



## Anatomy of a Replication Job

**B**BBj® replication allows users to have one or more up-to-date copies of their production data at all times. This feature is extremely useful for both backups and load balancing through read-only data access on target machines, and provides many benefits. Since BASIS built replication into BBj, developers do not have to update their applications to keep targets automatically current with changing source data. Replication overhead is extremely low, as BASIS designed it for minimal performance impact to writes on the source machine and no impact on read performance. It is also robust as it continues to replicate data flawlessly, regardless of how many times the source and/or target BBj Services are stopped and started. Administration is very flexible as users can maintain replication jobs easily in Enterprise Manager or in their own custom programs via the Admin API. Both methods allow replication job management on-premise and even in the cloud.

### Replication Overview

The basic architecture of BBj replication involves a source BBj Services with one or more replication jobs. Each replication

job describes a set of directories and files to replicate to a single target BBj Services. A user can create as many replication jobs as desired, each replicating any set of directories and/or files. In addition, users can either replicate the same data to many different targets or replicate different data to different targets.

### Creating the Initial Copy

When a user first creates a replication job, the target requires an initial copy of all of the files that the job specifies. While the replication process is capable of making initial copies of missing or out-of-date files by itself, this process will keep the file open in the background for however long it takes to copy the file. The faster the network connection, the better; large files on a very slow network connection could potentially take several hours.

Users may access files and make changes while the initial replication copy is in process, but they cannot delete the files. If users are willing to shut down their server for a short period, BASIS recommends restarting BBj Services in the “admin only” mode and manually [rsyncing](#) the data just before creating the replication job. This allows the administrator to control the impact of the required initial copy and minimize the effect on users of their application.

### Understanding the Replication Log

The replication job only needs to monitor changes to the source data, as read access does not modify the source data. Therefore, BBj Services tracks all modifications to the source files and records those changes to a replication log. The log then contains a continuous stream of file writes, removes, creations, and deletions for replicated files and directories without the replication job needing to actually open and read the source files or directories. Since this happens at the file level, the log reflects the data records exactly as they appear in the actual source file. >>



By **Chris Hardekopf**  
Software Engineer



Figure 1. BASIS data replicated across the world

Each replication job runs in the background and reads from the replication log in order to propagate changes to the target BBJ Services. Each target BBJ Services independently remembers its current position in the replication log for restart capability. If the replication is synchronous, any user change operation such as file or directory creation, revision, deletion, etc. will block the user application, waiting for the background target to process the change. However, if the job is asynchronous, as BASIS recommends, the user application can continue without waiting for targets to apply the change. As the targets process the file/directory changes registered in the replication log, BBJ Services removes those completed changes from the log to keep it as small as possible.

The replication log, therefore, allows for both the minimal performance impact of replication and safe and transparent handling of source and target BBJ Services downtime. The performance impact is minimal since the most that enabling replication will do is append a change record to the replication log, no matter how many targets are configured or how slow the communication with the target. Downtime for the source BBJ Services is easily handled since clearly there will be no new changes while it is down and all previous changes that have not yet been replicated to targets are preserved in the replication log. Downtime for the target BBJ Services is also managed automatically, since any changes happening on the source BBJ Services while the target is not available will be preserved in the replication log and sent to the target BBJ Services as soon as it is running once again.

### Seeing Replication in Action

BASIS is an example of a company who uses and relies on replication for continuous backups. They run their production BBJ Services on the Amazon cloud region on the West coast of the United States. BASIS continuously runs replication jobs to three other Amazon cloud regions: the United States (East coast), Japan, and Ireland. This means that the replication jobs provide complete backups of all BASIS' production BBJ data files in several places around the world. If necessary, BASIS could easily switch BBJ Services production to a different region as shown in **Figure 1**.

### Summary

Replication is such a powerful and easy-to-implement feature that every BASIS customer has reason to put it to work. It provides an invaluable safety net that anyone would pay dearly for if the situation ever required it. Files or directories, large or small...backups of any shape give extra security and a great deal of peace of mind. ■



Review these topics in the BASIS online help at

- [Replication Introduction](#)
- [Enterprise Manager - Replication Jobs](#)



**3DTek**

Innovative Search Solutions

## Finding the perfect match is an art, not an accident.

Technology is constantly evolving at BASIS International, and along with it, so are your staffing needs for BBJ, VPRO/5 or Barista professionals. At 3D Tek we are dedicated to helping you maintain your competitive edge by finding the right professionals for your company.

Our custom search solutions go beyond conventional recruiting, allowing us to locate, recruit, and bring employees on board who not only have the talent and skill set you need, but who share your goals and reflect your company culture.

Whether you need someone for a contract, contract-to-hire, direct hire, or a fixed price project, we can find the perfect match.

Improve your productivity and profitability with 3D Tek, your IT & Executive search and recruitment partner.



352-569-9203 or visit us at [www.3dtek.com](http://www.3dtek.com)