



Multi-site Data Synchronization with Tiger Bridge and the Cloud

Step-by-step guide for using multi-site sync with Tiger Bridge

Version: 1.0
February 2023



Table of Contents

Introduction	3
Typical Approaches	3
The Tiger Bridge Approach	4
Basic Configuration on the First Machine	5
Synchronization Configuration on the First Machine	8
Configuration on the Second Machine	10
Replication in Action	12
Modifying Files	12
Common Security Context	13
Keep in Mind	13
Conclusion	13

[Send Us Your Feedback](#)

[About Tiger Technology](#)



Introduction

Since more and more users and offices are located within different time zones, remote collaboration demands are on the rise. Organizations looking to provide long-term solutions often see **multi-site data synchronization** as the top choice.

However, anyone looking to implement this properly will need to provide sizeable initial investment as well as great attention to detail. Minor sync errors could reduce productivity while IT complexities might affect existing workflows, thus increasing the management strain.

Typical Approaches

There are two common ways to approach multi-site sync – through Virtual Private Networks (VPNs) and Distributed File Systems (DFS).

VPNs link client machines in a private network over the Internet in order to share data and applications. They may not require a large initial investment.

File transfers over VPN are generally considered secure as data is encrypted before entering the private network. However, an inevitable increase in latency due to new users being added will affect data transfer speeds. VPNs are also prone to connection breaks which prevent constant access to the data set. Setting up an effective VPN connection for businesses can be complex and costly, and configuration faults can expose data to external threats.

A **DFS** is a single file system distributed on multiple file servers or locations. It allows multiple remote users to access corporate data simultaneously. Despite its suitability for enterprises, it's not the go-to solution for SMBs and smaller teams due to associated costs and complexities.

DFS supports high availability – this allows them to continue working even in the event of a partial failure. On the other hand, they may require more work to ensure storage is secure – delays when dealing with large data sets are not uncommon. Managing a distributed system is overall a more complex process than managing a centralized storage infrastructure.

Check out our [YouTube channel](#) for more information.

The Tiger Bridge Approach

Tiger Bridge is a software-only data management solution that enables multi-site sync by using the cloud to store a copy of the data set that needs to be accessed from multiple locations. None of the complexities associated with VPN services are involved in the process.

Why Tiger Bridge?

Tiger Bridge blends on-premises and cloud/multi-cloud storage into a single space. While it enables hybrid workflows, it also preserves all Active Directory permissions and keeps your data in its native format. It supports any Microsoft, Amazon, Google or IBM Cloud tier, Backblaze, ERA Coeus, Hitachi HCP, Wasabi and cloud S3 compatible storage, plus local object and file storage.

With Tiger Bridge, you can sync multiple file servers in different offices, several client machines, or a file server and a client without connecting them over a common network or through VPN. This solution brings significant cost optimization - a single software license equals multiple clients. Based on simple data replication and sync policies, you can set the frequency of checking for updated contents to what matches your workflow needs and use optional file locking to prevent overwrite conflicts. Machines will also receive updates about any changes during periods when they were offline based on the settings configuration.

Multi-site sync with Tiger Bridge is based on strong eventual consistency which ensures a smooth user experience. Data synchronization through the cloud with Tiger Bridge also enables disaster recovery through off-site backups.

For the purposes of this guide, we have broken down the Tiger Bridge multi-site sync configuration and usage process into four separate steps:

1. [Basic configuration on the first machine](#)
2. [Synchronization configuration on the first machine](#)
3. [Configurations on the second machine](#)
4. [Replication in action](#)

Basic Configuration on the First Machine

Before you begin, you'll need to configure the program and simply leave it running in the background on your computer. The demonstrated synchronization scenario of Tiger Bridge will be used for syncing a file server and a desktop virtual machine situated in different time zones.

Note: Prior cloud account existence is required, along with a valid Tiger Bridge license for multi-site sync.


1. Install Tiger Bridge following the onscreen instructions and then open the program application. Initially, Tiger Bridge will be in paused state.

 Tiger Bridge Configuration

 Tiger Bridge (paused)


 Settings

2. Make the location of your enterprise class data a source folder through the Local Source button on the left side of the program. It should contain at least one file and/or folder.




On the right, select the respective folder and click OK to confirm your choice.


Select folder you want to replicate:



3. Set up a target cloud provider through the **Add Target** button to establish communication with a virtual storage space for your data. Tiger Bridge supports a variety of options, including AWS and Microsoft Azure.



To confirm your choice, click **OK**.

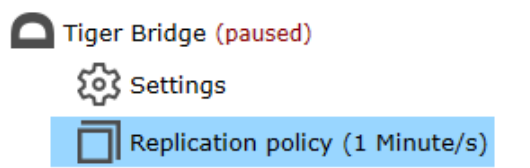


Proceed to fill out the requested information. You must have your target cloud provider account credentials prepared. If you have difficulties finding them on the cloud portal, an administrator may be able to fetch them instead.

Make sure you've filled out the fields correctly, and depending on the cloud target, select an appropriate container or bucket.


To conclude the process, click **Apply**.

Set up a **replication policy** by accessing the respective tab on the left half of Tiger Bridge.




This is done by specifying how much time (in seconds, minutes, hours, etc.) after a change you would like to wait before this change is replicated to the cloud. Use the dropdown menu on the right.

Tiger Bridge Replication policy:



Note: The lower the setting, the more real-time (close to instant) the replication will be. By default, it is set to one minute.


When done, **Apply** your setting.



Synchronization Configuration on the First Machine


Once you've covered the basics, it's time to apply some additional settings which are crucial to the Tiger Bridge multi-site synchronization setup.

1. From the main screen, click the Add sync button.



This will expand a space on the right where you need to specify how often the system will check for changes made to replicated files, as well as notify the systems involved in the synchronization process for these changes.

Tiger Bridge Synchronization:




The screenshot shows a configuration panel for Tiger Bridge Synchronization. It contains a checked checkbox labeled "Listen:", followed by a "Period:" label, a numeric input field with the value "1", a small up/down arrow control, and a dropdown menu currently showing "Minute/s". Below this is an unchecked checkbox labeled "Automatically restore files on the synchronized source".

■ **Listen** – sets how often a machine that participates in synchronization checks if there are active notifications for it. The lower the number, the faster the sync. Keep in mind that significantly faster syncs require more resources for communication with the cloud, which could add load to the machine and potentially generate cost. Extra collaborative workflows may require near-instant sync (5 to 10 seconds), while offline environments may be fine with much more.


■ **Automatically restore** files on the synchronized source – allows you to decide whether new data coming from other synchronized machines will be saved on your source or you want it to appear as empty stub files there which allow the user to download the data on-demand. Unchecked by default.

When you're done with the configurations, click **Apply**.



The screenshot shows a row of buttons. The "Apply" button is highlighted with a red border. Below it is a row of four buttons: "Resume" (with a play icon), "Refresh all", "Apply all", and "Close".

2. Make Tiger Bridge active by clicking the **Resume** button on the main screen.



Files and folders from the source folder will immediately begin transferring to the cloud in respect to the defined replication policy. Overlays on top of their icons allow you to see what's currently being synchronized.

Configuration on the Second Machine

The same steps need to be reproduced for the virtual machine in the multi-site sync process. You'll need to configure the program and leave it running on your computer.

1. Open a fresh build of Tiger Bridge (in **Paused** state) and add a source folder.

Note: If you add an empty source folder, it will get the files from the other machine populated. However, you also have the option to add a source which contains data you would like to get synchronized to the other machine since the synchronization works both ways.

2. Proceed to select the same target cloud provider you specified on the first machine, enter its respective credentials and select its bucket or container.

3. Confirm the target setup by clicking the **Apply** button.

You will be asked if you want to download the data already in the cloud or not. Configure this as desired. For our sync scenario you can choose **Import on demand** if you would like to see the files on the synced machine but without them occupying any space or restore all data to download all contents of the files locally.

Existing cloud data action ×

Choose action to be performed on any existing data in the target:

☒ No action	No data will be imported from the target. It can be later imported using explorer command.
☐ Import on demand	No data will be imported initially. Any folder that is accessed by an application will be imported automatically. This option will provide access to all data in the target without importing it first.
☐ Import all metadata	All metadata will be imported recreating the entire folder structure from the target. This may be a long operation based on the number of files and network speed.
☐ Restore all data	Import all metadata and restore all the data. This operation will copy the content of the target locally and require the necessary free space. May take very long time based on network speed.
☐ Link existing data	Links existing local files with corresponding object data on the target. Useful when data has already been uploaded to the target.

OK

4. Set the replication policy setting according to your preferences. As always, the lower the number, the faster the synchronization process will be.

5. Similarly, configure the Listen synchronization settings based on your needs.

6. Make Tiger Bridge active by clicking the **Resume** button.

Replication in Action

If you've successfully applied the instructions so far, the empty source you specified on the second machine should become populated with the exact same file and folder structure of the source you specified on the first machine. If it still happens to be empty, refresh it once to reflect the most recent synchronization.

Replicated data might or might not occupy space on the synchronized machine depending on your selection earlier. Check the files' icon overlay to verify.

Modifying Files

If you need to open a file and perform actions inside of it, this will trigger its download if it was previously empty.

Once you save the edited file, it effectively acquires a new version and updated time of modification, which will be visualized on the first machine according to the replication policy.

The same is valid the other way round - when you make a change to the contents of a file from the source folder of the first machine.


Modifications and updates will be reflected on the second machine according to the replication policy, so any time zone differences will be disregarded. You can rename or delete replicated files on one of the machines only to see them reflected on the other. The first source folder may get the changes slightly faster than the second one - it all depends on how the synchronization policy is defined.

Note: When you deploy Tiger Bridge for synchronization, it is advisable to change the default file deletion behavior and configure Tiger Bridge to keep the replica on the target even if the file is deleted on the source. In this case, to delete it from the target as well, you should access the target and manually delete the file.

Common Security Context

If the machines you're trying to synchronize are in a common security context (e.g., Active Directory Forest), you may want the same rights and restrictions to be applied to them.

For that purpose, open Tiger Bridge, access the **Settings** tab on the main screen and check the **Preserve security descriptor on sync** box on the left.



Keep in Mind

1. When a file server is synchronized with a virtual machine, it remains a fully operational server. The directory you defined there as a source will remain shared and its files and subfolders will be accessible over the network even without the use of Tiger Bridge.
2. The same source files and folders can also be accessed and downloaded externally from the respective cloud portal (e.g., AWS or Microsoft Azure).
3. A machine that doesn't have Tiger Bridge installed can access the file server through the network and be given the respective content.

Conclusion

Tiger Bridge is an ideal solution for cloud-enabled multi-site sync. It is transparent to users and applications and avoids vendor lock-in by preserving data in its native format.

Tiger Bridge is easy to set up and configure and does not introduce any admin complexities commonly associated with VPN or DFS approaches. It also delivers cost optimization: one subscription can be used to enable sync across multiple client machines.

More on the topic you can see in the following [video demos](#).

Send Us Your Feedback

Did you find the information you were looking for?
Email us your feedback at info@tiger-technology.com.

About Tiger Technology

Founded in 2003 and headquartered in Sofia, Bulgaria and Alpharetta, GA, USA, Tiger Technology specializes in the underlying technology of hybrid cloud workflows as well as data management software solutions designed to help customers of any size, scale and industry optimize their on-premises storage and enhance their workflows through cloud services. Tiger Technology solutions are admin-friendly, non-disruptive, transparent, and highly cost-effective.

Tiger Technology brings over 15 years of expertise in developing high-performance storage solutions for the most demanding workflows, which includes cross-platform NAS/SAN file system sharing, storage, user, project, and media management.

Throughout the years and the multitude of changes in IT and the digital landscape, Tiger Technology has chosen the path of growth and continual improvement, but it has never lost sight of what it all starts with - data. The company's current focus is enabling "on-premises-first" hybrid cloud workflows.

Find out more about Tiger Technology [here](#). Catch up on [our blog here](#).

Follow us on social media:

[LinkedIn](#) | [Twitter](#) | [YouTube](#) | [Facebook](#)