



Storage Assessment Report User's Guide

| | |
|---|----|
| <i>Introduction</i> | 3 |
| <i>How It Works</i> | 3 |
| <i>Getting Started</i> | 4 |
| <i>Reading the Reports</i> | 12 |
| <i>Analyzing the Report</i> | 13 |
| <i>Configure Tiger Bridge</i> | 19 |

This publication, or parts thereof, may not be reproduced in any form, by any method, for any purpose.

TIGER TECHNOLOGY MAKES NO WARRANTY, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES, OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, REGARDING THESE MATERIALS AND MAKES SUCH MATERIALS AVAILABLE SOLELY ON AN "AS-IS" BASIS.

IN NO EVENT SHALL TIGER TECHNOLOGY BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH OR ARISING OUT OF PURCHASE OR USE OF THESE MATERIALS. THE SOLE AND EXCLUSIVE LIABILITY TO TIGER TECHNOLOGY, REGARDLESS OF THE FORM OF ACTION, SHALL NOT EXCEED THE PURCHASE PRICE OF THE MATERIALS DESCRIBED HEREIN.

Tiger Technology reserves the right to revise and improve its products as it sees fit. This publication describes the state of this product at the time of its publication, and may not reflect the product at all times in the future.

THIRD-PARTY TRADEMARKS

All other brand names, product names, or trademarks belong to their respective holders.

Title: Storage Assessment Report User's Guide

Product version: 1.0

Date: April 27, 2022

Manual Revision and Control

Revision Record

| Date | Description | Page | Version |
|--------------|---------------|------|---------|
| 07 Feb. 2020 | Initial Draft | | |

Introduction

Tiger Technology's Storage Assessment Report utilizes Microsoft's Power BI engine to provide you with detailed reports about data usage on a locally mounted NTFS drive or an SMB/NFS share mapped on the computer. Statistics about the size and percentage of created, accessed and modified data in the selected path allow you to gain deeper insight into how your organization uses storage, to assess the approximate data growth as well as the gains of migrating to hybrid storage. The report can also be used as a starting point for configuring Tiger Bridge, Tiger Technology's data lifecycle management solution across storage tiers.

This manual will guide you in the process of downloading and installing the tool and generating reports.

How It Works

The Storage Assessment Report workflow goes through three simple steps:

1. Collect metadata from the path you want to analyze, using the Storage Assessment Data Collection tool.
2. Submit the collected metadata to the Microsoft Power BI engine for analysis, using the Storage Assessment Data Collection tool.
3. View the storage assessment report

To begin, you need to download and install the Storage Assessment Data Collection Tool on a computer running Microsoft Windows. You can then generate a Storage Assessment Report for any path accessible on the computer (a whole NTFS volume, a subfolder or a mapped SMB/NFS share). The tool collects data about the size, number and percentage of total size of created, accessed and modified files. It then uploads this metadata to a Microsoft Azure container, where the Microsoft Power BI engine analyzes it for the purposes of the Storage Assessment Report. Once the analysis is ready, you need to log on to the Tiger Technology licensing server to view the report. The Storage Assessment Report order is active for 30 days. After that the order and the associated data are automatically deleted.

Note: *Tiger Technology does not gain any access to your data.*

Number of Reports Per Order

When you register for a Storage Assessment Report order, you need to specify how many slots your order should be charged for. The number of slots in your order determines the maximum number of reports you can generate concurrently. For example, an order for two slots allows you to generate two concurrent reports for two drives either on the same computer or each on a separate computer. Even if you choose to create an order for just one slot, you can generate reports for as many drives/

SMB shares on the same computer as you wish, keeping in mind that a newly generated report overwrites the last one. You cannot overwrite a report generated on one computer with a report generated on another. To help you with discerning between slots and their respective reports, the Storage Assessment Data Collection Tool allows you to specify a name for each slot and to attach a UTC timestamp with the date and time of creation to the name. Additionally, the slots drop-down box also displays the serial number of the tool on the computer on which the slot has been used. You can get the tool's serial number on the computer following the steps below and this way discern between computers having used slots in your order.

To view the Storage Assessment Data Collection Tool serial number:

1. Start Command Prompt and change the directory to the folder containing the serial_tool.exe.
The serial_tool.exe is located in the "tools" sub-folder of the Storage Assessment Data Collection Tool installation folder. Thus, if the tool is installed in C:\Program Files\Tiger Technology\Storage Assessment, execute the following:
`cd "C:\Program Files\Tiger Technology\Storage Assessment\tools"`
2. To view the serial number of the tool on the computer, execute the following:
`serial_tool . exe`
The serial_tool.exe displays the serial number of the Storage Assessment Data Collection Tool on this computer.

Getting Started

Register an Order

Before you can download and install the tool, you need to register for access to the Tiger Technology licensing server and complete your payment for the Storage Assessment Report order, following the steps below.

To register for a Storage Assessment Report order:

1. In a web browser, go to:
<http://www.tiger-technology.com/services-and-support/storage-assessment-report>
2. Under Storage Assessment, click Get Started.

3. Fill in your name and email, and then click Submit.

Tiger Technology Storage Assessment Report

Please enter your name and email to get further instructions and finish your registration.

First name*

Last name*

Email*

protected by reCAPTCHA
Privacy - Terms



Submit

Note: If you can't find the Registration email from Tiger Technology in your **Inbox**, please check your **Spam** folder.

Note: Provide a valid email, on which to receive the link needed to complete your registration.

4. Follow the link received in the email from Tiger Technology and fill in the registration form:

| | | |
|------------------|---|--|
| Username | <input type="text" value="user@user.com"/> | |
| Password | <input type="password"/> | |
| Confirm Password | <input type="password"/> | |
| Full Name | <input type="text"/> | |
| Country | <input type="text" value="---"/> | |
| City | <input type="text"/> | |
| Postal/ZIP code | <input type="text"/> | |
| SAR Region: | <input type="text" value="(Europe) West Europe"/> | |
| SAR Reports: | <input type="text" value="1"/> | |
| <p>Register</p> | | |

Storage Assessment Report User's Guide

- Enter the password for your account on the Tiger Technology licensing server, your name and address in the corresponding fields.
- In SAR Region, select the region in which the container with the collected metadata to be created.
- In SAR Reports, enter the number of concurrent reports you want to be able to generate.

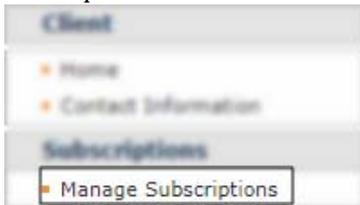
Note: The maximum number of concurrently generated reports per order is 30. In case you want to use more or to increase the number of reports after you have created your order, you need to create a new Storage Assessment Report order in your account for the additional number of reports, following the same steps.

5. Click Register.

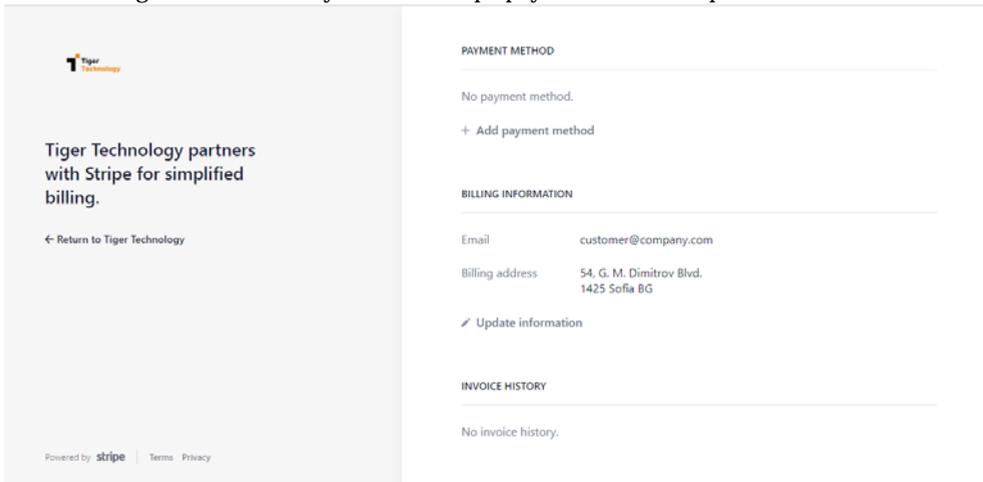
You are automatically redirected to the home page of your account on the Tiger Technology licensing server.

Tip: Click the Storage Assessment Report order name to view a sample report.

6. Click the name of the Storage Assessment Report order, then in the left pane, click Manage subscriptions.



The licensing server redirects you to the Stripe payment customer portal.



7. Click Add payment method.

The screenshot shows a section titled "PAYMENT METHOD" with a horizontal line below it. Below the line, the text "No payment method." is displayed. At the bottom of the section, there is a button with a plus sign and the text "Add payment method".

8. Enter the requested credit card information and select the “Use as default payment method” check box, then click Add.

The screenshot shows a form titled "Add payment method" with a breadcrumb "Billing > Payment method". Under the heading "Card information", there are two input fields: "Card number" and "MM / YY". Below these fields is a checkbox labeled "Use as default payment method" which is checked. At the bottom of the form are two buttons: "Add" (orange) and "Go back" (white).

By adding a payment method, you agree to Tiger Technology, Inc's Terms of Service and Privacy Policy.

Your payment information is saved and you can go back to your account on the Tiger Technology licensing server to download and install the tool, following the steps below.

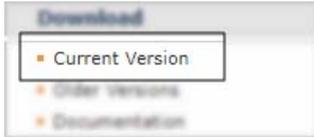
Software Digital Certificate

The Storage Assessment Data Collection Tool uses a digital certificate issued by GlobalSign certification authority. For the digital certificate to be verified upon installing Storage Assessment Data Collection Tool, GlobalSign's root certificate must be installed in the Trusted Root Certification Authorities of the Certificate Manager on the computer. On computers connected to the Internet and operating in less restrictive environments, this is done automatically. If the computer on which you want to install Storage Assessment Data Collection Tool operates in a more restrictive domain environment, you must manually download the root certificate from GlobalSign and install it yourself, before installing the product.

Download and Install the Tool

To download and install the Storage Assessment Data Collection Tool:

1. Log on to the Tiger Technology licensing server at:
<https://license.tiger-technology.com>
2. On the home page of your account, click the name of the Storage Assessment Report order, then in the left pane, click Current Version.



3. Click the Download button  next to the Storage Assessment Data Collection Tool.
4. Double-click the downloaded installation file to install the tool.
5. To open the tool, do one of the following:
 - Double-click the Storage Assessment Data Collection Tool shortcut on the desktop.
 - Navigate to the installation folder of the Storage Assessment Data Collection Tool and double-click the installation file.
6. When prompted, confirm that you allow the application to make changes to your device and if needed, enter the credentials of a user with administrative privileges.

Generate and View a Report

You can generate a report for any locally mounted NTFS volume or an SMB/NFS share accessible to the computer with the Storage Assessment Data Collection Tool installed.

Network Share Path Prerequisites

As the Storage Assessment Report operates at system level, to generate a report for a network share, you must first map it on the computer, executing the steps with elevated privileges:

1. Run command prompt as an administrator:
 - Click Start.
 - Type cmd.
 - Right-click Command Prompt and then click Run as administrator.

2. In command prompt, execute the following:

```
net use [drive letter] [path] /user: [username] [password]
```

where:

[drive letter] is the drive letter to which you want to map the network share

[path] is the full path to the folder you want to map as a network share

[username] is the name of the user with write access to the share on the computer exporting it

[password] is the password of the user with write access to the share on the computer exporting it

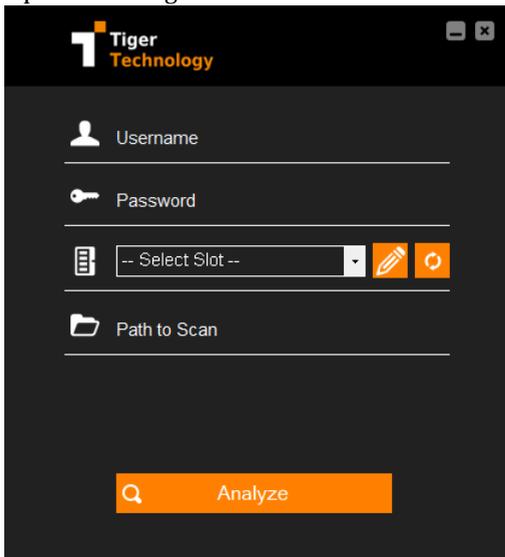
For example, to map a network path \\server\folder to drive letter F and to be able to access it with the privileges of user1 with password password1234, execute the following:

```
net use F: \\server\folder /user: user1 password1234
```

3. Repeat the step above for each network share the contents of which you want to analyze with the Storage Assessment Report.

To generate a report:

1. Open the Storage Assessment Data Collection tool.



The screenshot shows a dark-themed application window titled "Tiger Technology". It contains four input fields for user credentials and scan parameters:

- Username:** A text input field with a person icon on the left.
- Password:** A text input field with a key icon on the left.
- Slot Selection:** A dropdown menu with the text "-- Select Slot --" and a list icon on the left. To the right of the dropdown are two orange icons: a pencil (edit) and a refresh symbol.
- Path to Scan:** A text input field with a folder icon on the left.

At the bottom of the form is a large orange button with a magnifying glass icon and the text "Analyze".

2. Enter your order name and password in the respective fields.

Storage Assessment Report User's Guide

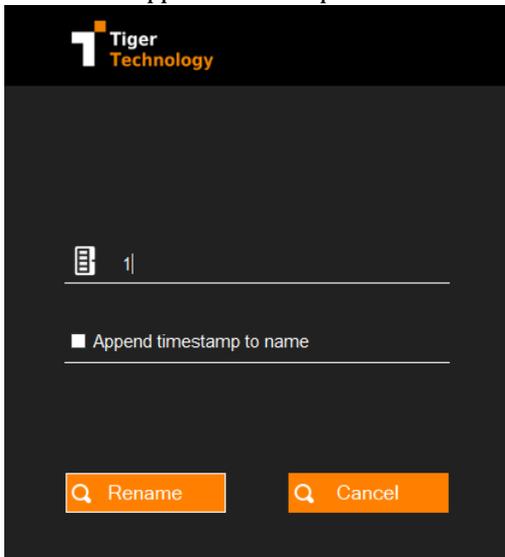
Note: Unless you have changed the password of your order, it uses the same password as the password of your account on the licensing server. You can find the order name in your profile on the Tiger Technology licensing server.



3. Select the slot to use for the report you are about to generate.

Note: If you select a slot, which has already been used on the same computer, you will overwrite the current report for the slot with the report for the newly analyzed path.

4. Click the edit button  and in the dialog that opens enter a name of the slot and optionally enable the “Append timestamp to name” check box, then click Rename.



5. In Path to Scan, click the folder icon and browse to and select a drive/mapped network share, or any subfolder, to analyze its data usage.

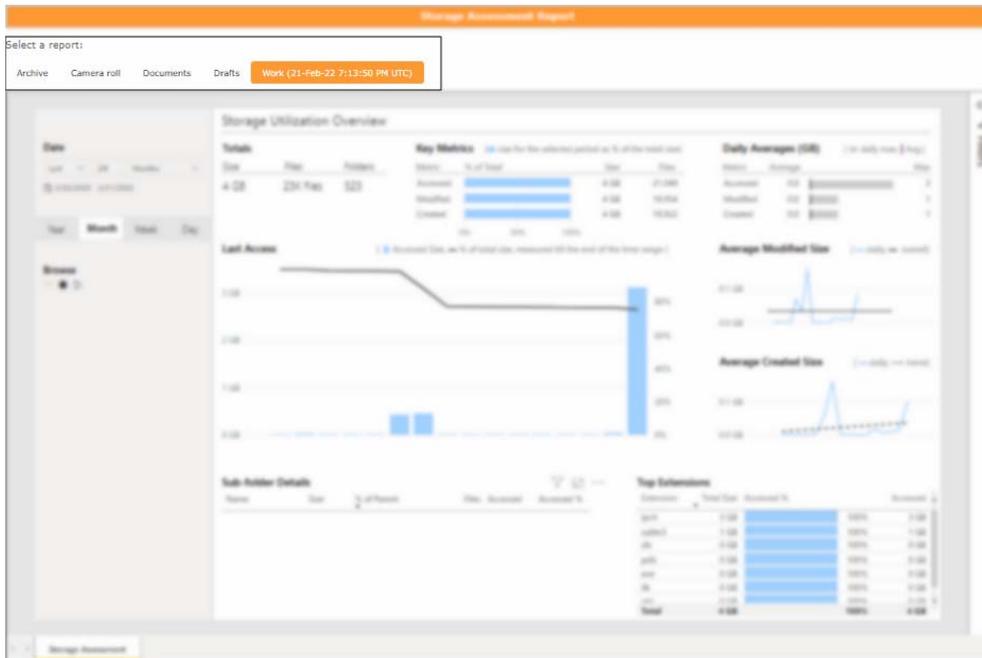
Note: You can narrow down the Storage Assessment Report to any subfolder of the selected path which is up to 5 levels deep in the folder hierarchy.

6. Click Analyze and when prompted, confirm that you agree to submit it for analysis.

To view the report:

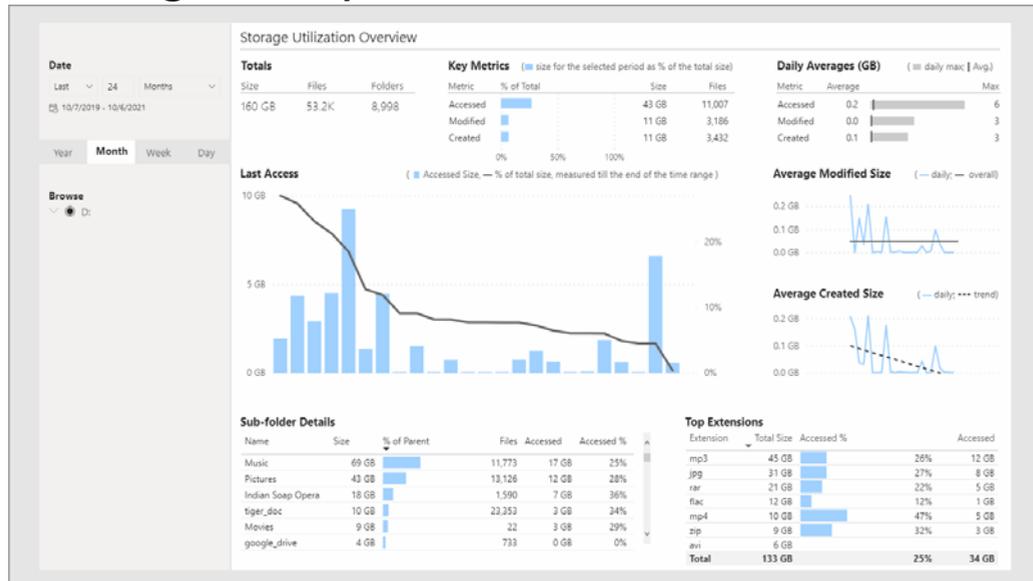
1. Log on to the Tiger Technology licensing server.
2. On the home page of your account, click the name of the Storage Assessment Report order.

- At the top of the page, click the name of the slot corresponding to the path whose report you want to view.



Note: If the page displays only the sample report, you have either not completed your registration with the required payment information or have not submitted any metadata for analysis, following the steps on page 9.

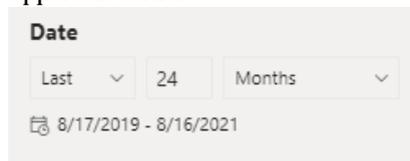
Reading the Reports



The web report about the selected drive gives you statistics about the number, size and % of total size of:

- created files
- accessed files
- modified files

The report presents the state of data usage for a selected time frame - any number of years, months, weeks or days. To refine the time frame for which you want to view a report, use the controls in the upper left corner:



You can choose to display the distribution of data for the selected time frame by year, month, week or day, clicking the respective tab below the time frame controls. This updates the Last Access, Average Modified and Average Created graphs. At the bottom of the report you can view statistics about each sub-folder of the analyzed path and about the top 7 file types by size.

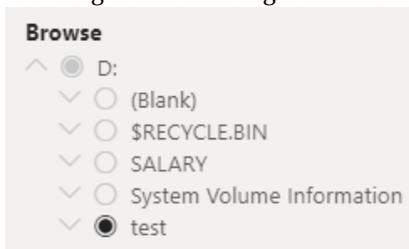
Hovering your mouse over any of the graphs displays additional information.

Choosing the Path to be Analyzed

While choosing the path to analyze depends greatly on your knowledge of your organization's workflow and data structure, you can adhere to the following rules for more consistent results:

- You can gain better insight into your organization's data usage, comparing statistics over longer periods of time. Thus, it is advisable to generate it for a storage device, which has been actively in use for at least one year.
- All sub-folders of the analyzed path should show a similar pattern of data usage over a longer period of time. Deviations from the general data usage pattern in one sub-folder can distort the conclusions you make about the parent folder.

With all of the above in mind, you can refine the report statistics to a specific path, by simply browsing to and selecting it under Browse in the lower left pane of the report:



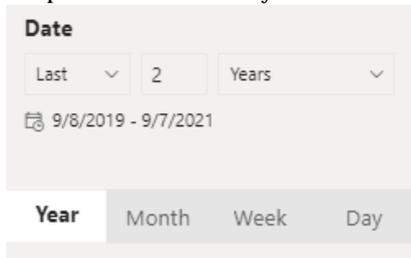
Analyzing the Report

Assessing Future Data Growth

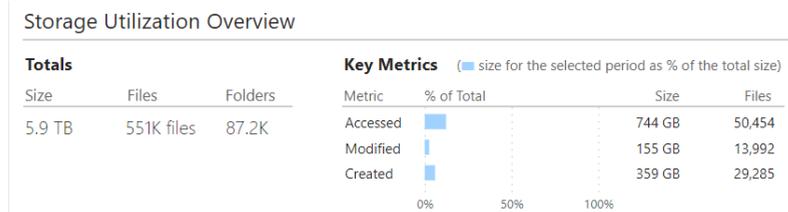
In today's digital world data is rarely decreasing. You can estimate by how much your organization's data is growing on a yearly basis by simply comparing the statistics about the size of all newly created data from several consecutive years. This, combined with calculations about planned wear and tear, allows you to estimate the approximate future investments needed to maintain your current storage workflow.

Sample steps to estimate your data growth:

1. Load the report for the desired slot and in the left pane, configure the controls above to display a report for the last two years.



2. Check the Created field of the Storage Utilization Overview.



3. Repeat the above two steps, changing the time criteria of the report to just one year.
4. Deduct the size of all newly created files on your storage for the last two years from the same value in the report for the last year only and you will get the approximate used space increase on your local storage.

Estimate the Benefits of Hybrid Storage

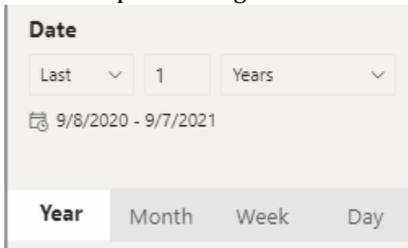
In the context of this document, hybrid storage refers to a storage environment made up of on-premises and a public/private cloud infrastructure. Users and applications perceive all components of the hybrid storage as one whole although it is comprised of the following tiers:

- local storage tier - this is the on-premises infrastructure, which provides immediate access to your data.
- nearline tier - this is the hot or cool tier (or their respective equivalents) of your public cloud storage provider. Both allow fast retrieval of data residing there to the local tier, where it can be previewed or modified. The difference between hot and cool lies in the ratio between costs for hosting data on the respective tier and the speed of retrieving it.
- archival tier - this is the cloud storage option intended for storing data that is hardly ever used. While the price for storing data in the archival tier is much lower compared to other cloud tiers, retrieving data from this tier takes considerable time and is more expensive.

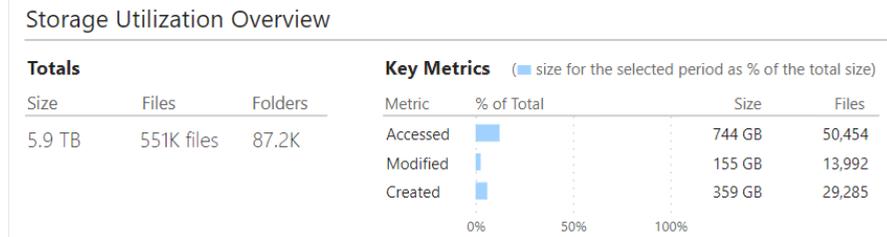
The main benefit of fast local storage is that all data is readily available for everyone to preview, access and modify. Deducting the amount of local data that has not been accessed in the last year, for example, allows you to easily assess how much dormant data is taking up valuable local storage resources. Moving such data to a nearline or archive tier of the hybrid storage will not only free local space, but will enormously ease the process of planned local storage replacement as there will be less data to migrate.

Sample steps to estimate the amount of data which can be moved to the cloud:

1. Load the report for the desired slot and if necessary, refine the analyzed path by selecting a sub-folder.
2. In the left pane, configure the Date controls to display a report for the last year.



3. Check the Storage Utilization Overview of the report.



4. Deduct the size of accessed files from the total size of files on your storage and you will get the approximate amount of space you can free on your local storage by moving data that has not been accessed to a nearline tier of hybrid storage. The sum will also allow you to get a rough estimate of the costs of hosting nearline or offline data in the cloud.

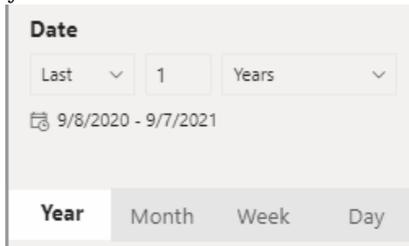
Estimate the Need for a Hybrid Storage Archival Tier

The main difference between cloud storage tiers is in the price and time needed to retrieve data from the respective tier. While the archival tier provides a much cheaper option for storing data, the price for retrieving data from it is higher and the time needed to accomplish it is longer. Thus, the archival tier of any cloud storage provider is most suitable for safely storing data that you do not access and that you may have to retrieve only to recover from disaster, for example.

An approximate evaluation of whether data in the analyzed path is suitable for archiving can be performed by comparing the graphs for last accessed data and average created size for a period of one year. If the graphs show similar patterns, data is suitable for archiving and vice versa.

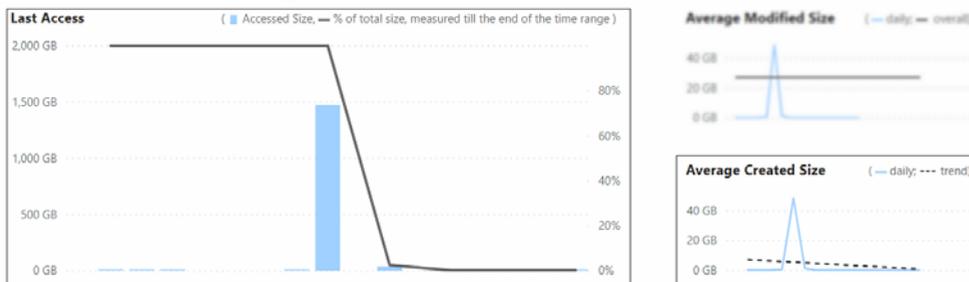
Sample steps to compare the “Last Access” and “Average Created” graphs:

1. Load the report for the desired slot and if necessary, refine the analyzed path by selecting a sub-folder.
2. In the left pane, click the year tab and configure the controls above to display a report for the last year.

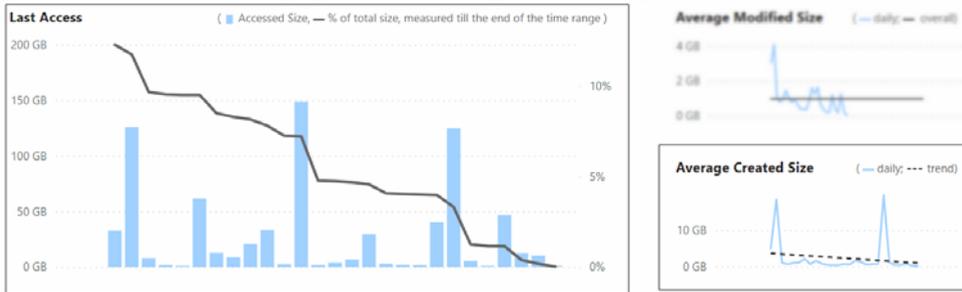


3. Compare the Last Access and Average Created Size graphs:

- an example of matching patterns - a peak in both file access and file creation is visible at the same point in time in both graphs:



- an example of patterns that do not match - a peak value in the middle of the analyzed period is present in the Last Access graph, but is missing in the Average Created Size graph:

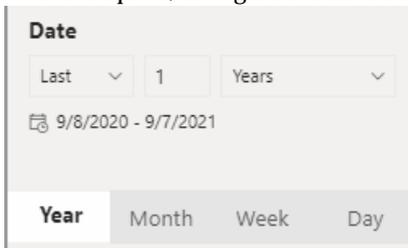


Estimating the Daily Cloud Upload Size

As traditional applications cannot write directly in the cloud, with hybrid storage data is first created on the local storage tier and then copied or moved to the respective cloud tier. Thus, aside from the amount of data that can be made accessible from the cloud tier, when planning the migration to a hybrid storage workflow, it is necessary to estimate approximately how much data will be uploaded to the cloud on a daily basis. By adding the average daily size of all created files to the average daily size of all modified files for a period of one year, for example, you can roughly estimate how much data will be replicated to the cloud daily. For example, if the daily average size of created files is 2.3GB and the daily average size of modified files is 1GB, you can expect to upload 3GB to the cloud on a daily basis.

To estimate the approximate daily upload size:

1. Load the report for the desired slot and if necessary, refine the analyzed path by selecting a sub-folder.
2. In the left pane, configure the controls to display a report for the last year.



3. Check the Daily Averages field of the report.

| Daily Averages (GB) (■ daily max; Avg.) | | | |
|---|---------|--|-----|
| Metric | Average | | Max |
| Accessed | 4.9 | | 114 |
| Modified | 1.0 | | 15 |
| Created | 2.3 | | 113 |

4. Add the average daily size of modified files to the average daily size of created files and you will get the approximate size of data which will be replicated daily.

Estimating the Amount of Versioned Data in the Cloud

When turning on versioning in the cloud, one of the main considerations is how much space versioned data would take up. You can easily estimate the costs of hosting each version of a replicated file in the cloud, by checking the size of modified files in the analyzed path for a period of one year, for example. You can also estimate approximately how much versioned data you will upload daily.

Sample steps to calculate the size of versioned data in the cloud:

1. Load the report for the desired slot and if necessary, refine the analyzed path by selecting a sub-folder.
2. In the left pane, configure the controls to display a report for the last year.

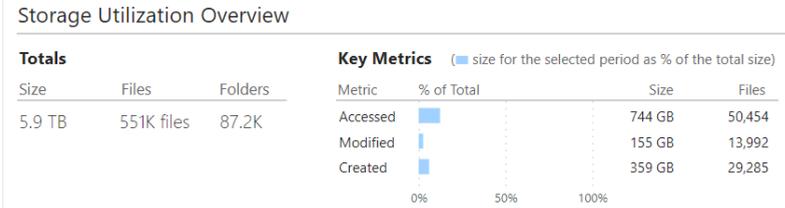
Date

Last ▾ 1 Years ▾

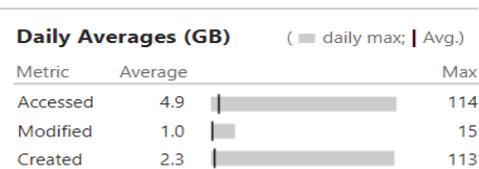
📅 9/8/2020 - 9/7/2021

Year | Month | Week | Day

3. Check the size of modified files for the selected time frame in the Storage Utilization Overview of the report - this value corresponds to the approximate size of versioned data that will be hosted on your cloud target on a yearly basis.



4. Check the average size of modified files in the Daily Averages section of the report to calculate approximately how much versioned data you will be uploading on a daily basis.



Configure Tiger Bridge

Tiger Bridge is Tiger Technology's data lifecycle manager for hybrid storage environments. It lets you pair a source (a locally mounted NTFS volume or an SMB/NFS network share) with a cloud storage target into seamless unity. Tiger Bridge automates the assignment of data to the source or the target tier, based on user-defined policies, thus addressing various workflow challenges - from alignment of data with storage costs, transparent data migration and synchronization between storage devices or geographically dispersed places, to extending your primary storage to the cloud and acting as a gateway to object storage.

Specify Tiger Bridge Sources

When configuring a source, you let Tiger Bridge automatically manage all data in the source's path including data in all sub-folders based on the policies for data replication, space reclaiming and data archiving you have specified. With this in mind, you can use the Storage Assessment Report to view statistics about a drive's folders and to decide whether to apply the same automatic data management policies to all of them or to configure separate sources depending on different data lifecycle patterns in different folders.

You can also use the report to decide whether a sub-folder of an analyzed path should be automatically managed by Tiger Bridge or added as an excluded location whose data you can manage only manually.

Configure Automatic Policies

Tiger Bridge utilizes user-defined policies to automatically manage data in your sources. These policies operate using parameters which instruct Tiger Bridge what data to process and when to process it. You can use the results of the Storage Assessment Report as a basis for specifying the parameters of the different policies depending on the specifics of your data usage workflow.

replication policy — this policy governs what data from the local storage needs to also be copied to the respective tier of the cloud storage target. Depending on the needs of your workflow, you can configure it to replicate all data as soon as possible (data backup and DR workflows) or synchronize the replication of data with its removal from the local storage (hierarchical storage and data archiving workflows). The replication policy uses just one parameter - Tiger Bridge replicates a file only when a specified amount of time has passed after its last modification, thus assuming the file will not undergo any further changes. To adjust it using the statistics of the report, compare values about the number and size of modified files on a source from different points in time.

space reclaiming policy — this policy governs which data needs to be moved to the nearline tier of the target, freeing space on your local storage. It operates using two parameters - for how long a file has not been accessed and how big the file is. Use the accessed files statistics from the report to decide what value would not interrupt your normal workflow by having to retrieve a file from the target in order to access it. The statistics about average file sizes can give you a clue how big a file should be in order to free the needed space on your local storage by moving it to cloud storage. In the space reclaiming policy you can also configure the used space and maximum used space thresholds based on the expected data growth, using the information in “Assessing Future Data Growth” on page 13.

data archiving policy — this policy governs what data needs to be copied to the archival tier of the target. The policy operates using two parameters - for how long a file has not been accessed and how big the file is. Use the accessed files statistics from the report to decide what files can be moved to the archival tier without any foreseeable need to keep them ready for retrieval unless you have to recover from disaster. When combined with the space reclaiming policy, all files copied to the archival tier can also be replaced by offline files on your source, thus freeing space on it. The statistics about average file sizes can give you a clue how big a file should be in order to free the needed space on your local storage by moving it to cloud storage.

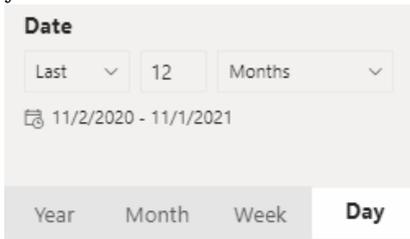
Configure Ransomware Protection

To prevent replication of files that have been encrypted on your source due to a ransomware attack, Tiger Bridge provides you with a fail-safe setting, which automatically pauses the replication once the number of files queued for re-replication exceeds a limit set by you. As files are queued for re-replication only if they are modified on the source, to configure this protection mechanism, you need to know what the maximum files modification rate of your organization is. Although the

maximum number of modified files may also include modified files that have not been replicated yet, it gives you a good starting point for calculating the limit above which the modification rate may signify that you are victim of a ransomware attack.

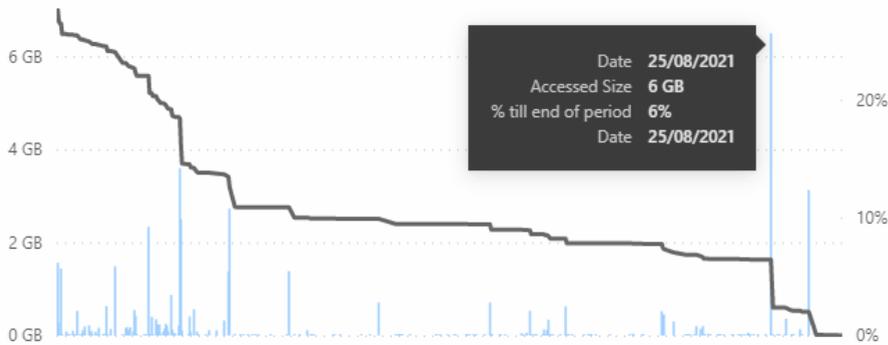
Sample steps to estimate the maximum number of modified files:

1. Load the report for the desired slot and navigate to path added as a source.
2. In the left pane, click the Day tab and configure the controls above to display a report for the last year.



3. In the Last Access graph, find peak value day, hover your mouse over it until the tooltip with additional details is displayed, and then click on the graph to update the statistics with data about just that day.

Last Access



4. Check the number of modified files for the date in the Storage Utilization Overview of the report - this is the approximate maximum modification rate, and any value above it may designate that your data is being encrypted due to a ransomware attack.

Storage Utilization Overview

| Totals | | | Key Metrics <small>(size for the selected period as % of the total size)</small> | | | |
|--------|------------|---------|--|------------|--------|--------|
| Size | Files | Folders | Metric | % of Total | Size | Files |
| 5.9 TB | 551K files | 87.2K | Accessed | | 744 GB | 50,454 |
| | | | Modified | | 155 GB | 13,992 |
| | | | Created | | 359 GB | 29,285 |

0% 50% 100%

Schedule Tiger Bridge Operations

Aside from configuring Tiger Bridge to automatically manage data in your sources, you can also use its shell-extension and command-line interface to manually perform data lifecycle management operations on selected data. The main advantage of using the command-line interface of Tiger Bridge is that you can schedule the execution of specific tasks by including the commands in a script. Using the statistics about peak access to data on a weekly basis, for example, can help you schedule Tiger Bridge to replicate files when this would not obstruct your organization's workflow.