

Spectra Logic StorCycle

Authors: Mitch Lewis, Randy Kerns

June 2020



Evaluator Group

Enabling you to make the best technology decisions



STORCYCLE[®]

Table of Contents

Executive Summary	2
Testing Environment.....	2
Identify	3
Testing and Results	3
Migrate	4
Testing and Results	4
Protect	6
Testing and Results	7
Access	7
Testing and Results	8
Additional Considerations	10
Cost Savings.....	10
Compatibility	11
Ease of Use	12
Performance.....	13
Future Features and Capabilities	14
Evaluator Group Assessment	16
Summary	16

Executive Summary

IT organizations around the world are consistently facing new data storage challenges as data volumes grow at explosive rates and the data storage systems and technologies continue to evolve. Many IT organizations are encountering immense amounts of data with little insight into the value of the data or how to best store and maintain it. While organizations may have a tiering strategy in place, without the proper tools they are often largely ineffective due to inactive data remaining undetected on the primary tier.

Far too often, primary storage systems, commonly consisting of the latest flash technology, are clogged and weighed down by inactive data and completed project data that is better suited for storage on a perpetual storage tier such as hard disk drives, tape, or the public cloud. Without properly identifying and managing this data, organizations are left wasting their expensive primary storage devices storing stale data that they are unaware of, leading to a loss of performance, and expensive acquisitions of additional primary storage.

To solve this problem, IT organizations need greater insight into their data to pinpoint what data is active and truly belongs on their fastest tier of storage, and what data would be better served moving to a perpetual storage location. Once inactive data has been located, it is crucial that it can be seamlessly migrated and restored as needed.

StorCycle, the latest solution from Spectra Logic, is a software defined storage lifecycle management product that is targeted at solving these data storage and management issues. StorCycle is capable of pinpointing inactive data by a variety of metrics, including age, size, and associated project, and migrating it to a more practical storage target such as a secondary NAS device, the public cloud, or an object storage system. With StorCycle, Spectra Logic identified four main functionalities that are crucial to properly archiving data: identifying, migrating, protecting, and accessing. Throughout a series of tests, Evaluator Group assessed StorCycle with regards to each of these categories in order to validate the product's capabilities.

Testing Environment

To execute the testing of the StorCycle software, StorCycle was deployed on a Windows server. An additional iteration of testing was also completed with StorCycle installed on a Linux server. Two NAS systems were connected to StorCycle, one as a primary storage system, and one as a secondary storage target. In addition, StorCycle's cloud compatibility was tested with a second target location set up in AWS. While StorCycle could also move data to tape or object storage using Black Pearl, that capability was not included in this test.

The data used in these tests was chosen to mimic that of a general IT environment, mostly using small files. Some larger files, such as imaging files from healthcare along with seismic and genomic data, were also tested to demonstrate StorCycle's effectiveness in handling various types of data.

Identify

Data identification is the first step in a successful data archiving strategy. Without a clear understanding of what data currently exists on a given storage system and how it is being used, it becomes overwhelmingly difficult to assign data to the most beneficial and logical tier.

Many organizations understand that files that have not been accessed in years, or belong to projects that have long been completed, do not need to reside on the highest performance, and often expensive, primary storage systems. The difficulty, however, is in visualizing the overall status of your storage and identifying the offending data.

StorCycle alleviates this issue with its data scanning capabilities. Scans may be run to identify data across the entire storage system, as well as individual directories, and reports are created to breakdown the data based on characteristics such as file size and age. This information provides critical insight into an organization's overall storage landscape and becomes the basis for a logical strategy of migrating data to a proper tier.

Testing and Results

To test the effectiveness of StorCycle's identification features, scans were run across the files stored in the test environment both as manual operations, as well as scheduled. Testing of the scanning process included scans across many different directories of test data, as well as full scans of the entire test data set.

Once scans were completed, the results were displayed. Scan results displayed basic information regarding the scan operation, such as showing the completion of the scan operation, the amount of data scanned, and the completion time, but the real value in the StorCycle scan process comes from a depiction of the stored data. As can be seen in Figure 1, StorCycle shows a breakdown of the scanned data distributed by file age and size.




Figure 1: StorCycle Data Scan

All StorCycle scan tests were completed successfully, with no errors or warnings, and properly displayed notifications once completed. For all directories scanned, data was identified by the age and size of scanned files. This visibility provides a valuable insight into what data is being stored and allows for an informed decision to be made regarding future data migration.

In July 2020, StorCycle will integrate a free scanner that will provide greater visibility into an organization's data. The tool will enable users to get an analysis of the cost savings they would realize by moving their inactive data from the primary tier, consisting of high-performance disk and flash, to the more cost-effective perpetual tier, which consists of spinning disk, object storage, tape and cloud.

Migrate

Once data's characteristics have been identified, data can then be migrated to the appropriate location. New or frequently used data will likely be best suited for staying on the primary tier where it is most likely accessed. Other data, however, such as older data that is infrequently used or very large files, may be stored more sensibly and efficiently on a different tier of storage.

When storing inactive data, economics and the need for capacity surpass the emphasis placed on performance that is necessary for primary data storage. This opens the door for IT organizations to leverage lower cost alternatives to flash storage such as HDDs or even tape. Very large files can become extremely expensive to store on primary storage systems and can be more effectively stored on a high capacity object storage system or in a public cloud, while freeing up primary storage resources.

StorCycle enables manual and automatic migration functionalities based on the age and size characteristics determined by its scans. In addition, data can be tagged by project and migrated accordingly. Unique to StorCycle, the Project Archive feature allows users to immediately migrate entire projects, such as completed scientific experiment output, finished videos and machine-generated data, to the perpetual tier for safekeeping on multiple storage platforms. The migration capabilities of StorCycle are capable of accommodating the many needs of various data sets and data storage environments by supporting migration to NAS systems and the public cloud along with object storage and tape via Spectra Logic's Black Pearl Object Storage Systems.

Testing and Results

To test the migration functionalities of StorCycle, data was migrated from a primary NAS device to a secondary NAS target as well as to the AWS public cloud. File migrations were conducted to both target environments based on file size and file age restrictions, as well as migrations which used a combination of the two parameters. Tests were conducted as one-time migrations, both run immediately and scheduled for specific times, as well as reoccurring scheduled processes.

Figure 2 shows the setup for a migration in which all data more than 90 days old and larger than 10 GB is to be migrated. As can be seen in figure 2, once migration parameters are selected, StorCycle displays a diagram of how much data will be archived based on the last run scan.




Figure 2: StorCycle Migration Setup

Testing of StorCycle's Project Archive functionality was also conducted. In testing this, entire project directories were migrated with metadata tags added to assist in identifying the project data.




Figure 3: Metadata Tagging During Migration

Tests found that StorCycle was able to successfully migrate data to both the target NAS and AWS, as well as migrating a copy to both systems simultaneously. Data was found to have been correctly migrated based on the defined parameters of age, size, and specified directory. Scheduled and recurring tests were found to run correctly and execute on schedule. All testing of various migration configurations, such as excluding specified file types or tagging project data, were found to have the expected migration outcomes.

The successful testing of StorCycle's migration functionality is indicative of its utility as an effective storage lifecycle management software platform. Through Evaluator Group's testing, large volumes of data were able to be relocated from primary storage to a perpetual storage solution. Migration criteria was easily specified by simple age or size ranges and displayed based on previously run scans. The migration features provided by StorCycle proved successful in offering a simplistic method to free up primary data storage by moving inactive data and entire projects to a more appropriate data tier.

Protect

A key aspect of data storage, whether it be primary or perpetual, is the protection of the data. Migrating or archiving data would be of little use if the data was vulnerable to attack during this process. Moving data between storage environments involves a few simple yet incredibly important requirements. Data moved from system A to system B must be guaranteed to actually arrive on system B. In addition, the data stored on system B must be completely unaltered from the original copy. It is also vital that valuable data cannot be obtained by a malicious actor during the migration process.

With data protection as a core value of StorCycle, many features are built in to ensure data migration can be achieved safely and correctly. Data can be encrypted during migration to prevent external access and checksums are implemented to ensure the validity of the migrated data. StorCycle also enables data protection practices via creating backup copies. While the original data can be removed from the source system during a migration, it can also be maintained while an additional copy is created on the target systems. Copies of data can be stored on up to three target systems, allowing organizations to protect data in multiple locations, including spinning disk, object storage, tape and cloud. With StorCycle, users can even use the software as a backup application to make a second copy on tape. Moreover, StorCycle enables users to store extra copies on tape for fast recovery in case of a disaster and for offline 'Air Gap' defense in case of ransomware attacks.

Testing and Results

When testing StorCycle's migration capabilities, Evaluator Group placed emphasis on ensuring the data protection requirements were upheld. Data that had been migrated was checked to ensure it was available and unaltered on the intended target system. This was done for data moved to the target NAS devices, data moved to AWS, and data moved to both systems.

StorCycle's copy functionality was also tested to ensure that data which was intended to be copied could be found unaltered on both the original primary source and the target systems.

By accessing the target systems, it was found that data was correctly migrated to both the NAS system and to AWS. Migrated files were found stored in the expected locations and remained unaltered from the original primary data. When testing the copy functionality, tests showed that the original data was left intact while copies were made on both target storage systems.

The protection of data is crucial to any organization's overall data storage needs and cannot be ignored in a data archiving strategy. StorCycle addresses these needs by performing its data migration functionality in a safe, trustworthy way. StorCycle ensures that data is migrated properly without any loss, alteration, or unwanted access affecting the data. In addition, StorCycle functions as a data protection tool that can create copies of important data on multiple storage systems to mitigate the risk of only storing a single copy of data.

Access

Moving old, inactive, or large data sets off of primary storage on to a cheaper or larger capacity storage tier can be a highly beneficial strategy to prioritize data, save money, and free up high performance storage for the data that needs it most. Certain situations may arise when previously archived data is needed, making it crucial that this data can be readily accessed.

StorCycle provides multiple options for handling the original data stored on the primary storage system as well as how to retrieve it when needed. When migrating data, the original data may remain on the primary system to create multiple copies of data, the original data may be removed from the source

completely, or the original data can be removed and replaced with an HTML link. StorCycle's HTML link functionality replaces the file with an HTML file of the same name which can then be used to restore the original file. File restoration can also be achieved via StorCycle's web interface which offers search functionality based on the migration project, file name, or metadata tags. These options provide flexibility for the varying migration and restoration needs of different data and remove the complexity involved in retrieving archived data.

Testing and Results

To test the access methods of StorCycle, migration jobs were completed using each of StorCycle's available options for maintaining the source data: retaining a copy, removing the source data entirely, and replacing the data with an HTML link. Data was then retrieved through a variety of methods.

For files in which the source data was intended to remain as a copy, the source system was checked to ensure that the files were correctly in place. When HTML links were chosen, those HTML links were tested by restoring single files as well as by restoring the entire associated data migration project. The following figures depict the process of migrating and retrieving data with HTML links during one such test.

Name	Date modified	Type	Size
 00e148aeea989ba56b	4/10/2020 4:26 PM	JPEG image	3,502 KB
 0a6bf3fa0a0d17aed4	4/10/2020 4:26 PM	JPEG image	3 KB
 0a37e2dba1c89ea4b7	4/10/2020 4:26 PM	JPEG image	13 KB
 0affaa5e5d0cd767f9	4/10/2020 4:26 PM	JPEG image	19 KB
 0b2f640b94a625e21e	4/10/2020 4:26 PM	JPEG image	58 KB
 0c9f5bd60c204a39a6	4/10/2020 4:26 PM	JPEG image	146 KB
 0c26b6cdb61fd8e8d9	4/10/2020 4:26 PM	JPEG image	314 KB
 0c79417e44b9523c60	4/10/2020 4:26 PM	JPEG image	223 KB
 0ca38898bfbf707dc5	4/10/2020 4:26 PM	JPEG image	1,778 KB
 0cfacfed56630d7875	4/10/2020 4:26 PM	JPEG image	92 KB
 0d11cf0159847e90be	4/10/2020 4:26 PM	JPEG image	48 KB
 0d61d5f95e66e535d1	4/10/2020 4:26 PM	JPEG image	66 KB
 0d794c49513ad0fdea	4/10/2020 4:26 PM	JPEG image	48 KB
 0da854bf125fd9e31b	4/10/2020 4:26 PM	JPEG image	36 KB
 0db71df03d631dfe94	4/10/2020 4:26 PM	JPEG image	32 KB
 0e56557c81927ebb	4/10/2020 4:26 PM	JPEG image	72 KB
 0e745380791e9686c0	4/10/2020 4:26 PM	JPEG image	13 KB
 0ef2b2fb71abab4e2e	4/10/2020 4:26 PM	PNG image	816 KB
 0ef148ffde892cb4fe	4/10/2020 4:26 PM	JPEG image	63 KB
 0f2b0f2107e928768b	4/10/2020 4:26 PM	JPEG image	358 KB
 0f31be490a342c1de0	4/10/2020 4:26 PM	JPEG image	98 KB
 0fb90613b639f93aa2	4/10/2020 4:26 PM	JPEG image	474 KB
 1a73c83d2aefcace85	4/10/2020 4:26 PM	JPEG image	84 KB

Figure 4: Migrated Image Files on Target

Figure 4 shows a series of image files that were migrated to a target system, while figure 5 shows the corresponding series of HTML files with the same names which replaced the images on the source system.

Name	Date modified	Type	Size
00e148aeaa989ba56b.JPG	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0a6bf3fa0a0d17aed4.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0a37e2dba1c89ea4b7.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0affaa5e5d0cd767f9.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0b2f640b94a625e21e.jpeg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0c9f5bd60c204a39a6.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0c26b6ddb61fd8e8d9.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0c79417e44b9523c60.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0ca38898bfbf707dk.5.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0cfacfed56630d7875.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0d11cf0159847e90be.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0d61d5f95e66e535d1.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0d794c49513ad0fdea.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0da854bf125fd9e31b.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0db71df03d631dfe94.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0e56557c1c81927ebb.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0e745380791e9686c0.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0ef2b2fb71abab4e2e.png	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0ef148fde892cb4fe.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0f2b0f2107e928768b.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
0f31be490a342c1de0.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB

Figure 5: HTML Links on Source

Figure 6 shows an example of an HTML file that is displayed when accessing one of the HTML links. The file shows information about the file and migration project as well as the option to restore the file.



This File has been Archived

Files have been moved from this device to archival storage, both for long-term preservation and more efficient use of primary media. The file has been replaced by this html file, but it can be recovered using the directions below.

File Information

Name: Copy1\Images\oregon-wildlife\oregon_wildlife\oregon_wildlife\bald_eagle\00e148aeaa989ba56b.JPG

Archived date: Friday, 17-Apr-20 18:44:39 UTC

Last accessed: Wednesday, 23-Oct-19 03:19:50 UTC

Size: 3.4 MiB (3,585,708 bytes)

Archived to: Target

Task: Migrate Oregon wildlife Images

Task Description: Migrating all Oregon Wildlife Images

Recovery Options

Restore Now

Restore *Copy1\Images\oregon-wildlife\oregon_wildlife\oregon_wildlife\bald_eagle\00e148aeaa989ba56b.JPG* to original location

Figure 6: File Restore from HTML Link

Figure 7 again shows the data on the source system, now after the entire project has been restored. It can be seen that the original files were restored and the html links remain.

Name	Date modified	Type	Size
 00e148aeaa989ba56b	4/10/2020 4:26 PM	JPEG image	3,502 KB
 00e148aeaa989ba56b.JPG	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0a6bf3fa0a0d17aed4	4/10/2020 4:26 PM	JPEG image	3 KB
 0a6bf3fa0a0d17aed4.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0a37e2dba1c89ea4b7	4/10/2020 4:26 PM	JPEG image	13 KB
 0a37e2dba1c89ea4b7.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0affaa5e5d0cd767f9	4/10/2020 4:26 PM	JPEG image	19 KB
 0affaa5e5d0cd767f9.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0b2f640b94a625e21e	4/10/2020 4:26 PM	JPEG image	58 KB
 0b2f640b94a625e21e.jpeg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0c9f5bd60c204a39a6	4/10/2020 4:26 PM	JPEG image	146 KB
 0c9f5bd60c204a39a6.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0c26b6ddb61fd8e8d9	4/10/2020 4:26 PM	JPEG image	314 KB
 0c26b6ddb61fd8e8d9.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0c79417e44b9523c60	4/10/2020 4:26 PM	JPEG image	223 KB
 0c79417e44b9523c60.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB
 0ca38898bfbf707dc5	4/10/2020 4:26 PM	JPEG image	1,778 KB
 0ca38898bfbf707dc5.jpg	4/10/2020 4:26 PM	Chrome HTML Docu...	18 KB

Figure 7: Restored Files and HTML Links on Source

For files that were removed and migrated without replacement by HTML links, data restoration was tested using the StorCycle web interface. The web interface was successful in restoring a specific file given its name and path as well as using the search functionality. Searches were conducted to restore files by the file name, migration project name, and assigned tags.

Successful testing of StorCycle's access methods demonstrate the system's ability to manage data throughout its full lifecycle. While previous tests demonstrated an ability to easily migrate data, an archiving solution would not be effective if the data could not be restored for further access just as easily. StorCycle provides a number of options to flexibly fit the migration and retrieval needs of different data, projects, and organizations.

Additional Considerations

Cost Savings

An important aspect of managing and archiving data is the associated cost. Primary storage systems are generally significantly more expensive than perpetual storage systems due to extra emphasis on performance. By identifying and migrating old and inactive data to lower cost perpetual storage tiers, StorCycle is able to offer significant cost savings and maximize the utilization of high-performance primary systems.

StorCycle features a built-in cost tracking tool that was enabled during Evaluator Group testing. The cost tracking tool allows a price per TB to be set for any configured storage location. Storage can also be assigned to departments for more granular cost tracking within an organization. To evaluate the cost

effectiveness of StorCycle while using a test data set of limited size, easily distinguishable costs of \$10,000/TB and \$1/TB were set for source and target sources respectively. After various testing that included migrating data based on a number of different size, age, and project criteria, a total of 1.4 TBs was archived. The cost tracking for this testing is shown to be \$14,456.62.

Evaluator Group Comment: It should be noted that the pricing numbers used in in this testing were intentionally selected values of \$10,000/TB and \$1/TB. These prices were chosen to demonstrate a large contrast in cost between storage tiers, and to easily display StorCycle's cost calculator functionality while using a small set of test data.




Figure 8: StorCycle Cost Tracking

There are a number of factors that can contribute to differing cost savings, including the criteria determined for migrations, the total amount of stored data, and the varying storage costs of source and target systems. While an in-depth cost analysis was not performed during this testing, it can be seen from the cost tracking example that StorCycle can be used as an effective tool to optimize data storage costs.

Compatibility

StorCycle was designed with compatibility in mind. The storage targets integrated with StorCycle support all industry-standard operating systems and major backup applications, such as Commvault and Veeam, allowing users the ability to run multiple applications and share the capacity of their storage targets with StorCycle.

Ease of Use

One of the key takeaways from this series of validation testing was the simplicity and ease of use of the StorCycle software. From installation and configuration through migration and data recall, StorCycle proved to be intuitive and simple to use.

Configuration of the source and target system was a simple process of providing the path to the system or, in the case of AWS information, access to the cloud storage. Few other configuration steps were necessary, all of which were equally simple.

Using StorCycle software was also found to be a straight forward process. Access to all of the main functionalities were easily located on a side panel that made it easy to move between scanning, migrating, restoring, and other information such as reports.




Figure 9: StorCycle Dashboard

For the creation of any new scan, migration, or restore job, StorCycle uses guided steps that lead the user through the appropriate settings, requirements and parameters. An example of this can be seen in the simple two step guide for a new scan shown below in figure 10.

The screenshot shows the 'New Scan' configuration window. At the top, there are two steps: '1 Scan Source' and '2 Schedule'. The 'Scan Source' step is active. The form includes the following fields and options:

- Project Name:** Scan_Copy1
- Scan Source:** Source (dropdown menu)
- Root Directory:** \Copy1
- Base or Working Directory:** A file browser showing a directory structure with folders: .. to SPECTRA2, AirlineData, and Audio.

Figure 10: Setup of New Scan Job

Upon completion of jobs, notifications of success, warnings, or errors were easily displayed from the StorCycle dashboard. In addition, SMTP was configured during testing and all notifications were received via email as well.

This testing of StorCycle involved the creation of two users on the system, an Administrator, having privileges for system configuration and running jobs, and a Storage Manager, with only the ability to run jobs. Both tested users were found to be successful in executing their assigned roles. StorCycle has additional functionality to integrate with Active Directory, however this was outside the scope of this testing project.

The overall user experience of StorCycle was found to be a simple one. System configuration can be done quickly and scanning, migrating, and restoring files can be easily executed by following the guided wizards. The ease of use of StorCycle allows it to transition seamlessly into an organization's data storage environment and quickly transform their data archival strategy.

Performance

This particular set of testing was not done with a focus on performance metrics; however, performance can be an impactful characteristic when considering data archival solutions. For example, Spectra Logic has stated that StorCycle can scan up to 10,000 files per second and transfer files at up to 2.75GB/s.

All jobs run in this testing were observed to complete in a timely manner, with some variance depending on the amount of data involved. When running a migration job, the performance was able to be increased by selecting to use the last run scan rather than running an entirely new scan.

StorCycle also has additional functionality that can alter the performance of the system. File transfers are performed in batches of either 10,000 files or 10 GB. Files are parallelized in streams within batches and this parallelization can be altered to fit specific performance needs.

An additional performance consideration offered by StorCycle is the ability to enable transfer throttling for peak hours. This can be configured per storage target and allows users to set a schedule of peak hours and desired performance limits. Job throttling configuration is depicted in figure 11.

The screenshot shows a configuration window titled 'ThrottleTransfer' with a green header. Below the header is a green box with the instruction: 'Select the day of the week and then drag the handles to select the times (all times UTC)'. The main area is titled 'Peak Hours' and contains two sections. The first section, 'Peak Hours Bandwidth', has a value of '40 MB/s'. The second section, 'Scan Objects per Second', has a value of '30'. Below these is a label 'Max MB per second' and a note 'Limit scan overhead: peak times'. A calendar view shows the days of the week (Su, Mo, Tu, We, Th, Fr, Sa) with horizontal bars indicating peak hours: Mo (08:30-14:00), Tu (07:00-18:00), and Fr (10:30-18:00).

Figure 11: Setup of New Scan Job

Future Features and Capabilities

Throughout this testing, StorCycle has proven to be an effective tool for managing and archiving data across storage tiers and has demonstrated a number of useful features. As Spectra Logic continues to develop the StorCycle software, additional functionality will be added to further enhance the solution.

Spectra Logic has adopted an agile approach to their software development with new features and functionalities being added on a quarterly basis to further improve the ease of use and technical capabilities of the software. In fact, Spectra performs more than 7,000 automated tests that run every

time they create a build of StorCycle which helps them provide the highest quality product in the most efficient manner possible.

Some of the announced roadmap capabilities include support for Amazon Glacier as a target location (available in June 2020), an enhanced cost savings calculator, and the ability to clone archive projects. Additional features have also been planned such as the ability to delete data manually or through a scheduled process (scheduled for September 2020), support for re-running migration jobs (scheduled for September 2020), file versioning capabilities (scheduled for December 2020), and support for using the cloud as a source location (scheduled for 2021). Spectra Logic also plans on adding advanced scanning capabilities to dive deeper into identifying inactive data (scheduled for 2021).

StorCycle's data access functionalities are also planned to receive new capabilities with the addition of symbolic links, which are scheduled for June 2020. In future versions, StorCycle will be able to replace an archived file with a symbolic link which allows users to access the data on a target system as if it were still on the source system.

The features included in StorCycle's roadmap, combined with the agile development and release cadence, will continue to enhance the effectiveness of StorCycle's data management and archival abilities. StorCycle has demonstrated value that can be critical to a comprehensive data storage strategy, and upcoming releases will add additional flexibility and functionality for StorCycle to further assist in data's storage lifecycle.




Figure 12: StorCycle Storage Management Software identifies and moves inactive data and finished projects from the costly Primary Tier to the cost-effective Perpetual Tier, made up of spinning disk, object storage, tape and cloud.

Evaluator Group Assessment

The testing documented in this report served as a functional validation of Spectra Logic StorCycle and its capabilities as a storage lifecycle management solution. After extensive testing of the StorCycle software, Evaluator Group has made the following assessment:

- StorCycle provides an effective scanning tool for visualizing stored data and identifying stale and inactive data. This identification is invaluable for organizations to develop a full understanding of the data they are storing.
- Once inactive data is identified, StorCycle's data migration capabilities allow it to be easily migrated to the most appropriate storage tier, which includes any combination of disk, object storage, tape and cloud, to provide greater storage utilization, enhanced storage planning, and large potential cost savings.
- StorCycle includes the ability to move data by two methods: Auto Migrate and Project Archive, giving organizations choices in how they store, track, manage and protect their data.
- The features provided with StorCycle provide functionality to facilitate a full data lifecycle from identification through retrieving archived data. The ability to quickly access and retrieve previously archived files and projects makes StorCycle a flexible solution for managing data as an organization's needs or resources may change.
- StorCycle maintains an intuitive user experience from installation through managing data across its full lifecycle process.
- Spectra Logic's commitment to continuous, agile development of the StorCycle software will further enhance the many storage management benefits that were found during this testing. Evaluator Group found current capabilities of StorCycle to be both simple to use and effective in managing data, and expects future releases to further build upon this.

Summary

As data growth continues to accelerate, the need to effectively manage data is becoming increasingly important. While simple in concept, data management requires a lifecycle of identifying, migrating, protecting, and accessing data, all of which become more challenging as the amount of data increases.

Spectra Logic's StorCycle is positioned to address the issues of data management during each step in this data lifecycle. The tests conducted were found to validate numerous StorCycle features including scanning, migrating, and retrieving data. StorCycle was also found to be secure and reliable while offering a simplistic interface and the potential for significant cost savings. These tests conclude that StorCycle is an effective tool for solving data management and archiving challenges.



About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2020 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, consequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.