



Simplifying Cloud Migration with Cirrus Data

Executive Summary

Over the last decade, cloud-based storage has gained momentum through Amazon Web Services (AWS), Microsoft's Azure, Google Cloud Platform (GCP), along with many others. The market growing at 20%+ has spurred all the primary storage manufacturers to develop solutions for the burgeoning opportunity, including:

- Cloud Block Store from Pure Storage
- Cloud Volumes ONTAP from NetApp
- HPE Cloud Volumes from Hewlett Packard Enterprise

One challenge faced by companies who want to use public cloud computing is migrating their existing applications and data to the cloud. For applications consisting in a set of relatively inactive data files (non-database) where a typical backup and restore process or the typical "free tools" provided by cloud vendors can be used for migration with acceptable amount of downtime, this may not be a significant issue. However, for mission critical applications that have larger, more complex data sets that are always in use and involving a large number of hosts, block migration is the only way to simplify the process and to reduce the cutover downtime to a manageable window. As such, block level cloud migration is becoming a growing need, and yet there are few software solutions available designed from the bottom up specifically for block-level data migration to the cloud. Today, data migration projects often

use disaster recovery or backup/restore software products. They rely on a mixed bag of tools for this purpose, resulting in unpredictable delivery times, unexpected outages, decreased quality of service for the application owners, and a substantial learning curve for the people who are expected to utilize the tools successfully.

Cirrus Data Solutions developed Cirrus Migrate Cloud - a web-based portal specifically built to migrate block-level storage data to and between clouds. The product design is based on Cirrus' patented Mini Transparent Datapath Intercept ("mTDI") technology which allows Cirrus Migrate Cloud web-based portal software to be used in cloud migrations with true zero downtime (except for remote migration to a new host where there is near-zero downtime due to host switching).

Additionally, Cirrus Migrate Cloud includes data reduction in-flight, checkpointing features, and in-flight encryption ensuring secure, efficient migration to the cloud. With the introduction of the cMotion™ (patent pending) feature, users can perform a storage level cutover from a source storage to a destination cloud without downtime to the cloud VM. cMotion™ is used to swing the workload from any source storage to any new destination storage, including native cloud block storage or the more advanced cloud virtual arrays at the cloud VM. For remote migrations such as on-prem to Cloud or Cloud to Cloud, cMotion ensures that all the updates are redirected to the destination prior to host switching to avoid the unpredictable transition time.

This document introduces the core features of Cirrus Migrate Cloud. It focuses on simplifying various steps in the migration process, reducing risk, and significantly reducing the overall timeline and cost of data migration projects.

In summary, Cirrus Migrate Cloud value proposition for cloud migration is:

- No downtime for applications throughout the migration project (except for cutover to a remote host)
- Host software-based migration (Cirrus Migrate Cloud) supports migration between systems residing in any public/private cloud
- Enables use of any source storage, including captive cloud storage to any cloud, including 3rd party cloud virtual array storage

Cirrus Migrate Cloud

Cirrus Migrate Cloud, or CMC, is a patented software solution created to help customers move data from any block storage to and between different vendors' clouds simply, efficiently, securely, and with minimal impact on production operations. In addition, CMC is block storage device agnostic, giving customers a single software tool for all their block data migration requirements. The right cloud migration software-based solution can significantly reduce the migration time, complexity, and cost, and can eliminate the typically inevitable "cutover downtime" for a host migrating from one storage to another.

For remote migration, the right solution can significantly reduce the cutover downtime by eliminating the "final sync" after the source hosts are shut down and prior to the new destination hosts are brought up.

The elimination or drastic reduction of downtime ensures that every migration project can be completed within the allocated time window for cutover.

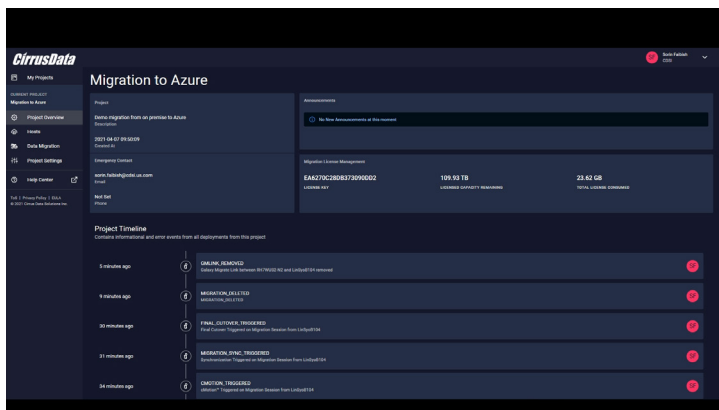
Cirrus Data Solutions developed Cirrus Migrate Cloud, a web-based portal specifically built to migrate block-level storage data to and between clouds.

CMC delivers a drastic reduction of downtime, reduced migration duration, inflight encryption, and a simple and efficient software management web portal.

This white paper presents in detail the methodology to successfully configure and execute the data migration using Cirrus Migrate Cloud. CMC simplifies the four phases of a typical storage cloud migration project: Planning, Deployment of Migration Solution, Migration, and Cutover. The paper will show how CMC delivers a drastic reduction of downtime, reduced migration duration, in-flight encryption, and a simple and efficient software management web portal. With the Cirrus Data Cloud portfolio of products, IT consumers now have several options for moving applications from on-premises to the cloud, between clouds, from captive cloud storage to 3rd party storage cloud arrays and back again, all without application disruption.

The Migration-as-a-Service Architecture

Cloud migration managers will find the Cirrus Migrate Cloud cloud-based portal easy to-use, since it requires only basic knowledge of the source or destination storage management. This is thanks to the simple and elegant architecture of CMC, which provides an easy to use Migration-as-a-Service offering that simplifies the migration process.



The CMC architecture is quite simple: a cloud-based portal discovers all hosts that need to migrate data in all the source/destination combinations among local, on-premises, cloud. Each host requires only running a single line curl Linux command or a single line Windows PowerShell command to register with the Migration Project defined in the portal followed by a download of the necessary CMC agent. For large scale projects, automation tools such as Terraform and Ansible can be used for massive deployment.

Given that each host is migrating its own data (under the central management of the CMC portal), the CMC solution can scale to meet the need of large-scale block storage migrations in the petabyte range. Unlike other solutions that require dedicated migration appliances that can be bottlenecks for the data flow, the

CMC architecture can have an unlimited number of hosts participate in the local or remote migration, as each host-to-host migration is independent. The migration leverages Cirrus Data's Intelligent Quality of Service, or iQoS, to ensure no impact on the production applications running at the source host.

With this simple architecture, there is no need to size the environment and plan for appropriately sized migration appliances as required by other solutions. You simply collect the list of hosts that need to migrate data, and add them to Migration Projects section of the CMC portal under your account. Once the hosts are registered, the Migration Project at the portal manages the rest of the migration process, starting with auto-remediation such as adding multipath drivers and iSCSI drivers if required by virtual cloud array destination storage. To further eliminate the need for any special skills or preparation work at the destination storage, CMC offers storage plugins for native cloud block storage (AWS EBS volumes, MS Azure Managed disks) and the many cloud virtual arrays (Pure CBS, NetApp CVO, HPE CVB, etc). The destination storage is allocated automatically without requiring any user interaction with the new storage. The block level data synchronization process then begins, with iQoS working in the background to ensure no impact to the host product I/O.

Once the Migration Project at the CMC Portal completes the initial block-level copy, the patent-pending cMotion™ process seamlessly moves the production IOs from the source storage to the destination storage. cMotion uses the new destination storage to impersonate the old source storage to avoid the need to shutdown the source filesystem for cutover, thereby eliminating the cutover downtime. As an example, for customers deploying a cloud virtual array to replace the native block storage, CMC can be deployed to migrate the data at the block level and the cMotion feature can swing the workload to the new storage all without downtime. The application continues to run throughout the migration process.

The following sections will provide details of the CMC process.

With CMC's simple architecture, there is no need to size the environment and plan for appropriately sized migration appliances as required by other solutions. You simply collect the list of hosts that need to migrate data, and add them to Migration Projects of the CMC portal under your account.

mTDI learns everything about the LUNs/disks: their identity and attributes, how each LUN/disk is presented to the hosts, initiator-target path information, LUN/disk sharing (clusters), and performance statistics such as IOPS, MB/s, latency, and queue depth.



Web-based Management Portal

The Cirrus Migrate Cloud Portal is a cloud management-as-a-service solution that allows users to manage migrations and protect any block storage anywhere. The Management Portal provides interfaces for all the Cirrus Migrate Cloud applications, configurations, management, and administrative tasks.

Another advantage of the web-based management application is its deployment in existing storage environments. Once installed, the agents are managed centrally by the Cirrus Migrate Cloud web-based management portal. Cirrus Migrate Cloud is the simplest, most straightforward data migration solution to achieve a successful result in the shortest time and with the lowest risk.

Automating the Discovery Process

The mTDI software automatically discovers the storage environment accessed by the host/VM. It learns everything about the LUNs/disks: their identity and attributes, how each LUN/disk is presented to the hosts, the initiator-target path information, LUN/disk sharing (clusters), and the performance statistics such as IOPS, MB/s, latency, and queue depth. Furthermore, the disk discovery process is fully automated and takes only seconds in the background to gather all the relevant information for the migration. The Cirrus Migrate Cloud agent then details this data on the web-based GUI.

Cirrus Migrate Cloud does not require manual entry of any additional details except selection of the devices to be migrated to the cloud. Since all of the source LUNs/disks are automatically identified and grouped by the host and storage, it is a straightforward process to isolate and select the disks/LUNs to be migrated using an on/off slider button in the Cirrus Migrate Cloud web-based portal.

Simplifying the Migration Process

Cirrus Migrate Cloud web-based portal executes the actual cloud data migration sessions. The web-based GUI of Cirrus Migrate Cloud is RESTful enabling command line or scripting tools to manage the migration.

The Cirrus Migrate Cloud web-based portal has many features that simplify the migration process and reduce the common risks of data migration projects.

a. Planning the Migration: Cirrus Migrate Cloud has simplified the Discovery Process to avoid having to run discovery processes on large numbers of hosts and collecting them for planning purposes. Once the CMC Agent is deployed using a project-specific code, the host automatically reaches out to the CMC portal (this is a secure, one-way connection that does NOT require opening any ports at the host) to report its system and storage details. All the source disks or LVM VG's are presented for selection for migration. In the case that a VG is selected, all the PV's that are part of the VG are automatically selected. This saves a lot of guesswork by the user and avoids potential mistakes in identifying PVs. Further, the disk signature or volume UUIDs are preserved on the migrated destination disks, thereby securing the intricate PV/VG/LV interdependency. As long as all the required disks are migrated for a set of LV's, the destination storage will continue to present the LV after the migration. No special preplanning work is needed to map out all the details of the PV/VG/LV. This drastically simplifies the pre-planning stage of the migration process if there are hundreds of hosts with large sets of LV's.

b. Preparing for Migration: Cirrus Migrate Cloud also has "plug-in" modules for popular brands of destination cloud storage and the native cloud block storage providers such as AWS EBS and Azure Managed Disks. CMC plug-ins create appropriately sized

The Cirrus Migrate Cloud web-based portal has many features that simplify the migration process and reduce the common risks of data migration projects.

CMC plug-ins create appropriately sized destination disks and matches up with the selected source disks. This means all the manual steps required by other migration solutions are eliminated.

destination disks and matches them up with the selected source disks. This means all the manual steps required by other migration solutions are eliminated, including creating a destination disk, as well as potentially having to create matching VG/LV's or filesystems in order to receive migrated data. Further, if the destination storage is one that requires other steps such as installing and configuring iSCSI drivers, multipath drivers, and other host remediation steps in order to access the new destination storage, the CMC plug-in will do this automatically. This provides a huge savings in time and a great simplification of the project in avoiding the requirement of a highly specialized storage management skillset. For remote migration such as on-premises to cloud, once the OS in the destination cloud VM is ready, the user needs to prepare the production application to run in the destination cloud VM and ensure that the correct version of the production applications/databases is available in the cloud with the appropriate licenses. Other steps unrelated to migration are the setup of networking, load-balancers, inbound and outbound security rules, etc.

c. Prepare the Cirrus Migrate Cloud Web Portal Account: To obtain a CMC account, follow the support page for the exact instructions on how to get an account at: <https://support.cirrusdata.cloud/portal/en/kb/articles/howdo-i-get-a-license-for-data-galaxy-complete>. Next, access the Cirrus Migrate Cloud portal using any browser to connect to and open a project for the migration on the Cirrus Migrate Cloud web-based Portal at <https://cloud.cirrusdata.com>. For remote migrations such as cloud to Cloud or On-premises to/from Cloud, ensure that either the source or the destination VM's have a specific inbound port open to allow connections with the other. If a storage plug-in needs to be used, obtain the credentials (or API Tokens) ahead of time to enable the auto-allocation of destination block storage devices.

d. Migration Project Creation: Each disk (example for Linux: /dev/sde, /dev/nvme2n1 /dev/dm-2, and for Windows: Disk2,

Disk3 ...) or each PV of LVM VG's is a unit of migration, usually defined by the host, application, or storage. Each set can be migrated and cut over as a group. Cirrus Migrate Cloud simplifies this process by utilizing a user-friendly web-based GUI that can select disks or even VG's at the Server or Storage level to be migrated with a single click. Selecting a VG will automatically select all its dependent PV's. Create a new project reflecting the specific migration characteristics, project name and description, type, owner of the migration, and any details needed to allow the operations.

e. Configure and Execute the Migration: Before starting the migration, we need to define the migration parameters including Migration Session Description, Data Synchronization Interval, and iQoS. The admin needs to decide the moderation level of the migration and Execute the Migration Session. Data reduction and encryption are used in-flight to reduce the bandwidth and time used for migration and to secure the data in transit.

f. Minimum Impact to Production: Cirrus Migrate Cloud accomplishes this by utilizing the iQoS intelligent Quality of Service functionality. This unique capability enables the migration session to be set to "Relentless," "Aggressive," "Moderate," or "Minimum" impact mode for each disk. Rather than artificially limiting the copy-process throughput or IOPS, Cirrus Migrate Cloud monitors the total I/O bandwidth consumption for each disk. If the maximum I/O bandwidth for a disk is exceeded, Cirrus Migrate Cloud will automatically reduce the copy process aggressiveness to yield to the host production I/O. This action is performed automatically and continuously, guaranteeing that the migration does not impact the business application performance while at the same time maximizing the usage of any unused I/O bandwidth for migration.

Overall, the Cirrus Migrate Cloud solution reduces risks and streamlines the migration process making it straightforward to create multiple migration sessions for disks and LV's with automatic iQoS control to

Cirrus Migrate Cloud integrates Cirrus Data's unique intelligent Quality of Service functionality, which enables migration administrators to configure the impact mode for each disk. If the specified maximum I/O bandwidth for a disk is exceeded, Cirrus Migrate Cloud will automatically reduce the copy process aggressiveness to yield to the host production I/O.

cMotion™ includes a revert feature that allows to workload to be moved back to the source storage so the can test the new storage without having to commit to the cutover.

ensure application owners do not experience any impact. The Cirrus Migrate Cloud web-based Portal can scale to thousands of migration sessions, hundreds of applications, and petabytes of data to be migrated.

Once the source and destination disks are synchronized and the migration sessions are in the “Pending Complete” state, Cirrus Migrate Cloud tracks the migrated delta blocks. It automatically performs a periodic resynchronization of the deltas, or the user can trigger the synchronization manually. This state is maintained indefinitely until the storage cutover is initiated. At this point, the systems are ready for cMotion™ migration cutover to be started. cMotion™ has introduced a new revert feature that allows to reverse the cMotion™ and swing the workload back to the source storage. This allows the user to try out the new storage without having to commit to the cutover.

Managing the Cutover Process

Once the source and destination disks have been synchronized, the process is ready to perform the cutover with cMotion™. Using cMotion™ cutover allows customers to migrate their data to any cloud safely. The production storage is retired, and the new cloud storage moves into production. The cutover process occurs as follows:

- a. Trigger synchronization to synchronize the new I/O's from source to destination.
- b. Trigger cMotion™ to swing the I/O workload from source to destination. I/O's from production storage are sent to the destination storage, and the source is quiesced.
- c. At any time, users can Revert cMotion™ and swing the I/O's back to the source production storage. This is possible until the cutover is finalized.
- d. Finalize the cutover. This locks in the cMotion™ and removes the ability to revert back to the source storage, thereby allowing the source storage to be removed to avoid further subscription

payments. After this step, the workload is moved to the destination cloud, and it is no longer possible to revert back to the source; the previous production disk(s) can be retired. No downtime is required, and the application continues to run. This is due to the fact that cMotion™ emulates the old storage using the new storage, thereby avoiding the need to remount the volume on the new storage, which requires downtime. This state can be maintained indefinitely.

The source hosts/VM's already reported the source storage details to the CMC portal. This knowledge is used by the plug-in to auto-allocate destination volumes. One-click and the new storage is set up by Cirrus Migrate Cloud.

Removing/Reducing Downtime for Cutover

Cirrus Migrate Cloud is efficient, and all the source storage devices are already synchronized. All that is left to be done is unmounting the old physical or cloud storage volumes and mount the new cloud storage at each of the cloud VMs. cMotion™ can be triggered forward, backward, and forward again, giving the user plenty of opportunities to fine-tune the destination cloud storage.

Cirrus Migrate Cloud allows all I/O's to be immediately redirected to the new cloud storage devices. The emulation of the old storage by the new storage eliminates downtime for cutover in the case of local (same host) migration between different block storage components (disks or arrays). In case of remote migration, downtime is minimized - reduced to a shutdown of the source application and a start of the application at the destination. The typical "final sync" required by other migration solutions is eliminated. This is because cMotion™ is redirecting all the production I/O to the destination storage prior to the host cutover, thereby eliminating any need for "final sync" altogether.

cMotion™ can be triggered forward, backward, and forward again, giving the user plenty of opportunities to fine-tune the destination cloud storage.



Cirrus Migrate Cloud passes on the new storage benefits of improved performance and reliability while removing the old physical or virtual storage and placing it offline immediately. No downtime is needed for the transition since the cloud host/VM already have access to the new cloud block devices. During all this time, cMotion™ tracks the dirty blocks from the source to the cloud destination and allows reversal of the I/O's back to source storage.

Cirrus Migrate Cloud web-based migration software portal is designed from the ground up to simplify and reduce the common risks associated with storage data migration projects.

Conclusion

Cirrus Data's portfolio of products is a solution to the issue of migrating data with limited or no application downtime. Cirrus Migrate Cloud web-based migration software portal is designed from the ground up to simplify and reduce the common risks associated with storage data migration projects. Compared with other data migration methods or tools, Cirrus Migrate Cloud has the least amount of risk, requires no special skills to manage cloud migrations, reduces the time and work to execute the migration to final cutover and, requires a minimal amount of specialized knowledge to deliver consistent, high-quality data migration services.

Cirrus Data's software products enable companies to use a single set of tools for all of their migration needs, including on-premises to a public cloud, between public cloud environments, and even from captive cloud storage to 3rd party cloud arrays.



About Cirrus Data

Cirrus Data Solutions Inc. (CDS) is a leading technology provider of next-generation solutions for data migration, data acceleration, and data protection. The company distributes its solutions through systems integrators, managed service providers, channel resellers, and partners including HPE, IBM, Dell/EMC, Datalink, Pure Storage, Infinidat, Align, SHI, CDI, Computacenter, Mainline, Sirius, WWT and many others. CDS is headquartered in Syosset, New York and has offices located in Dublin, Ireland, and Nanjing, China with sales and support offices in Boston, Chicago, and Denver. For more information, visit CDS online at www.cirrusdata.com.

Cirrus Data USA

6800 Jericho Turnpike
Suite 213W
Syosset, NY 11791

EMEA Headquarters

Unit 20A, Greenhills Indust. Estate
Greenhills Road, Dublin D12 E832
IRELAND

1.516.888.3389
info@cirrusdata.com
www.cirrusdata.com