



Feature Guide

Double-Take v4.5

Table of Contents

Product Description	3
Features	4
Intelligent Compression	4
E-mail Alerts	4
Management Console Server Filtering	5
Product Update Checking	5
Asynchronous, Byte-level, Continuous Replication	5
Hardware Independence	5
Simple Management Console	6
Uses Existing Networks	6
Task Command Processing	6
Unlimited Distance Replication	6
Many-to-One Failover	6
Automated Failover	7
Open-file Mirroring and Replication	7
Flow Control (Unlimited Disk Queueing)	7
Bandwidth Throttling	7
Resource Planning Tool	7
Active/Active MSCS Support	8
Serverless Backup Support	8
Failback/Restore	8
Block Checksum Re-Mirror	8
Replication Scheduling	8
Versioning Mirrors	8
Automatic Re-Mirror	8
File/Directory Selection	9
Support for Different File Types	9
Verification	9
SNMP Counters and Traps	9
Full Command-Line Control	9
Reports and Statistics	10

Product Description

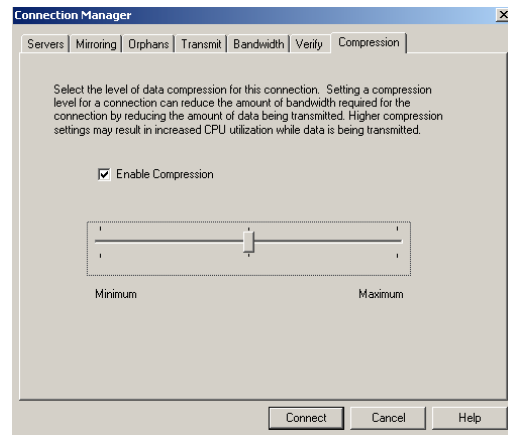
Double-Take® from Double-Take Software goes beyond periodic backup to provide accessible and affordable data protection, ensure minimal data loss, and enable immediate recovery from any disaster or system outage. Double-Take utilizes patented replication and failover technology that continuously captures byte-level changes as they happen and replicates those changes to one or more target servers at any location, locally or at a recovery site miles away. In the event of a disaster, Double-Take lets you recover from your target servers in minutes, if not seconds. What's more, Double-Take delivers better protection than many hardware-based solutions, and it costs tens of thousands less. Start with as few servers as you like, scale organically as your data requirements grow, and get a fast and continuing return on your investment.

Features

Intelligent Compression

Double-Take® is the only host-based replication solution that provides multilevel intelligent compression. Double-Take provides four options for data compression that allows custom configuration to suit each customer's individual requirements. These settings can be individually configured for each replication set defined within Double-Take, allowing even further customization by server, data, and/or network.

The first level is 'no compression'. Chose this setting for data that resides on application servers that require full availability of system resources or where network bandwidth is not a concern. Less critical data where latency is not a major concern may also be considered for no compression.



The next three levels enable compression, but to differing degrees. There is a tradeoff to consider when choosing compression levels. Higher compression levels require Double-Take to use additional system resources, including CPU and RAM, on both the source and target servers. Level one provides the least amount of compression, using the least amount of system resources, resulting in a smaller data transfer benefit while level three performs the highest degree of compression, resulting in the least amount of data being transmitted. However, additional CPU and RAM will be required. By providing multiple levels to choose from, and allowing each to be set on individual replication sets, customers can decide what level is necessary and what trade-offs are acceptable based on the actual data being protected.

As certain data types can actually increase in size when compressed, which will result in more data being transmitted rather than less, Double-Take compares the size of the compressed data to that of the uncompressed and will transmit the smaller of the two. This ensures that Double-Take compression does not increase the load on the network when intending to decrease it. Customers don't have to be concerned with first validating the data for its 'compressibility' prior to enabling compression. Double-Take will intelligently compress only the data from which a benefit will result.

Email Alerts

For optimal data protection and availability, Double-Take can provide event notifications via email for immediate awareness of possible breaks in service levels. Email alerts can be configured with different recipients for each Double-Take server, with each having its own event notification level (informational, warning, error). The email message will include useful information in the subject line, including server name where the event occurred and the error level and code. By using Double-Take, administrators can quickly and easily manage events via their email clients by sorting or filtering by these events. They can then choose which ones require immediate attention and which do not.

By providing immediate and at-a-glance updates of the Double-Take environment via email, administrators do not have to continually monitor the Double-Take management console to ensure everything is operating smoothly. This ability greatly increases the scalability and manageability of the Double-Take environment.

Management Console Server Filtering

Server Filtering allows individual administrators to create their own custom server views within the Double-Take management console. By hiding certain servers, administrators are able to see just their select set of servers, simplifying manageability of the Double-Take environment. Views are created by individual login IDs, allowing each administrator to only have to see the servers they are concerned with.

Through the use of customized server views, Double-Take allows for virtually unlimited scalability, as one can choose to view as few servers as desired without having to search through the tens or hundreds of other servers being managed by Double-Take.

Product Update Checking

To help ensure customers always install, or are at least aware of the latest version of Double-Take available, the installation process allows the user to check, via the web, for the latest product updates and releases. The user is given a list of updates available and can choose to download and install immediately. This update check is also available via the Double-Take GUI, allowing users to check for updates whenever they choose, without having to run the installation process.

As it is common for customers to install patches or custom enhancements to their Double-Take install, Double-Take will first determine what updates/patches the customer has installed and verify that these are all included in a particular update before presenting it to the customer as an option for download. Any update that does not include all the patches already installed by the customer are not listed for the customer to select.

Asynchronous, Byte-level, Continuous Replication

Asynchronous replication ensures that the replication process does not impact production applications, as can be the case with synchronous replication. Double-Take captures the data for processing but does not keep it from updating the local disk, whereas with synchronous replication the local disk is not updated until after the replicated data has been written/committed to the target disk and the corresponding acknowledgement has been sent back to the production server. This results in a performance decrease on the production application because it is waiting for confirmation that the replication has occurred. Double-Take allows the application to process data as it normally would and simply captures the changes for processing in the background.

Double-Take copies only the byte-level changes to files, not entire blocks (physical or logical) or entire files, thereby reducing the overhead on the servers and the network. Byte-level replication ensures that all transactions are captured and written in order on the target for maintaining exact replicas. These changes are captured as they are made, resulting in near real-time replication. The data is being replicated by Double-Take before it is even written to the source disk.

Hardware Independence

Whereas hardware-based synchronous replication is often proprietary for each storage vendor and provides no cross-vendor replication, Double-Take has no affinity to any hardware or storage provider or storage architecture (SCSI, FibreChannel, iSCSI). This independence allows Double-Take to be used on existing heterogeneous storage infrastructures and does not lock you in to a specific vendor for future purchases. Using a hardware-independent solution allows for a myriad of uses not possible with proprietary array-based solutions. Performing migrations to new hardware, implementation of a SAN or NAS, or any other project that requires moving data from one storage device to another is possible and efficient with Double-Take, and there are no distance limitations. Best of all, there is no installation or training required every time new hardware is added. Simply connect the storage to a host already running Double-Take, and define the data to be replicated.

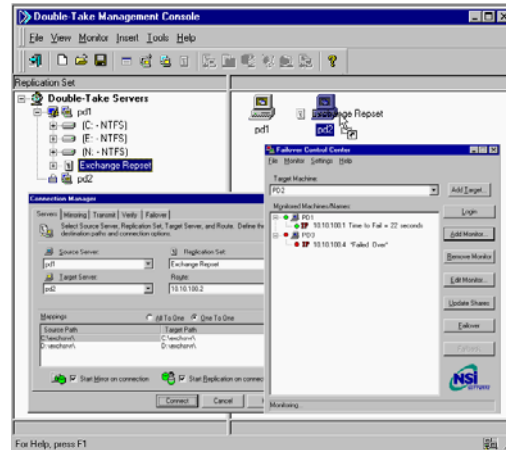
Double-Take 4.5 Feature Guide

Simple Management Console

The simplicity of navigating and configuring Double-Take provides for a much better overall user experience. This simplicity lends itself to the proficient and extended use of Double-Take in your environment, reducing overall cost-of-ownership and improving ROI. Double-Take provides an interface that is easy to use to ensure you can quickly and easily protect your data. The Double-Take UI automatically discovers servers running Double-Take and displays them in a single window. Managing the entire Double-Take deployment can be performed from this single console. Any transfers, reports, failures, and/or failbacks can be viewed from this single point.

Uses Existing Networks

Double-Take does not require its own private network for replicating data. Most often, existing networks are more than sufficient to facilitate this work, allowing you to implement lower-cost data replication and protection solutions. This reduces the overall cost-of-ownership by keeping the initial investment lower. This also removes any restrictions and additional costs for future network changes or purchases by utilizing any IP network in place. Double-Take can also be run on its own private network should you want to isolate replication traffic.



Task Command Processing

Double-Take allows insertion of system commands into the data stream for execution at different points during its regular processing. Tasks such as automatically initiating backup of your target servers are possible. Via in-band commands, you can ensure all files on the target are in sync with each other (i.e. log and data files), pause writing on the target (allowing source changes to still be transmitted and captured) and initiate the backup. Once the backup is complete, the task command can then enable writing on the target.

Unlimited Distance Replication

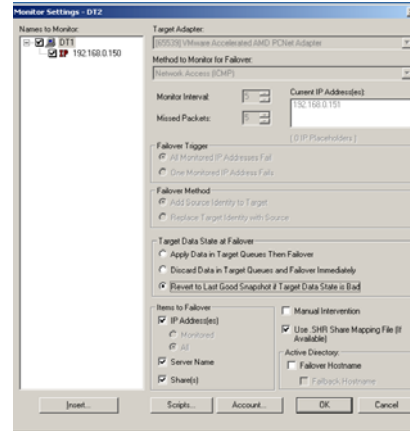
As Double-Take utilizes standard IP networks, there aren't the typical distance limitations as with some synchronous and array-based products. Being able to replicate over long distances lends itself to numerous uses that short distance products cannot accomplish. Solutions such as centralized backup to copy regional site data to a local server for backup is possible and practical regardless of geographic separations; protection against regional failures, replicating data across the country or across the globe to ensure the data is always available if needed.

Many-to-One Failover

Double-Take can be configured such that a single target server is used for numerous source servers. This can reduce the total cost of ownership as it is not necessary to have a one-to-one exact duplicate of the hardware for replication purposes, as is the case with many other replication solutions. This also helps to facilitate centralized backup since many local or remote servers can be backed up from a single server, further lowering costs by reducing the number of backup licenses needed.

Automated Failover

Whether running MSCS or not, Double-Take can provide high-availability failover of servers to ensure users remain online in case of a failure. While MSCS has distance limitations, Double-Take does not, thereby providing failover protection during local AND regional failures. This availability allows users and customers to quickly have access to their systems and data. For customer-facing systems, it ensures that customer satisfaction remains high and sales are not impacted.



Open-file Mirroring and Replication

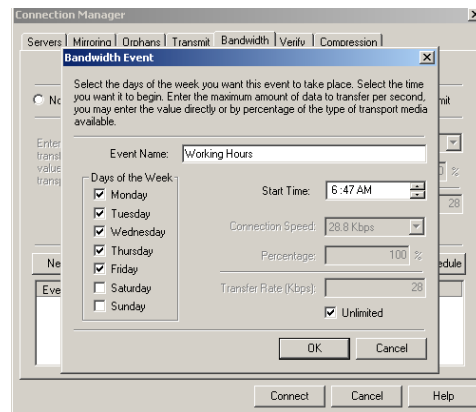
Double-Take doesn't require that applications be restarted each time additional files and/or directories are configured for replication. Users, customers and applications remain online and active while Double-Take is at work, never having a negative impact on your productivity. Unlike many products, Double-Take is able to process open files and ensure they are fully replicated without taking the files offline. Some backup products can provide open-file backup support but at an additional cost.

Flow Control (Unlimited Disk Queuing)

Double-Take is designed to handle spikes in the data rate-of-change even if the network connection to the target server is not sufficient to handle all the data at once. Double-Take will continue to filter all file changes and buffer these changes, while at the same time transmitting to the target as quickly as possible. Queuing ensures that all transactions are replicated to the target without loss of data. With the bandwidth throttling and scheduling features, users can now configure Double-Take around their production server and network requirements. If performing backups from the target sever, replication can be paused to ensure a complete point-in-time backup, while also ensuring that all changes on the Source server are still transmitted to the target and will be applied once the backup is complete.

Bandwidth Throttling

User-defined controls within Double-Take can limit the amount of the available network bandwidth it can use for data replication. This allows replication to occur real-time without impacting users on the same network. And, since Double-Take has the ability to queue data for transmission, all changes will be updated to the target and not lost due to network limits. Double-Take allows the administrator to define the actual type of network connection being used (T1, 128Kbps, etc.) and then define the amount that can be used for replication. Double-Take will not go beyond that limit, regardless of the amount of data it has to transfer. Double-Take also allows you to limit network use during busy work hours and increase or remove limits during non-peak hours.



Resource Planning Tool

To ensure an optimal replication environment, Double-Take includes a Resource Planning Tool which runs the Double-Take filter process without actually copying the data to a remote target. This gives information necessary to implement the appropriate networks and servers to meet your RPO and RTO. By running this tool before deployment, you can avoid costly overspends on hardware and networks and avoid an insufficient deployment that doesn't meet your needs.

Active/Active MSCS Support

If running in a Microsoft® clustered environment in an Active/Active configuration, Double-Take can be configured to replicate data from any or all members, to each other, or to remote systems. Double-Take can be configured to be aware of failovers and automatically adjust the sources and targets accordingly, removing the need for user intervention at the time of failure.

Serverless Backup Support

While Double-Take does not replace tape backup, it can greatly enhance your existing backup solutions. By replicating production data (application, file server, user, etc.) to a target server, backups can be performed from the target server without having to shut down applications or requiring users to log off the production server. The backup window, which is often too small and results in incomplete backups, is now not an issue, as the target server has virtually no time limits for the backup to complete. Users continue to run 24X7 on the production server while backup is offloaded to non-production target servers.

Failback/Restore

Should a failure occur, Double-Take can facilitate data restoration from the target back to the original source or to an alternate location. Through the Double-Take UI, users can easily restore data from the replicated disk back to the production disk once the failure is corrected. This greatly reduces the time to recover and restore, as it is not necessary to go offsite for tapes and then restore one at a time. This also ensures that you recover from the time the failure occurred, not from when the last backup was taken, which can result in a day or more of lost data and productivity. Unlike other solutions that make the users remember which files came from what location, the Double-Take restore process automatically reverses the direction of your original replication job.

Block Checksum Re-Mirror

Should a disconnect occur between the source and target, instead of doing a complete mirror of the entire replication set, Double-Take can perform a block-checksum re-mirror. This re-mirror will only replicate the file differences between the source and target, which will take much less time and resources to accomplish. This will ensure that the target is coordinated with the source. This feature greatly simplifies backup, recovery, and replication management.

Replication Scheduling

Double-Take provides the flexibility to schedule when replication occurs. Double-Take will continue to filter and queue up all appropriate data changes until the scheduled time has arrived, then replicates the queued changes to the target until the user-defined replication window closes. This allows users to fine-tune Double-Take around their business needs to ensure network and system resources are used efficiently and that replication does not impact production.

Versioning Mirrors

If multiple point-in-time versions are required, Double-Take can easily integrate with 3rd-party point-in-time applications for the creation of multiple copies of the source data from different times throughout the day.

Automatic Re-Mirror

Should a problem or even scheduled maintenance require that the connection between the source and target be broken, Double-Take will re-establish its connection when possible and will automatically re-mirror the source to the target(s), ensuring the target remains in sync with the source.

File/Directory Selection

With Double-Take you can select full file systems, directories or files for replication. This gives greater flexibility to configure efficient use of resources. Important files can be selected, while system and temporary files can be skipped. The use of wildcards and drag-n-drop facilitate quick and simple configuration.

Support for Different File Types

Double-Take has been successfully tested against different file types such as encrypted and compressed, and ultra-long file names, among others. This ensures that Double-Take will be sure to correctly replicate all the data in your environment.

Verification

Although Double-Take has numerous checks to validate that the data on the target is an exact replica with the source, there is also a verification option that can be run (either scheduled or immediately) that will verify the target is in sync with the source and create a report. This is useful when there is a brief outage, if services are stopped, or if someone has made direct updates to the data on the target. You can choose to just report on any differences or have all differences corrected.

SNMP Counters and Traps

Double-Take allows you to integrate it with your enterprise management frameworks. Double-Take will forward both replication statistics and events via SNMP allowing simplified management by administrators.

Full Command-Line Control

Double-Take allows all GUI functions to be controlled from the commandline, either via scripts or individual commands, giving you greater flexibility and automation capabilities.

Server Groups

Server groups and health status information within the Double-Take console allow you to easily manage multiple Double-Take servers and centrally monitor the health of real-time replication and failover.

Target Reboot Without Re-Mirror


You can gracefully shutdown the Double-Take Target server without requiring resynchronization of the entire protected dataset.

ICMP-less ping

Additional options for failover monitoring provide the ability to monitor the Double-Take service in place of monitoring the IP address of the production server.

MOM Integration Pack

Double-Take Management Pack for Microsoft Operations Manager allows you to monitor issues on all Double-Take servers across your enterprise and alerts you to only important events, allowing you to stay aware of, respond to, and demonstrate accountability for server and application services levels.

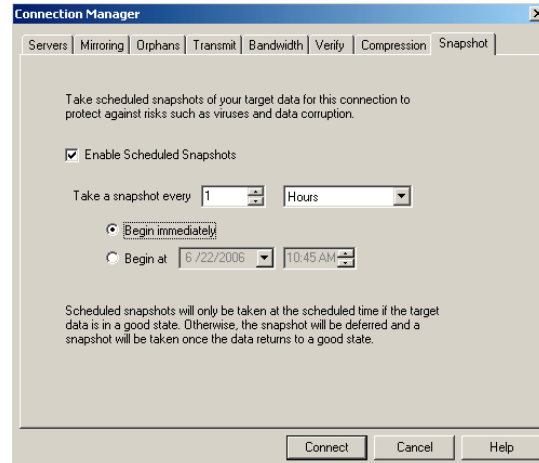


Volume Shadow Copy Service Snapshot Integration

Double-Take integrates with Microsoft Volume Shadow Copy Service to allow you to schedule and recover from up to 64 point-in-time copies of data on your Double-Take target. This ensures that you can easily recover from human error, corruption, viruses, and other unwanted changes.

Reports and Statistics

Double-Take provides reports and statistics about your replication environment which allow you to know at any time exactly what is happening on a Double-Take server, all from a single interface. Reports include information such as what has been replicated last, how much has been replicated, failovers, downed systems, and other useful statistics for ensuring adequate data protection.



64-bit Support

Double-Take provides 64-bit versions that run on Windows® Server x64 Edition (AMD64 and EMT64) or Windows Server 2003 for Itanium-based systems.

Enhanced Integration with Double-Take Application Manager

From within the standard Double-Take Management console, administrators can easily launch the Double-Take Application Manager to configure a new Exchange or SQL Server protection scenario or manage an existing one.

Target State Awareness

Double-Take can report to the user the “state” of replicated data. This provides the administrator additional status information about the replication process and can help them verify that their business-critical data is being protected by Double-Take or help them troubleshoot potential issues that may be affecting real-time replication.

Double-Take’s VSS Snapshot Integration feature leverages Target State Awareness to take automatic point-in-time snapshots of protected data as data moves from a “known” to an “unknown” state. This ensures that the administrator has at least a “known-good” copy of the protected data from which to recover.

Target Path Blocking

This feature allows the administrator to easily block any other application or process from writing to the replicated data on the Double-Take target server. This prevents unwanted changes from occurring inadvertently and ensures that the replicated data does not become corrupt due to these unwanted changes.

Summary

Double-Take® Software provides the world's most relied upon solution for accessible and affordable data protection for Microsoft® Windows® applications. The Double-Take product is the standard in data replication, enabling customers to protect business-critical data that resides throughout their enterprise. With its partner programs and professional services, Double-Take delivers unparalleled data protection, centralized back-up, high availability, and recoverability. It's the solution of choice for thousands of customers, from SMEs to the Fortune 500 in the banking, finance, legal services, retail, manufacturing, government, education and healthcare markets. Double-Take is an integral part of their disaster recovery, business continuity and overall storage strategies. For more information, please visit www.doubletake.com.

© Double-Take Software. All rights reserved. Double-Take, GeoCluster, and NSI are registered trademarks of Double-Take Software, Inc. Balance, Double-Take for Virtual Systems, and Double-Take for Virtual Servers are trademarks of Double-Take Software, Inc. Microsoft, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective companies.



Or send email to info@doubletake.com

Warning: No part of this document may be reproduced or transmitted in any form or by any means, electronic, or mechanical, for any reason, without the express written permission of NSI Software, Inc. The information in this document is subject to change without notice. Although we try to provide quality information, NSI makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained in this document. Companies, names and data used in examples herein are hypothetical and/or fictitious unless otherwise stated.