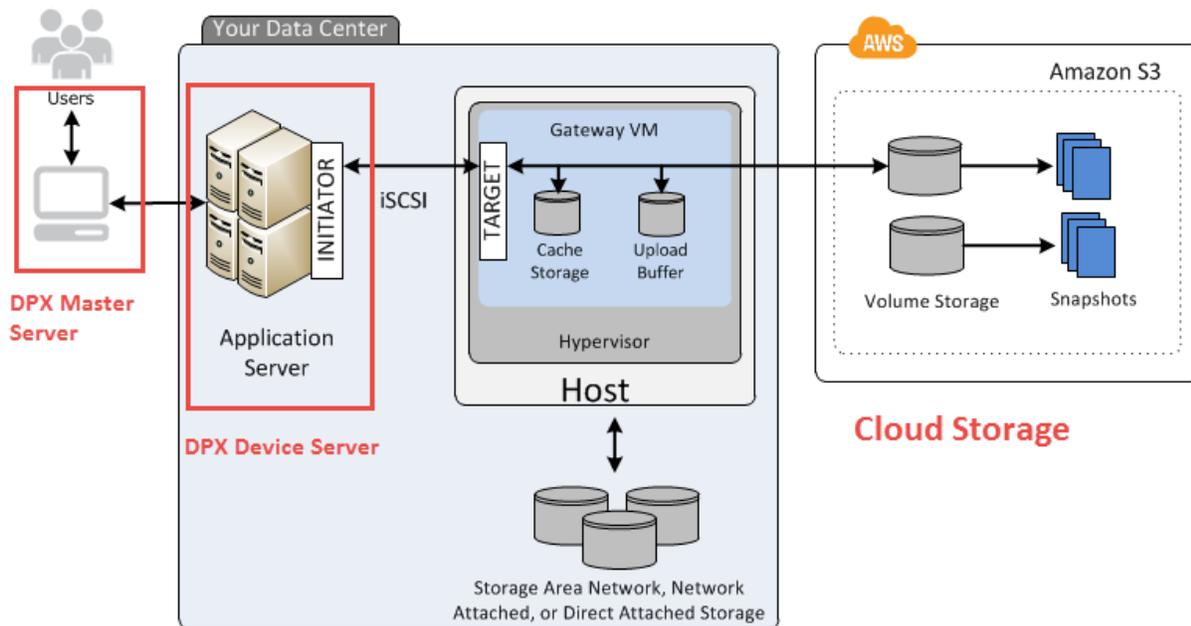
- Logical node name. This is a server with attached disk which is a part of DPX Enterprise, referred to as a “Device Server”
- Device ID. This is a predefined location on a disk volume attached to that server. Best results are achieved when the disk is a local array or iSCSI or FC connected directly to the server. NFS / CIFS share is also a supported configuration, however it must be defined as a UNC name and does not support all DPX backup types



AWS setup

Amazon Web Services (AWS) has developed mechanisms for accessing the cloud through a variety of methodologies including the use of iSCSI LUN. This iSCSI methodology presents a gateway to the cloud by allowing a LUN to be attached to a defined DPX Device Server, making it appear as a local drive, and enabling it to be used as a backup destination.

This methodology is described on the AWS documentation site for [AWS Storage Gateway](#). Amazon users must go through several steps (guided by a Wizard) to set up and activate AWS. In a nutshell, you will need to deploy a virtual machine on your Virtual environment. (In this example MS Hyper-V supervisor is supported as well). Virtual appliance (Gateway VM) provides you access to the Cloud and serves as a cash buffer for storing and forwarding data. After the appliance is set and configured, you will be able to define an iSCSI LUN that can be attached to the server of your choice. This process can be depicted as below:



Since deployment and setting up of AWS Storage gateway is an essential step for providing iSCSI LUN used as DPX destination device, following are certain individual actions describing the overall process.

Before starting, you have to plan and understand requirements; this information can be found [here](#).

Note that for this activity you deploy a **gateway on premises**.

- After obtaining the right credentials and permissions from your Amazon administrator, go the AWS Console (<http://console.aws.amazon.com/storagegateway>) and start deploying a new gateway



- Follow AWS documentation to download and deploy the AWS Storage Gateway to your host. Make sure you follow specific steps for AWS Storage Gateway and your platform – detailed

instructions can be found at

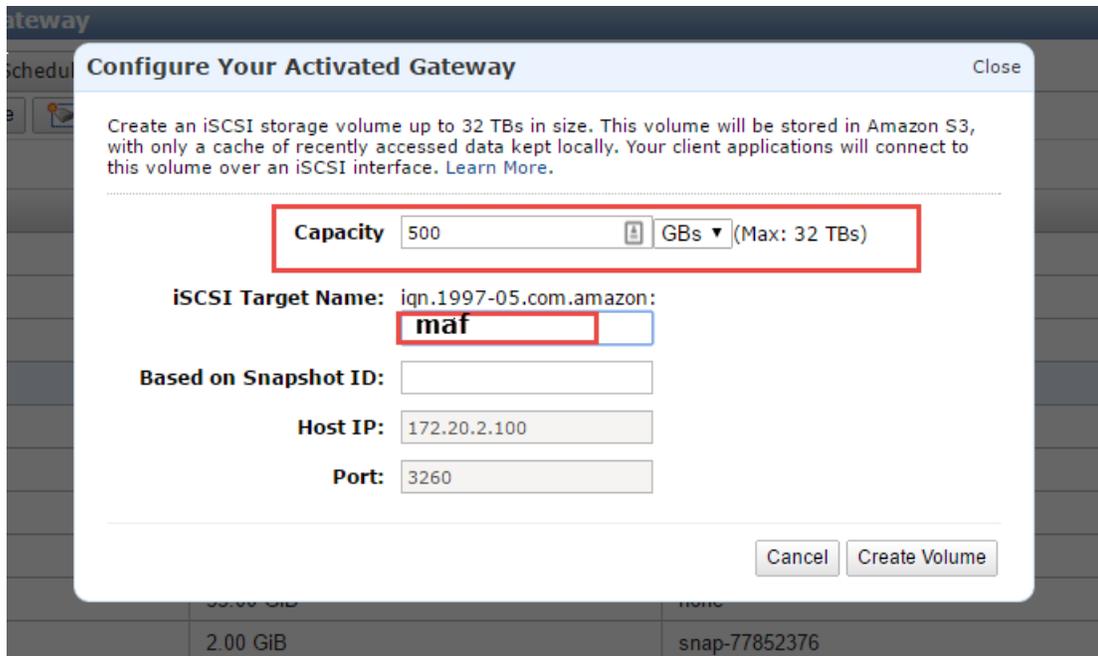
<http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedDownloadVM-common.html>

- During the deployment you will need to provision a local disk space for cache and upload buffer. Amazon provides you guideline to how to choose appropriate volumes. Those volumes can be readjusted after setup is complete.
- You also will need to assign an IP address for your appliance, which will be used for discovery of the iSCSI target. Ensure you remember it or get it from your administrator who installed the gateway
- After the gateway is configured, create a new volume which will be provisioned to the DPX Device server



Volume ID	Capacity
vol-13287678	500.00 GiB
vol-1760FE7C	500.00 GiB
vol-32E4F758	500.00 GiB
vol-39C73E53	20.00 GiB

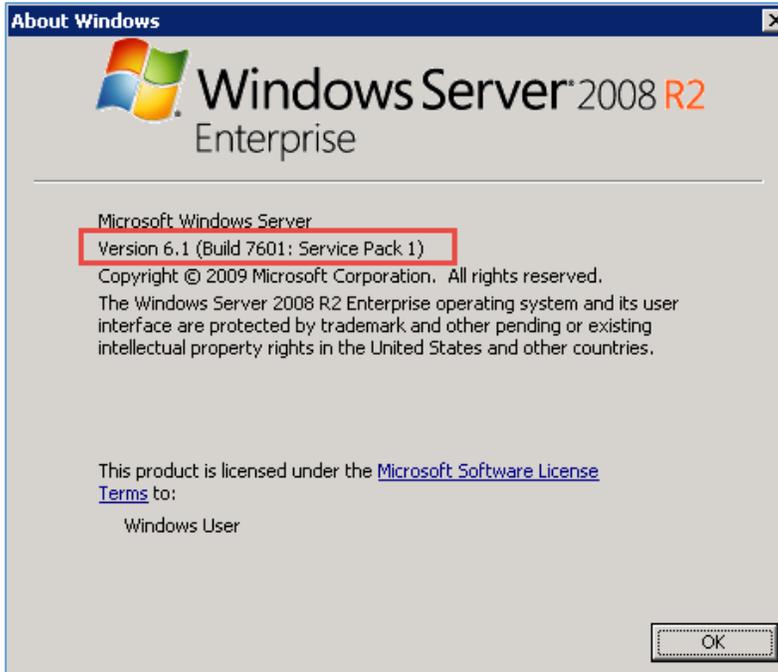
- You will need to indicate a volume capacity – this number is the required space to accommodate your backups. Depending on backup type, retention period, and other factors, this volume can vary considerably. Please consult with your Catalogic pre-sales engineer to calculate how much space you need. Volumes can be added later if your initial estimate is misjudged.
- After the volume is created, please note the iSCSI Target Name which consists of two parts: a name similar to iqn.1997-05.com.amazon and a suffix of your choice used to identify volume discovery during the next steps. In this example the suffix is “maf”.



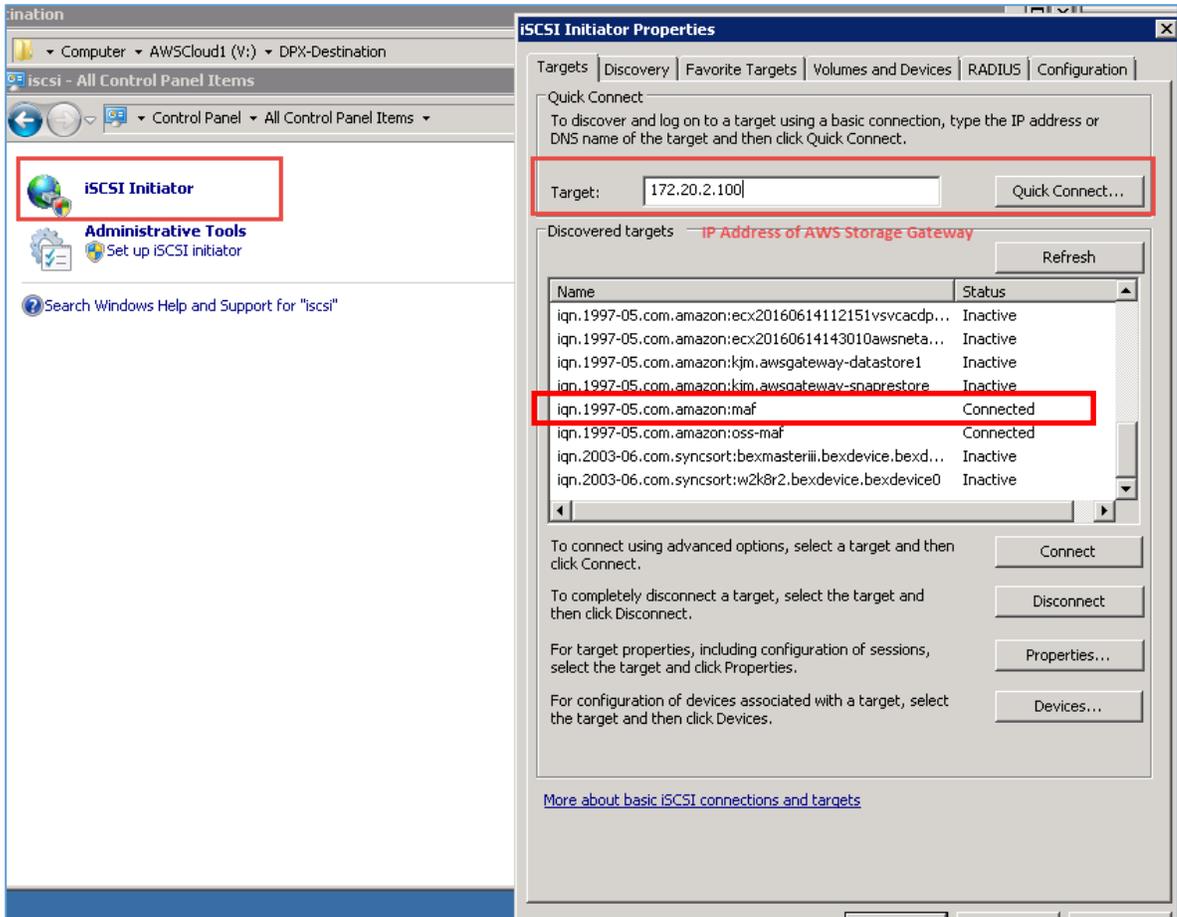
- After AWS is deployed, follow the following steps to provide the iSCSI LUN to the server, which is defined as a Device Server within the DPX Enterprise:

Device server setup

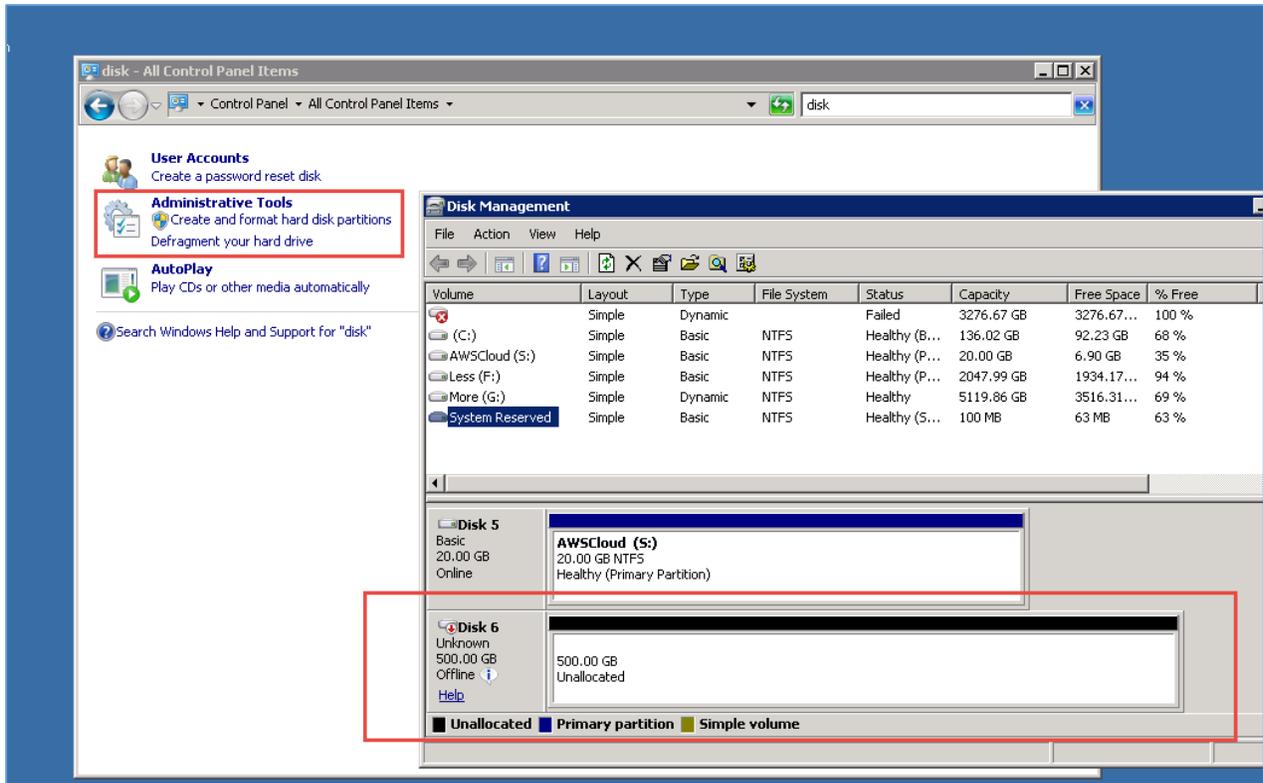
After the above procedure is completed, you will need to set up a Windows server which will be used as a device server in Catalogic DPX. Setting up of the Windows server is just the regular administrative procedure to attach and configure an iSCSI disk device to the server. If you want to use Linux server for this role, follow the guidelines for Linux regarding how to scan and attach an iSCSI disk to the server. Keep in mind following screenshots are taken from Windows 2008 R2;– they may look different for other versions of Windows server, however the substance of the procedure is the same.



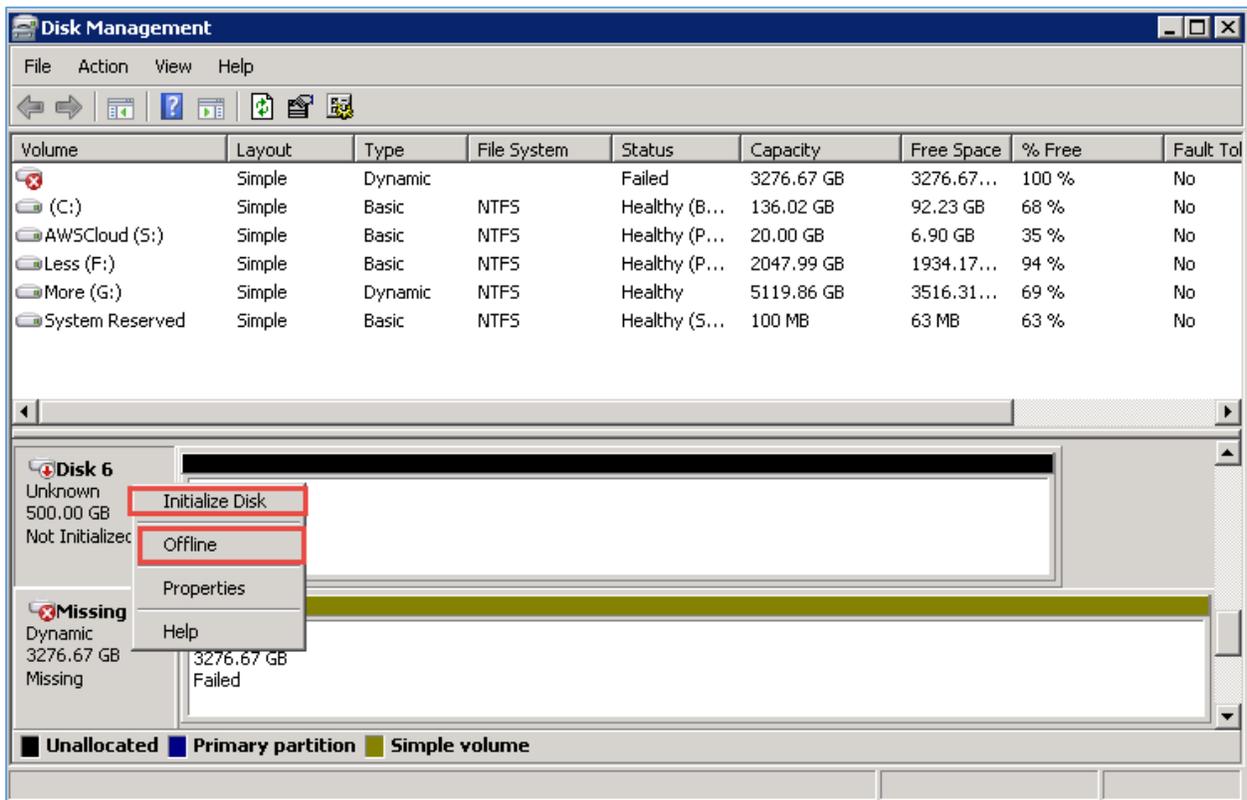
- Bring an iSCSI initiator on a Windows control panel to scan available targets using the IP address of the AWS Storage Appliance. Find the specific Target name from the previous step (iqn.1997-05.com.amazon:maf) and connect to it by pressing “Connect” button.



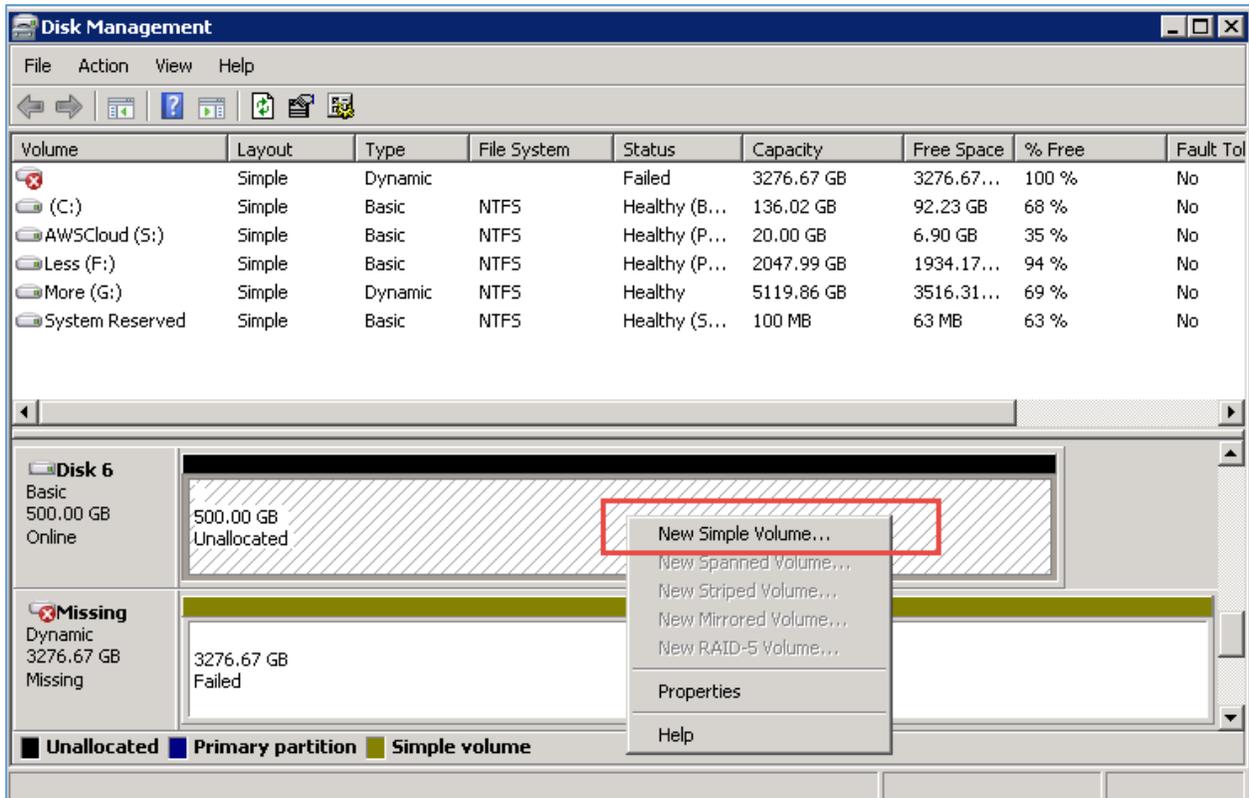
- After connection is established, the device will appear on a server as an unallocated device which can be observed from “Administrative Tools” on the Control Panel.

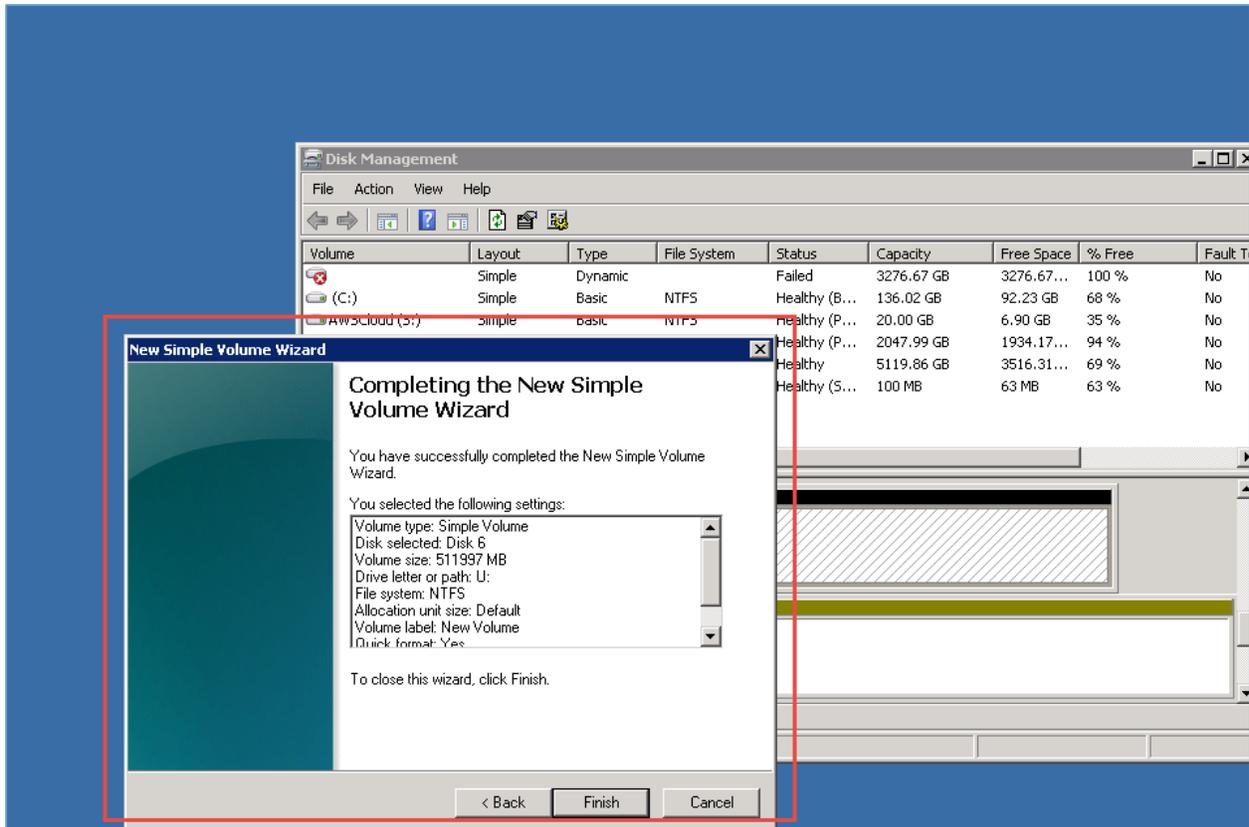


- On the Disk Management screen, select the disk, then right click to expose a context menu. Choose "Initialize Disk" and toggle "Offline" to bring the disk online.

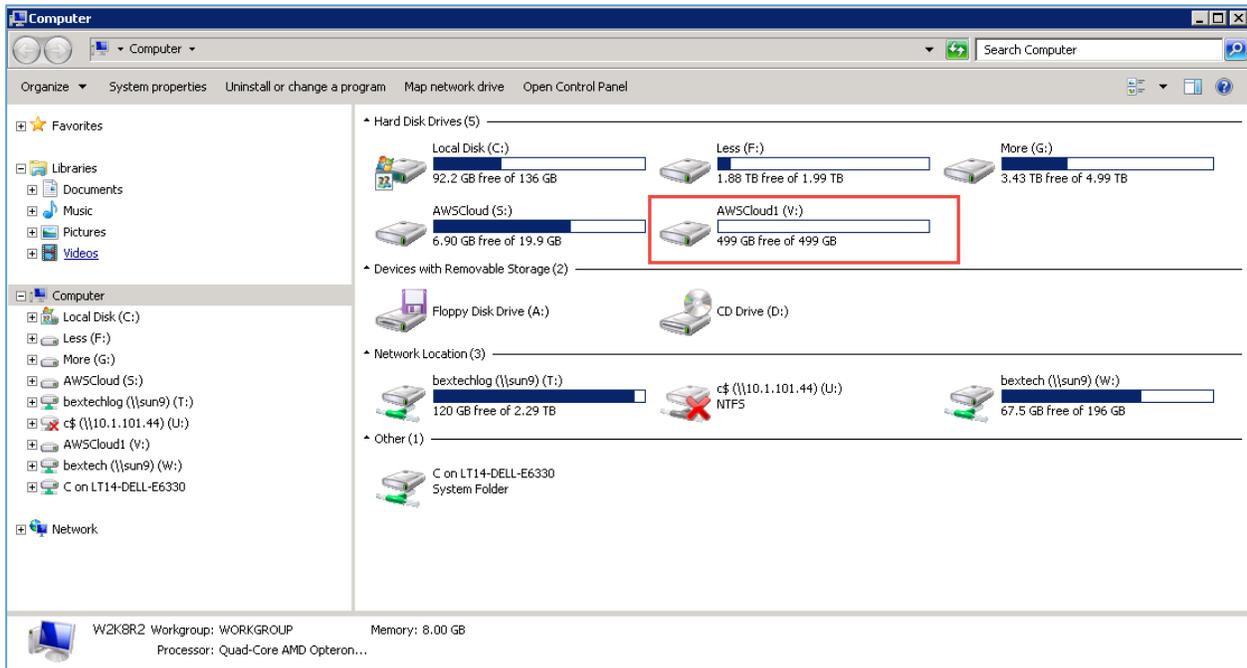
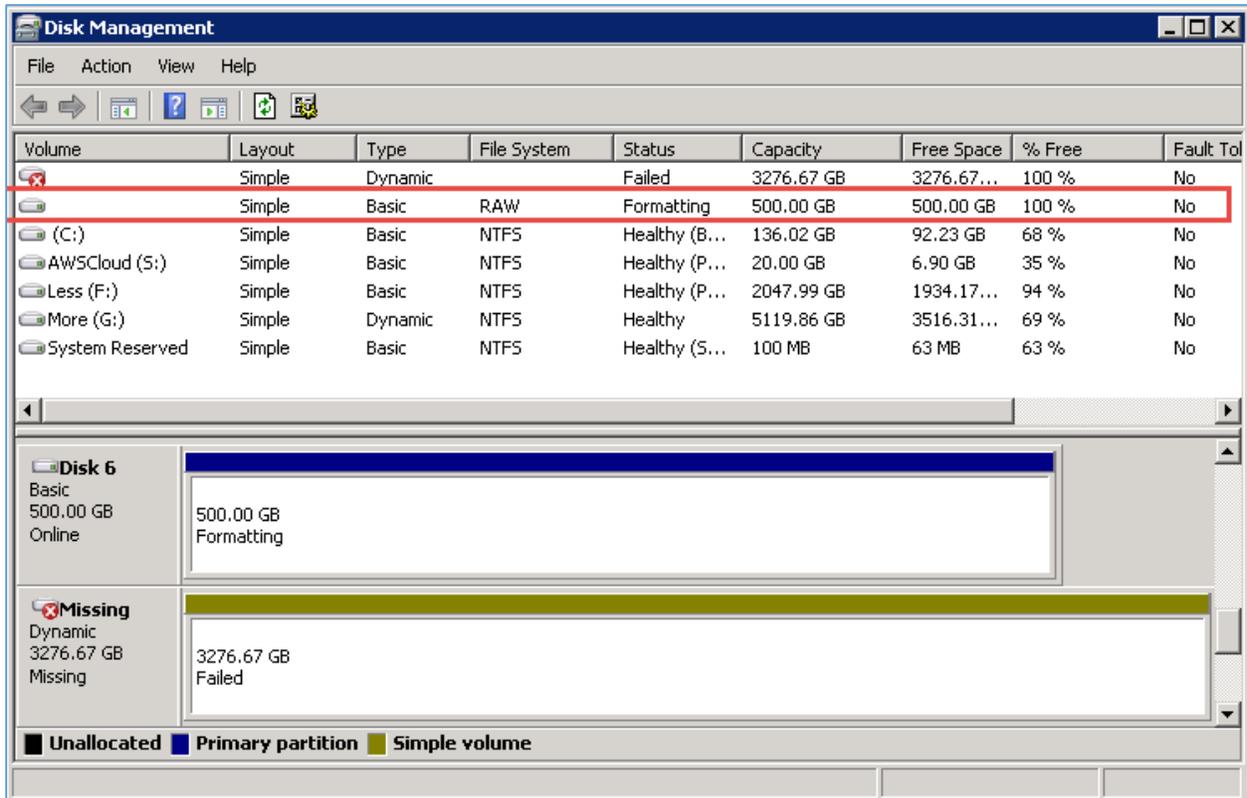


- Create a New Simple Volume by selecting a volume, then right-clicking to expose a context menu. Choose New Simple Volume, then proceed through the wizard and choose appropriate values, such as drive letter and file system.

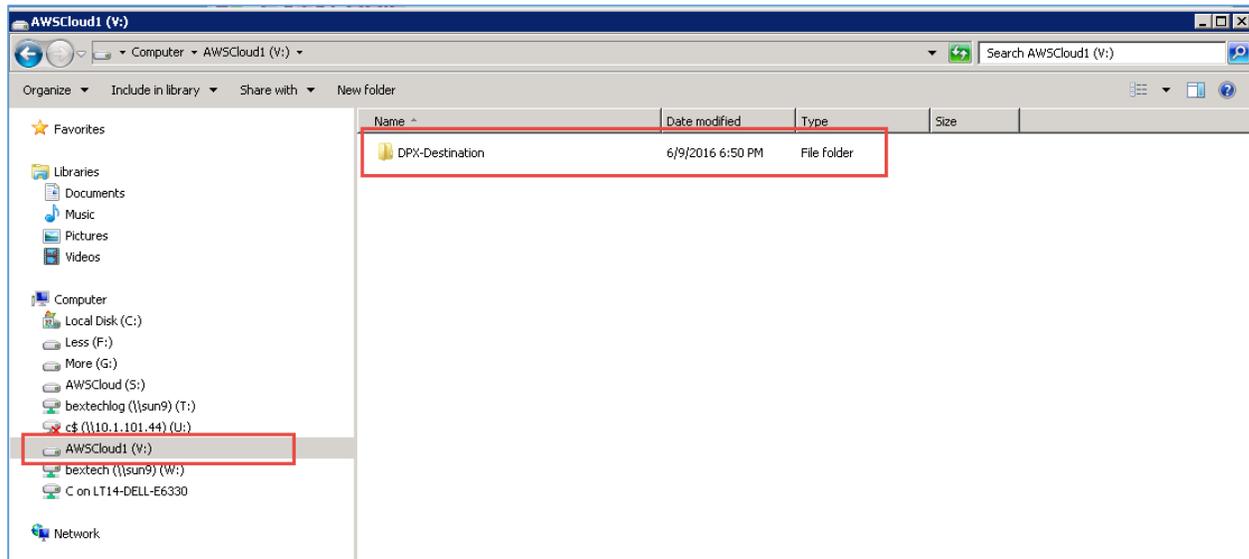




- Format the volume. While formatting, the volume appears in the top portion of the “Disk Management” window and will be visible as a local disk for this server



- Navigate to this volume and create the directory which is used in the tape drive definition in DPX



After the new volume is configured and visible as a local drive, it can be used as a destination for file level backup and NDMP backup, Check DPX documentation regarding setup and running of these job types. Keep in mind that the data saved on this drive is transferred to cloud and the performance of writing data on it depends on several factors which are external to DPX.

