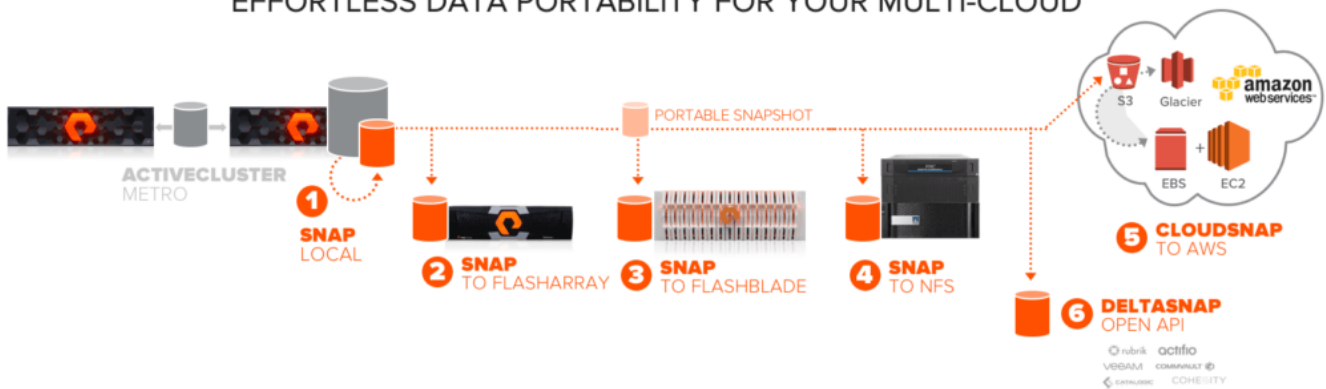


Moving Data Between Cloud and On-Premises Virtualized Environments

Purity **FA** Snap & CloudSnap EFFORTLESS DATA PORTABILITY FOR YOUR MULTI-CLOUD

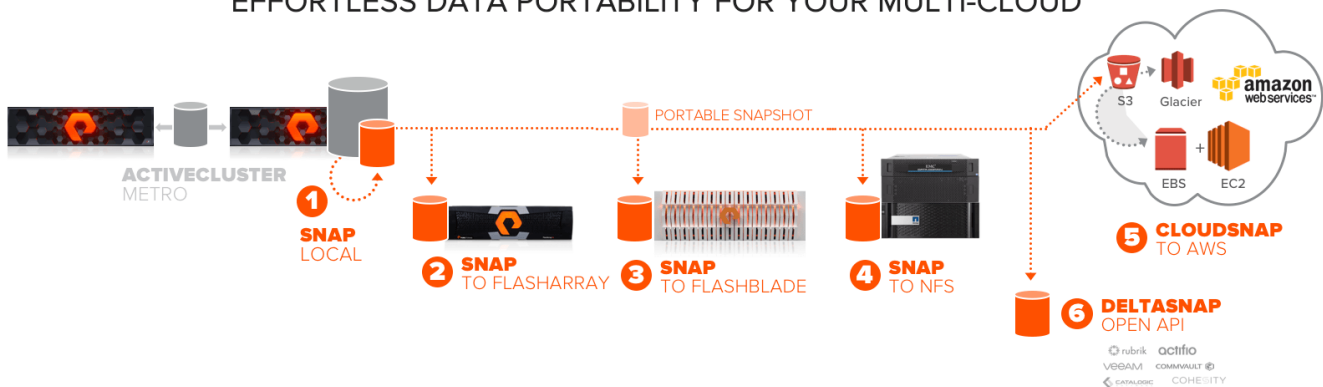


Note: CloudSnap to AWS GA expected in late 2018 with some features in 2019. Delivery dates and features subject to change.

[Cody Hosterman](#) and I delivered a session on “*Moving Data Between Cloud and On-Premises Virtualized Environments.*” The focus of this session was on data mobility and how Pure Storage customers and partners will be able to use different features and integrations to move data from on-premises, extending into the cloud, backing up using Snap To NFS (now) and converting Pure Storage volumes into cloud native format via our *Tech Preview* of CloudSnap.

Purity **FA** Snap & CloudSnap

EFFORTLESS DATA PORTABILITY FOR YOUR MULTI-CLOUD



Note: CloudSnap to AWS GA expected in late 2018 with some features in 2019. Delivery dates and features subject to change.

We will be doing a 5-part blog series that walks through each of the individual stages, which includes:

Stage 1 – VMware Virtual Volumes

In the first stage of the demo we set the scene with VVols. VVols allows you to non-disruptively convert VMFS-based virtual machines into the open format that VVols providers. Through a totally non-disruptive, VAAI-offloaded Storage vMotion, a VMFS VM will be converted into a VVol-based VM. Furthermore, using VM storage policies through the VVols Storage Policy Based Management engine, we can assign FlashArray replication to the specific VVols to replicate it to a remote FlashArray.

Stage 2 – Microsoft Hyper-V

Using the VVols from Stage 1 we use asynchronous replication of the SQL Server VVol data volume to our Equinix deployment. Once replicated, I connect to a Microsoft Hyper-V host and present that data volume to a Hyper-V virtual machine to perform additional data manipulation.

Stage 3 – AWS Direct Connect

In this stage we continue to use the same SQL Server data volume which we started with in Stage 1 by creating a crash consistent snapshot and connecting it to an EC2 instance over Direct Connect using iSCSI connectivity from the Pure Storage FlashArray.

Stage 4 – Snap To NFS

One of the critical steps with any workload is to create backups. In Stage 4 we use Snap To NFS to create a backup of the SQL Server data volume. After taking a backup we then restore that snapshot into an EC2 instance.

Stage 5 – CloudSnap

At Accelerate we said the Multi-Cloud model is important for data mobility regardless of what cloud provider someone may have chosen. Using our *Tech Preview* of CloudSnap we will take the SQL Server data volume and rehydrate into a cloud native format (EBS) that can be connected to an EC2. Think dev/test.